Microeconomics I

Personal Study Notes and Formal Derivations

Victor Alves

2026-06

Contents

Disclaimer

- **Incomplete Document:** This file constitutes supporting material under active development. Several sections and asymptotic derivations are still being revised, expanded, and supplemented.
- **AI-Assisted Construction:** This material was structured, reviewed, and expanded with the assistance of Artificial Intelligence models for didactic organization and Markdown/LaTeX formatting rigor.
- **Margin of Error:** Due to the technical nature of the matrix and asymptotic proofs, the text may contain typographical errors, algebraic omissions, or theoretical inaccuracies not yet reviewed by the author. It should not be used as the sole definitive bibliographic source.

Chapter 1: Consumer Theory

1.1 Budget Constraint and Elasticities

1.1.1 Motivation

The budget constraint addresses the fundamental problem of **resource scarcity**. Before defining what an individual *wants* (preferences), we map what he *can have*. The linear constraint assumes efficient markets and zero transaction costs, where the consumer acts as a price taker.

Practical Analogy: Imagine entering a supermarket with exactly R$100 and no credit. Prices are fixed and there are no volume discounts. These R$100 define your **budget set**, and the exact frontier of total expenditure is your **budget line**.

1.1.2 Welfare Mechanisms and Implications

Economic agents operate rationally within these limits based on two theoretical pillars:

- **Absence of Money Illusion (Homogeneity):** If all prices and income double simultaneously, real purchasing power remains strictly unchanged. The budget line does not shift and the optimal choice remains.
- **Income Exhaustion (Walras' Law):** Under the hypothesis of local non-satiation, the consumer chooses a point on the budget line, never below it, since spending all income maximizes feasible satisfaction.

Example: Wage Indexation If inflation raises prices by 10% and the government adjusts income by exactly 10%, the worker's budget constraint remains immobile in real terms. Welfare depends exclusively on *relative prices* and *real income*, not on nominal magnitudes.

1.1.3 Formal Definition

Let L be the number of available commodities. Define:

- **Price Vector:** $\mathbf{p} = (p_1, \dots, p_L) \in \mathbb{R}^L_{++}$ (strictly positive prices).
- **Consumption Vector:** $\mathbf{x} = (x_1, \dots, x_L) \in \mathbb{R}^L_{+}$ (non-negative quantities).
- **Income/Wealth:** $w \in \mathbb{R}_+$ (scalar available for expenditure).

Definition: Walrasian Budget Set

$$B_{\mathbf{p},w} = \{\mathbf{x} \in \mathbb{R}^L_+ : \mathbf{p} \cdot \mathbf{x} \leq w\}$$

Necessary Hypotheses: 1. **Price Linearity:** Total cost is the dot product $\mathbf{p} \cdot \mathbf{x} = \sum_{l=1}^{L} p_l x_l$. 2. **Infinite Divisibility:** Commodities can be acquired in any continuous fraction.

> **Model Limitations** - If prices depended on quantity (scale discounts), the constraint would be non-linear, invalidating homogeneity properties. - If zero or negative prices existed for desirable goods, the budget set would cease to be bounded (compact), making demand infinite and optimization unfeasible.

1.1.4 Properties and Elasticity Proofs

Auxiliary Theorem — Walras' Law Under locally non-satiated preferences, the demand function $\mathbf{x}(\mathbf{p}, w)$ satisfies $\mathbf{p} \cdot \mathbf{x}(\mathbf{p}, w) = w$ for all $(\mathbf{p}, w)$.

Property 1: Demand Homogeneity of Degree Zero **Statement:** $\mathbf{x}(\alpha\mathbf{p}, \alpha w) = \mathbf{x}(\mathbf{p}, w)$ for any $\alpha > 0$.

Proof:

1. Consider the set under new parameters:

$$B_{\alpha\mathbf{p},\alpha w} = \{\mathbf{x} \in \mathbb{R}^L_+ : (\alpha\mathbf{p}) \cdot \mathbf{x} \leq \alpha w\}$$

2. Since $\alpha > 0$, divide both sides of the inequality by α:

$$B_{\alpha\mathbf{p},\alpha w} = \{\mathbf{x} \in \mathbb{R}^L_+ : \mathbf{p} \cdot \mathbf{x} \leq w\} = B_{\mathbf{p},w}$$

3. If the set of alternatives and preferences remain identical, the unique maximizer (under strict convexity) must be the same. Therefore,

$$\mathbf{x}(\alpha\mathbf{p}, \alpha w) = \mathbf{x}(\mathbf{p}, w). \quad \blacksquare$$

Property 2: Engel Aggregation **Statement:** The sum of income elasticities η_l weighted by budget shares s_l equals 1:

$$\sum_{l=1}^{L} s_l \eta_l = 1$$

Proof:

1. Start from the total expenditure identity: $\sum_{l=1}^{L} p_l x_l(\mathbf{p}, w) = w$. Differentiate both sides with respect to wealth w:

$$\sum_{l=1}^{L} p_l \frac{\partial x_l}{\partial w} = 1$$

2. Multiply and divide each term l by $\frac{x_l}{x_l} \cdot \frac{w}{w}$:

$$\sum_{l=1}^{L} \left(\frac{p_l x_l}{w}\right) \left(\frac{\partial x_l}{\partial w} \frac{w}{x_l}\right) = 1$$

3. Identify the budget share $s_l = \frac{p_l x_l}{w}$ and the income elasticity $\eta_l = \frac{\partial x_l}{\partial w} \frac{w}{x_l}$. Substituting:

$$\sum_{l=1}^{L} s_l \eta_l = 1. \quad \blacksquare$$

Property 3: Cournot Aggregation **Statement:** For a price change in p_j, the sum of price elasticities weighted by budget shares is:

$$\sum_{l=1}^{L} s_l \epsilon_{lj} = -s_j$$

Proof:

1. Differentiate Walras' Law $\sum_{l=1}^{L} p_l x_l(\mathbf{p}, w) = w$ with respect to the specific price p_j. The right-hand side is zero since $\partial w / \partial p_j = 0$.
2. Applying the product rule on the left-hand side, isolate the term where $l = j$:

$$x_j(\mathbf{p}, w) + \sum_{l=1}^{L} p_l \frac{\partial x_l}{\partial p_j} = 0$$

3. Therefore:

$$\sum_{l=1}^{L} p_l \frac{\partial x_l}{\partial p_j} = -x_j$$

4. Multiply both sides by $\frac{p_j}{w}$:

$$\sum_{l=1}^{L} \left(\frac{p_l x_l}{w}\right) \left(\frac{\partial x_l}{\partial p_j} \frac{p_j}{x_l}\right) = -\frac{p_j x_j}{w}$$

5. Thus:

$$\sum_{l=1}^{L} s_l \epsilon_{lj} = -s_j. \quad \blacksquare$$

Property 4: Total Elasticity of Demand **Auxiliary Theorem (Euler):** If $f(\mathbf{z})$ is homogeneous of degree k, then $\sum_{i=1}^{n} \frac{\partial f}{\partial z_i} z_i = k \cdot f(\mathbf{z})$.

Statement: For any good i, the cross-price elasticities and the income elasticity sum to zero:

$$\sum_{j=1}^{L} \epsilon_{ij} + \eta_i = 0$$

Proof:

1. Since $x_i(\mathbf{p}, w)$ is homogeneous of degree zero ($k = 0$), expand its partial derivatives over the entire vector of arguments $(\mathbf{p}, w)$:

$$\sum_{j=1}^{L} \frac{\partial x_i}{\partial p_j} p_j + \frac{\partial x_i}{\partial w} w = 0$$

2. Divide the entire equation by the current consumption level $x_i(\mathbf{p}, w)$:

$$\sum_{j=1}^{L} \left(\frac{\partial x_i}{\partial p_j} \frac{p_j}{x_i} \right) + \left(\frac{\partial x_i}{\partial w} \frac{w}{x_i} \right) = 0$$

3. Therefore:

$$\sum_{j=1}^{L} \epsilon_{ij} + \eta_i = 0. \quad \blacksquare$$

Theoretical Connection This identity mathematically embodies the **absence of money illusion**. Under perfectly neutral and balanced inflation (where prices and income vary at the same rate), the net effect on real demand is zero.

1.1.5 Summary of Key Results

Concept	Definition	Key Formula
Budget Set	Set of affordable consumption bundles	$B_{\mathbf{p},w} = \{\mathbf{x} \in \mathbb{R}^L_+ : \mathbf{p} \cdot \mathbf{x} \leq w\}$
Walras' Law	Budget exhaustion under LNS	$\mathbf{p} \cdot \mathbf{x}(\mathbf{p}, w) = w$
Homogeneity of Degree Zero	Invariance to proportional price-income changes	$\mathbf{x}(\alpha\mathbf{p}, \alpha w) = \mathbf{x}(\mathbf{p}, w)$
Engel Aggregation	Weighted sum of income elasticities	$\sum s_l \eta_l = 1$
Cournot Aggregation	Weighted sum of price elasticities	$\sum s_l \epsilon_{lj} = -s_j$
Total Elasticity	Sum of cross-price and income elasticities	$\sum_j \epsilon_{ij} + \eta_i = 0$

1.2 Foundations of Choice Theory and Revealed Preference

1.2.1 Motivation

Axiomatic theory ensures that the economic agent has an internally consistent logical ordering of preferences, allowing welfare to be inferred from observable choices without relying on direct psychological measurements.

Tennis Tournament Analogy: - Completeness: Ensures that no match ends in "undetermined"; players are always comparable. - **Transitivity:** Prevents paradoxical cycles (where player A beats B, B beats C, but C eliminates A). - **Continuity:** Functions as "fine-tuning": an infinitesimal change in an athlete's preparation does not instantly transform him from champion to worst in the world.

1.2.2 Mechanisms and Axioms

- **Monotonicity:** "More is better." Ensures goods generate positive marginal utility (they are not treated as waste or satiation).
- **Convexity:** Valuing diversity. Consumers prefer convex linear combinations (averages) to extreme bundles. Equivalent to a diminishing Marginal Rate of Substitution (MRS).
- **Local Non-Satiation (LNS):** Guarantees that, for any chosen bundle, there will always be an infinitesimally close alternative that offers a higher level of satisfaction.

Impact of Axioms on Mathematical Structure:

- **Completeness** eliminates decision paralysis.

- **Transitivity** makes the *money pump* phenomenon impossible (infinite resource extraction through trade cycles).
- **Continuity** allows topological mapping into continuous utility functions, enabling the use of differential calculus.
- **Strict Convexity** ensures uniqueness and stability of the consumer equilibrium.

  Revealed Preference Utility is not directly observable. If an individual chooses bundle A when B was also affordable, he *reveals* that he prefers A to B. Choosing B the next day under the same conditions would constitute a logical inconsistency.

1.2.3 Formal Definitions

Let $X = \mathbb{R}^L_+$ be the available consumption space.

Axioms of Rationality and Continuity A preference relation $\succeq$ on X is **rational and continuous** if it satisfies:

- **Completeness:** $\forall x, y \in X$, we have $x \succeq y$ or $y \succeq x$ (or both).
- **Transitivity:** If $x \succeq y$ and $y \succeq z$, then $x \succeq z$.
- **Continuity:** For every sequence of pairs $(x^n, y^n) \to (x, y)$ such that $x^n \succeq y^n$ $\forall n$, then $x \succeq y$.

Geometric Properties of Preferences

- **Weak Monotonicity:** If $x \gg y$, then $x \succ y$.
- **Strong Monotonicity:** If $x \geq y$ and $x \neq y$, then $x \succ y$.
- **Convexity:** If $x \succeq z$ and $y \succeq z$, then $\alpha x + (1-\alpha)y \succeq z$, $\forall \alpha \in [0, 1]$.
- **Strict Convexity:** If $x \succeq z$, $y \succeq z$, and $x \neq y$, then $\alpha x + (1-\alpha)y \succ z$, $\forall \alpha \in (0, 1)$.

Local Non-Satiation (LNS) $\succeq$ satisfies LNS if, for every $x \in X$ and any $\epsilon > 0$, there exists a bundle $y \in X$ such that $\|y - x\| \leq \epsilon$ and $y \succ x$.

Revealed Preference Axioms (WARP and SARP) Consider a single-valued Walrasian demand function $\mathbf{x}(\mathbf{p}, w)$:

- **Weak Axiom (WARP):** If $\mathbf{p} \cdot \mathbf{x}(\mathbf{p}', w') \leq w$ and $\mathbf{x}(\mathbf{p}', w') \neq \mathbf{x}(\mathbf{p}, w)$, then $\mathbf{p}' \cdot \mathbf{x}(\mathbf{p}, w) > w'$.
- **Strong Axiom (SARP):** Prohibits any cyclical chain of revealed choices, imposing transitivity on revealed preference.

  Structural Failures Without completeness/transitivity: infinite cycles and empty optimal choice sets. Without continuity: non-existence of continuous utility function (Weierstrass fails). Without monotonicity/LNS: budget constraint may not be saturated, invalidating Walras' Law.

1.2.4 Proofs of Choice Theorems

Proof 1: Continuity of Utility Implies Continuity of Preference **Statement:** If the utility function $u(x)$ is continuous, the associated preference relation $\succeq$ is continuous.

Proof:

1. Let $u : X \to \mathbb{R}$ be a continuous function that represents $\succeq$.
2. Define the upper and lower contour sets for a fixed bundle x:

$$S(x) = \{y \in X : y \succeq x\}, \qquad I(x) = \{y \in X : x \succeq y\}$$

3. We can rewrite $S(x) = \{y \in X : u(y) \geq u(x)\} = u^{-1}([u(x), \infty))$.
4. Since the interval $[u(x), \infty)$ is closed in the real topology and u is continuous, the inverse image of a closed set is necessarily closed. Hence, $S(x)$ is closed.

5. Analogously, $I(x) = u^{-1}((-\infty, u(x)])$ is closed.
6. By the classical topological characterization, if the contour sets are closed, the relation $\succeq$ is continuous. ■

Proof 2: Strong Monotonicity Implies Budget Saturation **Statement:** Under strongly monotonic preferences, the optimal bundle x^* exhausts income: $\mathbf{p} \cdot x^* = w$.

Proof:

1. Suppose by contradiction that the optimum occurs in the interior of the set: $\mathbf{p} \cdot x^* < w$.
2. Given $\mathbf{p} \gg 0$, there exists a scalar radius $\epsilon > 0$ sufficiently small such that the expanded bundle $x' = x^* + \epsilon \mathbf{1}$ is still feasible: $\mathbf{p} \cdot x' \leq w$.
3. Since $x' \gg x^*$, the axiom of strong monotonicity imposes that $x' \succ x^*$.
4. This generates a direct contradiction with the assumption that x^* was the optimal choice within the feasible set. Therefore, $\mathbf{p} \cdot x^* = w$. ■

Proof 3: Strict Convexity Implies Unique Optimum **Statement:** If preferences exhibit strict convexity, the Utility Maximization Problem (UMP) generates at most one unique solution for each pair $(\mathbf{p}, w)$.

Proof:

1. Suppose by contradiction there exist two distinct optimal bundles x and y ($x \neq y$) achieving the same maximum utility level: $u(x) = u(y) = u^*$.
2. Both bundles satisfy the budget constraint: $\mathbf{p} \cdot x \leq w$ and $\mathbf{p} \cdot y \leq w$.
3. By the linearity of the budget space, any convex linear combination $x_\alpha = \alpha x + (1 - \alpha) y$ with $\alpha \in (0, 1)$ is also feasible: $\mathbf{p} \cdot x_\alpha \leq w$.
4. However, strict convexity of preferences dictates that $x_\alpha \succ x$, implying $u(x_\alpha) > u^*$.
5. We found a strictly superior and perfectly feasible bundle, breaking the optimality principle of u^*. Hence, the solution is unique. ■

Proof 4: LNS Implies Walras' Law **Statement:** If preferences are LNS and x^* is optimal for the UMP with $\mathbf{p} \gg 0$ and $w > 0$, then $\mathbf{p} \cdot x^* = w$.

Proof:

1. Suppose by contradiction that $\mathbf{p} \cdot x^* < w$. Define the slack $s = w - \mathbf{p} \cdot x^* > 0$.
2. By continuity of $f(x) = \mathbf{p} \cdot x$, there exists $\epsilon > 0$ such that $\forall y \in B_\epsilon(x^*)$: $\mathbf{p} \cdot y < w$.
3. By LNS, there exists $y \in B_\epsilon(x^*)$ with $y \succ x^*$.
4. This y is budget-feasible, contradicting the optimality of x^*.
5. Therefore $\mathbf{p} \cdot x^* \geq w$; combined with feasibility ($\mathbf{p} \cdot x^* \leq w$): $\mathbf{p} \cdot x^* = w$. ■

Proof 5: WARP Implies Homogeneity of Degree Zero **Statement:** If $x(\mathbf{p}, w)$ satisfies WARP and Walras' Law, then $x(\alpha \mathbf{p}, \alpha w) = x(\mathbf{p}, w)$ for all $\alpha > 0$.

Proof:

1. Let $x = x(\mathbf{p}, w)$ and $x_\alpha = x(\alpha \mathbf{p}, \alpha w)$.
2. By Walras' Law: $\mathbf{p} \cdot x = w$. Multiplying by α: $(\alpha \mathbf{p}) \cdot x = \alpha w$ — x is affordable under $(\alpha \mathbf{p}, \alpha w)$.

3. Analogously, $(\alpha \mathbf{p}) \cdot x_\alpha = \alpha w$, hence $\mathbf{p} \cdot x_\alpha = w$ — x_α is affordable under $(\mathbf{p}, w)$.
4. If $x \neq x_\alpha$, WARP applied to step 2 would require $(\alpha \mathbf{p}) \cdot x_\alpha > \alpha w$.
5. Contradiction with step 3. Therefore $x = x_\alpha$, i.e., $x(\alpha \mathbf{p}, \alpha w) = x(\mathbf{p}, w)$. ■

Proof 6: WARP Implies Slutsky Matrix NSD **Statement:** WARP implies that the Slutsky Substitution Matrix $S(\mathbf{p}, w)$ is negative semidefinite.

Proof:

1. Consider a Slutsky-compensated price change: $(\mathbf{p}, w) \to (\mathbf{p}', w')$ with $w' = \mathbf{p}' \cdot x(\mathbf{p}, w)$.
2. By construction, $x = x(\mathbf{p}, w)$ is just affordable under $(\mathbf{p}', w')$. By Walras' Law: $\mathbf{p}' \cdot x' = w' = \mathbf{p}' \cdot x$, hence $\mathbf{p}' \cdot (x' - x) = 0$.
3. If $x \neq x'$, WARP requires that x' was not affordable under $(\mathbf{p}, w)$: $\mathbf{p} \cdot x' > w = \mathbf{p} \cdot x$, hence $\mathbf{p} \cdot (x' - x) > 0$.
4. Adding the two inequalities: $(\mathbf{p}' - \mathbf{p}) \cdot (x' - x) = \mathbf{p}' \cdot (x' - x) - \mathbf{p} \cdot (x' - x) < 0$.
5. For infinitesimal changes ($dx = S\, d\mathbf{p}$): $d\mathbf{p}^\top S\, d\mathbf{p} \leq 0$, confirming that the Slutsky matrix S is NSD. ■

1.2.5 Summary of Key Results

Concept	Definition	Key Condition
Completeness	All bundles are comparable	$x \succeq y$ or $y \succeq x$
Transitivity	Consistent ordering	$x \succeq y,\ y \succeq z \implies x \succeq z$
Continuity	No jumps in preferences	$x^n \succeq y^n \to (x, y) \implies x \succeq y$
Monotonicity	More is better	$x \geq y, x \neq y \implies x \succ y$
Convexity	Diversity is valued	$\alpha x + (1-\alpha) y \succeq z$
LNS	Always a better bundle nearby	$\forall x, \epsilon, \exists y : \|y - x\| < \epsilon, y \succ x$
WARP	Consistency of revealed choice	$\mathbf{p} \cdot x(\mathbf{p}', w') \leq w \implies$ $\mathbf{p}' \cdot x(\mathbf{p}, w) > w'$
SARP	Transitive revealed preference	No cycles in revealed choices

---

1.3 Utility Representation, UMP, and Indirect Utility

1.3.1 Motivation

The utility function simplifies comparison of bundles by assigning a real number to each alternative, transforming qualitative comparison into numerical magnitude comparison. It does not measure a physical substance in the brain — it creates a scale that respects the *order* of preferences.

Topographic Map Analogy: Monotonicity ensures that walking northeast (more of all goods) goes up the mountain. Continuity ensures there are no infinite vertical cliffs, allowing continuous contour lines (indifference curves). Under these conditions, the "satisfaction" of any point is measured by comparing it with an equivalent point on the diagonal from the origin.

1.3.2 The Utility Maximization Problem (UMP)

The UMP analytically models rational choice by converting the qualitative ordering into a constrained optimization problem over a compact set:

$$\max_{x \in \mathbb{R}^L_+} u(x) \quad \text{s.t.} \quad \mathbf{p} \cdot x \leq w$$

The argument that solves this problem constitutes the **Marshallian Demand** $x(\mathbf{p}, w)$. Evaluating the utility function directly at the optimal demand point yields the **Indirect Utility Function** $v(\mathbf{p}, w)$:

$$v(\mathbf{p}, w) = u(x(\mathbf{p}, w))$$

Auxiliary Theorem 1 — Weierstrass If f is continuous on a non-empty compact set, then f attains a maximum and a minimum on that set.

Auxiliary Theorem 2 — Envelope (Restricted Version) Let $M(a) = \max_x f(x, a)$ s.t. $g(x, a) = 0$ with Lagrangian $\mathcal{L} = f - \lambda g$. Then:

$$\frac{dM(a)}{da} = \frac{\partial \mathcal{L}(x^*(a), \lambda(a), a)}{\partial a}$$

1.3.3 Proofs and Fundamental Identities

Proof 1: Invariance to Monotonic Transformations **Statement:** If $u(x)$ represents $\succeq$, then $v(x) = f(u(x))$ also represents $\succeq$ for any strictly increasing f.

Proof:

1. By definition: $x \succeq y \iff u(x) \geq u(y)$.
2. f strictly increasing preserves order: $a \geq b \iff f(a) \geq f(b)$.
3. Therefore: $u(x) \geq u(y) \iff f(u(x)) \geq f(u(y))$.
4. Combining: $x \succeq y \iff f(u(x)) \geq f(u(y))$, confirming that $v = f \circ u$ represents $\succeq$. ■

Proof 2: Quasi-Concavity of Utility **Statement:** The preference relation $\succeq$ is convex if and only if the utility function $u(x)$ is quasi-concave.

Proof:

- **Necessity ($\Rightarrow$):** Let two bundles be such that $u(x) \geq u(z)$ and $u(y) \geq u(z)$, implying $x \succeq z$ and $y \succeq z$. By convexity of preferences, the linear combination satisfies $\alpha x + (1 - \alpha)y \succeq z$. Translating to the numerical scale, we have $u(\alpha x + (1 - \alpha)y) \geq u(z)$, proving that the upper contour sets are convex (definition of quasi-concavity).
- **Sufficiency ($\Leftarrow$):** If u is quasi-concave, by definition $u(\alpha x + (1 - \alpha)y) \geq \min\{u(x), u(y)\}$. If we establish that $x \succeq y$, then $\min\{u(x), u(y)\} = u(y)$. Hence, $u(\alpha x + (1 - \alpha)y) \geq u(y)$, confirming that $\alpha x + (1 - \alpha)y \succeq y$. ■

Proof 3: Existence of Continuous Utility Function **Statement:** Under rational, continuous, and strictly monotonic preferences on $\mathbb{R}^L_+$, there exists a continuous $u(x)$ that represents them.

Proof:

1. Let $\mathbf{e} = (1, \dots, 1)^\top$. For each x, define:

$$A = \{\alpha \geq 0 : \alpha\mathbf{e} \succeq x\}, \qquad B = \{\alpha \geq 0 : x \succeq \alpha\mathbf{e}\}$$

2. By completeness, $A \cup B = [0, \infty)$. By continuity, A and B are closed. Both are non-empty by monotonicity.

3. Since $[0, \infty)$ is connected, $A \cap B \neq \emptyset$. Let $\alpha(x) \in A \cap B$; then $\alpha(x)\mathbf{e} \sim x$.
4. By strict monotonicity, $\alpha(x)$ is unique. Define $u(x) = \alpha(x)$.
5. **Representation:** $x \succeq y \Rightarrow u(x)\mathbf{e} \sim x \succeq y \sim u(y)\mathbf{e} \Rightarrow u(x) \geq u(y)$ (by monotonicity).
6. **Continuity:** $u^{-1}([a,b]) = \{x : a\mathbf{e} \preceq x \preceq b\mathbf{e}\} = S(a\mathbf{e}) \cap I(b\mathbf{e})$, intersection of closed sets — hence closed. ■

Proof 4: Homogeneity of Indirect Utility **Statement:** $x(\alpha\mathbf{p}, \alpha w) = x(\mathbf{p}, w)$ and $v(\alpha\mathbf{p}, \alpha w) = v(\mathbf{p}, w)$ for all $\alpha > 0$.

Proof:

1. $B_{\alpha\mathbf{p},\alpha w} = \{x : (\alpha\mathbf{p}) \cdot x \leq \alpha w\}$. Dividing by $\alpha > 0$, the set is identical to $B_{\mathbf{p},w}$.
2. Identical sets $\Rightarrow$ same maximizer: $x(\alpha\mathbf{p}, \alpha w) = x(\mathbf{p}, w)$.
3. Therefore: $v(\alpha\mathbf{p}, \alpha w) = u(x(\alpha\mathbf{p}, \alpha w)) = u(x(\mathbf{p}, w)) = v(\mathbf{p}, w)$. ■

Proof 5: First-Order Conditions (FOC) and Roy's Identity **Statement:** At the interior optimum of the UMP, the marginal rate of substitution equals the price ratio, and Roy's Identity holds: $x_i(\mathbf{p}, w) = -\frac{\partial v/\partial p_i}{\partial v/\partial w}$.

Proof:

1. Set up the Lagrangian for the UMP with multiplier λ:

$$\mathcal{L}(x, \lambda) = u(x) + \lambda(w - \mathbf{p} \cdot x)$$

2. The interior FOCs require the partial derivative to be zero for every good i:

$$\frac{\partial \mathcal{L}}{\partial x_i} = \frac{\partial u}{\partial x_i} - \lambda p_i = 0 \implies \frac{\partial u}{\partial x_i} = \lambda p_i$$

3. Taking the ratio of the FOCs for two distinct goods i and j, we isolate the multiplier and obtain the classical equality:

$$\frac{\partial u/\partial x_i}{\partial u/\partial x_j} = \frac{p_i}{p_j}$$

4. By the **Envelope Theorem**, the derivative of the value function $v(\mathbf{p}, w)$ with respect to the wealth parameter w is given by the partial derivative of the Lagrangian evaluated at the optimum:

$$\frac{\partial v}{\partial w} = \frac{\partial \mathcal{L}}{\partial w} = \lambda$$

5. Applying the Envelope Theorem again, now for the price of good i (p_i):

$$\frac{\partial v}{\partial p_i} = \frac{\partial \mathcal{L}}{\partial p_i} = -\lambda x_i$$

6. Substituting the value of λ obtained in step (4) into the equation from step (5), we isolate the consumption of good i:

$$\frac{\partial v}{\partial p_i} = -\left(\frac{\partial v}{\partial w}\right) x_i \implies x_i(\mathbf{p}, w) = -\frac{\partial v/\partial p_i}{\partial v/\partial w}$$

■

Proof 6: Quasiconvexity of Indirect Utility in p **Statement:** The set $V_{\bar{u}} = \{\mathbf{p} \mid v(\mathbf{p}, w) \leq \bar{u}\}$ is convex for any $\bar{u}$.

Proof:

1. Let $\mathbf{p}^1, \mathbf{p}^2$ with $v(\mathbf{p}^1, w) \leq \bar{u}$ and $v(\mathbf{p}^2, w) \leq \bar{u}$. Let $\mathbf{p}^t = t\mathbf{p}^1 + (1-t)\mathbf{p}^2$, and $x^t = x(\mathbf{p}^t, w)$.
2. By feasibility: $t\mathbf{p}^1 \cdot x^t + (1-t)\mathbf{p}^2 \cdot x^t \leq w$.
3. Hence, either $\mathbf{p}^1 \cdot x^t \leq w$ or $\mathbf{p}^2 \cdot x^t \leq w$ (otherwise the combination would violate the inequality).
4. In either case: $v(\mathbf{p}^t, w) = u(x^t) \leq v(\mathbf{p}^k, w) \leq \bar{u}$ for $k = 1$ or $k = 2$.
5. Therefore $V_{\bar{u}}$ is convex, confirming quasiconvexity. ■

1.3.4 Summary of Key Results

Concept	Definition	Key Formula
Utility Function	Numerical representation of preferences	$x \succeq y \iff u(x) \geq u(y)$
Marshallian Demand	Solution to UMP	$x(\mathbf{p}, w) = \arg\max_{p \cdot x \leq w} u(x)$
Indirect Utility	Maximum utility given prices and income	$v(\mathbf{p}, w) = u(x(\mathbf{p}, w))$
Roy's Identity	Recover demand from indirect utility	$x_i(\mathbf{p}, w) = -\frac{\partial v / \partial p_i}{\partial v / \partial w}$
Homogeneity of v	No money illusion	$v(\alpha\mathbf{p}, \alpha w) = v(\mathbf{p}, w)$
Quasiconvexity of v	Convex lower contour sets	$V_{\bar{u}} = \{\mathbf{p} : v(\mathbf{p}, w) \leq \bar{u}\}$ is convex

---

1.4 Duality: EMP, Hicksian Demand, and Slutsky & Roy Equations

1.4.1 Motivation

The Expenditure Minimization Problem (EMP) inverts the logic of the UMP: instead of "how much satisfaction can I get with my budget?", we ask "what is the minimum cost to attain a specific satisfaction level?" Duality provides "reverse engineering" tools for welfare analysis.

Dinner Party Analogy: You must organize a dinner with a required quality standard (target utility), buying ingredients at the lowest possible cost. The shopping list that minimizes cost is the **Hicksian (compensated) Demand**.

- **Roy's Identity:** Allows extraction of observable demand by analyzing how indirect utility varies with prices and income.
- **Slutsky Equation:** Isolates the *Substitution Effect* (the good became relatively cheaper) from the *Income Effect* (real purchasing power was altered).

1.4.2 Formal Definitions

Definition 1 — EMP

$$e(\mathbf{p}, u) = \min_{x \in \mathbb{R}^L_+} \mathbf{p} \cdot x \quad \text{s.t.} \quad u(x) \geq u$$

The minimizer is the **Hicksian demand** $h(\mathbf{p}, u)$; the minimum value is the **expenditure function** $e(\mathbf{p}, u)$.

Definition 2 — Duality Identities

$$h_i(\mathbf{p}, u) = x_i(\mathbf{p}, e(\mathbf{p}, u)) \qquad \text{and} \qquad v(\mathbf{p}, e(\mathbf{p}, u)) = u$$

Auxiliary Theorem — Young (Symmetry of Cross Derivatives) If f is twice continuously differentiable: $\dfrac{\partial^2 f}{\partial x_i \partial x_j} = \dfrac{\partial^2 f}{\partial x_j \partial x_i}$.

