Game Theory and Strategic Behavior

Comprehensive Course Notes

Table of Contents

---

1. Preliminaries

1.1 Introduction to Game Theory

Game theory is a mathematical framework for analyzing strategic interactions among multiple decision-makers (players) where the outcome depends on the choices of all players involved. Unlike single-person optimization problems, in game theory, the best decision for one player depends on the decisions made by others.

Definition 1.1 (Game): A game is a formal representation of a situation where:

1. Multiple decision-makers (players) interact
2. Each player has a set of possible actions
3. The outcome of the interaction depends on the actions chosen by all players
4. Each player has preferences over the possible outcomes

Key Characteristics of a Game:

1. Involvement of multiple players, each with individual objectives
2. Rules of the game known in advance (often a common knowledge assumption)
3. Players are assumed to be rational

4. The outcome depends on the joint choices of all involved participants

Applications of Game Theory:

- Economics: Trading, auctions, market behavior, oligopolies, monetary policy
- Engineering: Resource allocation, network formation, security protocols, congestion control
- Computer Science: Distributed algorithms, multi-agent systems, cryptography
- Biology: Evolution, animal behavior, ecological systems

Types of Games:

1. Static vs. Dynamic
2. Complete vs. Incomplete Information
3. Perfect vs. Imperfect Information
4. Zero-sum vs. Non-zero-sum
5. Cooperative vs. Non-cooperative

1.2 Decision Problems

Before diving into multi-player games, we first address single-person decision problems.

Definition 1.2 (Decision Problem): A decision problem consists of:

1. A set of actions A available to the decision-maker
2. A set of outcomes resulting from those actions
3. Preferences over those outcomes

Preferences

Definition 1.3 (Preference Relation): A binary relation $\succsim$ on set A where $a \succsim b$ means "a is preferred to or indifferent to b."

A preference relation $\succsim$ is rational if it satisfies:

1. **Completeness**: For all $a, b \in A$, either $a \succsim b$ or $b \succsim a$ or both
2. **Transitivity**: For all $a, b, c \in A$, if $a \succsim b$ and $b \succsim c$, then $a \succsim c$

From $\succsim$, we can derive:

- Strict preference: $a \succ b$ if $a \succsim b$ and not $b \succsim a$
- Indifference: $a \sim b$ if $a \succsim b$ and $b \succsim a$

Theorem 1.1: A preference relation $\succsim$ on a finite set A can be represented by a utility function $u : A \to \mathbb{R}$ if and only if $\succsim$ is rational.

A utility function u represents preference relation $\succsim$ if for all $a, b \in A$:

$$a \succsim b \iff u(a) \geq u(b)$$

Proof (Sketch):

- ($\Rightarrow$) If u represents $\succsim$, then completeness follows from the completeness of $\geq$ on $\mathbb{R}$. For transitivity, if $a \succsim b$ and $b \succsim c$, then $u(a) \geq u(b)$ and $u(b) \geq u(c)$, which implies $u(a) \geq u(c)$, thus $a \succsim c$.
- ($\Leftarrow$) If $\succsim$ is rational, define $u(a) = |b \in A : a \succsim b|$. This counts the number of elements "dominated by" a. If $a \succsim b$, then any element dominated by b is also dominated by a, so $u(a) \geq u(b)$. Conversely, if $u(a) \geq u(b)$, then either $a \succsim b$ or, by completeness, $b \succ a$. But if $b \succ a$, we would have $u(b) > u(a)$, a contradiction. Thus, $a \succsim b$.

Decision Trees: A graphical representation of sequential decision problems:

- Nodes represent decision points
- Branches represent possible actions
- Leaves represent final outcomes with associated utilities

1.3 Utility, Market, and Discount Factor

Utility Functions

Utilities (also called payoff functions) are an arbitrary quantification $u(q)$ of the goodness coming from some input q.

Properties of Utility Functions:

1. Ordinal representation: The exact values don't matter, only the ordering
2. Cardinal representation: When dealing with uncertainty, the exact values matter

Common Utility Functions:

1. Linear: $u(q) = aq + b$ where $a > 0$ (risk-neutral)
2. Logarithmic: $u(q) = \log(q)$ (risk-averse)
3. Quadratic: $u(q) = q - \alpha q^2$ where $\alpha > 0$ (risk-averse)
4. Power: $u(q) = q^\alpha$ where $0 < \alpha < 1$ (risk-averse) or $\alpha > 1$ (risk-loving)

Market

A market is a mechanism for allocating resources through the interaction of buyers and sellers.

Definition 1.4 (Market Clearing Price): The price at which quantity supplied equals quantity demanded.

Market Structures:

1. Perfect Competition: Many firms, homogeneous products, price-taking behavior
2. Monopoly: Single seller, price-setting behavior
3. Oligopoly: Few sellers, strategic interaction
   - Cournot: Firms compete on quantity
   - Bertrand: Firms compete on price
   - Stackelberg: Sequential quantity choice

Substitute Goods: Two goods x and y are substitutes if an increase in the price of one leads to an increase in demand for the other.

Definition 1.5 (Price Elasticity of Demand): The percentage change in quantity demanded in response to a one percent change in price:

$$\epsilon_d = \frac{\partial Q}{\partial P} \cdot \frac{P}{Q}$$

Discount Factor

Future payoffs are often valued less than immediate payoffs, captured by a discount factor.

Definition 1.6 (Discount Factor): A number $\delta \in (0, 1]$ that represents the present value of one unit of payoff in the next period.

The present value of a payoff x received t periods in the future is $\delta^t x$.

Interpretation of Discount Factor:

1. Time preference: Impatience
2. Risk of termination: Probability that the interaction continues
3. Interest rate: $\delta = \frac{1}{1+r}$ where r is the interest rate

1.4 Lotteries and Expected Utility

When outcomes are uncertain, decisions involve lotteries (probability distributions over outcomes).

Definition 1.7 (Lottery): A lottery p over set $X = x_1, x_2, \ldots, x_n$ is a probability distribution where $p(x_i) \geq 0$ for all i and $\sum_{i=1}^{n} p(x_i) = 1$.

In decision theory, we often represent uncertain events as choices made by "Nature," a non-strategic player that selects outcomes according to known probabilities.

Von Neumann-Morgenstern Expected Utility Theory

Definition 1.8 (von Neumann-Morgenstern Axioms): A preference relation $\succsim$ over lotteries satisfies the von Neumann-Morgenstern (vNM) axioms if:

1. **Completeness and Transitivity**: $\succsim$ is rational
2. **Continuity**: For any lotteries p, q, r with $p \succ q \succ r$, there exists $\alpha \in (0, 1)$ such that $\alpha p + (1 - \alpha)r \sim q$
3. **Independence**: For any lotteries p, q, r and $\alpha \in (0, 1)$, $p \succsim q \iff \alpha p + (1 - \alpha)r \succsim \alpha q + (1 - \alpha)r$

Theorem 1.2 (vNM Utility Representation): If $\succsim$ satisfies the vNM axioms, then there exists a utility function $u : X \to \mathbb{R}$ such that for any lotteries p and q:

$$p \succsim q \iff \mathbb{E}_p[u(x)] \geq \mathbb{E}_q[u(x)]$$

where $\mathbb{E}_p[u(x)] = \sum_{i=1}^{n} p(x_i)u(x_i)$ is the expected utility of lottery p.

Moreover, u is unique up to positive affine transformations.

Proof (Sketch):

1. For each outcome $x \in X$, define the degenerate lottery δ_x that gives x with probability 1.
2. Choose two outcomes x^* and x_* such that $\delta_{x^*} \succsim \delta_x \succsim \delta_{x_*}$ for all $x \in X$.
3. For each outcome x, find the unique $\alpha_x \in [0, 1]$ such that $\delta_x \sim \alpha_x \delta_{x^*} + (1 - \alpha_x)\delta_{x_*}$ (by continuity).
4. Define $u(x) = \alpha_x$.
5. Use the independence axiom to show that $p \succsim q \iff \mathbb{E}_p[u(x)] \geq \mathbb{E}_q[u(x)]$.
6. Uniqueness up to affine transformations follows from the fact that if $v(x) = au(x) + b$ where $a > 0$, then $\mathbb{E}_p[v(x)] = a\mathbb{E}_p[u(x)] + b$, which preserves the ordering.