1.4.3 Proofs

Proof 1: Roy's Identity **Statement:** $x_i(\mathbf{p}, w) = -\dfrac{\partial v/\partial p_i}{\partial v/\partial w}$

Proof:

1. Lagrangian of the UMP: $\mathcal{L} = u(x) + \lambda(w - \sum_j p_j x_j)$.
2. By the Envelope Theorem applied to the indirect utility function: $\partial v/\partial p_i = \partial \mathcal{L}/\partial p_i = -\lambda x_i$.
3. By the Envelope Theorem for income: $\partial v/\partial w = \partial \mathcal{L}/\partial w = \lambda > 0$.
4. Dividing (2) by (3) and canceling the Lagrange multiplier λ: $\dfrac{\partial v/\partial p_i}{\partial v/\partial w} = -x_i$.
5. Hence: $x_i(\mathbf{p}, w) = -\dfrac{\partial v/\partial p_i}{\partial v/\partial w}$. ■

Proof 2: Shephard's Lemma **Statement:** $h_i(\mathbf{p}, u) = \dfrac{\partial e(\mathbf{p}, u)}{\partial p_i}$

Proof:

1. Lagrangian of the EMP: $\mathcal{L} = \sum_j p_j x_j + \lambda(u - u(x))$.
2. By the Envelope Theorem applied to the expenditure function: $\partial e/\partial p_i = \partial \mathcal{L}/\partial p_i$.
3. Since the utility constraint $u(x)$ does not depend on the price vector p_i: $\partial \mathcal{L}/\partial p_i = x_i$.
4. Evaluating the derivative at the optimum where $x_i = h_i(\mathbf{p}, u)$: $\partial e/\partial p_i = h_i(\mathbf{p}, u)$. ■

Proof 3: Homogeneity of Degree Zero of Hicksian Demand **Statement:** $h(\alpha\mathbf{p}, u) = h(\mathbf{p}, u)$ for all $\alpha > 0$.

Proof:

1. The EMP under the scaled price vector $\alpha\mathbf{p}$ is: $\min(\alpha\mathbf{p}) \cdot x = \alpha(\mathbf{p} \cdot x)$ s.t. $u(x) \geq u$.
2. Since $\alpha > 0$, the scalar does not distort the slope of the price hyperplane. Minimizing $\alpha(\mathbf{p} \cdot x)$ generates exactly the same *arg min* as minimizing the original problem $\mathbf{p} \cdot x$.
3. Therefore, $h(\alpha\mathbf{p}, u) = h(\mathbf{p}, u)$. ■

Proof 4: Symmetry, NSD, and Slope of the Slutsky Matrix **Statement:** $S(\mathbf{p}, u) = D_p h(\mathbf{p}, u)$ is symmetric, negative semidefinite (NSD), and satisfies $\partial h_i/\partial p_i \leq 0$.

Proof:

- **Symmetry:** By Shephard's Lemma, $h_i = \partial e/\partial p_i$. Differentiating with respect to p_j, we have $\partial h_i/\partial p_j = \partial^2 e/(\partial p_j \partial p_i)$. By Young's Theorem on the continuity of second derivatives: $\partial^2 e/(\partial p_j \partial p_i) = \partial^2 e/(\partial p_i \partial p_j) = \partial h_j/\partial p_i$.
- **NSD:** The Slutsky matrix $S = D_p^2 e$ represents the Hessian matrix of the expenditure function with respect to prices. Since $e(\mathbf{p}, u)$ is defined as the infimum of linear functions in the price variable, it is geometrically concave. The Hessian of any concave function is guaranteed to be negative semidefinite: $v^\top S v \leq 0$ for all vectors v.

- **Slope:** The diagonal elements $s_{ii} = \partial h_i/\partial p_i$ represent the eigenvalues under canonical vectors. Being an NSD matrix, these terms are strictly non-positive, indicating that compensated demand is strictly downward-sloping in its own price. ■

Proof 5: Slutsky Equation **Statement:**

$$\frac{\partial x_i(\mathbf{p}, w)}{\partial p_j} = \frac{\partial h_i(\mathbf{p}, u)}{\partial p_j} - x_j(\mathbf{p}, w)\frac{\partial x_i(\mathbf{p}, w)}{\partial w}$$

Proof:

1. Start from the fundamental duality identity: $h_i(\mathbf{p}, u) = x_i(\mathbf{p}, e(\mathbf{p}, u))$.
2. Differentiate both sides with respect to price p_j. On the right-hand side, since the price affects Marshallian demand directly and also through the required income $e(\mathbf{p}, u)$, apply the chain rule:

$$\frac{\partial h_i(\mathbf{p}, u)}{\partial p_j} = \frac{\partial x_i(\mathbf{p}, e(\mathbf{p}, u))}{\partial p_j} + \frac{\partial x_i(\mathbf{p}, e(\mathbf{p}, u))}{\partial w} \cdot \frac{\partial e(\mathbf{p}, u)}{\partial p_j}$$

3. By Shephard's Lemma, substitute the derivative of the expenditure function: $\partial e(\mathbf{p}, u)/\partial p_j = h_j(\mathbf{p}, u)$.
4. Evaluating at the consistency point of equilibrium where nominal income equals minimum expenditure ($w = e(\mathbf{p}, u)$), the identity guarantees that $h_j(\mathbf{p}, u) = x_j(\mathbf{p}, w)$.
5. Substituting the terms and isolating the derivative of the observable demand:

$$\frac{\partial x_i(\mathbf{p}, w)}{\partial p_j} = \frac{\partial h_i(\mathbf{p}, u)}{\partial p_j} - x_j(\mathbf{p}, w)\frac{\partial x_i(\mathbf{p}, w)}{\partial w}$$

■

1.4.4 Summary of Key Results

Concept	Definition	Key Formula
Expenditure Function	Minimum cost to achieve utility u	$e(\mathbf{p}, u) = \min_{u(x)\geq u} \mathbf{p} \cdot x$
Hicksian Demand	Cost-minimizing bundle	$h(\mathbf{p}, u) = \arg\min_{u(x)\geq u} \mathbf{p} \cdot x$
Shephard's Lemma	Recover Hicksian demand from e	$h_i(\mathbf{p}, u) = \partial e/\partial p_i$
Duality Identities	Consistency between UMP and EMP	$h_i(\mathbf{p}, u) = x_i(\mathbf{p}, e(\mathbf{p}, u))$, $v(\mathbf{p}, e(\mathbf{p}, u)) = u$
Slutsky Equation	Decompose price effect	$\frac{\partial x_i}{\partial p_j} = \frac{\partial h_i}{\partial p_j} - x_j \frac{\partial x_i}{\partial w}$
Slutsky Matrix	Substitution effects	$S = D_p h(\mathbf{p}, u)$ is symmetric and NSD

1.5 Duality and Monetary Welfare Measures

1.5.1 Motivation

Duality theory establishes a stable mathematical mirror between abstract ordinal constructs (utility) and tangible monetary magnitudes (expenditure). This provides the analytical foundation for modern social welfare measurement, policy impact evaluation, and cost-of-living index design.

Wedding Budget Analogy: The client imposes a fixed standard of luxury and buffet (target utility), and the organizer selects the combination of inputs that meets the goal at the lowest market cost. The final amount on the invoice represents the **expenditure function**.

1.5.2 Welfare Measures

- **Compensating Variation (CV):** The monetary amount that must be given to or taken from the agent *after* the price shock to preserve exactly his initial welfare level u^0.
- **Equivalent Variation (EV):** The change in income that would generate an impact on utility *equivalent* to the actual impact of the shock, using the original prices as the reference vector.
- **Consumer Surplus (ΔCS):** A geometric and empirical approximation calculated under the observable Marshallian demand curve.

1.5.3 Formal Definitions

Definition — Expenditure Function

$$e(\mathbf{p}, u) = \min_{x \geq 0} \{\mathbf{p} \cdot x \mid u(x) \geq u\}$$

1.5.4 Proofs

Proof 1: Homogeneity of Degree 1 in Prices **Statement:** $e(\lambda \mathbf{p}, u) = \lambda e(\mathbf{p}, u)$ for all $\lambda > 0$.

Proof:

1. Write the definition for scaled prices:

$$e(\lambda \mathbf{p}, u) = \min_{x} \{(\lambda \mathbf{p}) \cdot x \mid u(x) \geq u\} = \min_{x} \{\lambda (\mathbf{p} \cdot x) \mid u(x) \geq u\}$$

2. Since the scale factor satisfies $\lambda > 0$, it can be extracted directly from the linear minimization operator without altering the ordering of the constraint set:

$$e(\lambda \mathbf{p}, u) = \lambda \min_{x} \{\mathbf{p} \cdot x \mid u(x) \geq u\} = \lambda e(\mathbf{p}, u)$$

■

Proof 2: Concavity of the Expenditure Function in Prices **Statement:** $e(\mathbf{p}, u)$ is concave in the price vector $\mathbf{p}$.

Proof:

1. Let two distinct price vectors $\mathbf{p}^1$ and $\mathbf{p}^2$, and a convex linear combination $\mathbf{p}^t = t\mathbf{p}^1 + (1-t)\mathbf{p}^2$ for $t \in [0, 1]$. Let x^t be the optimal argument that minimizes expenditure under the weighted price $\mathbf{p}^t$, so that $e(\mathbf{p}^t, u) = \mathbf{p}^t \cdot x^t$.

2. Since x^t meets the utility constraint ($u(x^t) \geq u$), it is classified as a feasible bundle (though not necessarily optimal) in the problems under the isolated prices $\mathbf{p}^1$ and $\mathbf{p}^2$. By the minimization property: $\mathbf{p}^1 \cdot x^t \geq e(\mathbf{p}^1, u)$ and $\mathbf{p}^2 \cdot x^t \geq e(\mathbf{p}^2, u)$.

3. Multiplying the inequalities by the weighting scalars t and $(1-t)$ and summing them:

$$(t\mathbf{p}^1 + (1-t)\mathbf{p}^2) \cdot x^t \geq te(\mathbf{p}^1, u) + (1-t)e(\mathbf{p}^2, u)$$

4. Substituting the left-hand side by the definition from step (1), we validate the fundamental geometric property of concavity:

$$e(\mathbf{p}^t, u) \geq te(\mathbf{p}^1, u) + (1-t)e(\mathbf{p}^2, u)$$

■

Proof 3: The Four Fundamental Duality Identities **Statements:** (1) $e(\mathbf{p}, v(\mathbf{p}, w)) = w$ (2) $v(\mathbf{p}, e(\mathbf{p}, u)) = u$ (3) $x_i(\mathbf{p}, w) = h_i(\mathbf{p}, v(\mathbf{p}, w))$ (4) $h_i(\mathbf{p}, u) = x_i(\mathbf{p}, e(\mathbf{p}, u))$

Proof of Identity (1):

1. Let $x^* = x(\mathbf{p}, w)$ be the optimal bundle that solves the UMP. By the local non-satiation (LNS) axiom, the budget constraint is exhausted with equality: $\mathbf{p} \cdot x^* = w$, and the maximum utility achieved is $u(x^*) = v(\mathbf{p}, w)$.
2. Since $u(x^*) = v(\mathbf{p}, w)$, the bundle x^* meets the minimum requirements of the EMP for a utility target equal to $v(\mathbf{p}, w)$, which establishes the upper bound: $e(\mathbf{p}, v(\mathbf{p}, w)) \leq \mathbf{p} \cdot x^* = w$.
3. Suppose by contradiction that the minimum expenditure were strictly less than the available income ($e(\mathbf{p}, v(\mathbf{p}, w)) < w$). This would imply the existence of an alternative bundle x' such that $\mathbf{p} \cdot x' < w$ with $u(x') \geq v(\mathbf{p}, w)$.
4. By the LNS hypothesis, if there is a budget surplus ($w - \mathbf{p} \cdot x' > 0$), we can construct a new local bundle x'' spending the remaining funds such that $u(x'') > u(x') \geq v(\mathbf{p}, w)$, violating the premise that $v(\mathbf{p}, w)$ was the maximum utility of the UMP problem.
5. Hence, by contradiction, the strict inequality is discarded: $e(\mathbf{p}, v(\mathbf{p}, w)) = w$. Identities (2), (3), and (4) derive analogously from the strict monotonicity of the functions and the uniqueness of optimal points guaranteed by strict quasi-concavity of preferences. ■

Proof 4: CV as a Hicksian Integral **Statement:** For a price change in good 1 from p_1^0 to p_1^1:

$$CV = \int_{p_1^1}^{p_1^0} h_1(\mathbf{p}, u^0)\, dp_1$$

Proof:

1. The compensating variation is the adjustment in endowment to restore the initial utility: $CV = w - e(\mathbf{p}^1, u^0)$.
2. Using Identity (1) proved above, substitute the initial income by the equivalent expenditure function: $w = e(\mathbf{p}^0, u^0)$, converting the measure into: $CV = e(\mathbf{p}^0, u^0) - e(\mathbf{p}^1, u^0)$.
3. Applying the Fundamental Theorem of Calculus to express the difference in values as the integral of its partial derivative:
$$CV = -\int_{p_1^0}^{p_1^1} \frac{\partial e(\mathbf{p}, u^0)}{\partial p_1}\, dp_1$$
4. Invoking Shephard's Lemma, substitute the derivative of the expenditure function by the Hicksian demand: $\partial e / \partial p_1 = h_1(\mathbf{p}, u^0)$.
5. Inverting the integration limits to eliminate the negative sign, we obtain the classical formula:
$$CV = \int_{p_1^1}^{p_1^0} h_1(\mathbf{p}, u^0)\, dp_1$$

■

Proof 5: Coincidence of Measures under Quasi-Linear Preferences **Statement:** If utility is quasi-linear in the numeraire good 1, for price shocks in other goods: $CV = EV = \Delta CS$.

Proof:

1. The quasi-linear preference structure ensures that second-order effects on wealth do not impact the demand for superior goods ($i > 1$), generating a zero marginal propensity to consume out of income: $\partial x_i / \partial w = 0$.

2. Analyzing the Slutsky Equation under this property: $\partial x_i/\partial p_j = \partial h_i/\partial p_j - x_j(0)$. The income effect term completely vanishes for all these goods.
3. Since the marginal slopes are identical and independent of financial endowment, the Marshallian and Hicksian demand functions are geometrically superimposed in space: $x_i(\mathbf{p}, w) = h_i(\mathbf{p}, u)$ for any reference utility level u.
4. This implies the direct equivalence of demand functions for different welfare levels: $h_i(\mathbf{p}, u^0) = h_i(\mathbf{p}, u^1) = x_i(\mathbf{p}, w)$.
5. Since the Hicksian and Marshallian integrands are perfectly identical, the three geometric areas calculated by the integrals collapse into the same monetary value: $CV = EV = \Delta CS$. ■

1.5.5 Summary of Key Results

Concept	Definition	Key Formula
Expenditure Function	Minimum cost to achieve utility u	$e(\mathbf{p}, u) = \min_{u(x) \geq u} \mathbf{p} \cdot x$
Compensating Variation (CV)	Income change to maintain utility after price change	$CV = e(\mathbf{p}^0, u^0) - e(\mathbf{p}^1, u^0)$
Equivalent Variation (EV)	Income change equivalent to price change at original prices	$EV = e(\mathbf{p}^0, u^1) - e(\mathbf{p}^0, u^0)$
Consumer Surplus (ΔCS)	Area under Marshallian demand	$\Delta CS = \int_{p^1}^{p^0} x_i(\mathbf{p}, w)\, dp_i$
CV as Hicksian Integral	CV equals area under Hicksian demand	$CV = \int_{p_1^1}^{p_1^0} h_1(\mathbf{p}, u^0)\, dp_1$
Quasi-Linear Equivalence	CV = EV = ΔCS	$\partial x_i/\partial w = 0$ for $i > 1$

1.6 The Structure of Preferences

1.6.1 Separability and Two-Stage Budgeting

Motivation The separability property postulates that the consumption bundle can be sliced into independent subgroups ("Food", "Clothing", "Leisure"). Under this hypothesis, choices within a specific group depend exclusively on the relative prices of that sector and the total budget allocated to it, without direct interference from the prices of goods in other categories.

Practical Analogy: The Section Budget Imagine you enter a bookstore with two separate budgets in envelopes: R$50 for "Fiction" and R$50 for "History". When evaluating which Fiction books to buy, you look only at the prices of fiction books on that shelf. If a biography in the History section suffers a price increase, this may change the initial division of your money between envelopes in the future, but it does not alter in any way your preference ordering or decision between the fiction novels today.

The Two-Stage Budgeting Mechanism

1. **First Stage:** The consumer allocates total wealth among macro-groups based on aggregate sectoral price indices.
2. **Second Stage:** Maximizes the sub-utility within each category using only the fractional budget assigned to that group, ignoring the individual prices of other goods in the economy.

Definition: Weak Separability A utility function $u(x)$ is **weakly separable** with respect to a partition of goods $\{N_1, \ldots, N_S\}$ if it can be expressed as:

$$u(x) = f(v_1(x^{(1)}), \ldots, v_S(x^{(S)}))$$

where f is a continuous and strictly increasing function, and each $v_s(x^{(s)})$ represents a sub-utility function that evaluates only the commodities belonging to group s.

> **Critical Model Failure** If preferences are not separable, the practical concept of estimating a "demand for food" as a function of a single price index loses theoretical validity. Without separability, Hicks' Composite Commodity Theorem would require all prices in the economy to vary in identical proportions to allow any legitimate aggregation.

Proof: Independence of the Marginal Rate of Substitution (MRS) **Statement:** The utility function $u(x)$ is weakly separable in group G if and only if the MRS between any pair of goods $i, j \in G$ is totally independent of the consumption level of any external good $k \notin G$.

Proof:

1. Write the function under the weak separability hypothesis: $u(x) = f(v_G(x_G), x_{\bar{G}})$, where $x_{\bar{G}}$ represents the vector of goods outside G.
2. Applying the classical chain rule to compute the marginal utilities of goods i and j belonging to G:

$$\frac{\partial u}{\partial x_i} = \frac{\partial f}{\partial v_G}\frac{\partial v_G}{\partial x_i}, \qquad \frac{\partial u}{\partial x_j} = \frac{\partial f}{\partial v_G}\frac{\partial v_G}{\partial x_j}$$

3. Calculate the structure of the Marginal Rate of Substitution (MRS):

$$MRS_{ij} = \frac{\partial u/\partial x_i}{\partial u/\partial x_j} = \frac{(\partial f/\partial v_G)(\partial v_G/\partial x_i)}{(\partial f/\partial v_G)(\partial v_G/\partial x_j)}$$

4. By the monotonicity axiom, the partial derivative of the macro-function is strictly positive ($\partial f/\partial v_G > 0$), allowing its direct algebraic cancellation:

$$MRS_{ij} = \frac{\partial v_G/\partial x_i}{\partial v_G/\partial x_j}$$

5. Since the sub-function v_G has as its argument only the vector x_G, the final expression for MRS_{ij} contains no external elements. Therefore, the cross-derivative with respect to any good outside the group is zero: $\partial MRS_{ij}/\partial x_k = 0$ for all $k \notin G$. ■

1.6.2 Integrability and Classical Consistency

Motivation The integrability problem addresses the reverse path of traditional theory: suppose a researcher empirically estimates a system of demand functions from market data. Is there any mathematical guarantee that this observed behavior stems from the maximizing choice of a rational consumer endowed with consistent and transitive preferences?

Practical Analogy: Footprints in the Forest Imagine finding deep footprints pressed into the forest floor. Analyzing the distance, depth, and shape of the steps, a biologist can reconstruct the size and anatomy of the animal that passed through there. Integrability works as the geometric structural test that determines whether the set of "footprints" left by consumption choices in the market actually has a geometry compatible with the anatomy of a stable rational maximizing agent.

Definition: The Integrability Conditions of Demand A system of observable demand functions $x(\mathbf{p}, w)$ is considered **integrable** if and only if it cumulatively satisfies four fundamental properties:

1. **Homogeneity of Degree Zero:** $x(\alpha\mathbf{p}, \alpha w) = x(\mathbf{p}, w)$ for all $\alpha > 0$.
2. **Walras' Law (Income Exhaustion):** $\mathbf{p} \cdot x(\mathbf{p}, w) = w$.
3. **Slutsky Symmetry:** $s_{ij}(\mathbf{p}, w) = s_{ji}(\mathbf{p}, w)$ for all pairs i, j.
4. **Negative Semidefiniteness:** The Slutsky matrix $S(\mathbf{p}, w)$ is NSD.

Where the terms are computed by:

$$s_{ij} = \frac{\partial x_i}{\partial p_j} + \frac{\partial x_i}{\partial w} x_j$$

Critical Model Failure If Slutsky symmetry is violated in the empirical data, the system of partial differential equations (PDEs) that reconstructs the expenditure function becomes overdetermined and inconsistent. Geometrically, the Frobenius integrability condition fails, meaning it is impossible to map these choices to any transitive preference relation. The economic agent would irremediably violate classical rationality.

Proof: Symmetry as an Integrability Condition **Statement:** Slutsky substitution effect symmetry ($s_{ij} = s_{ji}$) is a necessary condition for the integrability of the demand system.

Proof:

1. Resort to the fundamental duality identity: $h_i(\mathbf{p}, u) = x_i(\mathbf{p}, e(\mathbf{p}, u))$.
2. Substitute the Hicksian demand using the direct result from Shephard's Lemma ($h_i = \partial e/\partial p_i$):
$$\frac{\partial e(\mathbf{p}, u)}{\partial p_i} = x_i(\mathbf{p}, e(\mathbf{p}, u)) \quad \forall i = 1, \ldots, L$$
3. The previous step constitutes a complex system of partial differential equations where the expenditure function $e(\mathbf{p}, u)$ is the unknown to be integrated. For a unique solution to exist, the Integrability Theorem requires equality of second derivatives: $\partial^2 e/(\partial p_j \partial p_i) = \partial^2 e/(\partial p_i \partial p_j)$.
4. Differentiate the differential equation with respect to a generic price p_j, applying the chain rule on the right-hand side and maintaining consistency of the income argument ($w = e(\mathbf{p}, u)$):
$$\frac{\partial^2 e(\mathbf{p}, u)}{\partial p_j \partial p_i} = \frac{\partial x_i(\mathbf{p}, e(\mathbf{p}, u))}{\partial p_j} + \frac{\partial x_i(\mathbf{p}, e(\mathbf{p}, u))}{\partial w} \cdot \frac{\partial e(\mathbf{p}, u)}{\partial p_j}$$
5. Applying Shephard's Lemma again to substitute the remaining derivative ($\partial e(\mathbf{p}, u)/\partial p_j = h_j(\mathbf{p}, u) = x_j(\mathbf{p}, w)$), we reduce the expression to the classical Slutsky form:
$$\frac{\partial^2 e}{\partial p_j \partial p_i} = \frac{\partial x_i}{\partial p_j} + \frac{\partial x_i}{\partial w} x_j = s_{ij}$$
6. Executing the same procedure inverting the order of differentiation to obtain the symmetric cross-derivative:
$$\frac{\partial^2 e}{\partial p_i \partial p_j} = \frac{\partial x_j}{\partial p_i} + \frac{\partial x_j}{\partial w} x_i = s_{ji}$$
7. By Young's Theorem, second-order partial derivatives of a continuously differentiable function must be rigorously identical. Therefore, equating steps (5) and (6), we prove that Slutsky symmetry is mandatory for system consistency: $s_{ij} = s_{ji}$. ■

1.6.3 Homotheticity and Scale Properties

Motivation The homotheticity property establishes that the importance and relative proportion of goods in the chosen bundle remain totally static as the individual becomes wealthier. If a consumption bundle A is preferred to bundle B, proportionally expanding both alternatives by a scale factor preserves the same choice ordering. Geometrically, indifference curves maintain identical slopes along any linear ray projected from the origin of the consumption space.

Practical Analogy: The Architect's Blueprint If an architect expands the scale of a house drawing by ten times, the absolute size of the walls grows, but the exact proportion and geometry between rooms remain rigorously identical. Homothetic preferences replicate this spatial mechanics: the map of indifference curves is merely a scaled-up or scaled-down projection, without distortion of consumption proportions according to the size of the wallet.

Fundamental Characteristics

- **Constant MRS along Rays:** Moving along any straight line from the origin, the slope of indifference curves remains constant.
- **Unit Income Elasticity ($\eta_i = 1$):** A 10% increase in the individual's wealth results in an exact 10% increase in quantities demanded of all goods. Consequently, budget shares remain fixed.

Definition: Homothetic Preference A rational preference relation $\succeq$ on $\mathbb{R}^L_+$ is **homothetic** if it satisfies the following scale invariance property:

$$x \sim y \iff \alpha x \sim \alpha y \quad \forall \alpha > 0$$

Critical Model Failure If preferences violate homotheticity, Engel curves cease to be straight lines passing through the origin. Under non-homothetic scenarios, goods are dynamically differentiated into categories of **necessity goods** ($\eta_i < 1$) or **luxury goods** ($\eta_i > 1$), altering the qualitative structure of expenditure as the agent becomes wealthier.

Proof 1: Representation by Homogeneous Functions of Degree 1 **Statement:** A preference relation $\succeq$ is homothetic if and only if it can be represented by a utility function $u(x)$ that is homogeneous of degree 1, i.e., $u(\alpha x) = \alpha u(x)$ for all $\alpha > 0$.

Proof:

1. Consider the unit diagonal vector $\mathbf{e} = (1, \dots, 1)^\top$. For any generic bundle x, define the function $u(x)$ as the unique scalar multiplier $\alpha(x)$ that satisfies the indifference relation on the diagonal: $\alpha(x)\mathbf{e} \sim x$.
2. By direct construction of the scale, we have the fundamental identity: $x \sim u(x)\mathbf{e}$.
3. Applying the scale invariance axiom of homothetic preference, multiply both sides of the relation by a common scale factor $\alpha > 0$:
$$\alpha x \sim \alpha[u(x)\mathbf{e}] = [\alpha u(x)]\mathbf{e}$$
4. Now, apply the definition of the utility function to evaluate the new scaled argument αx:
$$\alpha x \sim u(\alpha x)\mathbf{e}$$
5. By transitivity between the indifference relations established in steps (3) and (4), we unite the ends:
$$u(\alpha x)\mathbf{e} \sim [\alpha u(x)]\mathbf{e}$$
6. By virtue of the strict monotonicity axiom applied along the diagonal straight line, a higher scale scalar level implies strict preference. Hence, to maintain the indifference in step (5), the numerical values must be identical: $u(\alpha x) = \alpha u(x)$, attesting to homogeneity of degree 1. ■

Proof 2: Linearity of Marshallian Demand in Income **Statement:** Under homothetic preferences, Marshallian demand is perfectly linear in income: $x(\mathbf{p}, \alpha w) = \alpha x(\mathbf{p}, w)$ for any $\alpha > 0$.

Proof:

1. The original UMP problem with income level w is $\max u(x)$ s.t. $\mathbf{p} \cdot x \leq w$, whose unique solution is $x(\mathbf{p}, w)$.
2. When raising income by factor α, the new UMP becomes: $\max u(x)$ s.t. $\mathbf{p} \cdot x \leq \alpha w$, generating the solution $x(\mathbf{p}, \alpha w)$.
3. Perform an analytical change of variable defining $z = x/\alpha$, which implies writing $x = \alpha z$. Substituting into the budget constraint inequality, we have: $\mathbf{p} \cdot (\alpha z) \leq \alpha w$. Dividing both sides by $\alpha > 0$, the clean constraint returns to: $\mathbf{p} \cdot z \leq w$.

4. Evaluate the impact of the variable change on the objective function. Invoking the homogeneity of degree 1 property proved earlier: $u(x) = u(\alpha z) = \alpha u(z)$.
5. Uniting the constraint and the modified objective, the optimization problem in variable z reconfigures as: $\max \alpha u(z)$ s.t. $\mathbf{p} \cdot z \leq w$.
6. Since the scalar multiplier α affects the entire objective function linearly, it has no capacity to alter the location of the maximizer. The optimal argument that solves the problem in z is identically the base solution: $z^* = x(\mathbf{p}, w)$.
7. Returning from the variable change ($z^* = x(\mathbf{p}, \alpha w)/\alpha$), we equate the expressions: $x(\mathbf{p}, \alpha w)/\alpha = x(\mathbf{p}, w)$, which isolates the linear result: $x(\mathbf{p}, \alpha w) = \alpha x(\mathbf{p}, w)$. ■

1.6.4 Extensions: Demand Aggregation and the Representative Consumer

Motivation In macroeconomics and general equilibrium models, the simplification of a single "Representative Consumer" is frequently employed. However, aggregate market demand depends crucially on how wealth is distributed among the population, not just on the aggregate amount. For us to be able to ignore inequality in modeling, a highly restrictive theoretical structure is required: all individuals must have rigorously identical marginal propensities to consume.

Theorem: The Gorman Polar Form A normative and positive representative consumer exists completely independently of the internal distribution of wealth if and only if the indirect utility functions of all individuals in the economy belong to the Gorman class:

$$v_i(\mathbf{p}, w_i) = a_i(\mathbf{p}) + b(\mathbf{p})w_i$$

where the wealth sensitivity coefficient $b(\mathbf{p})$ must be rigorously identical for all economic agents.

> **SMD Theorem (Sonnenschein-Mantel-Debreu)** If agents' preferences violate the restricted structure of the Gorman Form, the aggregate market demand function can systematically violate WARP and all revealed preference axioms, even if each individual is perfectly rational in isolation. The macroeconomic aggregate demand curve can take virtually any chaotic or bizarre geometric shape, making reduction to a stable representative agent impossible.

Proof 1: Necessary Condition for Distribution-Independent Aggregation **Statement:** The aggregate demand function for good l depends solely on the total sum of wealth ($w = \sum w_i$) if and only if the marginal propensity to consume the good is identical among all individuals in the economy: $\partial x_{li}/\partial w_i = \beta_l(\mathbf{p})$ for all i.

Proof:

1. Define the aggregate market demand function for commodity l as the direct sum of the individual demands of each of the I consumers:

$$x_l(\mathbf{p}, w_1, \ldots, w_I) = \sum_i x_{li}(\mathbf{p}, w_i)$$

2. For x_l to behave purely as a function of the linear sum $w = \sum w_i$, any marginal reallocation of wealth between two individuals (dw_i and dw_j) that keeps the macro-budget stable ($\sum_i dw_i = 0$) must generate a strictly zero net impact on summed demand ($dx_l = 0$).
3. Computing the total differential of aggregate demand with respect to individual wealth variations:

$$dx_l = \sum_{i=1}^{I} \left(\frac{\partial x_{li}}{\partial w_i} \right) dw_i = 0$$

4. By the fundamental linear algebra applicable to zero-sum hyperplane constraints, for the weighted sum in step (3) to result in zero for any displacement vector $\{dw_i\}$ satisfying $\sum dw_i = 0$, it is mathematically mandatory that all partial derivative weights be rigorously equal to a common value:
$$\frac{\partial x_{li}}{\partial w_i} = \beta_l(\mathbf{p}) \quad \forall i = 1, \ldots, I$$
5. Integrating the partial differential equation from step (4) with respect to each individual's wealth w_i, we obtain the linear structure of individual demand functions:
$$x_{li}(\mathbf{p}, w_i) = \alpha_{li}(\mathbf{p}) + \beta_l(\mathbf{p})w_i$$
This demonstrates geometrically that all agents' Engel curves must be perfectly parallel straight lines. ■

Proof 2: Aggregation of the Uncompensated Law of Demand (ULD) **Statement:** If each consumer i's demand function satisfies the Uncompensated Law of Demand (ULD) and each one's wealth is maintained as a fixed fraction of total wealth ($w_i = \alpha_i w$), then the aggregate market demand preserves the macroeconomic ULD.