Compound Lotteries

Definition 1.9 (Compound Lottery): A compound lottery is a lottery over lotteries.

Theorem 1.3 (Reduction of Compound Lotteries): Under the vNM axioms, a decision-maker is indifferent between a compound lottery and its reduced form.

If lottery p gives lottery q_j with probability α_j for $j = 1, \ldots, m$, then p is equivalent to the reduced lottery r where:

$$r(x_i) = \sum_{j=1}^{m} \alpha_j q_j(x_i)$$

The Value of Information

Information has value if it allows better decisions.

Definition 1.10 (Expected Value of Perfect Information): The difference between:

1. The expected utility with perfect information
2. The expected utility without additional information

Mathematically, if X is the set of states of the world with probability distribution p, A is the set of actions, and $u(a, x)$ is the utility of action a in state x, then:

$$EVPI = \sum_{x \in X} p(x) \max_{a \in A} u(a, x) - \max_{a \in A} \sum_{x \in X} p(x) u(a, x)$$

1.5 Risk Attitudes

Decision-makers may have different attitudes toward risk, categorized as:

Definition 1.11 (Risk Attitudes):

1. **Risk Neutral**: Values a lottery at its expected monetary value
   - Utility function is linear: $u(x) = ax + b$ where $a > 0$
   - $u(E[X]) = E[u(X)]$
2. **Risk Averse**: Prefers a certain outcome to a lottery with the same expected value
   - Utility function is concave: $u''(x) < 0$
   - $u(E[X]) > E[u(X)]$
   - Examples: $u(x) = \log(x)$ or $u(x) = \sqrt{x}$
3. **Risk Loving**: Prefers a lottery to a certain outcome with the same expected value
   - Utility function is convex: $u''(x) > 0$
   - $u(E[X]) < E[u(X)]$
   - Example: $u(x) = x^2$ for $x > 0$

Definition 1.12 (Arrow-Pratt Measure of Absolute Risk Aversion):

$$A(x) = -\frac{u''(x)}{u'(x)}$$

A higher value of $A(x)$ indicates greater risk aversion.

Definition 1.13 (Certainty Equivalent): For a lottery p, the certainty equivalent $CE(p)$ is the certain amount that makes the decision-maker indifferent between receiving $CE(p)$ for sure and facing the lottery p:

$$u(CE(p)) = \mathbb{E}_p[u(x)]$$

Definition 1.14 (Risk Premium): The difference between the expected value of a lottery and its certainty equivalent:

$$RP(p) = \mathbb{E}_p[x] - CE(p)$$

2. Static Games of Complete Information

2.1 Normal Form Games

Definition 2.1 (Normal Form Game): A static game of complete information consists of:

1. A set of players $N = 1, 2, \ldots, n$
2. A set of strategies S_i for each player $i \in N$
3. A utility function $u_i : S_1 \times S_2 \times \ldots \times S_n \rightarrow \mathbb{R}$ for each player $i \in N$

We denote such a game as $G = N, (S_i)_{i \in N}, (u_i)_{i \in N}$.

For simplicity, we often use $S = S_1 \times S_2 \times \ldots \times S_n$ to denote the set of strategy profiles.

For a strategy profile $s = (s_1, s_2, \ldots, s_n) \in S$, we use s_{-i} to denote the strategies of all players except i:

$$s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$$

Thus, $s = (s_i, s_{-i})$ for any player i.

Representation for Two-Player Games: For two-player games, we can represent the game as a bi-matrix where:

- Rows correspond to strategies of player 1
- Columns correspond to strategies of player 2
- Each cell contains a pair $(u_1(s_1, s_2), u_2(s_1, s_2))$

Example 2.1 (Prisoner's Dilemma): Two suspects are questioned separately. Each has two choices: cooperate (C) with the other or defect (D). The payoff matrix is:

	C	**D**
C	(-1,-1)	(-9,0)
D	(0,-9)	(-6,-6)

Definition 2.2 (Common Knowledge): A fact is common knowledge if:

1. All players know it
2. All players know that all players know it
3. All players know that all players know that all players know it, and so on ad infinitum

In static games of complete information, we assume that the structure of the game (players, strategies, and payoffs) is common knowledge.

2.2 Dominance and Rationality

A fundamental concept in game theory is the elimination of strategies that a rational player would never play.

Definition 2.3 (Strict Dominance): Strategy $s_i \in S_i$ is strictly dominated by strategy $s_i' \in S_i$ if:

$$u_i(s_i', s_{-i}) > u_i(s_i, s_{-i}) \quad \forall s_{-i} \in S_{-i}$$

Definition 2.4 (Weak Dominance): Strategy $s_i \in S_i$ is weakly dominated by strategy $s_i' \in S_i$ if:

$$u_i(s_i', s_{-i}) \geq u_i(s_i, s_{-i}) \quad \forall s_{-i} \in S_{-i}$$

with strict inequality for at least one s_{-i}.

Definition 2.5 (Dominant Strategy): Strategy $s_i^* \in S_i$ is a dominant strategy for player i if s_i^* strictly dominates all other strategies in S_i. That is, for all $s_i \in S_i$ with $s_i \neq s_i^*$:

$$u_i(s_i^*, s_{-i}) > u_i(s_i, s_{-i}) \quad \forall s_{-i} \in S_{-i}$$

Iterative Elimination Procedures:

Definition 2.6 (Iterated Elimination of Strictly Dominated Strategies, IESDS):

1. For each player, identify and eliminate all strictly dominated strategies
2. In the reduced game, identify and eliminate newly dominated strategies
3. Repeat until no more strategies can be eliminated

Let $S_i^0 = S_i$ for all $i \in N$. For each $k \geq 1$, define:

$$S_i^k = s_i \in S_i^{k-1} : s_i \text{ is not strictly dominated in the game } (S_1^{k-1}, \ldots, S_n^{k-1}, (u_i))$$

The set of strategies surviving IESDS is $S_i^\infty = \cap_{k=0}^\infty S_i^k$.

Theorem 2.1: The order of elimination in IESDS does not affect the final outcome.

Proof (Sketch): If s_i is strictly dominated by some strategy at some stage, it remains strictly dominated in all subsequent stages, regardless of the order of elimination.

Theorem 2.2: If IESDS yields a unique strategy profile $s^* = (s_1^*, s_2^*, \ldots, s_n^*)$, then s^* is the unique rational outcome of the game.

Example 2.2: Consider the following game:

	L	**C**	**R**
U	(3,3)	(5,1)	(8,3)
M	(4,1)	(0,5)	(7,3)
D	(4,0)	(3,7)	(9,9)

For player 1: M is dominated by a mixed strategy of $\frac{1}{2}$U + $\frac{1}{2}$D. For player 2: L is dominated by R.

After elimination:

	C	R
U	(5,1)	(8,3)
D	(3,7)	(9,9)

Now, C is dominated by R for player 2. After elimination:

	R
U	(8,3)
D	(9,9)

Finally, U is dominated by D for player 1, leaving (D, R) as the unique surviving strategy profile.

Definition 2.7 (Rationalizable Strategies): A strategy s_i is rationalizable if it survives the following procedure:

1. For each player, eliminate all strategies that are never best responses to some strategy profile of the other players
2. In the reduced game, eliminate all strategies that are never best responses
3. Repeat until no more strategies can be eliminated

Formally, let $R_i^0 = S_i$ for all $i \in N$. For each $k \geq 1$, define:

$$R_i^k = s_i \in R_i^{k-1} : \exists s_{-i} \in \times_{j \neq i} R_j^{k-1} \text{ such that } s_i \in BR_i(s_{-i})$$

where $BR_i(s_{-i})$ is the set of best responses of player i to s_{-i}. The set of rationalizable strategies is $R_i^\infty = \cap_{k=0}^\infty R_i^k$.