Proof:

1. The ULD property applied to individual consumer i dictates that price difference and Marshallian consumption difference vectors move in opposite directions:
$$(\mathbf{p}' - \mathbf{p}) \cdot [x_i(\mathbf{p}', w_i') - x_i(\mathbf{p}, w_i)] \leq 0 \quad \forall i$$
2. Substitute the stable endowment constraint by inserting the fixed participation proportion for initial income ($w_i = \alpha_i w$) and post-shock income ($w_i' = \alpha_i w'$). Then, perform the summation of all individual inequalities over the I consumers in the market:
$$\sum_{i=1}^{I} (\mathbf{p}' - \mathbf{p}) \cdot [x_i(\mathbf{p}', \alpha_i w') - x_i(\mathbf{p}, \alpha_i w)] \leq 0$$
3. Exploiting the distributive linearity property of the dot product, isolate the price difference vector $(\mathbf{p}' - \mathbf{p})$ outside the summation operator:
$$(\mathbf{p}' - \mathbf{p}) \cdot \sum_{i=1}^{I} [x_i(\mathbf{p}', \alpha_i w') - x_i(\mathbf{p}, \alpha_i w)] \leq 0$$
4. Substitute the sum of individual demands by the consolidated aggregate demand operator:
$$(\mathbf{p}' - \mathbf{p}) \cdot [x(\mathbf{p}', w') - x(\mathbf{p}, w)] \leq 0$$
5. The resulting inequality attests that the macroeconomic consolidated demand perfectly satisfies the aggregate ULD criteria. This ensures that the consolidated market substitution matrix will preserve the NSD property, shielding the macroeconomic model against aggregate behavioral anomalies. ■

1.6.5 Summary of Key Results

Concept	Definition	Key Property
Weak Separability	Utility partitions into independent sub-functions	$u(x) = f(v_1(x^{(1)}), \ldots, v_S(x^{(S)}))$

Concept	Definition	Key Property
MRS Independence	MRS within group independent of outside goods	$\partial MRS_{ij}/\partial x_k = 0$ for $k \notin G$
Integrability	Demand system derives from rational preferences	Satisfies homogeneity, Walras' Law, Slutsky symmetry, NSD
Slutsky Symmetry	Necessary integrability condition	$s_{ij} = s_{ji}\ \forall i, j$
Homotheticity	Scale-invariant preferences	$x \sim y \iff \alpha x \sim \alpha y$
Gorman Form	Linear indirect utility for aggregation	$v_i(\mathbf{p}, w_i) = a_i(\mathbf{p}) + b(\mathbf{p})w_i$
SMD Theorem	Aggregate demand may violate WARP	Without Gorman form, no representative consumer

---

Chapter 2: Production Theory and the Firm

2.1 The Production Set and the Technological Frontier

2.1.1 Motivation

The concept of technology in economics addresses the fundamental problem of **technical feasibility**: given current scientific knowledge, which combinations of inputs and outputs are physically possible? The production frontier is the outer limit of this capacity, representing the state in which the firm operates with maximum efficiency—that is, without resource waste.

Practical Analogy: The Professional Kitchen Imagine a professional kitchen. "Technology" is the set of recipes, equipment, and spatial layout. The production set includes all combinations of ingredients (inputs) and dishes (outputs) that are possible to produce. The frontier is reached when the kitchen delivers the maximum number of dishes without burning ingredients or leaving chefs idle.

2.1.2 Mechanisms and Implications

Economic agents, especially firms, behave as maximizers of objectives under technical constraints. While consumers seek to maximize utility, the firm seeks the production plan that generates the greatest market value (profit) within what technology allows.

The economic logic behind properties such as convexity of the production set lies in the idea that "averages are better than extremes" or in the presence of diminishing returns: if two distinct production plans are feasible, any mixture (convex combination) between them must also be feasible, reflecting the ability to balance the use of different resources.

In practical economics, this concept defines the limits of economic growth and firm efficiency. Public policies aimed at incentivizing innovation essentially seek to "push" this frontier outward, allowing more to be produced with fewer resources.

Example: Carbon Taxation By making a polluting input more expensive through taxation, the government forces firms to choose new points on the production frontier or to invest in new technologies (shifting the frontier itself) that substitute carbon for cleaner inputs to maintain economic viability.

2.1.3 Formal Definition

Let L be the number of commodities in the goods space $\mathbb{R}^L$. A **production plan** is a vector $\mathbf{y} = (y_1, \ldots, y_L) \in \mathbb{R}^L$, where we adopt the standard sign convention of firm theory: - $y_l > 0$ if commodity l is an output. - $y_l < 0$ if commodity l is an input.

Under this convention, the firm's profit for a price vector $\mathbf{p} \in \mathbb{R}^L_{++}$ is simply the dot product $\mathbf{p} \cdot \mathbf{y}$, eliminating the need to explicitly subtract costs.

Definition: Production Set and Transformation Function The **Production Set** $Y \subset \mathbb{R}^L$ contains all technologically feasible plans $\mathbf{y}$. We describe Y through a **transformation function** $F : \mathbb{R}^L \to \mathbb{R}$, such that:

$$Y = \{\mathbf{y} \in \mathbb{R}^L : F(\mathbf{y}) \leq 0\}$$

The **Production Frontier** is the set of plans with no technical waste, defined by $\{\mathbf{y} \in \mathbb{R}^L : F(\mathbf{y}) = 0\}$.

Regular Hypotheses for Y:

- **Y is non-empty and closed:** The firm always has some production option available, and the limits of the set belong to it (guaranteeing existence of solutions via Weierstrass' Theorem).
- **Inaction possibility ($\mathbf{0} \in Y$):** It is possible to shut down, producing nothing and consuming no inputs.
- **Free Disposal:** If $\mathbf{y} \in Y$ and $\mathbf{y}' \leq \mathbf{y}$ (component-wise), then $\mathbf{y}' \in Y$. It is possible to eliminate products or waste inputs at no additional cost.
- **Convexity:** If $\mathbf{y}, \mathbf{y}' \in Y$, then $\alpha\mathbf{y} + (1-\alpha)\mathbf{y}' \in Y$ for all $\alpha \in [0, 1]$.

  Structural Model Failures - Without the **convexity** hypothesis, the firm's technology may exhibit increasing returns to scale. This prevents the existence of a perfect competitive equilibrium, as the firm's profit would grow indefinitely with increased production scale. - Without the **closedness** hypothesis, the profit maximization problem may have no solution (the supremum of the profit set may not be attainable by any plan contained in Y).

2.1.4 Proofs of Technological Properties

Proof 1: The Marginal Rate of Transformation (MRT) as the Slope of the Frontier Statement: If $F(\mathbf{y}) = 0$ defines a differentiable surface and $\partial F/\partial y_k \neq 0$, then the technical rate of exchange between two goods l and k is given by the ratio of their partial derivatives.

Proof:

1. Consider two goods l and k on the production frontier, keeping the quantities of all other goods constant ($dy_m = 0$ for all $m \neq l, k$).
2. By the definition of the frontier, the transformation constraint must remain zero: $F(\mathbf{y}) = 0$.
3. Applying total differentiation to the transformation function $F(\mathbf{y})$:

$$\frac{\partial F}{\partial y_l} dy_l + \frac{\partial F}{\partial y_k} dy_k = 0$$

4. By the additive property of equality, isolate the terms on opposite sides:

$$\frac{\partial F}{\partial y_l} dy_l = -\frac{\partial F}{\partial y_k} dy_k$$

5. Through algebraic manipulation of the differentials, isolate the rate of change of y_k with respect to y_l:

$$\frac{dy_k}{dy_l} = -\frac{\partial F/\partial y_l}{\partial F/\partial y_k}$$

6. Define the Marginal Rate of Transformation (MRT_{lk}) as the absolute value of the slope of the production frontier:

$$MRT_{lk} = \frac{\partial F/\partial y_l}{\partial F/\partial y_k}$$

■

Proof 2: Profit Maximization Implies Technological Efficiency **Statement:** A production plan $\mathbf{y}^* \in Y$ is technologically efficient if there is no other alternative plan $\mathbf{y}' \in Y$ such that $\mathbf{y}' \geq \mathbf{y}^*$ and $\mathbf{y}' \neq \mathbf{y}^*$.

Proof:

1. Suppose by contradiction that plan $\mathbf{y}^*$ maximizes the firm's profit $\mathbf{p} \cdot \mathbf{y}$ for a strictly positive price vector $\mathbf{p} \gg 0$, but $\mathbf{y}^*$ is not technologically efficient.
2. If $\mathbf{y}^*$ is not efficient, by definition there must exist a feasible plan $\mathbf{y}' \in Y$ such that $\mathbf{y}' \geq \mathbf{y}^*$ and $\mathbf{y}' \neq \mathbf{y}^*$ (i.e., producing at least the same quantity of outputs with fewer inputs, or more outputs with the same inputs).
3. Since all market prices are strictly positive ($\mathbf{p} \gg 0$), the monotonicity property of the dot product guarantees:
$$\mathbf{p} \cdot \mathbf{y}' > \mathbf{p} \cdot \mathbf{y}^*$$
4. This directly contradicts the initial hypothesis that $\mathbf{y}^*$ is the global maximizer of profit in Y.
5. Therefore, by contradiction, the optimal plan $\mathbf{y}^*$ must be technologically efficient. ■

Proof 3: Profit under Constant Returns to Scale (CRS) **Statement:** A technology exhibits constant returns to scale (CRS) if the proportional expansion or reduction of any feasible plan is also feasible, i.e., if $\mathbf{y} \in Y \implies \alpha\mathbf{y} \in Y$ for all $\alpha \geq 0$.

Proof:

1. Let $\mathbf{y}^*$ be the plan that maximizes the firm's profit under prices $\mathbf{p}$, generating maximum profit $\pi^* = \mathbf{p} \cdot \mathbf{y}^*$.
2. Suppose by contradiction that this maximum profit is strictly positive, i.e., $\pi^* > 0$.
3. By the CRS hypothesis, if the firm doubles the scale of the production plan by choosing $\mathbf{y}' = 2\mathbf{y}^*$, we have $\mathbf{y}' \in Y$ (applying the scale parameter $\alpha = 2$).
4. By the linearity property of the dot product, the profit generated by this expanded production plan would be:
$$\pi' = \mathbf{p} \cdot (2\mathbf{y}^*) = 2(\mathbf{p} \cdot \mathbf{y}^*) = 2\pi^*$$
5. Since we assumed $\pi^* > 0$, it follows that $2\pi^* > \pi^*$, invalidating the fact that $\mathbf{y}^*$ is the global maximum profit plan in Y.
6. Hence, for a finite maximum profit to exist under CRS, we must have $\pi^* \leq 0$. Since the firm always has the inaction option ($\mathbf{0} \in Y$), which generates profit exactly equal to zero, the long-run equilibrium maximum profit under CRS must be exactly $\pi^* = 0$. ■

2.1.5 Summary of Key Results

Concept	Definition	Key Formula
Production Set	All technologically feasible plans	$Y = \{\mathbf{y} \in \mathbb{R}^L : F(\mathbf{y}) \leq 0\}$
Production Frontier	Plans with no technical waste	$F(\mathbf{y}) = 0$
Marginal Rate of Transformation	Slope of the production frontier	$MRT_{lk} = \frac{\partial F/\partial y_l}{\partial F/\partial y_k}$
Free Disposal	Waste at no cost	$\mathbf{y} \in Y, \mathbf{y}' \leq \mathbf{y} \implies \mathbf{y}' \in Y$
Convexity	Averages of feasible plans are feasible	$\alpha\mathbf{y} + (1-\alpha)\mathbf{y}' \in Y$
CRS Profit	Long-run profit under CRS	$\pi^* = 0$

2.2 The Production Function and Cost Minimization

2.2.1 Motivation

The production function serves as the "recipe book" of the economy. Before introducing monetary variables, prices, or profits, we need to map how available resources are physically transformed. It determines the maximum quantity of final output obtainable from a specific combination of inputs.

Practical Analogy: The Bakery The production function describes how many loaves of bread can be baked per hour given the number of bakers (labor) and ovens (capital). If we change only one input while keeping the other fixed (many bakers for only one oven), the marginal production of the last baker will be low: he will spend more time waiting for the oven to become available than actually operating production.

2.2.2 Mechanisms and Implications

Firms act as maximizing agents constrained by technology. While consumer theory focuses on utility maximization, the cost minimization approach focuses on allocative efficiency: achieving an exogenous production target at the lowest possible monetary cost.

The geometry of isoquants (level curves sharing the same output volume) reflects the technical substitutability between factors. If the wage of labor increases, the firm will respond by altering the intensity of factor use, moving along the isoquant and substituting labor for capital (automation), modifying the Marginal Rate of Technical Substitution ($MRTS$) until economic equilibrium with the new relative market prices is reestablished.

Example: Minimum Wage Impact If the labor factor becomes artificially more expensive relative to capital, the theory predicts a substitution of workers for automation. This underpins why industries in economies with expensive labor tend to operate with much higher capital technology density than those in developing countries.

2.2.3 Formal Definition

Let L be the number of productive inputs available. Define: - **Input Vector:** $\mathbf{x} = (x_1, \ldots, x_L) \in \mathbb{R}^L_+$ (non-negative quantities of factors). - **Factor Price Vector:** $\mathbf{w} = (w_1, \ldots, w_L) \in \mathbb{R}^L_{++}$ (strictly positive costs).

Definition: Production Function The **Production Function** $f : \mathbb{R}^L_+ \to \mathbb{R}_+$ associates to each input vector $\mathbf{x}$ the maximum quantity of a single output $y = f(\mathbf{x})$ technologically attainable.

Regular Hypotheses of the Production Function:

- **Inaction ($f(\mathbf{0}) = 0$):** No output generation without inputs (no "free lunch").
- **Monotonicity:** If $\mathbf{x}' \geq \mathbf{x}$, then $f(\mathbf{x}') \geq f(\mathbf{x})$, reflecting the free disposal hypothesis (additional inputs never reduce total output).
- **Quasiconcavity:** The input requirement set $V(y) = \{\mathbf{x} \in \mathbb{R}^L_+ : f(\mathbf{x}) \geq y\}$ is a convex set, indicating that convex combinations of extreme input plans are at least as productive as the extremes.

  Implications of Hypothesis Violations - Without the **quasiconcavity** hypothesis, isoquant curves lose convexity toward the origin. This makes the cost minimization problem ill-defined, generating corner solutions, discontinuities, and drastic jumps in input demand functions under marginal price variations. - Without the **monotonicity** hypothesis, the marginal productivity of a factor could take negative values, violating the principle that the firm operates strictly on its technical efficiency frontier.

2.2.4 Proofs of Production Theory

Proof 1: The MRTS as the Ratio of Marginal Productivities **Statement:** If the differentiable relation $f(x_1, x_2) = y$ implicitly defines x_2 as a function of x_1, the total derivative dx_2/dx_1 exists at any point where the marginal productivity of the second factor is non-zero ($\partial f/\partial x_2 \neq 0$).

Proof:

1. Fix the output level on a generic isoquant such that $f(x_1, x_2) = y$, where y is an exogenous constant.
2. Differentiate both sides with respect to x_1, applying the chain rule for total differentiation:
$$\frac{\partial f}{\partial x_1}\frac{dx_1}{dx_1} + \frac{\partial f}{\partial x_2}\frac{dx_2}{dx_1} = \frac{dy}{dx_1}$$
3. Since y is a constant, its derivative is zero ($dy/dx_1 = 0$). As $dx_1/dx_1 = 1$, the expression simplifies to:
$$\frac{\partial f}{\partial x_1} \cdot 1 + \frac{\partial f}{\partial x_2}\frac{dx_2}{dx_1} = 0$$
4. Isolate the differential derivative term using the additive property:
$$\frac{\partial f}{\partial x_2}\frac{dx_2}{dx_1} = -\frac{\partial f}{\partial x_1}$$
5. Perform the algebraic division of the partial derivatives (where $f_i(x) = \partial F/\partial x_i$):
$$\frac{dx_2}{dx_1} = -\frac{\partial f/\partial x_1}{\partial f/\partial x_2}$$
6. Since the Marginal Rate of Technical Substitution ($MRTS_{12}$) represents the absolute value of the isoquant slope:
$$MRTS_{12} = \left|\frac{dx_2}{dx_1}\right| = \frac{f_1(x)}{f_2(x)}$$

■

Proof 2: Optimality Condition of Cost Minimization **Statement:** To minimize a linear cost function subject to a technological constraint $f(\mathbf{x}) = y$, the gradient of the objective function must be proportional to the gradient of the constraint at the optimum point.

Proof:

1. The cost minimization problem of a firm operating with two inputs is expressed as:
$$\min_{x_1, x_2} w_1x_1 + w_2x_2 \quad \text{s.t.} \quad f(x_1, x_2) = y$$
2. Construct the Lagrangian function incorporating the Lagrange multiplier λ:
$$\mathcal{L}(x_1, x_2, \lambda) = w_1x_1 + w_2x_2 - \lambda(f(x_1, x_2) - y)$$
3. Calculate the First-Order Conditions (FOC) for inputs x_1 and x_2:
$$\frac{\partial \mathcal{L}}{\partial x_1} = w_1 - \lambda\frac{\partial f}{\partial x_1} = 0 \implies w_1 = \lambda f_1(x)$$
$$\frac{\partial \mathcal{L}}{\partial x_2} = w_2 - \lambda\frac{\partial f}{\partial x_2} = 0 \implies w_2 = \lambda f_2(x)$$
4. Divide the two FOC equations term by term:
$$\frac{w_1}{w_2} = \frac{\lambda f_1(x)}{\lambda f_2(x)}$$

5. Cancel the common scalar λ and substitute the definition of $MRTS_{12}$ proved earlier:

$$\frac{w_1}{w_2} = \frac{f_1(x)}{f_2(x)} \implies \frac{w_1}{w_2} = MRTS_{12}$$

6. We conclude that the optimal input plan requires that the technical substitution rate dictated by nature equals the price ratio dictated by the market. ■

2.2.5 Summary of Key Results

Concept	Definition	Key Formula
Production Function	Maximum output from inputs	$y = f(\mathbf{x})$
Input Requirement Set	Inputs producing at least y	$V(y) = \{\mathbf{x} : f(\mathbf{x}) \geq y\}$
Isoquant	Level curve of production	$f(x_1, x_2) = y$
Marginal Rate of Technical Substitution	Slope of the isoquant	$MRTS_{12} = \frac{f_1(x)}{f_2(x)}$
Cost Minimization Condition	MRTS equals price ratio	$MRTS_{12} = \frac{w_1}{w_2}$
Lagrangian Multiplier	Shadow price of output	$\lambda = \frac{w_i}{f_i(x)}$

2.3 Elasticity of Substitution and Functional Forms

2.3.1 Motivation

The elasticity of substitution (σ) measures the degree of flexibility and curvature of an isoquant. While the $MRTS$ expresses the slope of the curve at a given point, the elasticity indicates the speed with which this slope changes as we alter the proportion of input use.

Practical Analogy: Energy vs. Automobiles A thermoelectric plant operating interchangeably with natural gas or diesel oil has a high elasticity of substitution (isoquants close to straight lines). In contrast, automobile production requires strict proportions: four tires and one chassis. Additional tires without new chassis do not expand output. Here, $\sigma = 0$, characterizing perfect complements (L-shaped isoquants).

2.3.2 Mechanisms and Scale Elasticity

In response to price fluctuations, firms seek paths that minimize cost impact. If the relative cost of a production factor rises, the intensity of its substitution will depend directly on the technological rigidity described by σ: - **High σ:** Small oscillations in relative prices induce large reconfigurations in the proportion of inputs used. - **Low σ:** Technology is rigid; even if a factor suffers strong price shocks, the firm remains tied to the original combination to avoid scale losses.

Strict quasiconcavity of the production function ensures that $\sigma \geq 0$. This blocks uneconomic behaviors, ensuring that firms never increase the proportion of an input whose relative price has increased.

Example: Automation Policies In industries where the elasticity of substitution between labor and robots is high, imposing taxes on technological capital will have a strong effect on preserving jobs. If the technology has a low σ, taxation will only generate increases in production costs passed on to the final consumer, without altering the employment structure.

2.3.3 Formal Definition

Let $f(x_1, x_2)$ be a C^2 production function based on two factors. Define the input proportion as $r = \frac{x_2}{x_1}$ and the marginal rate of technical substitution as $MRTS = \frac{f_1}{f_2}$.

Definition: Elasticity of Substitution The **Elasticity of Substitution** σ measures the proportional variation in the factor ratio in response to a proportional variation in the marginal rate of technical substitution along an isoquant:

$$\sigma = \frac{d\ln(x_2/x_1)}{d\ln(MRTS)} = \frac{MRTS}{x_2/x_1} \cdot \frac{d(x_2/x_1)}{d(MRTS)}$$

Regularity Conditions:

- **Differentiability ($f \in C^2$):** Allows consistent calculation of second-order derivatives.
- **Strict Monotonicity ($f_i > 0$):** Ensures that the gradient and $MRTS$ are positive and well-defined quantities.
- **Strict Quasiconcavity:** Guarantees strictly convex isoquants relative to the origin, which imposes the derivative $\frac{d(x_2/x_1)}{d(MRTS)} > 0$, resulting in $\sigma \geq 0$.

2.3.4 Proofs of Functional Structures

Proof 1: Calculation of σ for the CES Functional Structure **Statement:** The functional specification $y = [x_1^\rho + x_2^\rho]^{1/\rho}$ is homogeneous of degree 1 and has an $MRTS$ that depends exclusively on the input utilization ratio r.

Proof:

1. Start from the classic definition of the symmetric CES function: $y = (x_1^\rho + x_2^\rho)^{1/\rho}$.
2. Applying the chain rule, partially differentiate the function with respect to input x_1:

$$f_1 = \frac{1}{\rho}(x_1^\rho + x_2^\rho)^{\frac{1}{\rho}-1} \cdot \rho x_1^{\rho-1} = (x_1^\rho + x_2^\rho)^{\frac{1-\rho}{\rho}} x_1^{\rho-1}$$

3. By algebraic symmetry, obtain the corresponding partial derivative for input x_2:

$$f_2 = (x_1^\rho + x_2^\rho)^{\frac{1-\rho}{\rho}} x_2^{\rho-1}$$

4. Determine the $MRTS$ by the ratio of marginal productivities (f_1/f_2):

$$MRTS = \frac{(x_1^\rho + x_2^\rho)^{\frac{1-\rho}{\rho}} x_1^{\rho-1}}{(x_1^\rho + x_2^\rho)^{\frac{1-\rho}{\rho}} x_2^{\rho-1}} = \frac{x_1^{\rho-1}}{x_2^{\rho-1}} = \left(\frac{x_2}{x_1}\right)^{1-\rho}$$

5. Performing the variable substitution by the input ratio $r = \frac{x_2}{x_1}$, we obtain $MRTS = r^{1-\rho}$.
6. Apply the natural logarithm to both sides of the equation:

$$\ln(MRTS) = (1-\rho)\ln(r)$$

7. Differentiate the linear expression in logarithms with respect to $\ln(r)$:

$$\frac{d\ln(MRTS)}{d\ln(r)} = 1 - \rho$$

8. Since elasticity is defined as the inverse of the calculated logarithmic derivative:

$$\sigma = \left[\frac{d\ln(MRTS)}{d\ln(r)}\right]^{-1} = \frac{1}{1-\rho}$$

■

Proof 2: Technological Limits of the Elasticity of Substitution **Statement:** Analyze the asymptotic behavior of $\sigma = \frac{1}{1-\rho}$ under three analytical regimes of the substitution parameter ρ.

Proof:

1. **Cobb-Douglas Regime ($\rho \to 0$):** When the parameter ρ approaches zero, the elasticity limit converges to $\sigma \to 1$. Applying the limit of the original CES function as $\rho \to 0$ (via L'Hôpital's rule on the exponent), the expression converges to the classical multiplicative form:
$$y = x_1^\alpha x_2^{1-\alpha}$$
2. **Perfect Substitutes Regime ($\rho \to 1$):** As the parameter ρ approaches unity, the elasticity tends to infinity ($\sigma \to \infty$). Substituting $\rho = 1$ into the functional structure, the technology collapses into an additive linear relationship:
$$y = x_1 + x_2$$
3. **Leontief / Fixed Proportions Regime ($\rho \to -\infty$):** When the parameter continuously decreases toward $-\infty$, the elasticity is nullified ($\sigma \to 0$). By the analytical theory of p-norm limits, the CES function converges to the minimum operator, characterizing perfectly complementary inputs:
$$y = \min\{x_1, x_2\}$$

■

2.3.5 Summary of Key Results

Concept	Definition	Key Formula
Elasticity of Substitution	Curvature of isoquant	$\sigma = \frac{d\ln(x_2/x_1)}{d\ln(MRTS)}$
CES Function	Constant elasticity of substitution	$y = [x_1^\rho + x_2^\rho]^{1/\rho}$
CES MRTS	MRTS for CES technology	$MRTS = r^{1-\rho}$
CES σ	Elasticity for CES	$\sigma = \frac{1}{1-\rho}$
Cobb-Douglas	$\sigma = 1$	$y = x_1^\alpha x_2^{1-\alpha}$
Perfect Substitutes	$\sigma = \infty$	$y = x_1 + x_2$
Leontief	$\sigma = 0$	$y = \min\{x_1, x_2\}$

2.4 Returns to Scale and Scale Elasticity

2.4.1 Motivation

The real economic problem that returns to scale attempt to solve is understanding what happens to productive efficiency when a firm decides to expand (or contract) its size in a strictly proportional manner. Unlike standard marginal analysis, we are not changing just one isolated factor, but rather multiplying the entire infrastructure by the same scalar. The central question becomes whether the final output volume will grow identically, more, or less than this expansion.

Real-World Analogy Think of a pizza chain. When opening a second unit strictly identical to the first (same number of ovens, inputs, and pizzaiolos), the logic of pure replication suggests that production volume will double (**constant returns**). However, if this expanded scale enables the purchase of flour at much lower wholesale prices or the implementation of a centralized automated delivery system that would only be economically viable for large volumes, the output generated will grow proportionally more (**increasing returns**).

2.4.2 Mechanisms and Implications for Market Structure

Economic agents react directly to the curvature of technology. In scenarios with **Increasing Returns to Scale (IRS)**, large productive structures operate at efficiency levels superior to small units, acting as a natural inducer of industrial concentration and formation of natural monopolies.

On the other hand, **Decreasing Returns to Scale (DRS)** manifest when the coordination and control complexity of a giant corporation begins to generate informational bottlenecks and bureaucratic inefficiencies, establishing an organizational limit to horizontal firm growth. The theoretical logic behind these properties rests on the premises of divisibility, labor specialization, and fixed cost dilution.

In practical economics, these regimes shape the conditions of competition in a sector: - Under **Constant Returns to Scale (CRS)**, the optimal firm size remains indeterminate, allowing the balanced coexistence of small and large firms. - Under strong **Increasing Returns (IRS)**, antitrust policies become imperative, given that the market organically converges toward consolidation.

Example: Aviation Industry The development of a new commercial aircraft involves astronomical initial fixed costs. Once the model is designed, the production of hundreds of units does not replicate the initial research and development cost. This massive IRS regime justifies why the global market is dominated by a structural duopoly, as the minimum efficient scale acts as an insurmountable barrier to smaller entrants.

2.4.3 Formal Definition

Let $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n_+$ be the input vector and $f(\mathbf{x})$ the production function mapping these factors to the scalar output $y \in \mathbb{R}_+$.

Definition: Global Scale Regimes We evaluate global returns to scale by introducing a scalar multiplier $t > 1$: - **Constant Returns (CRS):** $f(t\mathbf{x}) = tf(\mathbf{x}), \quad \forall t > 0$ (homogeneous function of degree 1). - **Increasing Returns (IRS):** $f(t\mathbf{x}) > tf(\mathbf{x}), \quad \forall t > 1$. - **Decreasing Returns (DRS):** $f(t\mathbf{x}) < tf(\mathbf{x}), \quad \forall t > 1$.

Locally, we define the **Scale Elasticity** $\mu(\mathbf{x})$ through the limit of the percentage variation in output resulting from a 1% infinitesimal variation in the scale of all inputs simultaneously:

$$\mu(\mathbf{x}) = \lim_{t \to 1} \frac{\partial f(t\mathbf{x})}{\partial t} \frac{t}{f(\mathbf{x})}$$

> **Incompatibility of IRS with Perfect Competition** If the production technology exhibits global increasing returns to scale (IRS), the profit maximization problem under perfect competition admits no finite solution. Since average cost continuously decreases with output volume, the firm would find incentives to expand its supply indefinitely toward infinity, collapsing the hypothesis of price-taking agents.

2.4.4 Proofs

Proof 1: Scale Elasticity as the Sum of Input Elasticities

> **Euler's Theorem for Homogeneous Functions** If a differentiable function $f(\mathbf{x})$ is homogeneous of degree k with respect to its arguments, then the sum of its partial derivatives weighted by their respective variables satisfies:
>
> $$\sum_{i=1}^{n} \frac{\partial f(\mathbf{x})}{\partial x_i} x_i = kf(\mathbf{x})$$

Proof:

1. Structure the output level conditioned on a generic scale t by $y(t) = f(tx_1, \dots, tx_n)$.

2. By definition, the local scale elasticity coefficient corresponds to:

$$\mu(\mathbf{x}) = \left.\frac{dy/y}{dt/t}\right|_{t=1} = \left.\frac{dy(t)}{dt}\frac{t}{y(t)}\right|_{t=1}$$

3. Applying the chain rule for the multivariate gradient, differentiate $y(t)$ with respect to t:

$$\frac{dy(t)}{dt} = \sum_{i=1}^{n} \frac{\partial f(t\mathbf{x})}{\partial(tx_i)}\frac{d(tx_i)}{dt}$$

4. Since the derivative of each linear component $\frac{d(tx_i)}{dt} = x_i$, the expression simplifies to:

$$\frac{dy(t)}{dt} = \sum_{i=1}^{n} \frac{\partial f(t\mathbf{x})}{\partial(tx_i)} x_i$$

5. Evaluating the marginal variation at the initial scale point ($t = 1$), we obtain:

$$\frac{dy(1)}{dt} = \sum_{i=1}^{n} \frac{\partial f(\mathbf{x})}{\partial x_i} x_i$$

6. Substituting this derivative into the elasticity equation structured in step 2, noting that $y(1) = f(\mathbf{x})$:

$$\mu(\mathbf{x}) = \left(\sum_{i=1}^{n} \frac{\partial f(\mathbf{x})}{\partial x_i} x_i\right) \frac{1}{f(\mathbf{x})}$$

7. Using the distributive property of division over summation:

$$\mu(\mathbf{x}) = \sum_{i=1}^{n} \frac{\frac{\partial f(\mathbf{x})}{\partial x_i} x_i}{f(\mathbf{x})} = \sum_{i=1}^{n} \epsilon_i$$

   where ϵ_i represents the partial input elasticity of production factor i. ■

Proof 2: Fundamental Relationship between Costs (AC/MC) and Scale Elasticity

> **Envelope Theorem Applied to Costs** At the point that minimizes the firm's expenditure, the derivative of the total cost value function with respect to the output level y is equal to the Lagrange multiplier associated with the technical constraint ($\frac{\partial C}{\partial y} = \lambda = MC$).