Theorem 2.3: In a finite game, a strategy is rationalizable if and only if it survives iterated elimination of never best responses.

2.3 Best Response and Nash Equilibrium

Definition 2.8 (Best Response): A strategy $s_i^* \in S_i$ is a best response to strategies s_{-i} if:

$$u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i}) \quad \forall s_i \in S_i$$

The set of best responses for player i to s_{-i} is denoted as $BR_i(s_{-i})$:

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

Definition 2.9 (Nash Equilibrium): A strategy profile $s^* = (s_1^*, s_2^*, \dots, s_n^*)$ is a Nash equilibrium if for every player $i \in N$:

$$s_i^* \in BR_i(s_{-i}^*)$$

Equivalently, for all $i \in N$ and all $s_i \in S_i$:

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

Properties of Nash Equilibrium:

1. A Nash equilibrium is self-enforcing: no player has an incentive to deviate unilaterally.
2. Not all games have a pure-strategy Nash equilibrium.
3. A game may have multiple Nash equilibria.
4. Nash equilibria may be Pareto inefficient.

Theorem 2.4: If s^* is the unique strategy profile that survives IESDS, then s^* is a Nash equilibrium.

Proof: If s^* is the unique survivor of IESDS, then s_i^* is the only strategy in S_i^∞ for all i. Since any strategy that is not a best response to some strategy profile is strictly dominated, s_i^* must be a best response to s_{-i}^*. Thus, s^* is a Nash equilibrium.

Theorem 2.5: Every strictly dominated strategy is never a best response to any strategy profile of the other players.

Proof: If s_i is strictly dominated by s_i', then for all s_{-i}, $u_i(s_i', s_{-i}) > u_i(s_i, s_{-i})$. Thus, s_i cannot be a best response to any s_{-i}.

Example 2.3 (Battle of the Sexes): Two players want to meet but have different preferences over two locations. The payoff matrix is:

	Opera (O)	**Football (F)**
Opera (O)	(2,1)	(0,0)
Football (F)	(0,0)	(1,2)

This game has two pure-strategy Nash equilibria: (O, O) and (F, F).

2.4 Applications of Nash Equilibrium

Cournot Duopoly

Two firms (1 and 2) simultaneously choose production quantities q_1 and q_2. The market price is $P(Q) = a - Q$ where $Q = q_1 + q_2$ and $a > 0$. Each firm has constant marginal cost c where $0 < c < a$.

The profit function for firm i is:

$$\pi_i(q_1, q_2) = q_i \cdot (P(Q) - c) = q_i \cdot (a - q_1 - q_2 - c)$$

To find the Nash equilibrium, we compute the best response functions:

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

Solving, we get:

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

The Nash equilibrium is the intersection of the best response functions:

$$q_1^* = q_2^* = \frac{a - c}{3}$$

With profits:

$$\pi_1^* = \pi_2^* = \frac{(a - c)^2}{9}$$

Bertrand Duopoly

Two firms simultaneously choose prices p_1 and p_2. Consumers buy from the firm with the lower price, or split equally if prices are the same.

If $p_i < p_j$, firm i captures the entire market: $q_i = a - p_i$ and $q_j = 0$.

If $p_i = p_j = p$, the firms split the market: $q_i = q_j = \frac{a-p}{2}$.

The profit function for firm i is:

$$\pi_i(p_i, p_j) = \{(p_i - c)(a - p_i) \quad \text{if } p_i < p_j \; (p_i - c)\frac{a-p_i}{2} \quad \text{if } p_i = p_j \; 0 \quad \text{if } p_i > p_j$$

The Nash equilibrium is:

$$p_1^* = p_2^* = c$$

With profits:

$$\pi_1^* = \pi_2^* = 0$$

This is the "Bertrand paradox": with just two firms, the competitive outcome is achieved.

Tragedy of the Commons

Consider n farmers sharing a common resource. Each farmer i decides how many units g_i to extract. The total extraction is $G = \sum_{i=1}^{n} g_i$.

The value per unit extracted is $v(G)$, which is decreasing in G. The cost per unit is constant at c.