Proof:

1. The cost minimization problem expresses the value function as $C(\mathbf{w}, y) = \min \sum w_i x_i$ subject to $f(\mathbf{x}) = y$.
2. Structure the respective Lagrangian operator: $\mathcal{L} = \sum w_i x_i - \lambda(f(\mathbf{x}) - y)$.
3. The first-order conditions for the input vector require that:

$$w_i = \lambda \frac{\partial f(\mathbf{x}^*)}{\partial x_i}, \quad \forall i$$

4. Multiplying both sides of each equality by the respective optimal input x_i^* and applying summation:

$$\sum_{i=1}^{n} w_i x_i^* = \lambda \sum_{i=1}^{n} \frac{\partial f(\mathbf{x}^*)}{\partial x_i} x_i^*$$

5. By the definition of the cost function, the left-hand side represents the optimal total cost $C(\mathbf{w}, y)$:

$$C(\mathbf{w}, y) = \lambda \sum_{i=1}^{n} f_i(\mathbf{x}^*) x_i^*$$

6. Incorporating the Envelope Theorem identity ($\lambda = MC$):

$$C(\mathbf{w}, y) = MC \cdot \sum_{i=1}^{n} f_i(\mathbf{x}^*) x_i^*$$

7. Dividing both sides of the equation by the product between output y and Marginal Cost (MC):

$$\frac{C(\mathbf{w}, y)}{y \cdot MC} = \frac{\sum_{i=1}^{n} f_i(\mathbf{x}^*) x_i^*}{y}$$

8. Identifying that $\frac{C}{y} = AC$ (Average Cost) and the term on the right maps the scale elasticity $\mu(\mathbf{x})$ calculated in Proof 1, we conclude:

$$\frac{AC}{MC} = \mu(\mathbf{x})$$

It follows directly that if $AC > MC \implies \mu(\mathbf{x}) > 1$ (**IRS**); if $AC = MC \implies \mu(\mathbf{x}) = 1$ (**CRS**). ■

2.4.5 Summary of Key Results

Concept	Definition	Key Formula
Constant Returns (CRS)	Output scales proportionally	$f(t\mathbf{x}) = tf(\mathbf{x})$
Increasing Returns (IRS)	Output scales more than proportionally	$f(t\mathbf{x}) > tf(\mathbf{x})$
Decreasing Returns (DRS)	Output scales less than proportionally	$f(t\mathbf{x}) < tf(\mathbf{x})$
Scale Elasticity	Local returns to scale	$\mu(\mathbf{x}) = \sum_{i=1}^{n} \epsilon_i$
AC-MC Relationship	Returns from cost structure	$\frac{AC}{MC} = \mu(\mathbf{x})$
Euler's Theorem	Homogeneity identity	$\sum f_i x_i = kf(\mathbf{x})$

2.5 The Cost Minimization Problem (CMP)

2.5.1 Motivation

The Cost Minimization Problem (CMP) isolates the firm's internal allocative efficiency decision: assuming a fixed production target, which input choice will demand the least financial expenditure? While the profit decision dictates the ideal market volume, the CMP focuses strictly on the choice of production technique, seeking the optimal factor vector within an isoquant.

Buffet Analogy Imagine a buffet service tasked with delivering exactly 500 meals (output target y). The manager can hire additional cooks (labor) or rent high-performance combination ovens (capital). The cost minimization point will be reached when the exact balance between labor and capital that generates the lowest total monetary outlay is established, given current wages and rental rates.

2.5.2 Substitutability and Analytical Duality

The decomposition of firm behavior through the CMP allows the engineering of production to be split into two independent stages: *how to produce* and, subsequently, *how much to produce.* When the market price of a given factor rises, the firm reacts by substituting it for inputs whose relative prices remained lower, moving along the isoquant.

The optimality condition requires that the technical substitution rate permitted by nature ($MRTS$) be numerically identical to the financial exchange rate imposed by the market (factor price ratio). If these magnitudes diverge, there will be room for input rearrangements that reduce total cost without compromising the final output volume. The minimum cost function derived from this problem constitutes the foundation for determining the marginal cost curve, a guiding variable for supply decisions under competition and for the exercise of market power under monopolistic structures.

2.5.3 Formal Definition

Consider n inputs with quantity vector $\mathbf{x} = (x_1, \ldots, x_n) \in \mathbb{R}^n_+$ and a factor price vector denoted by $\mathbf{w} = (w_1, \ldots, w_n) \in \mathbb{R}^n_{++}$.

Definition: Minimum Cost Function Given a production function $f(\mathbf{x})$ and an exogenous output target y, the **Cost Minimization Problem** is mathematically expressed as:

$$c(\mathbf{w}, y) = \min_{\mathbf{x} \geq 0} \{\mathbf{w} \cdot \mathbf{x} : f(\mathbf{x}) \geq y\}$$

Regularity Hypotheses:

- **Strictly positive price vector ($\mathbf{w} \gg 0$):** Prevents infinite demands and guarantees the existence of a finite bounded minimum.
- **Continuity and strict growth of $f(\mathbf{x})$:** Ensures that the technological constraint operates actively on the frontier of the set, i.e., $f(\mathbf{x}) = y$.
- **Quasiconcavity of $f(\mathbf{x})$:** Ensures that the input requirement set $V(y) = \{\mathbf{x} : f(\mathbf{x}) \geq y\}$ is convex, making the first-order conditions necessary and sufficient to characterize a global minimum.

  Implications of Absence of Quasiconcavity If the production function $f(\mathbf{x})$ does not exhibit quasi-concave behavior, the CMP admits multiple local solutions. The corresponding cost function will lose differentiability properties at critical points, provoking discontinuities and abrupt jumps in input choices in response to marginal variations in market prices.

2.5.4 Proofs

Proof 1: Equivalence between MRTS and the Price Ratio at the Optimum **Proof:**

1. Construct the Lagrangian equation associated with the constrained minimization problem:

$$\mathcal{L}(\mathbf{x}, \lambda) = \sum_{i=1}^{n} w_i x_i - \lambda(f(\mathbf{x}) - y)$$

2. Compute the first-order conditions (FOC) with respect to any two generic inputs i and j:

$$\frac{\partial \mathcal{L}}{\partial x_i} = w_i - \lambda \frac{\partial f(\mathbf{x}^*)}{\partial x_i} = 0 \implies w_i = \lambda f_i(\mathbf{x}^*)$$

$$\frac{\partial \mathcal{L}}{\partial x_j} = w_j - \lambda \frac{\partial f(\mathbf{x}^*)}{\partial x_j} = 0 \implies w_j = \lambda f_j(\mathbf{x}^*)$$

3. Divide the FOC equations term by term to eliminate the multiplier λ:

$$\frac{w_i}{w_j} = \frac{\lambda f_i(\mathbf{x}^*)}{\lambda f_j(\mathbf{x}^*)}$$

4. Simplify the constant λ and recall the analytical definition of $MRTS_{ij}$ as the ratio between the marginal productivities of the factors:

$$\frac{w_i}{w_j} = \frac{f_i(\mathbf{x}^*)}{f_j(\mathbf{x}^*)} \implies MRTS_{ij} = \frac{w_i}{w_j}$$

■

Proof 2: Concavity Property of the Cost Function with Respect to Prices (w)

Analytical Definition of Concavity A value function $g(\mathbf{w})$ is said to be concave if, and only if, for any convex combination parameterized by a scalar $t \in [0, 1]$, the following inequality holds:

$$g(t\mathbf{w} + (1-t)\mathbf{w}') \geq tg(\mathbf{w}) + (1-t)g(\mathbf{w}')$$

Proof:

1. Let $\mathbf{w}$ and $\mathbf{w}'$ be two distinct factor price vectors, and define $\mathbf{w}_t = t\mathbf{w} + (1-t)\mathbf{w}'$ as the convex combination between them.
2. Let $\mathbf{x}_t$ be the optimal input vector that minimizes total cost under the combined price vector $\mathbf{w}_t$ to achieve the output target y.
3. Expand the minimum cost function associated with $\mathbf{w}_t$ through the dot product:

$$c(\mathbf{w}_t, y) = \mathbf{w}_t \cdot \mathbf{x}_t = [t\mathbf{w} + (1-t)\mathbf{w}'] \cdot \mathbf{x}_t$$

4. Applying the distributive linear property of vector products:

$$c(\mathbf{w}_t, y) = t(\mathbf{w} \cdot \mathbf{x}_t) + (1-t)(\mathbf{w}' \cdot \mathbf{x}_t)$$

5. Since $\mathbf{x}_t$ is technologically feasible to produce y, but is not necessarily the optimal choice when prices are isolated as $\mathbf{w}$ or $\mathbf{w}'$, by definition of minimum it follows that:

$$\mathbf{w} \cdot \mathbf{x}_t \geq c(\mathbf{w}, y) \quad \text{and} \quad \mathbf{w}' \cdot \mathbf{x}_t \geq c(\mathbf{w}', y)$$

6. Multiplying the inequalities by the positive scalars t and $(1-t)$ and summing the expressions:

$$t(\mathbf{w} \cdot \mathbf{x}_t) + (1-t)(\mathbf{w}' \cdot \mathbf{x}_t) \geq tc(\mathbf{w}, y) + (1-t)c(\mathbf{w}', y)$$

7. Substituting the left-hand term with the equality obtained in step 4, we validate the property:

$$c(t\mathbf{w} + (1-t)\mathbf{w}', y) \geq tc(\mathbf{w}, y) + (1-t)c(\mathbf{w}', y)$$

■

2.5.5 Summary of Key Results

Concept	Definition	Key Formula
Cost Minimization Problem	Minimize cost for given output	$c(\mathbf{w}, y) = \min_{\mathbf{x} \geq 0}\{\mathbf{w} \cdot \mathbf{x} : f(\mathbf{x}) \geq y\}$
Cost Function	Minimum cost to produce y	$c(\mathbf{w}, y)$
Optimality Condition	MRTS equals price ratio	$MRTS_{ij} = \frac{w_i}{w_j}$
Concavity of c in w	Cost increases sublinearly with prices	$c(t\mathbf{w} + (1-t)\mathbf{w}', y) \geq tc(\mathbf{w}, y) + (1-t)c(\mathbf{w}', y)$

2.6 Conditional Factor Demand and Shephard's Lemma

2.6.1 Motivation

Conditional factor demand establishes the mapping of the firm's optimal hiring channels, strictly isolating the effect of production scale. The nomenclature "conditional" derives from the problem's submission to an exogenous final product level (q), allowing analysis of pure technical choices without interference from market demand behavior for the product.

Confectionery Analogy Imagine a confectioner receiving a fixed order for a five-tier wedding cake (output target q). She has flexibility to balance the use of industrial mixers (capital) or kitchen assistants (labor). If a spike occurs in electricity prices, her conditional demand for mixers will be reduced, while the conditional demand for labor will expand, ensuring the delivery of the same cake at the lowest possible cost.

2.6.2 Structural Implications and the Input Demand Law

The mapping of conditional demands gives the analyst the power to anticipate reconfigurations in the factor matrix of an economic sector resulting from exogenous price shocks, independent of final consumption behavior.

One of the pillars of this modeling is the **Input Demand Law**, derived directly from Shephard's Lemma and the concavity properties of the value function. It imposes that the conditional demand curve for any production factor must be necessarily negatively (or zero) sloped with respect to its own price. Positive shocks in the hiring cost of an input induce, necessarily, its saving or substitution in the production plant, as long as the intended scale level remains constant.

Example: Industrial Automation Subsidies Public policies aimed at reducing the acquisition cost of technological capital goods reduce the relative price of this factor. The CMP demonstrates that, even if aggregate industrial activity remains stagnant (maintenance of the reference q), industries will adjust their conditional demands, reducing jobs and increasing robot density to optimize their expenditure allocations along isoquants.

2.6.3 Formal Definition

Let L be the number of production factors available, $\mathbf{z} = (z_1, \ldots, z_L) \in \mathbb{R}^L_+$ the input vector, and $\mathbf{w} = (w_1, \ldots, w_L) \in \mathbb{R}^L_{++}$ the corresponding hiring price vector.

Definition: Conditional Factor Demand **Conditional Factor Demand** $\mathbf{z}(\mathbf{w}, q)$ corresponds to the single-valued optimal argument that solves the firm's constrained minimization problem:

$$\mathbf{z}(\mathbf{w}, q) = \arg\min_{\mathbf{z}\geq 0}\{\mathbf{w} \cdot \mathbf{z} : f(\mathbf{z}) \geq q\}$$

Necessary Hypotheses:

- **Strict growth and continuity of $f(\mathbf{z})$:** Makes the technical constraint strictly active in equality ($f(\mathbf{z}) = q$).
- **Strict quasiconcavity of $f(\mathbf{z})$:** Guarantees strict convexity of level curves, ensuring that the optimal argument is a unique point (well-defined function) and not a set of alternatives.

2.6.4 Proofs

Proof 1: Shephard's Lemma **Proof:**

1. Structure the Lagrangian operator associated with the firm's cost problem:

$$\mathcal{L}(\mathbf{z}, \lambda) = \sum_{j=1}^{L} w_j z_j - \lambda(f(\mathbf{z}) - q)$$

2. At the optimum point, the minimum cost value function coincides with the value assumed by the Lagrangian:
$$c(\mathbf{w}, q) = \mathcal{L}(\mathbf{z}^*(\mathbf{w}, q), \lambda^*(\mathbf{w}, q), \mathbf{w}, q)$$
3. Applying the Envelope Theorem, the total derivative of the value function with respect to the input price w_i equals the partial derivative of the Lagrangian operator evaluated at the optimal vector $\mathbf{z}^*$:
$$\frac{\partial c(\mathbf{w}, q)}{\partial w_i} = \left.\frac{\partial \mathcal{L}}{\partial w_i}\right|_{(\mathbf{z}^*, \lambda^*)}$$
4. Expanding the partial derivative of the Lagrangian's linear structure through the differential operator:
$$\frac{\partial \mathcal{L}}{\partial w_i} = \frac{\partial(\sum w_j z_j)}{\partial w_i} - \frac{\partial[\lambda(f(\mathbf{z}) - q)]}{\partial w_i}$$
5. Since the variables λ, $f(\mathbf{z})$, and q have no explicit functional dependence on market prices w_i, the second term vanishes, leaving:
$$\frac{\partial \mathcal{L}}{\partial w_i} = z_i$$
6. Substituting the result into the identity obtained via the Envelope Theorem, we validate the Lemma:
$$\frac{\partial c(\mathbf{w}, q)}{\partial w_i} = z_i(\mathbf{w}, q)$$

■

Proof 2: Homogeneity of Degree Zero of Conditional Demand with Respect to Prices
Proof:

1. Take the definition of the CMP under a scalar price multiplier $\alpha > 0$:
$$c(\alpha\mathbf{w}, q) = \min_{\mathbf{z}\geq 0}(\alpha\mathbf{w}) \cdot \mathbf{z} \quad \text{s.t.} \quad f(\mathbf{z}) \geq q$$
2. Using the linearity property of the dot product, extract the multiplicative factor:
$$c(\alpha\mathbf{w}, q) = \min_{\mathbf{z}\geq 0} \alpha(\mathbf{w} \cdot \mathbf{z}) \quad \text{s.t.} \quad f(\mathbf{z}) \geq q$$
3. Since α is a fixed positive constant, it does not alter the minimizing argument $\mathbf{z}^*$. Hence, it can be extracted from the optimization operator:
$$c(\alpha\mathbf{w}, q) = \alpha \min_{\mathbf{z}\geq 0}\{\mathbf{w} \cdot \mathbf{z} : f(\mathbf{z}) \geq q\} = \alpha c(\mathbf{w}, q)$$
The minimum cost function is, therefore, homogeneous of degree 1 with respect to vector $\mathbf{w}$.
4. By Shephard's Lemma proved earlier, we know that $z_i(\mathbf{w}, q) = \frac{\partial c(\mathbf{w},q)}{\partial w_i}$.
5. Applying the classical property of differentiation of homogeneous functions (the derivative of a homogeneous function of degree k constitutes a new homogeneous function of degree $k - 1$):
$$z_i(\alpha\mathbf{w}, q) \quad \text{is homogeneous of degree} \quad 1 - 1 = 0$$
6. By the algebraic definition of homogeneity of degree zero, we conclude that:
$$z_i(\alpha\mathbf{w}, q) = z_i(\mathbf{w}, q)$$

■

Proof 3: Input Demand Law (Negative Slope of Conditional Demand) **Proof:**

1. Structure the input substitution matrix through the Jacobian of conditional demand, denoted by $D_{\mathbf{w}}\mathbf{z}(\mathbf{w}, q)$.
2. Substituting Shephard's Lemma identity ($\mathbf{z}(\mathbf{w}, q) = \nabla_{\mathbf{w}} c(\mathbf{w}, q)$) into the Jacobian matrix:

$$D_{\mathbf{w}}\mathbf{z}(\mathbf{w}, q) = D_{\mathbf{w}}[\nabla_{\mathbf{w}} c(\mathbf{w}, q)] = \nabla^2_{\mathbf{w}} c(\mathbf{w}, q)$$

   The Jacobian of conditional factor demand corresponds rigorously to the Hessian matrix of the minimum cost function with respect to $\mathbf{w}$.
3. In Proof 2 of the previous section, it was analytically demonstrated that the cost function $c(\mathbf{w}, q)$ is concave with respect to the price vector $\mathbf{w}$.
4. Linear algebra imposes that the Hessian matrix of any concave C^2 function be **Negative Semidefinite (NSD)**.
5. By definition of NSD matrix properties, all elements arranged along its main diagonal must be non-positive:

$$\frac{\partial z_i(\mathbf{w}, q)}{\partial w_i} \leq 0, \quad \forall i$$

   It is demonstrated that conditional demand for an input has a strictly non-positive response to variations in its own price. ■

2.6.5 Summary of Key Results

Concept	Definition	Key Formula
Conditional Factor Demand	Cost-minimizing input choice	$\mathbf{z}(\mathbf{w}, q) = \arg\min_{\mathbf{z}}\{\mathbf{w} \cdot \mathbf{z} : f(\mathbf{z}) \geq q\}$
Shephard's Lemma	Recover demand from cost function	$\frac{\partial c(\mathbf{w},q)}{\partial w_i} = z_i(\mathbf{w}, q)$
Homogeneity of z	Invariant to proportional price changes	$\mathbf{z}(\alpha\mathbf{w}, q) = \mathbf{z}(\mathbf{w}, q)$
Input Demand Law	Own-price effect is non-positive	$\frac{\partial z_i}{\partial w_i} \leq 0$
Substitution Matrix	Hessian of cost function	$D_{\mathbf{w}}\mathbf{z} = \nabla^2_{\mathbf{w}} c$ (NSD)

2.7 The Cost Function and the Cost Minimization Problem

2.7.1 Motivation

The cost function synthesizes the firm's economic efficiency, elegantly isolating technological constraints from goods market decisions. The fundamental problem it proposes to solve consists of internal productive efficiency: given an exogenous physical output target (q), what is the input arrangement that demands the least total financial outlay? While the profit maximization decision establishes *how much* to sell, the cost function determines *how* to produce.

Real-World Analogy Consider a shoe factory. To fulfill a fixed order of 1,000 pairs of shoes (production target), management faces a technical choice: adopt a manual assembly line with high artisan density or implement a robotized system based on high-performance machinery. The cost function acts as the operational optimization manual that points to which of these techniques will generate the least financial impact, indexing the decision strictly to current wages and prevailing energy tariffs.

2.7.2 Mechanisms and Economic Properties

Firms operate as optimizing agents linked to technological frontiers. Upon observing factor prices in the market, companies calibrate their production plants: if the cost of a specific input suffers a positive shock (such as a readjustment in electricity tariffs), economic rationality imposes the substitution of this input for relatively cheaper alternative vectors, generating a continuous shift along the production isoquant.

The central property governing this behavior is **concavity in input prices**. Under the occurrence of a price increase in a production factor, the firm's total cost expands at a strictly non-explosive (sublinear or, at the limit, linear) rate. This occurs because the firm readjusts conditional demands (the input mix), attenuating the cost increase through internal technical substitution.

2.7.3 Formal Definition

Let L be the number of available inputs. Define the input quantity vector by $\mathbf{z} = (z_1, \ldots, z_L) \in \mathbb{R}^L_+$ and the respective price vector by $\mathbf{w} = (w_1, \ldots, w_L) \in \mathbb{R}^L_{++}$.

Definition: The Cost Function as a Value Function Given a production function $f(\mathbf{z})$ and a final output target $q \in \mathbb{R}_+$, the **Cost Function** $c(\mathbf{w}, q)$ corresponds to the value function derived from the Cost Minimization Problem (CMP):

$$c(\mathbf{w}, q) = \min_{\mathbf{z} \geq 0} \{\mathbf{w} \cdot \mathbf{z} : f(\mathbf{z}) \geq q\}$$

Structural Hypotheses:

- **Strictly positive prices ($\mathbf{w} \gg 0$):** Ensures that cost level sets are compact (bounded), guaranteeing the mathematical existence of a global minimum point.
- **Continuity and strict monotonicity of $f(\mathbf{z})$:** Ensures that the production constraint operates actively in equality, i.e., $f(\mathbf{z}) = q$ at the optimum.
- **Quasiconcavity of $f(\mathbf{z})$:** Imposes that the input requirement set $V(q) = \{\mathbf{z} \in \mathbb{R}^L_+ : f(\mathbf{z}) \geq q\}$ constitutes a convex set, making the first-order conditions (FOC) sufficient for determining the minimum.

  Implications of Non-Quasiconcavity If the technology expressed by $f(\mathbf{z})$ is not quasiconcave, the CMP will admit multiple local solutions. As a direct result, the cost function will lose differentiability at transition points, provoking discontinuities and abrupt jumps in the input demand vector in response to infinitesimal variations in market prices.

2.7.4 Proofs

Proof 1: Homogeneity of Degree 1 with Respect to the Price Vector (w) **Proof:**

1. Evaluate the cost function under a price vector inflated by a generic scalar $\alpha > 0$:

$$c(\alpha\mathbf{w}, q) = \min_{\mathbf{z} \geq 0} \{(\alpha\mathbf{w}) \cdot \mathbf{z} : f(\mathbf{z}) \geq q\}$$

2. Applying the linearity property of vector dot products, isolate the parameter α:

$$c(\alpha\mathbf{w}, q) = \min_{\mathbf{z} \geq 0} \{\alpha(\mathbf{w} \cdot \mathbf{z}) : f(\mathbf{z}) \geq q\}$$

3. Since α constitutes a strictly positive constant, its presence does not alter the minimizing argument $\mathbf{z}^*$. By the algebra of infimum and minimum operators, the scalar can be extracted from the operator:

$$c(\alpha\mathbf{w}, q) = \alpha \min_{\mathbf{z} \geq 0} \{\mathbf{w} \cdot \mathbf{z} : f(\mathbf{z}) \geq q\}$$

4. Substituting the definition of the CMP contained in the core of the operator, we obtain the homogeneity of degree 1 identity:
$$c(\alpha \mathbf{w}, q) = \alpha c(\mathbf{w}, q)$$

■

Proof 2: Concavity of the Cost Function with Respect to Prices (w) Proof:

1. Let $\mathbf{w}$ and $\mathbf{w}'$ be two arbitrary input price vectors, and define $\mathbf{w}_\lambda = \lambda \mathbf{w} + (1-\lambda)\mathbf{w}'$ as their convex combination, where $\lambda \in [0, 1]$.
2. Let $\mathbf{z}_\lambda$ be the optimal input vector that minimizes total cost under the mixed price regime $\mathbf{w}_\lambda$ to meet the output target q.
3. Expand the minimum cost function evaluated at $\mathbf{w}_\lambda$ through the dot product:
$$c(\mathbf{w}_\lambda, q) = \mathbf{w}_\lambda \cdot \mathbf{z}_\lambda = [\lambda \mathbf{w} + (1-\lambda)\mathbf{w}'] \cdot \mathbf{z}_\lambda$$
4. Distributing the input vector $\mathbf{z}_\lambda$ over the components of the linear sum:
$$c(\mathbf{w}_\lambda, q) = \lambda(\mathbf{w} \cdot \mathbf{z}_\lambda) + (1-\lambda)(\mathbf{w}' \cdot \mathbf{z}_\lambda)$$
5. Since $\mathbf{z}_\lambda$ is a technologically feasible input plan to achieve output q, but is not necessarily the optimal choice under the isolated price vectors $\mathbf{w}$ or $\mathbf{w}'$, it follows from the minimum conditions that:
$$\mathbf{w} \cdot \mathbf{z}_\lambda \geq c(\mathbf{w}, q) \quad \text{and} \quad \mathbf{w}' \cdot \mathbf{z}_\lambda \geq c(\mathbf{w}', q)$$
6. Weighting the respective inequalities by the non-negative scalars λ and $(1-\lambda)$, and performing their sum:
$$\lambda(\mathbf{w} \cdot \mathbf{z}_\lambda) + (1-\lambda)(\mathbf{w}' \cdot \mathbf{z}_\lambda) \geq \lambda c(\mathbf{w}, q) + (1-\lambda)c(\mathbf{w}', q)$$
7. Substituting the left-hand member by the equality deduced in step 4, we confirm concavity:
$$c(\lambda \mathbf{w} + (1-\lambda)\mathbf{w}', q) \geq \lambda c(\mathbf{w}, q) + (1-\lambda)c(\mathbf{w}', q)$$

■

Proof 3: Shephard's Lemma via the Envelope Theorem Proof:

1. Formulate the Lagrangian operator attached to the firm's CMP:
$$\mathcal{L}(\mathbf{z}, \lambda) = \sum_{j=1}^{L} w_j z_j - \lambda(f(\mathbf{z}) - q)$$
2. At the optimal choice point $\mathbf{z}^*(\mathbf{w}, q)$, the productive constraint operates in active regime $(f(\mathbf{z}^*) = q)$, equating the value function to the value assumed by the Lagrangian:
$$c(\mathbf{w}, q) = \mathcal{L}(\mathbf{z}^*(\mathbf{w}, q), \lambda^*(\mathbf{w}, q), \mathbf{w}, q)$$
3. Applying the Envelope Theorem, the total derivative of the value function with respect to the price parameter w_i corresponds to the partial derivative of the Lagrangian function:
$$\frac{\partial c(\mathbf{w}, q)}{\partial w_i} = \left. \frac{\partial \mathcal{L}}{\partial w_i} \right|_{(\mathbf{z}^*, \lambda^*)}$$

4. Expanding the differential operator over the additive terms of the Lagrangian:

$$\frac{\partial \mathcal{L}}{\partial w_i} = \frac{\partial}{\partial w_i}\left[\sum_{j=1}^{L} w_j z_j\right] - \frac{\partial}{\partial w_i}[\lambda(f(\mathbf{z}) - q)]$$

5. Since the variables z_j, λ, $f(\mathbf{z})$, and q have no explicit functional dependence on the price w_i within the scope of the Lagrangian, the constraint term vanishes, leaving only the linear coefficient:

$$\frac{\partial \mathcal{L}}{\partial w_i} = z_i - 0$$

6. Uniting the results, we establish that the partial derivative of minimum cost maps the conditional demand for the production factor:

$$\frac{\partial c(\mathbf{w}, q)}{\partial w_i} = z_i(\mathbf{w}, q)$$

■

2.7.5 Summary of Key Results

Concept	Definition	Key Formula
Cost Function	Minimum cost to produce q	$c(\mathbf{w}, q) = \min_{\mathbf{z}\geq 0}\{\mathbf{w}\cdot\mathbf{z} : f(\mathbf{z}) \geq q\}$
Homogeneity of c	Linear in prices	$c(\alpha\mathbf{w}, q) = \alpha c(\mathbf{w}, q)$
Concavity of c	Sublinear response to price increases	$c(\lambda\mathbf{w} + (1-\lambda)\mathbf{w}', q) \geq \lambda c(\mathbf{w}, q) + (1-\lambda)c(\mathbf{w}', q)$
Shephard's Lemma	Derivative yields demand	$\frac{\partial c(\mathbf{w},q)}{\partial w_i} = z_i(\mathbf{w}, q)$

2.8 Cost Structures under Different Time Horizons

2.8.1 Motivation

The analysis of costs under different time horizons investigates the flexibility of production factor adaptation in response to variations in activity levels. The concepts of short and long run are not tied to a rigid chronological calendar, but rather to the legal, contractual, or physical nature of input immobilization. In the short run, the firm operates under rigidity constraints, possessing fixed factors that cannot be changed. In the long run, all input vectors become fully malleable.

Logistics Fleet Analogy Imagine a transportation company. In the short run, the volume of available trucks is fixed due to existing leasing contracts. If a sudden expansion in freight demand occurs, the only technical way to increase production is through the intensification of the variable factor, via payment of overtime to drivers. In the long run, the firm acquires full freedom to resize the fleet or invest in more efficient modes, re-optimizing the cost architecture.

2.8.2 The Dynamics of Short-Run and Long-Run Costs

Since the long run encompasses all operational combinations executable in the short run, in addition to freeing the variation of previously constrained factors, the long-run total cost acts as the **absolute lower bound** of short-run structures. The long run fundamentally constitutes the strategic choice of the most efficient short-run arrangement for each targeted output volume.

This distinction underpins the formulation of investment decisions and measures industrial inertia. Regulatory vectors (such as carbon emission taxation) may exhibit low responsiveness in pollutant reduction in the short run, a period in which the machinery park is fixed. However, they trigger profound technological reconfigurations in the long run, a horizon in which fixed assets undergo depreciation and obsolescence, being replaced by cleaner matrices.

2.8.3 Formal Definition

We divide the factor space into two sub-vectors: $\mathbf{x}_v \in \mathbb{R}^k$ (variable inputs) and $\mathbf{x}_f \in \mathbb{R}^{L-k}$ (fixed inputs), associated with the price vectors $\mathbf{w}_v \gg 0$ and $\mathbf{w}_f \gg 0$.

Definition: Short-Run and Long-Run Cost Functions

- **Long-Run Cost Function (c):** Full factor flexibility.

$$c(\mathbf{w}, y) = \min_{\mathbf{x} \geq 0} \{\mathbf{w} \cdot \mathbf{x} : f(\mathbf{x}) \geq y\}$$

- **Short-Run Cost Function (sc):** Restricted to fixed vector $\bar{\mathbf{x}}_f$.

$$sc(\mathbf{w}, y; \bar{\mathbf{x}}_f) = \min_{\mathbf{x}_v \geq 0} \{\mathbf{w}_v \cdot \mathbf{x}_v + \mathbf{w}_f \cdot \bar{\mathbf{x}}_f : f(\mathbf{x}_v, \bar{\mathbf{x}}_f) \geq y\}$$

2.8.4 Proofs

Proof 1: Long-Run Cost as the Envelope Frontier of Short-Run Cost **Proof:**

1. Let $\mathbf{x}^*(\mathbf{w}, y) = (\mathbf{x}_v^*(\mathbf{w}, y), \mathbf{x}_f^*(\mathbf{w}, y))$ be the choice vector that solves the long-run CMP for an output y.
2. By definition of the value function, the long-run minimum cost is expressed as:

$$c(\mathbf{w}, y) = \mathbf{w} \cdot \mathbf{x}^*(\mathbf{w}, y) = \mathbf{w}_v \cdot \mathbf{x}_v^*(\mathbf{w}, y) + \mathbf{w}_f \cdot \mathbf{x}_f^*(\mathbf{w}, y)$$

3. Take the constraint of the short-run problem for any fixed vector $\bar{\mathbf{x}}_f$:

$$\Omega_{SR} = \{\mathbf{x}_v \in \mathbb{R}^k : f(\mathbf{x}_v, \bar{\mathbf{x}}_f) \geq y\}$$

4. The set of feasible short-run choices Ω_{SR} constitutes a subset contained in the unrestricted long-run possibility space $\Omega_{LR} = \{\mathbf{x} \in \mathbb{R}^L : f(\mathbf{x}) \geq y\}$.
5. The imposition of additional constraints reduces, or at most maintains, the search space. By the variational principle, the minimum computed over a larger set is less than or equal to the minimum restricted to a subset:

$$c(\mathbf{w}, y) \leq sc(\mathbf{w}, y; \bar{\mathbf{x}}_f), \quad \forall \bar{\mathbf{x}}_f, y$$

6. If we evaluate the specific scenario where the short-run fixed vector coincides with the long-run optimal choice, i.e., $\bar{\mathbf{x}}_f = \mathbf{x}_f^*(\mathbf{w}, y)$, the long-run optimal plan becomes feasible in the short run.
7. Being feasible and minimizing, the short-run cost collapses to equality with the long-run cost:

$$c(\mathbf{w}, y) = sc(\mathbf{w}, y; \mathbf{x}_f^*(\mathbf{w}, y))$$

■

Proof 2: Tangency Condition and Equivalence of Marginal Costs

Envelope Theorem Identity Given a parameterized minimization structure $V(a) = \min_x g(x, a)$ subject to $h(x, a) = 0$, the total derivative of the value function with respect to the parameter satisfies $\frac{dV(a)}{da} = \frac{\partial \mathcal{L}}{\partial a}\Big|_{x^*}$.

Proof:

1. Retrieve the fundamental identity established in Proof 1:

$$c(\mathbf{w}, y) = sc(\mathbf{w}, y; \mathbf{x}_f^*(\mathbf{w}, y))$$

2. Differentiate both sides of the equation with respect to the output target y, applying the chain rule for total differentiation on the right-hand member:

$$\frac{\partial c(\mathbf{w}, y)}{\partial y} = \frac{\partial sc(\mathbf{w}, y; \mathbf{x}_f^*)}{\partial y} + \sum_{i \in \text{fixed}} \frac{\partial sc(\mathbf{w}, y; \mathbf{x}_f^*)}{\partial x_{f,i}} \cdot \frac{\partial x_{f,i}^*(\mathbf{w}, y)}{\partial y}$$

3. At the point that tunes the long-run equilibrium, the firm made the choice of $\mathbf{x}_f$ focusing on the global analytical optimization of the CMP.
4. Since the vector $\mathbf{x}_f^*(\mathbf{w}, y)$ minimizes total cost, the first-order conditions for the immobilization variables in the long run require the cancellation of their partial derivatives in the restricted cost function:

$$\frac{\partial sc(\mathbf{w}, y; \mathbf{x}_f^*)}{\partial x_{f,i}} = 0, \quad \forall i$$

5. Substituting the optimality condition into the differentiation equation obtained in step 2, the summation term vanishes:

$$\frac{\partial c(\mathbf{w}, y)}{\partial y} = \frac{\partial sc(\mathbf{w}, y; \mathbf{x}_f^*)}{\partial y} + 0$$

6. Since the partial derivatives of cost with respect to output define the Long-Run Marginal Cost (LMC) and Short-Run Marginal Cost (SMC) functions, we confirm tangency:

$$LMC(y) = SMC(y; \mathbf{x}_f^*)$$

■

2.8.5 Summary of Key Results

Concept	Definition	Key Formula
Long-Run Cost	All factors variable	$c(\mathbf{w}, y) = \min_{\mathbf{x}}\{\mathbf{w} \cdot \mathbf{x} : f(\mathbf{x}) \geq y\}$
Short-Run Cost	Fixed factors constrained	$sc(\mathbf{w}, y; \bar{\mathbf{x}}_f) = \min_{\mathbf{x}_v}\{\mathbf{w}_v \cdot \mathbf{x}_v + \mathbf{w}_f \cdot \bar{\mathbf{x}}_f : f(\mathbf{x}_v, \bar{\mathbf{x}}_f) \geq y\}$
Envelope Property	LR cost is lower bound of SR cost	$c(\mathbf{w}, y) \leq sc(\mathbf{w}, y; \bar{\mathbf{x}}_f)$
Tangency Condition	LR and SR marginal costs equal at optimum	$LMC(y) = SMC(y; \mathbf{x}_f^*)$
Long-Run Choice	Optimal fixed factor level	$\mathbf{x}_f^*(\mathbf{w}, y) = \arg\min_{\mathbf{x}_f} sc(\mathbf{w}, y; \mathbf{x}_f)$

Chapter 3: Market Structures

3.1 Monopoly

3.1.1 Motivation

The concept of monopoly addresses how a firm behaves when it is the **sole supplier** of a good with no close substitutes, protected by barriers to entry. In perfect competition, the firm is a "price taker"; in monopoly, it is the "price maker."

Practical Analogy: Imagine the only ferry crossing a wide river between two isolated cities. If you want to cross, you must pay whatever price the ferry owner sets. He knows that if he raises the price, some people may give up crossing (the quantity effect), but he will earn much more per person who still decides to go (the price effect). The monopolist's problem is finding the "sweet spot" in this balance to maximize profit.

3.1.2 Mechanisms

- **Consumers:** Behave according to their willingness to pay, reflected in the demand curve. They suffer a welfare loss because the monopoly price is always above the marginal cost of production.
- **The Firm:** Unlike the competitive firm, the monopolist perceives that its production decision affects the market price. To sell an extra unit, it must lower the price not only for that unit but for all previous units it intended to sell.
- **Economic Logic:** The monopolist stops expanding production when the additional gain from selling one more unit (**Marginal Revenue**) exactly equals the cost of producing it (**Marginal Cost**). Producing beyond that point would reduce total profit.

3.1.3 Implications

The main practical change is **allocative inefficiency**. The monopoly produces less and charges more than would be socially optimal, creating the so-called "deadweight loss" (welfare loss that goes to no one).

Concrete Example: The regulation of **natural monopolies**, such as water or electricity distribution companies. Since infrastructure costs are extremely high, it makes no sense to have two companies digging holes to lay pipes on the same street. The government intervenes by setting tariffs or subsidizing production so that the price approaches marginal cost, trying to mimic the efficiency of competition and reduce social loss.

3.1.4 Formal Definition

Let there be a market with a single producer facing an inverse demand function $p(q)$ and having a cost structure $c(q)$.

- **Variables and Parameters:**
  - $q \in [0, \bar{q}]$: Quantity produced (non-negative reals).
  - $p(q) : [0, \bar{q}] \to \mathbb{R}_+$: Inverse demand function (price the market pays for q).
  - $c(q) : [0, \bar{q}] \to \mathbb{R}_+$: Total cost function.
  - $\pi(q)$: Monopolist's profit function.
- **Hypotheses:**
  1. **Differentiability:** $p(q)$ and $c(q)$ are continuous and differentiable.
  2. **Decreasing Demand:** $p'(q) < 0$ (the demand curve is downward sloping).
  3. **Non-decreasing Costs:** $c'(q) \geq 0$ (marginal cost is positive).
  4. **Profit Concavity (SOC):** $\pi''(q) < 0$ to guarantee a global maximum.
- **What "breaks" without the hypotheses:** If $p'(q) > 0$, the good would be a "Giffen good," invalidating the standard market power analysis. If $\pi''(q) > 0$, the First-Order Condition would find a **minimum** profit point, not a maximum.

3.1.5 Proofs

Property 1: The First-Order Condition (FOC) The monopolist's objective is to maximize $\pi(q) = p(q) \cdot q - c(q)$.

1. $\pi(q) = p(q) \cdot q - c(q)$ (by definition of profit as Total Revenue minus Total Cost).
2. $\frac{d\pi}{dq} = \frac{d}{dq}[p(q) \cdot q] - \frac{dc(q)}{dq} = 0$ (by the necessary condition for an extremum of differentiable functions).
3. $[p(q) \cdot 1 + q \cdot p'(q)] - c'(q) = 0$ (by applying the **Product Rule** to the first part of the derivative).
4. $p(q) + q \cdot p'(q) = c'(q)$ (by transposition of terms).
5. $MR(q) = MC(q)$ (by definition of Marginal Revenue and Marginal Cost). ■

Property 2: Relationship with Price Elasticity of Demand (ϵ) Elasticity is defined as $\epsilon = \frac{dq}{dp} \cdot \frac{p}{q}$.

1. $MR = p(q) + q \cdot \frac{dp}{dq}$ (from the FOC proved above).
2. $MR = p(q)\left[1 + \frac{q}{p(q)} \cdot \frac{dp}{dq}\right]$ (factoring out $p(q)$).
3. $MR = p(q)\left[1 + \frac{1}{\frac{p}{q} \cdot \frac{dq}{dp}}\right]$ (by the inverse derivative property: $\frac{dp}{dq} = \frac{1}{dq/dp}$).
4. $MR = p(q)\left[1 + \frac{1}{\epsilon}\right]$ (by substituting the formal definition of price elasticity). ■

Property 3: The Markup (Lerner Index) The markup is defined as the margin of price over marginal cost relative to price.

1. $p\left(1 + \frac{1}{\epsilon}\right) = MC$ (by the FOC and the elasticity relationship proved above).
2. $p + \frac{p}{\epsilon} = MC$ (by applying the distributive property).
3. $p - MC = -\frac{p}{\epsilon}$ (by algebraic transposition of terms).
4. $\frac{p-MC}{p} = \frac{1}{|\epsilon|}$ (dividing both sides by p and using the absolute value of elasticity, since $\epsilon < 0$). ■

3.1.6 Formal Connection: Intuition → Proof

Mechanism from Intuition	Corresponding Variable/Condition in Proof	Step/Equation
Choosing the "Sweet Spot" of profit	FOC: $MR(q^*) = MC(q^*)$	Property 1, Step 5
Balance between Price and Quantity Effects	Term $p(q) + q \cdot p'(q)$	Property 1, Step 4
Dependence on consumer reaction	Elasticity ϵ in the MR formula	Property 2, Step 4
Power to charge above cost (Markup)	Lerner Index: $\frac{p-MC}{p} = \frac{1}{\lvert\epsilon\rvert}$	Property 3, Step 4
Inefficiency (Price > MC)	$p^* = \frac{MC}{1+1/\epsilon} > MC$ since $\epsilon < -1$ in equilibrium	Property 3 derivation

---

3.2 First-Degree Price Discrimination

First-degree price discrimination, also known as perfect discrimination, represents the theoretical limit of market power, where a firm manages to extract the entirety of consumer surplus. Under this structure, the monopolist operates to transform every bit of willingness to pay into pure profit, eliminating allocative inefficiencies but drastically altering welfare distribution.

3.2.1 Motivation

In a uniform-price monopoly, the firm faces a dilemma: to sell more units, it must reduce the price of all units sold, generating "deadweight loss" — mutually beneficial transactions that do not occur. First-degree price discrimination attempts to solve this "waste" of potential profit by treating each consumer and each unit individually.

Practical Analogy: Imagine an omniscient antique auctioneer who knows exactly the maximum value each person in the room would pay for an item. He does not announce a fixed price; he approaches each person and charges exactly their willingness to pay. If you would pay \$100, the price for you is \$100. If your neighbor would pay \$10, the price for him is \$10.

3.2.2 Mechanisms

- **Firm Behavior:** The monopolist acts as a marginal value capturer. It continues producing and selling as long as the next customer's willingness to pay is greater than or equal to the cost of producing that extra unit ($P \geq MC$).
- **Consumer Behavior:** The consumer is reduced to his "reservation price." Since the price charged is exactly equal to his perceived benefit, the consumer is indifferent between buying or not, resulting in zero net surplus.
- **Economic Logic:** The central property is that **marginal revenue becomes equal to price** (the demand curve itself), because selling an extra unit at a lower price does not require reducing the price of previous units.

3.2.3 Implications

In practice, the concept illuminates that monopoly inefficiency does not come from profit itself, but from the quantity restriction. Ironically, perfect monopoly is **allocatively efficient** (i.e., produces the same quantity as perfect competition), but it is socially controversial because total social welfare is composed entirely of producer profit.

Concrete Example: Custom enterprise software negotiations or high-level consulting. The supplier analyzes the client's balance sheet and expected productivity gain to set a fee that captures almost all the value generated by the solution, varying drastically from client to client.

3.2.4 Formal Definition

Let there be a market with a single producer facing a continuous inverse demand function $P(q)$ and a total cost function $C(q)$.

- **Variables and Parameters:**
  - $q \in [0, \infty)$: Quantity produced.
  - $P(q) : \mathbb{R}_+ \to \mathbb{R}_+$: Maximum price the market pays for the q-th unit.
  - $C(q) : \mathbb{R}_+ \to \mathbb{R}_+$: Total production cost.
  - $W(q) = \int_0^q P(x)dx$: Total willingness to pay (gross revenue under perfect discrimination).
- **Hypotheses:**
  1. **Perfect Information:** The monopolist knows $P(q)$ for each unit and individual.
  2. **No Arbitrage:** Consumers cannot resell the good among themselves.
  3. **Differentiability:** $P(q)$ and $C(q)$ are differentiable.
- **What "breaks" formally:** Without Hypothesis 1, we fall into second or third-degree discrimination (asymmetric information). Without Hypothesis 2, the price converges to a uniform price due to competition among resellers.

3.2.5 Proofs

Auxiliary Theorem: Fundamental Theorem of Calculus (Part 1) *If f is continuous on $[a, b]$, then the function $F(x) = \int_a^x f(t)dt$ is differentiable on (a, b) and $F'(x) = f(x)$.*

Property 1: Perfect Discrimination Equilibrium The monopolist's objective is to maximize total profit $\pi(q)$.

1. $\pi(q) = \text{Total Revenue} - C(q)$ (by the accounting definition of profit).
2. $\pi(q) = \int_0^q P(x)dx - C(q)$ (by the hypothesis that the revenue from the x-th unit is $P(x)$, accumulated by the monopolist).
3. $\frac{d\pi}{dq} = \frac{d}{dq}\left[\int_0^q P(x)dx\right] - \frac{dC(q)}{dq}$ (by the necessary first-order condition for optimization).
4. $\frac{d\pi}{dq} = P(q) - C'(q)$ (by applying the Fundamental Theorem of Calculus to the integral and the definition of marginal cost).
5. $P(q) - C'(q) = 0$ (setting the derivative to zero to find the critical point).

6. $P(q) = MC(q)$ (by definition of marginal cost $C'(q) = MC$). ■

Property 2: Allocative Efficiency and Zero Surplus

1. $S(q) = \int_0^q P(x)dx -$ Total Payment (definition of consumer surplus).
2. Total Payment $= \int_0^q P(x)dx$ (by the definition of first-degree discrimination, where the price of each unit x is $P(x)$).
3. $S(q) = 0$ (direct algebraic subtraction of identical terms).
4. Social Welfare $= S(q) + \pi(q) = 0 + \pi(q) = \int_0^q P(x)dx - C(q)$ (by definition of social welfare as the sum of surpluses).
5. Since $P(q^*) = MC(q^*)$, the quantity produced q^* is the same as in perfect competition, where the market price equals marginal cost. ■

3.2.6 Formal Connection: Intuition → Proof

Mechanism from Intuition	Corresponding Variable/Condition in Proof	Step/Equation
Capture of all willingness to pay	Total Revenue: $\int_0^q P(x)dx$	Property 1, Step 2
Production decision on the last unit	FOC: $P(q) - C'(q) = 0$	Property 1, Step 5
Price equals marginal cost	$P(q^*) = MC(q^*)$	Property 1, Step 6
Zero consumer surplus	$S(q) = 0$	Property 2, Step 3
Allocative Efficiency (Social)	Quantity q^* equal to competition	Property 2, Step 5

3.3 Second-Degree Price Discrimination

Second-degree price discrimination, often called non-linear pricing or menu of packages, occurs when a firm has market power but faces **asymmetric information**: it knows the distribution of consumer preferences but cannot individually identify each customer's willingness to pay at the time of purchase. To solve this problem, the firm designs a "menu" of options (different quantities or qualities at different prices) and induces consumers to "self-identify" through their choices.

3.3.1 Motivation

Unlike first-degree discrimination (where the seller is "omniscient"), here the monopolist knows that there are customers who highly value the product and others who value it little, but he does not know who is who. If he charges a single high price, he loses low-income customers; if he charges a single low price, he fails to profit from high-income customers. Second-degree discrimination attempts to capture surplus by offering volume discounts or different quality levels.

Practical Analogy: Consider cell phone data plans. The operator does not know exactly how much internet you use, so it offers a menu: a basic 2GB plan for $30 and a "premium" 20GB plan for $100. By choosing the larger plan, you signal to the company that you have a high willingness to pay, allowing it to capture more profit from you than if it offered only a medium plan for everyone.

3.3.2 Mechanisms

- **Consumers:** They maximize their utility by choosing the package that offers them the greatest net surplus. The firm's challenge is to ensure that the high-demand customer does not prefer the package intended for the low-demand customer (a phenomenon called **personal arbitrage**).
- **Firm:** The monopolist acts as a mechanism designer. It must satisfy two critical conditions:
  1. **Participation (Individual Rationality):** The price cannot be higher than the value the consumer attributes to the good, otherwise he will not buy.

2. **Incentive Compatibility:** The package designed for the "High" type must give him more satisfaction than the "Low" type's package, so that he does not try to "pretend" to be a low-income consumer.

- **Economic Logic:** To prevent the rich from choosing the poor's package, the firm deliberately **worsens** the poor's offer (reducing quantity or quality), making it unattractive to those with more money.

3.3.3 Implications

The main change is that the market becomes more segmented. This can be socially beneficial if it allows low-income groups to access the good (which they would not have under a single high price), but it generates inefficiency: low-demand consumers end up consuming less than socially optimal.

Concrete Example: Electricity or water tariffs that vary according to consumption blocks. Large consumers may pay different marginal rates, allowing the utility to recover high fixed costs through a price structure that segments industrial from residential users without requiring complex supporting documents.

3.3.4 Formal Definition

Let there be a monopolist facing two types of consumers, θ_1 (low demand) and θ_2 (high demand), with $\theta_2 > \theta_1 > 0$.

- **Variables and Parameters:**
  - $q_i \in [0, \infty)$: Quantity offered to type i.
  - $T_i \in \mathbb{R}_+$: Total tariff paid by type i.
  - $U(q, \theta) = \theta V(q) - T$: Consumer utility, where $V'(q) > 0$ and $V''(q) < 0$ (diminishing marginal utility).
  - $C(q) = cq$: Total production cost with constant marginal cost c.
  - $\lambda \in (0, 1)$: Proportion of consumers of type θ_1 in the population.
- **Hypotheses:**
  1. **Single-Crossing Property (Spence-Mirrlees):** $\frac{\partial^2 U}{\partial q \partial \theta} > 0$, meaning that the high-demand consumer values each additional unit more than the low-demand consumer.
  2. **Asymmetric Information:** The firm knows $\lambda, \theta_1, \theta_2$, but not the individual type at the time of transaction.
  3. **Monotonicity:** $q_2 \geq q_1$ must be satisfied in equilibrium.
- **What "breaks" formally:** If the single-crossing hypothesis fails, the firm cannot separate the types (pooling equilibrium), since indifference curves could cross multiple times, preventing the ordering of choices.

3.3.5 Proofs

Objective: Maximize monopolist profit $\Pi = \lambda(T_1 - cq_1) + (1 - \lambda)(T_2 - cq_2)$.

Constraints: - (IR1) $\theta_1 V(q_1) - T_1 \geq 0$ (Type 1 must want to buy). - (IC2) $\theta_2 V(q_2) - T_2 \geq \theta_2 V(q_1) - T_1$ (Type 2 must prefer its package to type 1's).

Property: Inefficiency at the Bottom and Efficiency at the Top

1. **Identification of Active Constraints:** The monopolist will increase T_1 and T_2 until the constraints bind. In equilibrium, $T_1 = \theta_1 V(q_1)$ (extracts all surplus from type 1) and $T_2 = \theta_2 V(q_2) - \theta_2 V(q_1) + T_1$ (by manipulating IC2).
2. **Tariff Substitution:** Substitute T_1 into T_2: $T_2 = \theta_2 V(q_2) - (\theta_2 - \theta_1)V(q_1)$ (by algebraic substitution).
3. **Substituted Profit Function:**

$$\Pi = \lambda[\theta_1 V(q_1) - cq_1] + (1 - \lambda)[\theta_2 V(q_2) - (\theta_2 - \theta_1)V(q_1) - cq_2]$$

(by substituting the values of T_i into profit).

4. **FOC for q_2:** $\frac{\partial \Pi}{\partial q_2} = (1-\lambda)[\theta_2 V'(q_2) - c] = 0$ (by the first-order condition for optimization).
5. **Efficiency at the Top:** $\theta_2 V'(q_2) = c$ (by simplification). This implies that the high-demand type consumes the **efficient quantity** ($MRS = MC$).
6. **FOC for q_1:**

$$\frac{\partial \Pi}{\partial q_1} = \lambda\theta_1 V'(q_1) - \lambda c - (1-\lambda)(\theta_2 - \theta_1)V'(q_1) = 0$$

(by the partial derivative with respect to q_1).

7. **Isolate Marginal Utility of q_1:** $V'(q_1)[\lambda\theta_1 - (1-\lambda)(\theta_2 - \theta_1)] = \lambda c$ (by factoring $V'(q_1)$).
8. **Inefficiency at the Bottom:**

$$\theta_1 V'(q_1) = c + \frac{1-\lambda}{\lambda}(\theta_2 - \theta_1)V'(q_1)$$

(by algebraic manipulation). Since the term on the right is positive, we have $\theta_1 V'(q_1) > c$, proving that q_1 **is smaller than the socially optimal level**. ■

3.3.6 Formal Connection: Intuition → Proof

Mechanism from Intuition	Corresponding Variable/Condition in Proof	Step/Equation
Ensuring the "poor" customer buys	Participation Constraint (IR1)	Step 1 (Constraints)
Preventing the "rich" customer from pretending to be poor	Incentive Compatibility Constraint (IC2)	Step 1 (Constraints)
The "Rich" consumes what is ideal for him	$\theta_2 V'(q_2) = c$ (Efficiency at the Top)	Step 5 of Proof
The "Poor" suffers a reduction in supply (inefficiency)	$\theta_1 V'(q_1) > c$ (Downward distortion)	Step 8 of Proof
The "Rich" retains part of his value	Informational rent: $(\theta_2 - \theta_1)V(q_1)$	Step 2 (Substitution)

3.4 Third-Degree Price Discrimination

Third-degree price discrimination, commonly referred to as multi-market pricing, represents the situation where a monopolist manages to segment its consumer base into distinct groups and charge different linear prices to each group. Unlike second-degree discrimination (where the consumer self-assigns to a package), here the firm uses **exogenous observable signals** to perform the separation.

3.4.1 Motivation

The central problem this topic solves is the **inefficiency of a single price** when facing groups with drastically different price sensitivities. If a monopolist sets a high price, it loses the low-income market; if it sets a low price, it fails to capture the surplus from the high-income group. Third-degree discrimination allows the firm to "extract" more value by adapting the price to each segment's elasticity.

Practical Analogy: Imagine a publisher releasing an awaited book. It knows that avid fans will pay a high price for the deluxe hardcover edition (low elasticity), while students will prefer the paperback economy version months later (high elasticity). By offering the same basic content in distinct formats or times to identifiable groups, it maximizes its return in both niches.

3.4.2 Mechanisms

- **The Firm:** Acts as an elasticity arbitrator. It not only seeks the point where **Marginal Revenue (MR) equals Marginal Cost (MC)**, but must ensure that **MR is identical in all served markets**. If MR were higher in market A than in B, the firm would gain more by diverting a unit of sale from B to A.
- **Consumers:** Are classified by attributes they cannot easily change (age, location, occupation). They do not choose the price; they face a specific offer based on the group they belong to.
- **Economic Logic:** The firm charges a higher price to the group that has "fewer options" or more urgent need, reflected in a more rigid (inelastic) demand curve.

3.4.3 Implications

The main practical change is **surplus redistribution** and, potentially, market expansion. In some cases, a good is only produced if discrimination is allowed, since revenue from a single market may not cover fixed costs.

Concrete Example: Airline tariffs. Airlines charge more from business travelers (inelastic, buying last minute) and less from tourists (elastic, planning ahead). Without this differentiation, promotional tickets might not exist, or direct business flights would be economically unviable.

3.4.4 Formal Definition

Let there be a monopolist producing a good with total cost $C(Q)$, where Q is total production distributed across k independent markets.

- **Variables and Parameters:**
  - q_i: Quantity sold in market i, where $i = 1, \ldots, k$.
  - $p_i(q_i)$: Inverse demand function in market i.
  - $Q = \sum_{i=1}^{k} q_i$: Total production.
  - $\epsilon_i = -\frac{dq_i}{dp_i}\frac{p_i}{q_i}$: Price elasticity of demand in market i (defined as a positive value).
- **Hypotheses:**
  1. **Market Independence:** Demand in market i depends only on p_i.
  2. **No Arbitrage:** Consumers in the low-price market cannot resell to the high-price market.
  3. **Differentiability:** Demand and cost functions are continuous and differentiable.
- **What "breaks" formally:** If there is arbitrage between groups, prices must converge to a single price, invalidating segmentation.

3.4.5 Proofs

Property 1: Equality of Marginal Revenues to Marginal Cost The monopolist solves: $\max_{q_1,\ldots,q_k} \Pi = \sum_{i=1}^{k} p_i(q_i)q_i - C(\sum_{i=1}^{k} q_i)$.

1. $\frac{\partial \Pi}{\partial q_i} = 0, \forall i$ (by the necessary first-order optimality condition).
2. $\frac{d[p_i(q_i)q_i]}{dq_i} - \frac{dC(Q)}{dQ}\frac{\partial Q}{\partial q_i} = 0$ (by the linearity rule and chain rule).
3. $\frac{\partial Q}{\partial q_i} = 1$ (by the derivative of the sum $\sum q_j$ with respect to one of its components q_i).
4. $MR_i(q_i) - MC(Q) = 0$ (by definition of Marginal Revenue and Marginal Cost).
5. $MR_1(q_1) = MR_2(q_2) = \cdots = MR_k(q_k) = MC(Q)$ (by applying the result for all i). ■

Property 2: The Inverse Elasticity Rule (Lerner Markup) From the FOC $MR_i = MC$:

1. $p_i + q_i \frac{dp_i}{dq_i} = MC$ (by expanding the derivative of the product $p_i(q_i) \cdot q_i$).
2. $p_i \left[1 + \frac{q_i}{p_i}\frac{dp_i}{dq_i}\right] = MC$ (by factoring p_i).
3. $p_i \left[1 - \frac{1}{\epsilon_i}\right] = MC$ (by substituting the definition of elasticity ϵ_i).
4. $p_i - \frac{p_i}{\epsilon_i} = MC$ (by the distributive property).

5. $p_i - MC = \frac{p_i}{\epsilon_i}$ (by transposition of terms).
6. $\frac{p_i - MC}{p_i} = \frac{1}{\epsilon_i}$ (dividing both sides by p_i). ■

Formal Conclusion: Since MC is common to all markets, if $\epsilon_i < \epsilon_j$, then necessarily $p_i > p_j$. **More inelastic markets suffer larger markups.**

3.4.6 Formal Connection: Intuition → Proof

Mechanism from Intuition	Corresponding Variable/Condition in Proof	Step/Equation
Exogenous Identity Signal	Segmentation into k independent demands $p_i(q_i)$	Definition 3.4.4
Optimal Allocation between Groups	Equality of Marginal Revenues: $MR_i = MR_j$	Property 1, Step 5
Focus on Demand Rigidity	Dependence on Elasticity ϵ_i	Property 2, Step 3
Adapted Price (Markup)	Multi-market Lerner Index: $\frac{p_i - MC}{p_i} = \frac{1}{\epsilon_i}$	Property 2, Step 6
Centralized Production	MC dependent on total sum $Q = \sum q_i$	Property 1, Step 4

3.5 Collusion

3.5.1 Motivation

The central problem collusion attempts to solve is the **erosion of profits** caused by strategic interdependence. In a standard oligopoly (like Cournot or Bertrand), the individual desire to capture market share pushes prices down, often to marginal cost, which is optimal for consumers but suboptimal for producers. Collusion emerges as an attempt to "fix" this coordination failure among firms so they can share a more generous "pie" of profits.

Practical Analogy: Imagine two neighbors sharing a fishing lake. If both fish aggressively (competition), the fish stock is quickly depleted and both lose in the long run. If they agree to fish only what is necessary to maintain equilibrium (collusion), the stock is preserved and both profit. However, there is always the temptation for one of them to fish secretly at night to gain a little more, hoping the other still respects the agreement.

3.5.2 Mechanisms

- **The Firm and the Dilemma:** Firm behavior is paradoxical. Collectively, they want to maximize the sum of their profits, acting like a **multi-plant monopoly**. However, individually, each firm faces a marginal incentive to "cheat" on the agreement, because at the high collusive price, selling an extra unit generates an individual gain above marginal cost.
- **Logic of Repetition:** Since collusion is rarely a legally binding contract (due to antitrust laws), its stability depends on **repeated interaction**. Firms use the threat of future punishments (like a price war) to discipline present behavior.
- **The Role of Patience:** The viability of cooperation depends on how valuable the future is to firms. If they are "patient" (low discount rate), the loss of future profits due to punishment will outweigh the immediate gain from cheating.