The payoff for farmer i is:

$$u_i(g_1, \ldots, g_n) = g_i \cdot (v(G) - c)$$

To find the Nash equilibrium, we compute the first-order condition:

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

In a symmetric equilibrium, $g_i = \frac{G}{n}$, so:

$$v(G^*) + \frac{G^*}{n} \cdot v'(G^*) = c$$

The socially optimal extraction would satisfy:

$$v(G^{opt}) + G^{opt} \cdot v'(G^{opt}) = c$$

Since $v'(G) < 0$ and $\frac{G^*}{n} < G^*$, we have $G^* > G^{opt}$, leading to over-extraction of the resource.

2.5 Efficiency and Price of Anarchy

Definition 2.10 (Pareto Dominance): A strategy profile s' Pareto dominates another profile s if:

$$u_i(s') \geq u_i(s) \quad \forall i \in N$$

with strict inequality for at least one player.

Definition 2.11 (Pareto Efficiency): A strategy profile s is Pareto efficient if there is no other strategy profile s' that Pareto dominates s.

Definition 2.12 (Social Welfare): The social welfare of a strategy profile s is the sum of all players' utilities:

$$SW(s) = \sum_{i=1}^{n} u_i(s)$$

Definition 2.13 (Socially Optimal Outcome): A strategy profile s^{opt} is socially optimal if it maximizes social welfare:

$$s^{opt} \in \arg\max_{s \in S} SW(s)$$

Definition 2.14 (Price of Anarchy): The price of anarchy (PoA) is the ratio of the social welfare at the socially optimal outcome to the worst-case social welfare at any Nash equilibrium:

$$PoA = \frac{SW(s^{opt})}{\min_{s \in NE} SW(s)}$$

where NE is the set of Nash equilibria.

Definition 2.15 (Price of Stability): The price of stability (PoS) is the ratio of the social welfare at the socially optimal outcome to the best-case social welfare at any Nash equilibrium:

$$PoS = \frac{SW(s^{opt})}{\max_{s \in NE} SW(s)}$$

Example 2.4 (Prisoner's Dilemma): In the Prisoner's Dilemma, the unique Nash equilibrium (D, D) with payoffs (-6, -6) is Pareto dominated by the profile (C, C) with payoffs (-1, -1). The price of anarchy is:

$$PoA = \frac{SW((C, C))}{SW((D, D))} = \frac{-1 + (-1)}{-6 + (-6)} = \frac{-2}{-12} = \frac{1}{6}$$

Example 2.5 (Pigou's Example): Consider a routing game with two parallel edges from s to t. The upper edge has constant latency 1, and the lower edge has latency x, where x is the flow on that edge.

The socially optimal solution is to route half the traffic on each edge, giving an average latency of $\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \frac{1}{2} = \frac{3}{4}$.

However, in the Nash equilibrium, all traffic uses the lower edge, giving a latency of 1. The price of anarchy is:

$$PoA = \frac{1}{\frac{3}{4}} = \frac{4}{3}$$

2.6 Mixed Strategies

When players randomize over their pure strategies, we enter the domain of mixed strategies.

Definition 2.16 (Mixed Strategy): A mixed strategy σ_i for player i is a probability distribution over S_i. The set of all mixed strategies for player i is denoted by $\Delta(S_i)$.

If $S_i = s_i^1, s_i^2, \ldots, s_i^k$, then $\sigma_i = (\sigma_i(s_i^1), \sigma_i(s_i^2), \ldots, \sigma_i(s_i^k))$ where $\sigma_i(s_i^j) \geq 0$ for all j and $\sum_{j=1}^{k} \sigma_i(s_i^j) = 1$.