3.5.3 Implications

Collusion generates severe **allocative inefficiency**, creating social "deadweight loss" by maintaining prices above marginal cost and restricting production. In practice, this motivates the existence of regulatory bodies and laws like the Sherman Act in the US, which prohibit conspiracies to restrain trade.

Concrete Example: The **OPEC** cartel. Member countries try to fix production quotas to keep the price of oil high. However, we often observe episodes where members exceed their quotas (cheating) to increase their national revenue, which can trigger drastic drops in world prices when the "collusion" temporarily collapses.

3.5.4 Formal Definition

Let there be a symmetric oligopoly with n identical firms operating in an infinite time horizon $t = 1, 2, \ldots$.

- **Variables and Parameters:**
  - q_i^t: Quantity produced by firm i in period t.
  - $P(Q)$: Inverse demand function, where $Q = \sum q_i$.
  - c: Constant marginal cost.
  - $\delta \in (0, 1)$: Intertemporal discount factor.
  - $\pi_i(q_i, q_{-i})$: Profit function of firm i.
- **Hypotheses:**
  1. **Perfect Information:** Firms observe the past actions of all others.
  2. **Infinite Horizon:** Necessary to avoid "unraveling" by backward induction that occurs in finite games.
  3. **Grim Trigger Strategy:** Firms cooperate if all cooperated in the past; otherwise, they revert to the static Nash equilibrium forever.
- **What "breaks" formally:** If the horizon is finite (T), the firm will cheat in T (since there is no future to punish), leading to cheating in $T - 1$, and collusion collapses in all periods.

3.5.5 Proofs

Property 1: The Instability of Static Collusion *Statement: In a one-period game, collusion is not a Nash equilibrium.*

1. $\pi_{total}(Q) = P(Q)Q - cQ$ (by definition of joint monopoly profit).
2. $\frac{d\pi_{total}}{dQ} = P'(Q)Q + P(Q) - c = 0$ (first-order condition for the joint maximum).
3. Let Q^m be the quantity satisfying (2). Each firm produces $q_i^c = Q^m/n$ and earns $\pi_i^c = \pi_{total}(Q^m)/n$.
4. Consider the individual incentive of firm i to deviate: $\frac{\partial \pi_i}{\partial q_i} = P'(Q)q_i + P(Q) - c$.
5. Substituting $P(Q^m) - c = -P'(Q^m)Q^m$ from Eq. (2) into Eq. (4): $\frac{\partial \pi_i}{\partial q_i} = P'(Q^m)q_i - P'(Q^m)Q^m = P'(Q^m)(q_i - Q^m)$ (by algebraic substitution).
6. Since $P'(Q) < 0$ and $q_i = Q^m/n < Q^m$ for $n > 1$, then $\frac{\partial \pi_i}{\partial q_i} > 0$ (by sign analysis).
7. **Conclusion:** The firm always increases its profit by producing slightly more than the collusive quota, proving that static collusion is unstable. ■

Property 2: Sustainability Condition (Discount Factor) *Statement: Collusion is a Subgame Perfect Nash Equilibrium if and only if $\delta \geq \bar{\delta}$.*

1. $V_c = \sum_{t=0}^{\infty} \delta^t \pi^c = \frac{\pi^c}{1-\delta}$ (present value of perpetual cooperation, by geometric series).
2. $V_d = \pi^d + \sum_{t=1}^{\infty} \delta^t \pi^n = \pi^d + \frac{\delta \pi^n}{1-\delta}$ (present value of cheating today and suffering perpetual Nash punishment).
   - π^d: optimal deviation profit.
   - π^n: profit in the static Nash equilibrium.
3. For sustainability, we require $V_c \geq V_d$ (no-deviation condition).
4. $\frac{\pi^c}{1-\delta} \geq \pi^d + \frac{\delta \pi^n}{1-\delta}$ (substituting the definitions).
5. $\pi^c \geq \pi^d(1 - \delta) + \delta \pi^n$ (multiplying both sides by $(1 - \delta)$).
6. $\pi^c - \pi^d \geq -\delta \pi^d + \delta \pi^n$ (by distributivity and transposition).
7. $\delta(\pi^d - \pi^n) \geq \pi^d - \pi^c$ (isolating terms with δ).
8. $\delta \geq \frac{\pi^d - \pi^c}{\pi^d - \pi^n}$ (dividing by the profit differential, which is positive). ■

3.5.6 Formal Connection: Intuition → Proof

Mechanism from Intuition	Corresponding Variable/Condition in Proof	Step/Equation
Attempt to act as a Monopoly	Maximization of π_{total} and Q^m	Property 1, Step 2
Individual incentive to cheat	$\frac{\partial \pi_i}{\partial q_i} > 0$ at Q^m	Property 1, Step 6
Future Punishment (Stick)	Term $\frac{\delta \pi^n}{1-\delta}$ (perpetual loss)	Property 2, Step 2
Immediate Gain (Cheating)	π^d (optimal deviation profit)	Property 2, Step 2
Required Patience Level	Threshold $\bar{\delta} = \frac{\pi^d - \pi^c}{\pi^d - \pi^n}$	Property 2, Step 8

3.6 Oligopoly

3.6.1 Bertrand Duopoly

Bertrand duopoly challenges our intuition about market power by suggesting that competition between just two firms can be sufficient to replicate the efficiency of a perfectly competitive market.

Motivation The Bertrand model, proposed in 1883, attempts to solve the problem of how firms interact when their main competitive weapon is **price**, not quantity produced. While in the Cournot model firms decide how much to produce and the market determines the price, Bertrand argued that, in reality, firms announce prices and consumers decide from whom to buy.

Practical Analogy: Imagine two identical popcorn carts side by side in a square. If both charge $5.00, they split the customers. But if one lowers the price to $4.95, all customers in the square will flock to it, leaving the other with no sales. This incentive to "bid lower" is the engine of the model.

Mechanisms

- **Consumers:** Behave strictly rationally, seeking the lowest available price. Since products are perfect substitutes, brand loyalty is non-existent; the lowest price captures 100% of demand.
- **Firms:** Perceive that their demand is extremely price-sensitive (perfectly elastic). They anticipate that if they charge any value above production cost, the competitor will have an irresistible incentive to charge one cent less and steal the entire market.
- **Economic Logic:** The continuous undercutting process only stops when there is no more profit to sacrifice. This occurs when the price equals **marginal cost**, at which point economic profit is zero.

Implications The result is the so-called **Bertrand Paradox**: a market configuration with only two competitors produces the same outcome as a market with infinite firms (perfect competition).

Concrete Example: This illuminates the aggressiveness in **reverse auctions** or public tenders. When the government wants to buy supplies and the only criterion is the lowest price, two firms with similar costs end up offering bids so close to their manufacturing costs that the profit margin disappears, maximizing consumer surplus (in this case, the State).

Formal Definition Let there be a market with two firms, $i \in \{1, 2\}$, producing homogeneous goods.

- **Variables and Parameters:**
  - $p_i \in [0, \infty)$: Price chosen by firm i.
  - $c \geq 0$: Constant and identical marginal cost for both firms.
  - $D(p)$: Total market demand function, where $D'(p) < 0$.
  - $\pi_i(p_i, p_j)$: Profit function of firm i.

- **Individual Demand Function (D_i):**

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

- **Hypotheses:**
  1. **Simultaneity:** Firms choose prices at the same time without observing the rival's choice.
  2. **Perfect Substitutes:** Goods are indistinguishable to consumers.
  3. **Perfect Information:** Firms know each other's cost structure.
- **What "breaks" formally:** If there is **capacity constraint** ($q_i \leq K_i$), the price may not fall to marginal cost (Edgeworth Paradox). If products are **differentiated**, profit functions become continuous and equilibrium involves prices above marginal cost.

Proofs **Auxiliary Theorem: Definition of Nash Equilibrium (NE)** *A pair of prices (p_1^*, p_2^*) is a NE if $\pi_i(p_i^*, p_j^*) \geq \pi_i(p_i, p_j^*)$ for all $p_i \geq 0$.*

Property: The Unique Bertrand Equilibrium *Statement: In a symmetric Bertrand duopoly, the unique Nash Equilibrium is $p_1^* = p_2^* = c$.*

Proof by Exclusion (Cases):

1. **Case $p_1 > p_2 > c$:**
   - Firm 1 has profit $\pi_1 = 0$ (by the demand definition where $p_1 > p_j$).
   - If firm 1 chooses $p_1 = p_2 - \epsilon$ (where ϵ is infinitesimally small):
   - $\pi_1 = (p_2 - \epsilon - c)D(p_2 - \epsilon) > 0$ (by simple algebra, since $p_2 > c$).
   - **Conclusion:** This case is not a NE, since firm 1 has an incentive to deviate.
2. **Case $p_1 = p_2 = \bar{p} > c$:**
   - Each firm earns $\pi_i = \frac{1}{2}(\bar{p} - c)D(\bar{p})$ (by the market split rule).
   - If firm 1 deviates to $p_1 = \bar{p} - \epsilon$:
   - $\pi_1' = (\bar{p} - \epsilon - c)D(\bar{p} - \epsilon)$ (captures the entire market).
   - Since $\lim_{\epsilon \to 0} \pi_1' = (\bar{p} - c)D(\bar{p})$, and $(\bar{p} - c)D(\bar{p}) > \frac{1}{2}(\bar{p} - c)D(\bar{p})$:
   - $\pi_1' > \pi_1$ for ϵ sufficiently small (by term comparison).
   - **Conclusion:** Not a NE.
3. **Case $p_1 = p_2 = c$:**
   - Both have profit $\pi_i = (c - c)D(c) = 0$.
   - If firm i increases price ($p_i > c$): its profit remains 0 (sales fall to zero).
   - If firm i decreases price ($p_i < c$): its profit becomes negative (sells below cost).
   - **Conclusion:** No firm wants to deviate. It is a Nash Equilibrium. ■

Formal Connection: Intuition → Proof

Mechanism from Intuition	Corresponding Variable/Condition in Proof	Step/Equation
Price War (Undercutting)	Deviation $p_i = p_j - \epsilon$ yields higher profit	Property, Cases 1 and 2
Total market capture	Demand $D_i(p_i, p_j) = D(p_i)$ if $p_i < p_j$	Definition (Demand)
Zero Profit in Equilibrium	$\pi_i = (p_i^* - c) \cdot \text{Sales} = 0$	Property, Case 3
Inevitability of Marginal Cost	Implicit FOC: deviation incentive ceases at $p = c$	Uniqueness Proof
Perfect Substitutes	Market split $\frac{1}{2}D(p)$ in case of tie	Definition (Demand)

3.6.2 Cournot Duopoly

Cournot duopoly, one of the pillars of modern Industrial Organization, was proposed by Augustin Cournot in 1838. It grounds our understanding of how strategic competition in quantities shapes market outcomes.

Motivation The Cournot model addresses the problem of **strategic interdependence**: how should a firm decide its production level when the final price of the product depends not only on its own decision but also on the quantity launched into the market by its competitor? It solves the dilemma of firms that need to commit production capacities before the equilibrium price is determined.

Practical Analogy: Imagine two mineral water producers extracting water from natural springs at almost zero cost. They must decide, simultaneously and without knowing the other's choice, how many gallons to bring to the town square in the morning. If both bring too much, the water will be left over and the price will drop drastically. If they bring too little, the price will be high, but they may regret not having sold more. Equilibrium occurs when each brings a quantity that is the "best response" to the (correctly anticipated) quantity of the other.

Mechanisms

- **The Firm:** Unlike a price taker, the Cournot firm recognizes that it has **market power**. It perceives that by increasing its production (q_i), it causes a negative effect on the market price (P), which reduces the revenue from all units it already planned to sell.
- **Consumers:** Behave passively according to a downward-sloping aggregate demand curve. They buy the total production ($Q = q_1 + q_2$) at the price that clears the market.
- **Economic Logic:** The central mechanism is the **reaction function** (or best response). If Firm 2 produces a lot, there is little "residual demand" left for Firm 1, inducing it to produce less to avoid a price collapse. Equilibrium is the point where no one wants to alter their production, given the rival's actions.

Implications The model illuminates the **inefficiency of oligopoly**. The Cournot result is a middle ground: more is produced and less is charged than in a monopoly, but less is produced and more is charged than in perfect competition.

Concrete Example: The model is used to analyze **mergers and acquisitions**. If two Cournot firms merge, they start acting as a monopoly, increasing price by contracting total quantity. Antitrust regulators use this logic to predict whether a merger will excessively increase market power and harm consumer welfare.

Formal Definition Let there be a market with two players (firms), $i \in \{1, 2\}$.

- **Variables and Parameters:**
  - $q_i \in [0, \infty)$: Quantity strategy of firm i.
  - $P(Q) = a - bQ$: Inverse demand function, where $Q = q_1 + q_2$ and $a, b > 0$.
  - $C_i(q_i) = c \cdot q_i$: Total cost function with constant marginal cost c, where $a > c \geq 0$.
  - $\pi_i(q_i, q_j)$: Profit (payoff) function of firm i.
- **Hypotheses:**
  1. **Simultaneity:** Firms choose q_i without observing the rival's choice.
  2. **Homogeneity:** Products are perfect substitutes.
  3. **Complete Information:** Demand and cost structure is common knowledge.
  4. **Rationality:** Firms maximize profit non-cooperatively.
- **What "breaks" formally:** If the game is sequential, it becomes the **Stackelberg** model, where the leader produces more. If marginal costs are decreasing or demand is highly convex, profit may not be concave, and equilibrium may not exist or be multiple.

Proofs

Property 1: Derivation of Reaction Functions Firm 1's objective is to maximize its profit $\pi_1(q_1, q_2)$.

1. $\pi_1 = P(q_1 + q_2)q_1 - C(q_1)$ (by definition of profit as revenue minus cost).
2. $\pi_1 = [a - b(q_1 + q_2)]q_1 - cq_1$ (by substituting the demand and cost functions).
3. $\pi_1 = aq_1 - bq_1^2 - bq_1q_2 - cq_1$ (by algebraic distributivity).
4. $\frac{\partial \pi_1}{\partial q_1} = a - 2bq_1 - bq_2 - c = 0$ (by the necessary first-order condition for the maximum).
5. $2bq_1 = a - c - bq_2$ (by transposition of terms).
6. $q_1(q_2) = \frac{a-c}{2b} - \frac{q_2}{2}$ (by dividing both sides by $2b$; this is the reaction function). ■

Property 2: The Cournot Nash Equilibrium (q_1^*, q_2^*) Equilibrium occurs where the reaction functions intersect: $q_1 = q_1(q_2)$ and $q_2 = q_2(q_1)$.

1. $q_1 = \frac{a-c}{2b} - \frac{q_2}{2}$ (from Property 1, Step 6).
2. $q_2 = \frac{a-c}{2b} - \frac{q_1}{2}$ (by symmetry between firms).
3. $q_1 = \frac{a-c}{2b} - \frac{1}{2}\left[\frac{a-c}{2b} - \frac{q_1}{2}\right]$ (by substituting (2) into (1)).
4. $q_1 = \frac{a-c}{2b} - \frac{a-c}{4b} + \frac{q_1}{4}$ (by scalar distributivity).
5. $q_1 - \frac{q_1}{4} = \frac{2(a-c)-(a-c)}{4b}$ (by transposition and LCM).
6. $\frac{3q_1}{4} = \frac{a-c}{4b}$ (by term simplification).
7. $q_1^* = q_2^* = \frac{a-c}{3b}$ (multiplying both sides by 4/3 and applying symmetry). ■

Property 3: Equilibrium Price and Profit

1. $Q^* = q_1^* + q_2^* = \frac{2(a-c)}{3b}$ (by definition of aggregate production).
2. $P^* = a - b\left[\frac{2(a-c)}{3b}\right]$ (substituting Q^* into inverse demand).
3. $P^* = \frac{3a-2a+2c}{3} = \frac{a+2c}{3}$ (by algebraic simplification).
4. $\pi_i^* = (P^* - c)q_i^*$ (by definition of unit profit times quantity).
5. $\pi_i^* = \left(\frac{a+2c}{3} - c\right)\left(\frac{a-c}{3b}\right) = \frac{(a-c)^2}{9b}$ (by direct algebra). ■

Formal Connection: Intuition → Proof

Mechanism from Intuition	Corresponding Variable/Condition in Proof	Step/Equation
Market Power (Markup)	Term $-2bq_1$ in the revenue derivative	Property 1, Step 4
Interdependence	Dependence on q_j in firm i's FOC	Property 1, Step 4
Best Response	Reaction Function $q_i(q_j)$	Property 1, Step 6
Equilibrium Stability	Intercept of functions q_1^* and q_2^*	Property 2, Step 7
Allocative Inefficiency	$P^* = \frac{a+2c}{3} > c$ (Price above Marginal Cost)	Property 3, Step 3

3.7 Summary of Key Results

Concept	Definition	Key Formula
Monopoly FOC	Marginal revenue equals marginal cost	$MR(q^*) = MC(q^*)$
Lerner Index	Markup over marginal cost	$\frac{p-MC}{p} = \frac{1}{\lvert\epsilon\rvert}$
First-Degree Discrimination	Price equals willingness to pay per unit	$P(q) = MC(q)$ at equilibrium; $CS = 0$
Second-Degree Discrimination	Menu design with self-selection	$\theta_2 V'(q_2) = c$ (top efficiency); $\theta_1 V'(q_1) > c$ (bottom distortion)

Concept	Definition	Key Formula
Third-Degree Discrimination	Multi-market pricing	$MR_i = MC = MR_j$; $\frac{p_i - MC}{p_i} = \frac{1}{\epsilon_i}$
Collusion Sustainability	Grim trigger strategy	$\delta \geq \frac{\pi^d - \pi^c}{\pi^d - \pi^n}$
Bertrand Equilibrium	Price equals marginal cost	$p_1^* = p_2^* = c$
Cournot Equilibrium	Quantity competition	$q_i^* = \frac{a-c}{3b}$, $P^* = \frac{a+2c}{3}$, $\pi_i^* = \frac{(a-c)^2}{9b}$

Chapter 4: Social Welfare Theory and Preference Aggregation

4.1 Individual Preferences and Social Welfare Functions

4.1.1 Motivation

The topic of individual preferences and social welfare functions addresses the fundamental problem of **collective aggregation**: how to transform the preference orderings of diverse individuals into a single institutional choice that represents the "will of society"? While a single economic agent can easily order their choices, a society composed of multiple individuals with divergent tastes and endowments faces severe logical constraints to consistently define whether one social state is superior to another.

Practical Analogy: Imagine a group of friends trying to choose a pizza flavor to share. Each has their own personal ranking (individual preference). The method of choice—whether majority voting, imposition by a leader, or summing everyone's happiness—represents the Social Welfare Function (SWF). The conflict arises when the aggregation rule fails to preserve basic rationality properties, allowing the group to prefer flavor A over B, B over C, but inconsistently prefer C over A.

4.1.2 Mechanisms

Economic agents operate under the following axiomatic structure:

- **Consumers:** Are endowed with a binary preference relation ($\succeq$) that allows them to compare alternative consumption bundles. Rational behavior presupposes that they choose the best alternative contained in their possibility set.
- **Social Planner:** Acts as a central arbiter who adopts an SWF to evaluate alternative distributions of resources and utilities in the economy.

The microeconomic logic behind the properties of completeness and transitivity aims to guarantee the consistency of choices. Without transitivity, social choices could enter infinite decision cycles, preventing the determination of a stable optimal point. At the aggregate level, axioms such as Independence of Irrelevant Alternatives (IIA) seek to ensure that social choice between two options depends strictly on how individuals evaluate those two options, without suffering distortions from the introduction of a third secondary alternative.

4.1.3 Implications

In public economics, this concept establishes the formal limits of central planning and democratic voting regimes, mathematically demonstrating that no purely ordinal aggregation rule can simultaneously satisfy basic criteria of distributive justice and logical rationality.

Concrete Example: When faced with the choice of budget allocation between building a hospital or a school, if the social planner aggregates citizens' preferences via simple majority, the collective choice can become inconsistent. To circumvent this aggregation failure, governments adopt specific ethical criteria translated into mathematical functions, such as the *Utilitarian SWF* (focused on maximizing total surplus) or the *Rawlsian SWF* (focused on maximizing the welfare of the worst-off agent).

4.1.4 Formal Definition

I. Individual Preferences Let X be the set of all feasible alternatives (social states or consumption bundles). The preference relation of individual i is a binary relation $\succeq_i$ defined over the domain $X \times X$.

- **Hypothesis 1 (Completeness):** $\forall x, y \in X$, we have $x \succeq_i y$ or $y \succeq_i x$.
- **Hypothesis 2 (Transitivity):** $\forall x, y, z \in X$, if $x \succeq_i y$ and $y \succeq_i z$, then $x \succeq_i z$.

If the relation $\succeq_i$ is continuous under the standard topology, there exists a continuous utility function $u_i : X \to \mathbb{R}$ such that:

$$x \succeq_i y \iff u_i(x) \geq u_i(y)$$

II. Social Welfare Function (SWF) A Bergson–Samuelson SWF is an aggregator function $W : \mathbb{R}^I \to \mathbb{R}$ that maps the vector of individual utilities $(u_1, \ldots, u_I)$ into a social scalar.

- **Domain:** The set of all realizable utility profiles in the economy.
- **Parameters:** Distributive ethical weights $\alpha_i > 0$ associated with each individual.

  What breaks without the hypotheses: If the transitivity property fails, the set of social choices ceases to have a maximal element, generating decision paralysis. In the case of the SWF, the absence of the IIA axiom allows the inclusion of an irrelevant alternative to alter the social ordering of the main options, enabling strategic manipulation of electoral agendas.

4.1.5 Proofs

Auxiliary Theorem: Monotonicity and Pareto Efficiency *If the social welfare function $W(u)$ is strictly increasing in each component u_i, the argument that maximizes W over the utility possibility set is necessarily Pareto-efficient.*

Proof: Social Optimum Implies Pareto Efficiency

1. Let $x^* \in X$ be the allocation that maximizes the social planner's objective function:

$$\max_{x \in X} W\big(u_1(x), \ldots, u_I(x)\big)$$

2. Suppose, by contradiction, that the optimal allocation x^* is not Pareto-efficient.
3. By the definition of Pareto dominance, there must exist a feasible alternative allocation $x' \in X$ such that:

$$u_i(x') \geq u_i(x^*) \quad \forall i \in \{1, \ldots, I\}$$
$$u_j(x') > u_j(x^*) \quad \exists j \in \{1, \ldots, I\}$$

4. Since the function $W(\cdot)$ is strictly increasing in all its inputs, apply the strict monotonicity operator to the modified utility vector:

$$W\big(u_1(x'), \ldots, u_j(x'), \ldots, u_I(x')\big) > W\big(u_1(x^*), \ldots, u_j(x^*), \ldots, u_I(x^*)\big)$$

5. The strict inequality obtained in the previous step establishes that:

$$W(u(x')) > W(u(x^*))$$

6. **Conclusion:** This generates a direct contradiction with the initial hypothesis that x^* is the global maximizer of the function W. Therefore, by *reductio ad absurdum*, the socially optimal allocation x^* is necessarily Pareto-efficient. ■

4.1.6 Formal Connection: Intuition → Proof

Mechanism from Intuition	Corresponding Variable/Condition in Proof	Reference in Proof
Individual Maximization	Ordering via utility $u_i(x) \geq u_i(y)$	Section 4.1.4 - Subsection I
Collective Rationality	Preservation of Transitivity in social mapping	Section 4.1.4 - Hypothesis 2
Allocative Efficiency	Maximization of W makes Pareto dominance impossible	Section 4.1.5 - Steps 4 and 5
Distributive Justice	Introduction of linear ethical weights α_i in the SWF	Section 4.1.4 - Subsection II

4.2 The Condorcet Paradox and Majority Rule

4.2.1 Motivation

The Condorcet Paradox exposes a structural logical inconsistency in democratic collective choice processes: the possibility that a group of individuals, each with strictly rational, transitive, and consistent preferences, generates a social preference ordering that is cyclical and irrational. The central problem lies in the aggregation mechanism: simple majority rule applied to binary comparisons is unable to preserve the mathematical property of transitivity from the individual to the collective level.

Practical Analogy: Consider three friends trying to choose a travel destination among Paris (x), Rome (y), and London (z). If they hold pairwise votes, a majority may prefer Paris over Rome, another majority may prefer Rome over London, and paradoxically, a third majority may prefer London over Paris. The collective ordering takes on the behavior of a "rock-paper-scissors" game, where no alternative is stable or immune to majority coalitions.

4.2.2 Mechanisms

The paradox operates through the application of simple majority rule to sequential binary choices:

- **Individual Agents:** Are defined as standard microeconomic maximizers. Their preferences are complete and transitive, making internal cycles impossible in their private decisions.
- **Collective Aggregator:** Society uses majority vote counting to construct a social ranking.

The failure in social transitivity occurs because the majorities that win in each binary round have different individual compositions. The "will of the majority" ceases to be a direct and unified reflection of the population's preferences, becoming an artificial result that intrinsically depends on the design of the voting process.

4.2.3 Implications

The main macroeconomic consequence of the paradox is the permanent instability of collective decisions and the frequent absence of a *Condorcet Winner* (the alternative that defeats all others in pairwise contests). This scenario opens space for the mechanism of agenda manipulation.

Concrete Example: In a legislative vote to define public spending ceilings among three options—High (x), Medium (y), and Low (z)—under cyclical preferences, the legislative committee chair can determine which option will emerge victorious simply by controlling the order of voting. If he wants to approve the High spending target (x), he simply structures the first round between Medium (y) and Low (z). The Low option will be eliminated, and in the final confrontation, the majority will vote for the High option, revealing how institutional stability can be artificial.

4.2.4 Formal Definition

Let $X = \{x, y, z\}$ be the finite set of alternatives and $N = \{1, 2, 3\}$ the set of agents. Each individual $i \in N$ has a strict preference relation $\succ_i$, assuming the absence of indifference to isolate the structure of the paradox.

- **Axiom of Individual Rationality:** For all $i \in N$, the relation $\succ_i$ is strictly transitive and complete.
- **Simple Majority Aggregation Rule:** The strict social preference relation P is governed by the vote counting rule:

$$xPy \iff \big|\{i \in N : x \succ_i y\}\big| > \big|\{i \in N : y \succ_i x\}\big|$$

  What breaks formally: By allowing an unrestricted domain of preference profiles, the social relation P violates the transitivity property, making it impossible for the collective ordering to be modeled as a standard rational preference order.

4.2.5 Proofs

Theorem: Transitivity Failure of Majority Rule

1. Define the preference profile of the three voters over the set of alternatives X:

$$\begin{aligned} \text{Individual 1} &: \quad x \succ_1 y \succ_1 z \\ \text{Individual 2} &: \quad y \succ_2 z \succ_2 x \\ \text{Individual 3} &: \quad z \succ_3 x \succ_3 y \end{aligned}$$

2. Evaluate the social binary contest between alternatives x and y:
   - Individual 1 has $x \succ_1 y$.
   - Individual 2 has $y \succ_2 x$.
   - Individual 3 has $x \succ_3 y$.
   - Calculating cardinalities: $\big|\{i : x \succ_i y\}\big| = 2$ and $\big|\{i : y \succ_i x\}\big| = 1$.
   - By simple majority rule, we obtain the social relation: xPy.
3. Evaluate the social binary contest between alternatives y and z:
   - Individual 1 has $y \succ_1 z$.
   - Individual 2 has $y \succ_2 z$.
   - Individual 3 has $z \succ_3 y$.
   - Calculating cardinalities: $\big|\{i : y \succ_i z\}\big| = 2$ and $\big|\{i : z \succ_i y\}\big| = 1$.
   - By simple majority rule, we obtain the social relation: yPz.
4. Evaluate the social binary contest between alternatives z and x:
   - Individual 1 has $x \succ_1 z$ (from the transitivity of $\succ_1$).
   - Individual 2 has $z \succ_2 x$.
   - Individual 3 has $z \succ_3 x$.
   - Calculating cardinalities: $\big|\{i : z \succ_i x\}\big| = 2$ and $\big|\{i : x \succ_i z\}\big| = 1$.
   - By simple majority rule, we obtain the social relation: zPx.
5. Consolidate the social preference relation P constructed in the previous steps:

$$xPy \quad \text{and} \quad yPz \implies xPz \quad \text{(if } P \text{ were transitive)}$$
$$\text{However, Step 4 proved that:} \quad zPx \iff \neg(xPz)$$

6. **Conclusion:** The collective preference relation P results in the continuous cycle $x \to y \to z \to x$. It is proven that majority rule fails to generate a transitive social ordering. ■

4.2.6 Formal Connection: Intuition → Proof

Mechanism from Intuition	Corresponding Variable/Condition in Proof	Reference in Proof
Individual Rationality	Strict transitivity of $\succ_i$	Section 4.2.4 and Section 4.2.5, Step 1
Binary Voting	Pairwise contests ordered in disjoint sets	Section 4.2.5, Steps 2, 3, and 4
Will of the Majority	Simple majority condition by cardinality counting	Section 4.2.4 - Equation of the Rule
Collective Cyclicality	Simultaneous coexistence of xPy, yPz, and zPx	Section 4.2.5, Steps 5 and 6

4.3 Arrow's Impossibility Theorem

4.3.1 Motivation

Arrow's Impossibility Theorem establishes an insurmountable mathematical barrier for collective choice theory. It investigates how to convert profiles of individual ordinal preferences into a single, stable, and coherent social preference ordering. Arrow proved that if the universe of alternatives contains three or more options, it is impossible to construct a Social Welfare Function (SWF) that simultaneously satisfies a set of basic axioms of distributive justice and logical consistency without resorting to a dictatorial structure.

Practical Analogy: Imagine an electoral committee developing a voting system for a multi-candidate election. The committee demands four basic rules: the system must accept any voter ballot; if everyone prefers candidate A to B, society must choose A; the contest between the main candidates cannot suffer interference from the entry or exit of a candidate with no chance; and finally, the system cannot be based on a "leader" whose vote decides everything. Arrow's theorem proves that this perfect system is a mathematical impossibility; one of the rules will have to be violated.

4.3.2 Mechanisms

The model operates under four fundamental axioms that define a reasonable decision rule:

- **Unrestricted Domain (U):** The function must be able to process any logically possible combination of individual rankings.
- **Weak Pareto Principle (WP):** If absolutely all agents prefer alternative A to alternative B, the social ordering must place A above B.
- **Independence of Irrelevant Alternatives (IIA):** The social choice or ordering between options A and B must depend strictly on the relative positioning that individuals give to A and B, remaining immune to changes in a third alternative C.
- **Non-Dictatorship (D):** There must not exist a central individual whose private preferences determine social choice independently of the rest of the population's vote.

The logical gear of the theorem demonstrates that the axioms U, WP, and IIA force the concentration of all decision-making power of society into the hands of a single economic agent, directly violating the Non-Dictatorship axiom.

4.3.3 Implications

The theorem demonstrates the impossibility of finding a perfect ordinal voting mechanism. This forces the formulation of economic policies toward second-best solutions, requiring the relaxation of one of the Arrovian axioms—whether by imposing restrictions on the format of agents' preferences (such as single-peaked preferences) or by introducing data on the intensity of utilities (cardinal analysis), prohibited in Arrow's pure ordinal modeling.

Concrete Example: In public resource allocation processes and parliamentary amendments, the introduction of a new unviable proposal can completely alter which of the original proposals would obtain the majority of votes within parliament, generating instability in budget approval, unless a centralized leadership (a dictator in Arrow's definition) previously defines the voting agenda.

4.3.4 Formal Definition

Let X be the set of social states with $|X| \geq 3$, and $N = \{1, \ldots, n\}$ the set of individuals with $n \geq 2$. Denote by $\mathcal{R}$ the set of all rational individual preference relations (complete and transitive) over X.

A Social Welfare Function is a mapping $f : \mathcal{R}^N \to \mathcal{R}$ that transforms each profile of individual preferences $(R_1, \ldots, R_n)$ into a stable social preference relation $R = f(R_1, \ldots, R_n)$.

- **Unrestricted Domain (U):** The domain of definition of f is the complete Cartesian product $\mathcal{R}^N$.
- **Weak Pareto (WP):** If xP_iy for all $i \in N$, then xPy, where P denotes strict social preference.
- **Independence of Irrelevant Alternatives (IIA):** Let two profiles (R_i) and (R'_i). If $\forall i \in N$, $xR_iy \iff xR'_iy$, then $xf(R_i)y \iff xf(R'_i)y$.
- **Non-Dictatorship (D):** There does not exist an agent $i^* \in N$ such that $\forall (x, y) \in X^2, xP_{i^*}y \implies xPy$.

4.3.5 Proofs

The demonstration uses the classical concept of **Decisive Coalitions**.

> **Theoretical Definition:** A group of agents $G \subseteq N$ is classified as decisive for alternative x over alternative y (denoted by $D_G(x, y)$) if, whenever all members of G have xP_iy, the resulting social response is necessarily xPy.

Lemma 1: Expansion of the Decisive Field *If a coalition G is decisive for a specific pair (a, b), it is necessarily decisive for any other ordered pair within X.*

1. Suppose coalition G is decisive for the pair of alternatives (a, b). We want to demonstrate that G is also decisive for the generic pair (a, c), where $c \neq a, b$.
2. Invoking the Unrestricted Domain Axiom (U), construct a specific individual preference profile where:
$$\forall i \in G: \quad aP_ibP_ic$$
$$\forall j \notin G: \quad bP_jcP_ja$$
3. Evaluating the subset of alternatives (a, b): observe that all members belonging to G prefer a to b. Since, by hypothesis, G is decisive for this pair, the society's response must register:
$$aPb$$
4. Evaluating the subset of alternatives (b, c): note that both agents in G and agents outside G prefer alternative b to alternative c. By the Weak Pareto Axiom (WP), society must order:
$$bPc$$
5. Combining the results obtained in steps 3 and 4 via the transitivity property of the social preference relation $R \in \mathcal{R}$:
$$aPb \quad \text{and} \quad bPc \implies aPc$$
6. By the Independence of Irrelevant Alternatives Axiom (IIA), the resulting social ordering between options a and c depends solely on the positions of a and c in the individual rankings. In the constructed profile, the members of G are the only ones who prefer a to c.

7. **Conclusion of Lemma 1:** Therefore, the collective ordering aPc was determined exclusively by the preference of G, which formally proves that coalition G is decisive for the pair (a, c). By algebraic symmetry, the same applies to any pair in X. ■

Lemma 2: Contraction to the Dictator Agent

1. By the Weak Pareto Axiom (WP), the universal set containing all individuals in the economy (N) is a decisive coalition.
2. Let $G \subseteq N$ be a decisive coalition of minimum size. Suppose G is not unitary, i.e., contains at least two agents ($|G| \geq 2$).
3. Partition the minimum coalition G into two disjoint and non-empty subgroups, G_1 and G_2, such that $G_1 \cup G_2 = G$.
4. Construct a preference profile allowed by Axiom U over three distinct alternatives $\{a, b, c\}$:

$$\begin{aligned} \text{Agents in } G_1: &\quad aP_ibP_ic \\ \text{Agents in } G_2: &\quad bP_icP_ia \\ \text{Agents outside } G: &\quad cP_iaP_ib \end{aligned}$$

5. Since $G = G_1 \cup G_2$ is a global decisive coalition and all its members prefer b to c, by definition society chooses bPc.
6. Now analyze the social positioning of alternative a relative to c:
   - If society decides aPc, then, given that only the members of G_1 prefer a to c in this profile, the subgroup G_1 would be classified as a decisive coalition (via Lemma 1). This would generate a contradiction with the hypothesis that G was the decisive coalition of minimum size.
   - If society decides $c \succeq a$, combining with the stable condition from step 5 (bPc), we would have by social transitivity that bPa. However, inspecting the profile from step 4, the only agents who sustain bP_ia are the components of G_2. This would imply that G_2 is the decisive coalition, again breaking the minimality property of G.
7. **General Conclusion:** Since the logical structure generates contradiction in all possible branches, the initial hypothesis that $|G| \geq 2$ is false. The minimum decisive coalition must be composed of a single individual, $G = \{i^*\}$. This agent i^* holds the power to determine all social preferences, formally characterizing itself as a dictator and proving Arrow's Theorem. ■

4.3.6 Formal Connection: Intuition → Proof

Mechanism from Intuition	Corresponding Variable/Condition in Proof	Reference in Proof
Power of Unanimity	Weak Pareto Axiom (WP)	Section 4.3.4 and Lemma 1, Step 4
Consistency of Ordering	Transitivity of the social relation R	Lemma 1, Step 5
Isolation of Stable Options	Independence of Irrelevant Alternatives Axiom (IIA)	Lemma 1, Step 6
Collective Decisive Power	Decisive Coalition Property $D_G(x, y)$	Section 4.3.5 - Auxiliary Definition
Non-Existence of Fair Rule	Mathematical derivation of the unitary dictator agent i^*	Lemma 2 - General Conclusion

---

4.4 The Utility Possibility Set and Linear SWFs

4.4.1 Motivation

The Utility Possibility Set (UPS) functions as the geometric "menu" of all feasible combinations of welfare and happiness that a society can achieve from its finite endowments of resources and technology. In an economy with scarcity constraint, the expansion of one individual's utility necessarily requires the contraction of another agent's utility along the frontier of the set.

Since there are infinite efficient combinations on the UPS frontier (where there is no resource waste), the social planner uses the Linear Social Welfare Function (SWF) to apply weighted weights to each agent, serving as an ethical scale designed to elect a single optimal point as the socially ideal allocation.

Practical Analogy: Imagine the UPS as the physical frontier of a piece of land that a government can transform into different configurations of public parks. Linear SWFs act as the guidelines of an ethics committee that decides, for example, that the welfare of a low-income citizen has a weighted weight equal to 2, while that of a high-income citizen has weight 1. The planner's goal is to find the exact coordinates on the land's frontier that maximize this specific weighted sum.

4.4.2 Mechanisms

The allocative processes of the economy are guided by two components:

- **Consumers:** Seek to individually maximize their own utilities given market budget constraints.
- **Benevolent Social Planner:** Optimizes the aggregator function $W = \sum \alpha_h u_h$, where the parameter α_h mirrors the ethical weight associated with individual h.

The economic logic dictates that if the mathematical formulation of the SWF is strictly increasing (strictly positive weights), the point that maximizes social welfare will necessarily be a point belonging to the Pareto efficiency frontier. The First-Order Conditions (FOC) of this problem guarantee that the Marginal Rates of Substitution (MRS) between any pair of goods are equalized among all consumers in the economy, eliminating remaining gains from trade.

4.4.3 Implications

This theoretical framework provides the mathematical foundation for the Two Fundamental Theorems of Welfare Economics. It demonstrates that if society elects a specific welfare distribution guided by criteria of distributive justice, the State does not need to intervene directly by destroying the free market price mechanism. Instead, the government should operate through non-distortionary initial wealth transfers (*lump-sum*) and allow competition to autonomously lead the economy to the chosen optimal efficiency point.

Concrete Example: Evaluation of a large-scale logistics infrastructure project that raises corporate income but generates negative local externalities for neighboring communities. A purely utilitarian SWF (identical ethical weights for everyone) will approve the project's execution if the summed gain exceeds the loss of those affected. However, the FOC instrument demonstrates that if, after project implementation, the groups' MRS diverge, there will be room for compensatory transfers that guarantee a transition based on Pareto improvement.

4.4.4 Formal Definition

I. Utility Possibility Set (UPS) Let X be the set of all feasible allocations and $u_h : X \to \mathbb{R}$ the continuous utility function of individual $h \in \{1, \ldots, H\}$. The UPS (denoted by $\mathcal{U}$) is formally defined as the set:

$$\mathcal{U} = \left\{(v_1, \ldots, v_H) \in \mathbb{R}^H : \exists x \in X \text{ such that } v_h \leq u_h(x), \forall h \in \{1, \ldots, H\}\right\}$$

- **Convexity Property:** If the set of feasible allocations X is convex and all utility functions $u_h(\cdot)$ are strictly concave, the set $\mathcal{U}$ will be strictly convex.

II. Linear Social Welfare Function A Generalized Utilitarian SWF maps the utility vector into a scalar indicator through the weighted sum:

$$W(u_1, \ldots, u_H) = \sum_{h=1}^{H} \alpha_h u_h, \quad \text{where } \alpha_h \geq 0 \quad \forall h$$

What breaks formally: If the social planner defines null ethical weights ($\alpha_h = 0$) for a subgroup of agents, the maximization process of W can result in optimal points that violate Pareto efficiency, since the social algorithm becomes totally indifferent to welfare losses imposed on individuals disregarded by the weighting.

4.4.5 Proofs

Theorem 1: Maximization of Linear SWF with Positive Weights Implies Pareto Efficiency

1. Let $x^* \in X$ be the allocation resulting from the global maximization of the linear social welfare function, under the condition of strictly positive ethical weights:

$$\max_{x \in X} \sum_{h=1}^{H} \alpha_h u_h(x) \quad \text{with } \alpha_h > 0 \quad \forall h \in \{1, \ldots, H\}$$

2. Suppose, by contradiction, that the optimal allocation x^* is not Pareto-efficient.
3. By the formal definition of Pareto dominance, there must exist a feasible alternative allocation $x' \in X$ such that:

$$u_h(x') \geq u_h(x^*) \quad \forall h \in \{1, \ldots, H\}$$
$$u_j(x') > u_j(x^*) \quad \exists j \in \{1, \ldots, H\}$$

4. Multiplying each individual inequality by its respective ethical weight $\alpha_h > 0$, we preserve the direction of the mathematical operators:

$$\alpha_h u_h(x') \geq \alpha_h u_h(x^*) \quad \forall h \in \{1, \ldots, H\}$$
$$\alpha_j u_j(x') > \alpha_j u_j(x^*) \quad \text{(for individual } j\text{)}$$

5. Applying the linear property of the summation operator over the vector of combined inequalities:

$$\sum_{h=1}^{H} \alpha_h u_h(x') > \sum_{h=1}^{H} \alpha_h u_h(x^*)$$

6. By the mathematical definition of the social planner's objective function (W), the inequality from step 5 translates to:

$$W(u(x')) > W(u(x^*))$$

7. **Conclusion of Theorem 1:** The result directly contradicts the initial premise that x^* is the global maximizing argument of W. Hence, by contradiction structure, the allocation x^* is necessarily Pareto-efficient. ■

Theorem 2: First-Order Conditions and the Equalization of MRS

1. To map a Pareto efficiency point, the planner maximizes the utility of individual 1 subject to reservation utility guarantees for the other agents ($\bar{u}_i$) and under the physical resource endowment constraint (ω):

$$\max_{x} u_1(x_1) \quad \text{s.t. } u_i(x_i) \geq \bar{u}_i \quad \forall i \in \{2, \ldots, n\} \quad \text{and} \quad \sum_{i=1}^{n} x_i = \omega$$

2. Construct the corresponding Lagrangian operator for the social optimization problem:

$$\mathcal{L} = u_1(x_1) + \sum_{i=2}^{n} \kappa_i \big(u_i(x_i) - \bar{u}_i\big) + \sum_{l=1}^{L} \mu_l \left(\omega_l - \sum_{i=1}^{n} x_{il}\right)$$

3. Extract the First-Order Conditions (FOC) by partially differentiating with respect to the choice variables of goods l and h for a generic individual i:

$$\frac{\partial \mathcal{L}}{\partial x_{il}} = \kappa_i \frac{\partial u_i}{\partial x_{il}} - \mu_l = 0$$
$$\frac{\partial \mathcal{L}}{\partial x_{ih}} = \kappa_i \frac{\partial u_i}{\partial x_{ih}} - \mu_h = 0$$

4. Isolate the marginal utility terms by basic algebraic manipulation:

$$\frac{\partial u_i}{\partial x_{il}} = \frac{\mu_l}{\kappa_i} \quad \text{and} \quad \frac{\partial u_i}{\partial x_{ih}} = \frac{\mu_h}{\kappa_i}$$

5. Dividing the two expressions to eliminate the individual Lagrangian multiplicative multiplier κ_i:

$$\frac{\partial u_i / \partial x_{il}}{\partial u_i / \partial x_{ih}} = \frac{\mu_l / \kappa_i}{\mu_h / \kappa_i} = \frac{\mu_l}{\mu_h}$$

6. The term on the left side of the equation represents, by microeconomic definition, the Marginal Rate of Substitution of individual i between goods l and h (MRS^i_{lh}).
7. **Conclusion of Theorem 2:** Since the shadow price ratio on the right side ($\frac{\mu_l}{\mu_h}$) is independent of the identity of agent i, the equilibrium requires the simultaneous equalization of exchange rates for the entire population:

$$MRS^1_{lh} = MRS^2_{lh} = \cdots = MRS^n_{lh} = \frac{\mu_l}{\mu_h}$$

■

4.4.6 Formal Connection: Intuition → Proof

Mechanism from Intuition	Corresponding Variable/Condition in Proof	Reference in Proof
Social Welfare Menu	Frontier of the geometric set $\mathcal{U}$	Section 4.4.4 - Subsection I
Ethical Weighting	Multiplicative weight parameter α_h	Theorem 1, Step 1
Allocative Efficiency	Deduction of contradiction $W(u(x')) > W(u(x^*))$	Theorem 1, Step 6
Exchange Equilibrium	Ratio of marginal utilities equals $\frac{\mu_l}{\mu_h}$	Theorem 2, Step 5
Physical Scarcity of Goods	Resource Lagrange multiplier μ_l	Theorem 2, Step 2

4.5 Maximization of Linear SWF and Wealth Allocation

4.5.1 Motivation

The maximization of the linear social welfare function under wealth constraint formally describes the process by which a benevolent centralized planner operationalizes distributive criteria. The planner's choice problem occurs when the economy is already positioned on the Pareto Frontier. Since all points on the frontier eliminate physical waste, the planner needs an ethical indexer to select the ideal distribution of monetary wealth among citizens.

Practical Analogy: Imagine a father with a fixed amount of monetary resources to distribute as an allowance among his children. He knows that any totalizing division avoids wasting money (allocative efficiency). If the father values everyone's welfare symmetrically, the weights in his linear function will be identical. However, if one of the children faces unforeseen medical expenses and needs greater support, the father alters the ethical weights of the function to channel a larger fraction of income to this child, maximizing the family's consolidated welfare.

4.5.2 Mechanisms

The model operates by integrating two optimization behaviors:

- **Individuals:** Privately maximize their indirect utilities $v_i(p, w_i)$, reacting to market prices p and the wealth endowment w_i received.
- **Social Planner:** Solves the optimization problem by allocating the total budget through the equalization of marginal utilities of wealth weighted by ethical weights.

The macroeconomic logic dictates that to achieve maximum linear social welfare, the planner must transfer portions of wealth to agents who demonstrate greater capacity to convert monetary units into weighted social utility. Planning equilibrium is achieved when the last unit of currency distributed generates exactly the same social welfare impact across all population axes.

4.5.3 Implications

This instrument offers the framework for structuring tax systems and redistributive public policies, such as the design of progressive income taxes and focused income transfer programs. The model analytically demonstrates that government choices about equity can be integrated into the competitive price system, validating the Second Theorem of Welfare Economics.

Concrete Example: In formulating a progressive tax reform, the State's action is equivalent to assuming that the marginal utility of wealth of lower-income citizens has a higher ethical weight (α_i) in the calculation of the social objective function. By taxing high-income strata and redistributing resources via direct transfers, the planner seeks to equalize the marginal social gain obtained at the base of the pyramid with the marginal social cost generated at the top.

4.5.4 Formal Definition

Consider an exchange economy with I individuals, stable market prices given by vector $p \in \mathbb{R}^L$, and aggregate wealth endowment equal to w_{total}.

- **Individual Indirect Utility:** $v_i(p, w_i)$ represents the maximum utility achieved by agent i when possessing an individual wealth level w_i.
- **Aggregator Mechanism (Linear SWF):** The Bergson–Samuelson linear function is structured as:
$$W(w_1, \dots, w_I) = \sum_{i=1}^{I} \alpha_i v_i(p, w_i), \quad \text{where } \alpha_i > 0 \quad \forall i \in \{1, \dots, I\}$$
- **Central Budget Constraint:** The resource distribution is limited by the economy's global endowment:
$$\sum_{i=1}^{I} w_i = w_{\text{total}}$$

What breaks formally: If the individual utility functions are not strictly concave, the Utility Possibility Set loses the convexity property. Under non-convexity, the linear optimization algorithm based on Lagrange multipliers may present discontinuities (jumps), making it impossible for certain Pareto-efficient allocations to be achieved through the use of fixed linear weights.

4.5.5 Proofs

Proof 1: First-Order Conditions and the Equality of Social Marginal Utility of Wealth

1. The benevolent social planner solves the wealth allocation problem to maximize the indicator W:

$$\max_{w_1,\dots,w_I} \sum_{i=1}^{I} \alpha_i v_i(p, w_i) \quad \text{s.t.} \quad \sum_{i=1}^{I} w_i = w_{\text{total}}$$

2. Structure the classical problem through the Lagrangian operator method, defining μ as the multiplier associated with the global budget endowment constraint:

$$\mathcal{L}(w_1, \dots, w_I, \mu) = \sum_{i=1}^{I} \alpha_i v_i(p, w_i) + \mu \left(w_{\text{total}} - \sum_{i=1}^{I} w_i \right)$$

3. Determine the First-Order Conditions (FOC) by calculating the partial derivative of $\mathcal{L}$ with respect to the wealth allocation variable w_i for each individual in the economy:

$$\frac{\partial \mathcal{L}}{\partial w_i} = \alpha_i \frac{\partial v_i(p, w_i)}{\partial w_i} - \mu = 0 \quad \forall i \in \{1, \dots, I\}$$

4. Substitute the term of the partial derivative of indirect utility with respect to wealth ($\frac{\partial v_i}{\partial w_i}$) with its corresponding microeconomic concept: the individual marginal utility of wealth, denoted by λ_i:

$$\alpha_i \lambda_i - \mu = 0 \implies \alpha_i \lambda_i = \mu$$

5. Since the Lagrange multiplier μ is an identical constant for all equations of the FOC system, apply the transitive equality property to consolidate market equilibrium:

$$\alpha_1 \lambda_1 = \alpha_2 \lambda_2 = \dots = \alpha_I \lambda_I = \mu$$

6. **Conclusion of Proof 1:** At the social optimum point, the marginal utility of wealth of all individuals, when weighted by the respective ethical weight assigned by the planner ($\alpha_i \lambda_i$), must be rigorously identical. Isolating the ethical parameter, we deduce that:

$$\alpha_i = \frac{\mu}{\lambda_i}$$

   This formally proves that the planner's ethical weights must be inversely proportional to the individual marginal utilities of wealth at the equilibrium point. ■

4.5.6 Formal Connection: Intuition → Proof

Mechanism from Intuition	Corresponding Variable/Condition in Proof	Reference in Proof
Weights of Ethical Importance	Multiplier parameter α_i	Proof 1, Steps 1 and 4
Redistribution Instrument	Endogenous wealth choice variable w_i	Proof 1, Steps 1 and 3
Shadow Price of Budget	Institutional Lagrange multiplier μ	Proof 1, Steps 2 and 4
Marginal Welfare Impact	Marginal utility of income variable λ_i	Proof 1, Steps 4 and 5
Distributive Equilibrium	Weighted equality condition $\alpha_i \lambda_i = \mu$	Proof 1, Steps 5 and 6

4.6 Summary of Key Results

Concept	Definition	Key Formula
Individual Preference	Complete and transitive binary relation	$x \succeq_i y \iff u_i(x) \geq u_i(y)$
Social Welfare Function	Maps utility profiles to social scalar	$W(u_1, \ldots, u_I)$
Pareto Efficiency	No one can be made better off without harming another	$u_i(x') \geq u_i(x^*)$, with at least one strict
Condorcet Paradox	Majority rule can generate cycles	xPy, yPz, zPx simultaneously
Arrow's Impossibility	No SWF satisfies U, WP, IIA, and D	Any SWF with $\|X\| \geq 3$ is dictatorial
Decisive Coalition	Group whose unanimous preference determines social choice	$D_G(x, y)$
Utility Possibility Set	All feasible utility combinations	$\mathcal{U} = \{v \in \mathbb{R}^H : v_h \leq u_h(x)\}$
Linear SWF	Weighted sum of utilities	$W = \sum_{h=1}^{H} \alpha_h u_h$
MRS Equalization	Efficiency condition	$MRS^1_{lh} = \cdots = MRS^n_{lh} = \frac{\mu_l}{\mu_h}$
Optimal Wealth Allocation	Weighted marginal utilities equalized	$\alpha_i \lambda_i = \mu$ for all i

---

Chapter 5: Externalities and Market Failures

5.1 Missing Markets and Production Externalities

5.1.1 Motivation

The central problem this topic addresses is the **coordination failure** when one agent's actions affect another's welfare or production possibilities without any price-mediated payment or compensation. In a world of complete markets, everything that affects agents' welfare would have a price. If a factory pollutes the air an individual breathes, that individual would buy the right to clean air or the firm would buy from him the right to pollute. The impasse occurs because, in practice, many of these assets (like clean air or silence) lack markets. When a market for a side effect does not exist, we call it a **missing market**.

Practical Analogy: Imagine an all-you-can-eat buffet where customers do not pay for waste, only for entry. If a customer takes much more food than they can eat and throws it away, they generate a cost for the restaurant and potentially increase the price for future customers or reduce the quality of available food. Since there is no price charged per gram wasted (missing market), the social cost of waste does not enter the individual's private decision.

5.1.2 Mechanisms

- **Firms and Consumers:** Behave as maximizers, but only look at prevailing market prices. If the firm does not pay for the damage it causes (or does not receive for the benefit it generates for third parties), it simply ignores this effect in its profit calculation.
- **Economic Logic:** The price mechanism serves to equalize marginal benefit to marginal cost. In the presence of externalities, the Private Marginal Cost (what the firm actually spends) diverges from the Social Marginal Cost (the real cost to society, including the damage generated to third parties).

5.1.3 Implications

The concept illuminates that total market freedom does not always lead to the best social outcome. The main implication is **allocative inefficiency**: if the externality is negative (pollution), there will be overproduction of the good; if it is positive (research and development that benefits other firms), there will be underproduction.

Concrete Example: The carbon credit market. Previously, companies could emit greenhouse gases at no direct cost (missing market). By creating the carbon credit, the government designs and implements the missing market. Now, emitting pollution has a market price, forcing firms to internalize the social cost in their private balance sheets.

5.1.4 Formal Definition

Let there be an economy with two firms, 1 and 2. Firm 1 produces good q_1 and, in this process, generates a side effect that affects Firm 2.

- **Variables and Parameters:**
  - $q_1, q_2 \in [0, \infty)$: Quantities produced by firms 1 and 2.
  - $p_1, p_2 > 0$: Market prices of the goods, taken as given (competitive environment).
  - $c_1(q_1)$: Cost function of Firm 1.
  - $c_2(q_2, q_1)$: Cost function of Firm 2, which explicitly depends on the production of Firm 1 (production externality).
- **Hypotheses:**
  1. **Differentiability:** The cost functions are continuous and twice continuously differentiable.
  2. **Convexity:** c_i is strictly convex in its own quantity ($\frac{\partial^2 c_i}{\partial q_i^2} > 0$) to guarantee the existence of a unique global maximum.
  3. **Negative Externality:** $\frac{\partial c_2}{\partial q_1} > 0$ (an increase in q_1 raises Firm 2's total cost).
- **What "breaks" formally:** If we remove the dependence of c_2 on q_1, i.e., make $\frac{\partial c_2}{\partial q_1} = 0$, the externality disappears and the conditions of the First Welfare Theorem are restored, guaranteeing allocative efficiency via competitive prices.

5.1.5 Proofs

Property 1: Divergence between Private Equilibrium and Social Optimum A. Market Equilibrium (Firm 1's Private Decision)

1. Firm 1 maximizes its profit π_1 ignoring the effect on Firm 2:

$$\pi_1(q_1) = p_1 q_1 - c_1(q_1)$$

2. Taking the first-order condition (FOC) for the maximum:

$$\frac{d\pi_1}{dq_1} = p_1 - \frac{dc_1(q_1)}{dq_1} = 0$$

3. Resulting in:

$$p_1 = c_1'(q_1)$$

B. Social Optimum (Benevolent Planner's Decision)

1. The planner maximizes the sum of the economy's profits (Social Welfare):

$$W(q_1, q_2) = [p_1 q_1 - c_1(q_1)] + [p_2 q_2 - c_2(q_2, q_1)]$$

2. Taking the necessary first-order optimality condition for variable q_1:

$$\frac{\partial W}{\partial q_1} = p_1 - c_1'(q_1) - \frac{\partial c_2(q_2, q_1)}{\partial q_1} = 0$$

3. Isolating the price:

$$p_1 = c_1'(q_1) + \frac{\partial c_2}{\partial q_1}$$

4. Since $\frac{\partial c_2}{\partial q_1} > 0$ by hypothesis, then $p_1 > c_1'(q_1)$ at the social optimum.

   Conclusion: Comparing the private result with the social one, we see that in private equilibrium Firm 1 produces at a point where $c_1'(q_1)$ is smaller than at the social optimum. Given that marginal cost is strictly increasing ($c_1'' > 0$), this rigorously proves that $q_1^{private} > q_1^{social}$, confirming the overproduction of the polluting good. ■

5.1.6 Formal Connection: Intuition → Proof

Mechanism from Intuition	Corresponding Variable/Condition in Proof	Step/Equation
Missing Market (No price for pollution)	Absence of revenue/cost term for q_1 in Firm 2	Section 5.1.4
Private Profit Maximization	Firm 1 equates $p_1 = c_1'(q_1)$	Property 1, A.3
Social Marginal Cost	Sum $c_1'(q_1) + \frac{\partial c_2}{\partial q_1}$	Property 1, B.3
Inefficiency (Overproduction)	Result $q_1^{private} > q_1^{social}$	Property 1, Conclusion
Production Externality	Term $\frac{\partial c_2}{\partial q_1} \neq 0$	Section 5.1.4 and Property 1, B.2

5.2 The Centralized Correction Mechanism: Pigouvian Tax

5.2.1 Motivation

Imagine that every economic transaction occurs within a delimited space called the market, where everything produced or consumed carries a price that reflects its total opportunity cost to society. The real economic problem arises when one agent's actions spill over outside this space, affecting third parties without any corresponding payment.

We say the market is absent because there is no pricing for this side effect. In the absence of price, the generating agent treats the impact as if it cost zero (if it is damage) or were worth zero (if it is a benefit), misaligning private equilibrium with social welfare.

Practical Analogy: Think of an apartment dweller who decides to play drums at midnight. He reaps the full benefit of leisure, but the marginal cost (the neighbors' loss of sleep) does not enter his personal optimization calculation, as there is no structured market where neighbors could sell their right to silence to him.

5.2.2 Mechanisms

- **Economic Agents:** In the absence of intervention, firms and consumers maximize their own objectives based strictly on prevailing prices. If a factory pollutes a river to produce steel, it ignores the damage caused to fishermen downstream because the use of the river as a waste disposal site constitutes an absent market.
- **Economic Logic:** The fundamental consequence is that the Competitive Equilibrium ceases to be Pareto-efficient. The price system fails to signal the real scarcity of affected resources. The Pigouvian tax acts as an artificial price set by the government to replace the missing market, forcing the agent to internalize the costs of the externality.
- **Government:** Intervenes by taxing the activity to restore the exact equality between private marginal cost and social marginal cost.

5.2.3 Implications

In practice, this shifts public policy formulation from a purely prohibitive approach ("ban pollution") to an incentive-based view ("pollute only if the marginal value of your product exceeds the marginal cost of environmental damage").

Concrete Example: Carbon emission taxation. Instead of the State dictating the operational technologies of each factory, it imposes a fixed rate per ton of CO2 emitted. This mimics a virtual market for clean air, where only firms that value production above the social cost of carbon will continue emitting, optimizing aggregate resource allocation.

5.2.4 Formal Definition

Consider an economy with two agents. Agent 1 chooses an activity level $h \in \mathbb{R}_+$ that directly affects Agent 2's welfare.

- **Variables and Parameters:**
  - $\pi(h)$: Profit or utility function of Agent 1.
  - $\phi(h)$: Damage function (negative externality) suffered by Agent 2.
  - t: Pigouvian tax rate charged per unit of activity h.
- **Hypotheses:**
  1. **Strict Concavity:** The function $\pi(h)$ is assumed strictly concave ($\pi''(h) < 0$) and continuously differentiable.
  2. **Increasing Marginal Damage:** The damage function satisfies $\phi'(h) > 0$ and $\phi''(h) > 0$.
  3. **Absent Markets:** There is no market arrangement by which Agent 2 can charge Agent 1 for the level of h.
  4. **Planner's Information:** The government has perfect information about the marginal benefit and damage functions.
  5. **Interior Solutions:** We assume the optimal level occurs at $h > 0$.
- **What "breaks" without the hypotheses:** If property rights were well-defined and transaction costs were zero, the assumptions of the Coase Theorem would be satisfied. The market would cease to be absent and private bargaining would achieve Pareto efficiency without the need for government taxation.

5.2.5 Proofs

Property 1: The Market Equilibrium is Inefficient

1. Agent 1 maximizes its private profit:
$$\max_h \pi(h)$$
2. The first-order condition for the maximum requires:
$$\frac{d\pi(h)}{dh} = 0$$
Let h^* be the unique solution to this equation.
3. Social Welfare (W) is defined as the sum of benefits minus damages generated:
$$W(h) = \pi(h) - \phi(h)$$
4. Differentiating W with respect to h:
$$\frac{dW}{dh} = \frac{d\pi}{dh} - \frac{d\phi}{dh}$$
5. Evaluating the social impact at the private equilibrium point h^*:
$$\frac{dW(h^*)}{dh} = 0 - \frac{d\phi(h^*)}{dh} < 0$$

(substituting the private FOC and applying the hypothesis $\phi' > 0$)

Conclusion: Since the derivative of social welfare is strictly negative at h^*, a marginal reduction in h will cause a net gain in social welfare. Hence, the private level h^* is socially excessive and inefficient. ■

Property 2: The Pigouvian Tax Restores Efficiency

Auxiliary Theorem (Pigou's Golden Rule): The optimal tax t must be exactly equal to the marginal damage evaluated at the social optimum point.