Expected Utility: Given mixed strategy profile $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_n)$, the expected utility for player i is:

$$u_i(\sigma) = \sum_{s \in S} \left(\prod_{j=1}^{n} \sigma_j(s_j) \right) u_i(s)$$

Definition 2.17 (Support): The support of a mixed strategy σ_i, denoted $supp(\sigma_i)$, is the set of pure strategies played with positive probability:

$$supp(\sigma_i) = s_i \in S_i : \sigma_i(s_i) > 0$$

Definition 2.18 (Mixed Strategy Nash Equilibrium): A mixed strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*, \ldots, \sigma_n^*)$ is a Nash equilibrium if for every player $i \in N$:

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*) \quad \forall \sigma_i \in \Delta(S_i)$$

Theorem 2.6 (Characterization of Mixed NE): A mixed strategy profile σ^* is a Nash equilibrium if and only if for all $i \in N$:

1. For all $s_i \in supp(\sigma_i^*)$: $u_i(s_i, \sigma_{-i}^*) = u_i(\sigma_i^*, \sigma_{-i}^*)$
2. For all $s_i \notin supp(\sigma_i^*)$: $u_i(s_i, \sigma_{-i}^*) \leq u_i(\sigma_i^*, \sigma_{-i}^*)$

Proof:

- ($\Rightarrow$) If σ^* is a Nash equilibrium and there exists $s_i \in supp(\sigma_i^*)$ with $u_i(s_i, \sigma_{-i}^*) < u_i(\sigma_i^*, \sigma_{-i}^*)$, then player i could increase their expected utility by shifting probability from s_i to other strategies, contradicting that σ^* is a Nash equilibrium. Similarly, if there exists $s_i \notin supp(\sigma_i^*)$ with $u_i(s_i, \sigma_{-i}^*) > u_i(\sigma_i^*, \sigma_{-i}^*)$, player i could increase their expected utility by shifting probability to s_i.
- ($\Leftarrow$) If conditions 1 and 2 hold, then for any mixed strategy σ_i:

$$u_i(\sigma_i, \sigma_{-i}^*) = \sum_{s_i \in S_i} \sigma_i(s_i) u_i(s_i, \sigma_{-i}^*) \leq \sum_{s_i \in S_i} \sigma_i(s_i) u_i(\sigma_i^*, \sigma_{-i}^*) = u_i(\sigma_i^*, \sigma_{-i}^*)$$

  Thus, σ^* is a Nash equilibrium.

Example 2.6 (Matching Pennies): Players simultaneously choose Heads (H) or Tails (T). Player 1 wins if the choices match; player 2 wins if they differ.

	H	T
H	(1,-1)	(-1,1)
T	(-1,1)	(1,-1)

The unique Nash equilibrium is for both players to play $\sigma_1^* = \sigma_2^* = (0.5, 0.5)$.

Example 2.7 (Battle of the Sexes - Mixed NE): In the Battle of the Sexes game from Example 2.3, there is also a mixed-strategy Nash equilibrium where:

$$\sigma_1^* = \left(\frac{2}{3}, \frac{1}{3}\right) \quad \text{and} \quad \sigma_2^* = \left(\frac{1}{3}, \frac{2}{3}\right)$$

To verify, we check that each player is indifferent between their pure strategies given the other player's mixed strategy:

For player 1:

- $u_1(O, \sigma_2^*) = \frac{1}{3} \cdot 2 + \frac{2}{3} \cdot 0 = \frac{2}{3}$
- $u_1(F, \sigma_2^*) = \frac{1}{3} \cdot 0 + \frac{2}{3} \cdot 1 = \frac{2}{3}$

For player 2:

- $u_2(\sigma_1^*, O) = \frac{2}{3} \cdot 1 + \frac{1}{3} \cdot 0 = \frac{2}{3}$
- $u_2(\sigma_1^*, F) = \frac{2}{3} \cdot 0 + \frac{1}{3} \cdot 2 = \frac{2}{3}$

Since both players are indifferent between their pure strategies, this mixed strategy profile is a Nash equilibrium.

2.7 Nash Theorem

Theorem 2.7 (Nash, 1950): Every finite strategic-form game has at least one Nash equilibrium (possibly in mixed strategies).

Proof (Sketch): The proof uses the Kakutani fixed-point theorem, applied to the best-response correspondence.