1. The Social Planner defines the optimal level h^o by maximizing social welfare:

$$\max_h [\pi(h) - \phi(h)]$$

2. The planner's FOC establishes:

$$\pi'(h^o) - \phi'(h^o) = 0 \implies \pi'(h^o) = \phi'(h^o)$$

3. The government introduces a tax t on Agent 1. The firm's new maximization problem becomes:

$$\max_h [\pi(h) - t \cdot h]$$

4. FOC of the firm internalizing the tax:

$$\pi'(h) - t = 0 \implies \pi'(h) = t$$

5. For the decentralized firm's choice (h) to coincide strictly with the social optimum (h^o), the first-order conditions must be functionally identical.
6. Equating both FOCs, we deduce the optimal tax rate:

$$t^* = \phi'(h^o)$$

■

5.2.6 Formal Connection: Intuition → Proof

Mechanism from Intuition	Corresponding Variable/Condition in Proof	Reference in Proof
Private Profit Maximization	$\max \pi(h)$	Property 1, Step 1
Absent Market (ignores damage)	Absence of $\phi(h)$ in the firm's problem	Property 1, Step 2
Internalization of Externality	Term $-t \cdot h$ in firm's profit	Property 2, Step 3
Price of Marginal Damage	$t = \phi'(h^o)$	Property 2, Step 6
Pareto Efficiency	$\pi'(h^o) - \phi'(h^o) = 0$	Property 2, Step 2

5.3 Decentralized Institutional Solutions: The Coase Theorem

5.3.1 Motivation

The central problem this topic addresses is the externality, which creates a structural mismatch between private cost and social cost. The essence of Ronald Coase's thinking is that externalities are not necessarily market failures inherent to technology or goods, but rather the direct result of **poorly defined property**

rights. If it is not legally clear who owns common assets like clean air or silence, the price system loses its pricing capacity.

The Coase Theorem states that if property rights are well-defined by legal institutions and **transaction costs are zero**, agents will bargain privately to achieve an efficient allocative outcome, **regardless of who initially received the legal right**.

Practical Analogy: Imagine a doctor and a baker who share a commercial wall. The noise from the baker's machines prevents the doctor from properly auscultating his patients' hearts. If the law grants the doctor the "right to silence," the baker can financially compensate the doctor to let him work (if the bakery's profit exceeds this compensation). If the law grants the baker the "right to noise," the doctor can pay the baker to change his hours or install soundproofing (if the medical practice is more profitable). In the end, the decision to produce noise will depend on which activity generates greater aggregate social value, and not on which agent the law initially favored.

5.3.2 Mechanisms

- **Agent Behavior:** Agents act as maximizers of economic surplus. They perceive that there is an unexploited "gain from trade." If the polluter values production at \$100 and the victim suffers damage valued at \$60, there is a social surplus of \$40 to be appropriated through contractual bargaining.
- **Economic Logic:** The fundamental property here is the **neutrality of resource allocation**. Property rights define only the direction of payment flows (who compensates whom), but do not alter the final level of real economic activity (pollution or noise), which will always be at the level that maximizes the sum of utilities, given that income does not affect marginal demands (absence of wealth effect).

5.3.3 Implications

In practice, the concept shifts the focus of government regulatory action: the State's primary role ceases to be direct taxation or prohibition and becomes the design of clear, swift property rights and the active reduction of transaction costs (search costs, legal negotiation, and contract enforcement costs).

Concrete Example: The creation of decentralized carbon credit markets. Instead of the government setting restrictive operational caps for each industrial plant, it establishes tradable emission rights. Companies that efficiently reduce pollution sell their legal surplus rights to plants with prohibitive mitigation costs. Efficiency arises from the direct negotiation of these property endowments, exactly as modeled by Coase.

5.3.4 Formal Definition

Consider two economic agents, indexed by $h = 1$ (polluting agent) and $h = 2$ (victim agent).

- **Variables and Parameters:**
  - $x \in \mathbb{R}_+$: Level of activity generating the externality (e.g., emission level).
  - $v_1(x)$: Utility or profit function of Agent 1.
  - $v_2(x)$: Utility function of Agent 2.
  - m_i: Initial monetary wealth endowment (*numeraire*).
  - T: Compensatory monetary transfer made between agents.
- **Hypotheses:**
  1. **Strict Convexity/Concavity:** The functions satisfy $v_1''(x) < 0$, $v_2'(x) < 0$ (characterizing the negative externality), and $v_2''(x) < 0$.
  2. **Quasilinear Preferences:** The total utility function takes the form $U_i(x, m_i) = v_i(x) + m_i$, which neutralizes any income effects on the optimal demand for the externality.
  3. **Zero Transaction Costs:** Agents can draft, monitor, and enforce contracts at zero cost.
  4. **Perfect Information:** Both parties have complete knowledge of the structural functions $v_i(x)$ involved.

- **What "breaks" formally:** If we introduce information asymmetry (Agent 2 privately hiding their real level of damage suffered), the Myerson-Satterthwaite Theorem formally proves that no decentralized voluntary exchange mechanism can simultaneously guarantee Pareto efficiency and individual rationality of the parties.

5.3.5 Proofs

Property 1: The Social Optimum (x^o)

1. The Social Planner maximizes the simple sum of the quasilinear utility functions:

$$\max_x W(x) = v_1(x) + v_2(x)$$

2. The necessary first-order condition for the global maximum determines:

$$\frac{dW}{dx} = v_1'(x) + v_2'(x) = 0$$

3. Resulting in the following efficient allocation rule:

$$v_1'(x^o) = -v_2'(x^o)$$

Property 2: Independence of Rights Allocation (Coase Theorem) Case A: Property Right Belongs to the Victim (Agent 2)

The victim has legal backing to demand $x = 0$. To exercise any level of polluting activity, Agent 1 must transfer a monetary compensation T to Agent 2.

1. Agent 1 models its problem by proposing values of x and T:

$$\max_{x,T} [v_1(x) - T]$$

2. Subject to Agent 2's Participation Constraint (PC), given that it holds the legal reservation point $x = 0$:

$$v_2(x) + T \geq v_2(0)$$

3. At the bargaining optimum, the participation constraint operates strictly actively:

$$T = v_2(0) - v_2(x)$$

4. Substituting the function T directly into Agent 1's objective function:

$$\max_x [v_1(x) - (v_2(0) - v_2(x))]$$

5. Differentiating the final expression with respect to the choice variable x:

$$v_1'(x^A) - (-v_2'(x^A)) = 0 \implies v_1'(x^A) + v_2'(x^A) = 0 \implies x^A = x^o$$

Case B: Property Right Belongs to the Polluter (Agent 1)

The polluter has legal support to pollute up to its maximum pure private optimization level, denoted by x_{max} (the point where $v_1'(x_{max}) = 0$). The victim (Agent 2) offers a payment T to induce him to reduce x.

1. Agent 2 models its problem by choosing the activity scale and the compensation offer:

$$\max_{x,T} [v_2(x) - T]$$

2. Subject to Agent 1's Participation Constraint (PC), whose reservation utility is assured at x_{max}:

$$v_1(x) + T \geq v_1(x_{max})$$

3. In perfect bargaining equilibrium, the constraint assumes equality:

$$T = v_1(x_{max}) - v_1(x)$$

4. Substituting the constraint T into Agent 2's utility function under optimization:

$$\max_x [v_2(x) - (v_1(x_{max}) - v_1(x))]$$

5. Differentiating the equation with respect to the control variable x:

$$v_2'(x^B) - (-v_1'(x^B)) = 0 \implies v_2'(x^B) + v_1'(x^B) = 0 \implies x^B = x^o$$

Conclusion of the Proof: The comparison of institutional equilibria analytically demonstrates that the initial legal allocation of property endowments does not alter the final equilibrium level of the economy's real activity ($x^A = x^B = x^o$). The allocation of rights impacts exclusively the final monetary wealth balances (m_i) through the compensatory transfer vector T. ■

5.3.6 Formal Connection: Intuition → Proof

Mechanism from Intuition	Corresponding Variable/Condition in Proof	Reference in Proof
Maximization of Total Value	$\max W(x) = v_1(x) + v_2(x)$	Social Optimum, Step 1
Bargaining / Compensation	Monetary transfer T	Case A and B, Step 1
Veto Power (Reservation Right)	Reservation value in constraint ($v_2(0)$ or $v_1(x_{max})$)	Case A and B, Step 2
Internalization of Damage	Inclusion of $v_2(x)$ in the profit objective after substitution	Case A, Step 4
Independence of Efficiency	Final result $x = x^o$ in both arrangements	Case A and B, Step 5

5.4 Common Property Goods and the Tragedy of the Commons

5.4.1 Motivation

The *commons* problem refers to the conflict over scarce resources that arises from the inevitable tension between individual selfish incentives and collective welfare. The central challenge this topic addresses is the chronic overuse of public or open-access resources, where the absence of well-defined property rights prevents the real cost of resource depletion from being charged or internalized by those who use it.

Practical Analogy: Imagine a pasture shared by several village herders. Each herder seeks to maximize their yield by increasing the size of their herd. For the herder, adding a goat to the herd brings the full benefit of the animal's sale. However, the cost—the degradation of pasture quality—is a burden distributed among all herders in the village. Since the individual retains all the gain but shares only a fraction of the loss, the incentive is to expand the herd until the pasture is destroyed, harming the collective.

5.4.2 Mechanisms

- **Agent Behavior:** Economic agents (farmers, fishermen, firms) act non-cooperatively, responding strictly to private incentives. They seek to maximize their own utility or individual profit, ignoring the negative effect (negative externality) that their extraction actions impose on other users of the resource.
- **Economic Logic:** The fundamental property is that market equilibrium (or Nash equilibrium) does not guarantee Pareto optimality. Since each agent shares the cost of degradation with the group but retains the marginal benefit of use, the common resource is exploited beyond the level that would be socially desirable or sustainable.

5.4.3 Implications

This concept refutes the intuition that "more market freedom" always leads to allocative efficiency, evidencing the critical need for government intervention or institutional arrangements (such as privatization or regulation) to manage common-access resources.

Concrete Example: The management of fish stocks in the oceans. Without a regulatory authority, each fishing vessel has a dominant incentive to catch as much as possible in the present. The aggregate result is population collapse of the species. Economic policy based on this theory introduces catch quotas or fishing licenses, forcing fishermen to internalize the marginal cost of reducing the fish population for future generations.

5.4.4 Formal Definition

Consider a village with n farmers who take goats to graze on a common pasture.

- **Variables and Parameters:**
  - $g_i \in [0, \infty)$: Number of goats of farmer i.
  - $G = \sum_{j=1}^{n} g_j$: Total goats in the village.
  - $v(G)$: Value to the farmer of grazing a goat when the aggregate level is G.
  - c: Cost of acquisition and care per unit of goat (constant).
- **Hypotheses:**
  1. **Resource Saturation:** $v'(G) < 0$ for $G < G_{max}$ (the value per goat falls as the pasture becomes denser).
  2. **Increasing Marginal Degradation:** $v''(G) < 0$ (diminishing marginal returns in pasture productivity).
  3. **Commons Collapse:** $v(G) = 0$ for $G \geq G_{max}$ (exhausted pasture).
  4. **Game Structure:** Strictly rational and non-cooperative behavior, resulting in a Nash Equilibrium.
- **What "breaks" formally:** If property rights were privatized (dividing the pasture into n exclusive and excludable parcels), each farmer would become a monopolist over their own resource, internalizing the term $g_i v'(G)$ and achieving allocative efficiency automatically.

5.4.5 Proofs

Property 1: The Equilibrium Utilization Level (G^*) is Inefficient ($G^* > G^{**}$)

1. Each farmer i solves their profit maximization problem in a decentralized manner:

$$\max_{g_i} \pi_i(g_i, g_{-i}) = g_i \cdot v(g_i + g_{-i}) - c \cdot g_i$$

2. Taking the First-Order Condition (FOC) for farmer i:

$$\frac{\partial \pi_i}{\partial g_i} = v(G) + g_i v'(G) - c = 0$$

(by the product rule)

3. Assuming a symmetric Nash equilibrium where $g_i = g^*$ for all i, we have $G^* = ng^*$, or $g_i = \frac{G^*}{n}$. Substituting into the FOC:
$$v(G^*) + \frac{G^*}{n}v'(G^*) - c = 0$$
4. The Social Planner maximizes the aggregate surplus of the economy (Social Welfare):
$$\max_G W(G) = G \cdot v(G) - c \cdot G$$
5. The social planner's FOC for the social optimum (G^{**}) is:
$$v(G^{**}) + G^{**}v'(G^{**}) - c = 0$$
(by the product rule)
6. Comparing (3) and (5), note that the term distorting the private FOC is $\frac{G^*}{n}v'(G^*)$, while in the social optimum it is $G^{**}v'(G^{**})$.
7. **Conclusion:** Since $n > 1$, the weight given to the negative effect (marginal damage) by the individual ($\frac{1}{n}$) is strictly smaller than the weight considered by the social planner (1). Given that $v'(G) < 0$, the farmer underestimates the negative impact of their activity. This formally proves that $G^* > G^{**}$ (overuse). ■

Property 2: The Tragedy Increases with the Number of Agents (n)

1. We retake the FOC of the symmetric Nash Equilibrium:
$$v(G^*) + \frac{G^*}{n}v'(G^*) - c = 0$$
2. We apply implicit differentiation with respect to n to find the behavior of G^*:
$$\left[v'(G^*)\frac{dG^*}{dn}\right] + \left[\frac{1}{n}v'(G^*)\frac{dG^*}{dn} + \frac{G^*}{n}v''(G^*)\frac{dG^*}{dn} - \frac{G^*}{n^2}v'(G^*)\right] = 0$$
(by the chain rule and quotient rule)
3. Grouping the terms containing $\frac{dG^*}{dn}$:
$$\frac{dG^*}{dn}\left[v'(G^*)\left(1 + \frac{1}{n}\right) + \frac{G^*}{n}v''(G^*)\right] = \frac{G^*}{n^2}v'(G^*)$$
(by algebraic manipulation)
4. Analyzing the signs (under the hypothesis of $v' < 0$ and $v'' < 0$):
   - The term in brackets is negative.
   - The term on the right-hand side is negative.
5. **Conclusion:** Therefore, isolating the gradient:
$$\frac{dG^*}{dn} = \frac{\frac{G^*}{n^2}v'(G^*)}{v'(G^*)(1 + 1/n) + \frac{G^*}{n}v''(G^*)} > 0$$
It is analytically proven that the aggregate level of overuse grows monotonically with the number of agents. ■

5.4.6 Formal Connection: Intuition → Proof

Mechanism from Intuition	Corresponding Variable/Condition in Proof	Reference in Proof
Private Incentive (Gain)	Own revenue term $g_i \cdot v(G)$	Property 1, Step 1
Social Cost (Degradation)	Hypothesis $v'(G) < 0$	Property 1, Step 2
Ignored Externality	Fractional factor $\frac{1}{n}$ in the individual FOC	Property 1, Step 3
Free-Rider Logic	Divergence between Equations (3) and (5)	Property 1, Step 6
Tragedy (Overuse)	Inequality result $G^* > G^{**}$	Property 1, Conclusion

---

5.5 Summary of Key Results

Concept	Definition	Key Formula
Production Externality	One firm's output affects another's cost	$\frac{\partial c_2}{\partial q_1} \neq 0$
Private Equilibrium	Firm ignores externality	$p_1 = c_1'(q_1)$
Social Optimum	Planner internalizes externality	$p_1 = c_1'(q_1) + \frac{\partial c_2}{\partial q_1}$
Pigouvian Tax	Tax equals marginal damage	$t^* = \phi'(h^o)$
Coase Theorem	Efficiency independent of rights allocation	$x^A = x^B = x^o$ (with zero transaction costs)
Tragedy of the Commons	Overuse of common resources	$G^* > G^{**}$
Nash Equilibrium (Commons)	Individual FOC	$v(G^*) + \frac{G^*}{n} v'(G^*) - c = 0$
Social Optimum (Commons)	Planner's FOC	$v(G^{**}) + G^{**} v'(G^{**}) - c = 0$
Scale Effect	Overuse increases with n	$\frac{dG^*}{dn} > 0$

Chapter 6: Public Goods

6.1 Public Goods Provision and the Samuelson Rule

Public goods theory is one of the most refined topics in microeconomics, where the properties of excludability and rivalry challenge the market's ability to allocate resources efficiently.

6.1.1 Motivation

A public good is defined by two fundamental properties: **non-rivalry** (one person's consumption does not diminish the amount available to another) and **non-excludability** (it is impossible or prohibitively costly to prevent someone from consuming the good). The real economic problem is that these goods generate benefits for everyone, but the cost of producing them falls on specific individuals.

The market fails here because if I cannot be excluded from enjoying the benefit, I have a rational incentive not to pay and wait for others to pay. This is called the **"free-rider" problem**.

Practical Analogy: Imagine a lighthouse on a dangerous coast. Once lit, all passing ships benefit from the light (non-rivalry) and the lighthouse has no way to "turn off the light" only for ships that did not pay the fee (non-excludability). The likely result is that no individual captain will voluntarily pay, and the lighthouse will never be built, despite its total value to the fleet far exceeding the construction cost.

6.1.2 Mechanisms

- **Agent Behavior:** In private provision, each agent decides how much to contribute to the public good based solely on their private marginal benefit. They ignore the benefit their contribution generates for all other members of society.
- **Economic Logic:** Efficiency requires us to take everyone's utility into account. While a private good should be produced until the last consumer's benefit equals the cost ($MB = MC$), the public good should be produced until the sum of all consumers' utilities equals the marginal cost.

6.1.3 Implications

This concept illuminates the need for collective action or state intervention. Since the private market will always produce a suboptimal quantity (or zero), the government intervenes via mandatory taxation to finance efficient provision.

Concrete Example: National Defense. If military protection were sold as a market service, many citizens would not pay, hoping to be protected by the "security bubble" paid for by neighbors. Without tax collection (public provision), the country would be vulnerable, as voluntary revenue would be insufficient to maintain the armed forces.

6.1.4 Formal Definition

Consider an economy with n agents and two goods: a private good y (used as *numeraire*) and a public good x.

- **Variables and Parameters:**
  - $x \in [0, \infty)$: Quantity of the public good.
  - $y_i \in \mathbb{R}$: Quantity of the private good consumed by agent i.
  - $v_i(x)$: Utility agent i obtains from the public good.
  - $C(x)$: Cost function to produce x.
- **Hypotheses:**
  1. **Strict Concavity:** $v_i''(x) < 0$ to guarantee the existence of a maximum.
  2. **Production Technology:** $C'(x) > 0$ and $C''(x) \geq 0$.
  3. **Quasilinear Preferences:** Agent i's utility is given by $U_i(x, y_i) = v_i(x) + y_i$.
  4. **Non-Excludability and Non-Rivalry:** All agents simultaneously consume the same total quantity x.
  5. **Differentiability:** All functions are continuously differentiable.
- **What "breaks" without the hypotheses:** If utility were not quasilinear, efficient provision would depend on the initial wealth distribution (income effects). If the good were rival, the efficiency condition would collapse back to the conventional private good equality ($v_i'(x) = C'(x)$).

6.1.5 Proofs

Property 1: The Samuelson Rule for Efficient Provision (x^o) *Auxiliary Theorem (Sum of Benefits): In an environment with quasilinear preferences, an allocation is Pareto-efficient if and only if it maximizes the economy's total surplus.*

1. The Social Planner maximizes Social Welfare (W), defined as the sum of utilities net of cost:

$$\max_x W(x) = \sum_{i=1}^{n} v_i(x) - C(x)$$

2. Applying the derivative operator with respect to quantity x:

$$\frac{dW(x)}{dx} = \frac{d}{dx}\left[\sum_{i=1}^{n} v_i(x)\right] - \frac{dC(x)}{dx}$$

3. By the linearity of the derivative operator, we obtain the First-Order Condition (FOC):

$$\sum_{i=1}^{n} v_i'(x^o) - C'(x^o) = 0$$

4. Isolating the marginal cost by basic algebra, we arrive at the **Samuelson Rule**:

$$\sum_{i=1}^{n} v_i'(x^o) = C'(x^o)$$

This proves that the social optimum occurs when the sum of all individuals' marginal utilities equals the marginal cost of production. ■

Property 2: The Inefficiency of Private Provision (Nash Equilibrium, x^*)

1. Suppose the marginal cost of x is constant and equal to c. Each agent i chooses their own voluntary contribution $g_i \geq 0$, such that the total quantity provided is $x = g_i + \sum_{j \neq i} g_j$.
2. Agent i maximizes their individual payoff non-cooperatively:

$$\max_{g_i} \pi_i(g_i, g_{-i}) = v_i\left(g_i + \sum_{j \neq i} g_j\right) - c \cdot g_i$$

3. Taking the FOC for an agent making a strictly positive contribution ($g_i > 0$):

$$v_i'(x^*) = c$$

4. Rewriting the Samuelson Rule (Social Optimum x^o) under the same constant cost structure c:

$$\sum_{j=1}^{n} v_j'(x^o) = c$$

5. Since $v_j'(x) > 0$ for all agents, the sum of all marginal utilities is strictly greater than any individual marginal utility evaluated at the same point:

$$\sum_{j=1}^{n} v_j'(x^o) > v_i'(x^o) \implies c > v_i'(x^o)$$

6. For the private FOC from step (3) to be satisfied ($v_i'(x^*) = c$) under the concavity hypothesis ($v_i'' < 0$), the decreasing marginal derivative requires that:

$$x^* < x^o$$

This formally proves that decentralized private provision always results in an equilibrium quantity strictly below the social optimum (underprovision). ■

6.1.6 Formal Connection: Intuition → Proof

Mechanism from Intuition	Corresponding Variable/Condition in Proof	Reference in Proof
Non-Rivalry	Use of the same argument x for all functions v_i	Property 1, Step 1

Mechanism from Intuition	Corresponding Variable/Condition in Proof	Reference in Proof
Free-Rider Logic	Agent decides looking only at individual return $v_i'(x^*) = c$	Property 2, Step 3
Sum of Benefits	The summation operator $\sum v_i'(x^o)$	Property 1, Step 4
Internalization of Externality	Divergence between social FOC $(\sum v_i')$ and private FOC (v_i')	Property 2, Step 5
Market Underprovision	Final inequality demonstrated: $x^* < x^o$	Property 2, Step 6

6.2 Demand Revelation Mechanisms: The VCG Mechanism

Demand revelation and decentralized implementation. We will explore how to design "rules of the game" that force self-interested agents to reveal the truth, allowing society to reach the social optimum even when critical information is dispersed and hidden.

6.2.1 Motivation

The real economic problem this topic addresses is **asymmetric information** in collective decisions. Often, a planner wants to make a decision affecting everyone (e.g., building a bridge), but the value each individual assigns to that decision is private information. If the planner simply asks "how much is this worth to you?", people have strategic incentives to lie: exaggerate the value to ensure the project or underestimate it to avoid paying for it.

The essence of the solution is to create a mechanism where the best individual strategy coincides exactly with honesty. This is called **Incentive Compatibility**.

Practical Analogy: Imagine a group of friends deciding which pizza to order. If the rules say that whoever chooses the most expensive pizza pays the difference alone, no one will choose their favorite just because of the cost. The demand revelation mechanism is like an intelligent "bill-splitting" system where the price you pay depends only on the inconvenience (externality) you cause others by imposing your preference.

6.2.2 Mechanisms

- **Agent Behavior:** Agents are treated as rational utility maximizers in a game environment. They observe the mechanism's rules and calculate which "message" (reported value) maximizes their personal welfare, regardless of what others will do.
- **Economic Logic:** The central property is the **Revelation Principle**, which states that any social outcome achievable by a complex game can be replicated by a "direct revelation" mechanism, where agents simply tell the truth. The VCG (Vickrey-Clarke-Groves) mechanism uses monetary transfers to make each agent internalize the externality that their report imposes on the rest of the group.

6.2.3 Implications

This concept allows modern economies to allocate scarce resources efficiently without the need for an omniscient planner. It replaces bureaucratic control with precise financial incentives.

Concrete Example: Spectrum Auctions. Governments use these mechanisms to sell cell phone licenses. Instead of arbitrarily choosing a company, the government designs a mechanism where the company that truly values the license the most (and is therefore likely the most efficient) wins, paying a value that reflects the opportunity cost to other bidders.

6.2.4 Formal Definition

Let there be an economy with a set of n agents, denoted by $N = \{1, \ldots, n\}$, and a set of discrete or continuous public alternatives X.

- **Variables and Parameters:**
  - $\theta_i \in \Theta_i$: The type (private information) of agent i, representing their true valuation.
  - $u_i(x, \theta_i)$: Utility agent i obtains from alternative $x \in X$ given their true type.
  - $\hat{\theta}_i$: The publicly reported type (message) by agent i.
  - $m_i \in \mathbb{R}$: Monetary transfer received ($m_i > 0$) or paid ($m_i < 0$) by the agent.
- **Hypotheses:**
  1. **Quasilinear Preferences:** Agent i's total utility function is given by $V_i(x, m_i, \theta_i) = u_i(x, \theta_i) + m_i$.
  2. **Independent Private Values:** Agent i's valuation depends only on their own type θ_i.
  3. **No Strict Budget Constraint:** Money acts as linear free compensation.
- **What "breaks" without the hypotheses:** If we remove the quasilinearity hypothesis, the Gibbard-Satterthwaite Theorem proves that any mechanism that is simultaneously efficient and immune to strategic manipulation will necessarily be dictatorial.

6.2.5 Proofs

Property 1: Efficiency of the Social Optimum (x^*)

1. A state is Pareto-efficient in a quasilinear environment if and only if the choice of alternative maximizes the sum of agents' utilities. The planner seeks a choice rule $x^*(\cdot)$ such that:

$$x^*(\theta) \in \arg\max_{x \in X} \sum_{i=1}^{n} u_i(x, \theta_i)$$

2. By definition of the maximum argument, for any other alternative $x' \in X$:

$$\sum_{i=1}^{n} u_i(x^*(\theta), \theta_i) \geq \sum_{i=1}^{n} u_i(x', \theta_i)$$

3. Since the marginal utility of money is constant and equal to 1 for all due to quasilinearity, any redistribution of monetary endowments through transfers where $\sum m_i \leq 0$ will preserve the Pareto efficiency of the choice of x^*. ■

Property 2: Incentive Compatibility in Dominant Strategies of the VCG Mechanism
Theorem (Truthfulness as a Dominant Strategy): In the VCG mechanism, reporting the truth ($\hat{\theta}_i = \theta_i$) is a dominant strategy for every agent i.

1. Define the classic VCG transfer rule for agent i based on the collective reports $\hat{\theta}$:

$$m_i(\hat{\theta}) = \sum_{j \neq i} u_j(x^*(\hat{\theta}), \hat{\theta}_j) + h_i(\hat{\theta}_{-i})$$

   where $h_i(\hat{\theta}_{-i})$ is an arbitrary function that depends exclusively on the reports of other agents, being outside i's control.

2. Agent i chooses their report $\hat{\theta}_i$ to maximize their expected real utility:

$$\max_{\hat{\theta}_i} \left[u_i(x^*(\hat{\theta}_i, \hat{\theta}_{-i}), \theta_i) + m_i(\hat{\theta}_i, \hat{\theta}_{-i}) \right]$$

3. Substituting the VCG mechanism transfer rule established in step (1) directly into the agent's objective function:

$$\max_{\hat{\theta}_i} \left[u_i(x^*(\hat{\theta}_i, \hat{\theta}_{-i}), \theta_i) + \sum_{j \neq i} u_j(x^*(\hat{\theta}_i, \hat{\theta}_{-i}), \hat{\theta}_j) + h_i(\hat{\theta}_{-i}) \right]$$

4. Since the term $h_i(\hat{\theta}_{-i})$ does not contain $\hat{\theta}_i$, agent i has no ability to affect its value through their message. Hence, it is irrelevant to the optimization problem and can be omitted:

$$\max_{\hat{\theta}_i} \left[u_i(x^*(\hat{\theta}_i, \hat{\theta}_{-i}), \theta_i) + \sum_{j \neq i} u_j(x^*(\hat{\theta}_i, \hat{\theta}_{-i}), \hat{\theta}_j) \right]$$

5. Observe the mathematical structure: agent i's objective has become to choose a report $\hat{\theta}_i$ that induces the planner to select an alternative x that maximizes the weighted expression above.
6. By definition (see Property 1), the planner's decision function $x^*(\cdot)$ has already been designed to compute the maximum of the sum of reported utilities:

$$x^*(\hat{\theta}_i, \hat{\theta}_{-i}) \in \arg\max_{x \in X} \left[u_i(x, \hat{\theta}_i) + \sum_{j \neq i} u_j(x, \hat{\theta}_j) \right]$$

7. Comparing the two expressions, if agent i chooses to report the exact truth about their preferences, setting $\hat{\theta}_i = \theta_i$, the planner will solve exactly the problem from step (4).
8. **Conclusion:** Since this identity is independent of the values contained in the vector of other agents $(\hat{\theta}_{-i})$, the strategy of being perfectly honest is optimal under any game scenario. It is proven that truthfulness is a dominant strategy in the VCG mechanism. ■

6.2.6 Formal Connection: Intuition → Proof

Mechanism from Intuition	Corresponding Variable/Condition in Proof	Reference in Proof
Maximization of Welfare	The social choice function $\arg\max_x \sum u_i$	Property 1, Step 1
Private Information (Type)	The hidden exogenous parameter θ_i	Section 6.2.4
Internalization of Externality	The summation of third parties $\sum_{j \neq i} u_j$ injected into the payoff	Property 2, Step 3
Honesty as Strategy	The message sending choice where $\hat{\theta}_i = \theta_i$	Property 2, Step 7
Independence from Others	Validity of maximization for any generic vector $\hat{\theta}_{-i}$	Property 2, Step 8

6.3 Summary of Key Results

Concept	Definition	Key Formula
Public Good	Non-rival and non-excludable good	x consumed equally by all agents
Samuelson Rule	Efficient provision condition	$\sum_{i=1}^n v_i'(x^o) = C'(x^o)$
Private Provision (Nash)	Individual contribution equilibrium	$v_i'(x^*) = c$ (with constant MC)
Free-Rider Problem	Underprovision in private market	$x^* < x^o$
Quasilinear Utility	No income effects on public good demand	$U_i(x, y_i) = v_i(x) + y_i$
VCG Mechanism	Truthful demand revelation	$m_i(\hat{\theta}) = \sum_{j \neq i} u_j(x^*(\hat{\theta}), \hat{\theta}_j) + h_i(\hat{\theta}_{-i})$
Dominant Strategy	Truth-telling is optimal	$\hat{\theta}_i = \theta_i$ for all i

Concept	Definition	Key Formula
Revelation Principle	Any outcome can be replicated by direct revelation	Truthful reporting as equilibrium