Define the best-response correspondence $BR : \Delta(S) \to 2^{\Delta(S)}$ where:
$BR(\sigma) = BR_1(\sigma_{-1}) \times BR_2(\sigma_{-2}) \times \ldots \times BR_n(\sigma_{-n})$

A Nash equilibrium is a fixed point of this correspondence: $\sigma^* \in BR(\sigma^*)$.

1. $\Delta(S)$ is non-empty, compact, and convex.
2. $BR(\sigma)$ is non-empty for all $\sigma \in \Delta(S)$.
3. $BR(\sigma)$ is convex for all $\sigma \in \Delta(S)$.
4. BR has a closed graph.

By the Kakutani fixed-point theorem, BR has a fixed point, which is a Nash equilibrium.

Computing Nash Equilibria: Finding Nash equilibria is computationally hard (PPAD-complete). For two-player games with two strategies each, we can find mixed equilibria analytically:

Let p and q be the probabilities with which players 1 and 2 play their first strategies. In equilibrium, each player must be indifferent between their pure strategies, leading to equations: $u_1(s_1^1, q) = u_1(s_1^2, q)$ $u_2(p, s_2^1) = u_2(p, s_2^2)$

Solving these equations gives the mixed Nash equilibrium.

Theorem 2.8 (Uniqueness of Completely Mixed NE): In a two-player game, if there is a completely mixed Nash equilibrium (both players randomize over all their pure strategies), then it is the unique completely mixed Nash equilibrium.

Proof (Sketch): In a completely mixed Nash equilibrium, each player is indifferent between all their pure strategies. This gives a system of linear equations whose solution, if it exists, is unique.

2.8 Potential Games

Definition 2.19 (Exact Potential Game): A game G is an exact potential game if there exists a function $\Phi : S \to \mathbb{R}$ such that for every player i, every strategy profile s_{-i}, and every pair of strategies $s_i, s_i' \in S_i$:

$$u_i(s_i, s_{-i}) - u_i(s_i', s_{-i}) = \Phi(s_i, s_{-i}) - \Phi(s_i', s_{-i})$$

Definition 2.20 (Weighted Potential Game): A game G is a weighted potential game if there exists a function $\Phi : S \to \mathbb{R}$ and weights $w_i > 0$ for each player i such that for every player i, every strategy profile s_{-i}, and every pair of strategies $s_i, s_i' \in S_i$:

$$u_i(s_i, s_{-i}) - u_i(s_i', s_{-i}) = w_i(\Phi(s_i, s_{-i}) - \Phi(s_i', s_{-i}))$$

Definition 2.21 (Ordinal Potential Game): A game G is an ordinal potential game if there exists a function $\Phi : S \to \mathbb{R}$ such that for every player i, every strategy profile s_{-i}, and every pair of strategies $s_i, s_i' \in S_i$:

$$u_i(s_i, s_{-i}) - u_i(s_i', s_{-i}) > 0 \iff \Phi(s_i, s_{-i}) - \Phi(s_i', s_{-i}) > 0$$

Properties of Potential Games:

1. Every potential game has at least one pure-strategy Nash equilibrium (the strategy profile that maximizes the potential function).
2. The Nash equilibria of a potential game are the local maxima of the potential function.
3. Fictitious play converges to a Nash equilibrium in potential games.
4. Best-response dynamics converge to a Nash equilibrium in finite potential games.

Theorem 2.9: In a finite potential game, any sequence of strict best-response improvements converges to a Nash equilibrium in a finite number of steps.

Proof (Sketch): Each strict best-response improvement increases the potential function. Since the game is finite, the potential function has a maximum value, so the sequence must terminate at a Nash equilibrium.

Examples of Potential Games:

1. **Congestion Games**: Consider a set of resources E and players choosing subsets of resources. The cost of using a resource depends on the number of players using it. For player i choosing resources $s_i \subseteq E$, the cost is:

   $$c_i(s_1, \ldots, s_n) = \sum_{e \in s_i} c_e(n_e(s))$$

   where $n_e(s)$ is the number of players using resource e in profile s.

   The potential function is:

   $$\Phi(s) = \sum_{e \in E} \sum_{j=1}^{n_e(s)} c_e(j)$$

2. **Coordination Games**: Players receive higher payoffs when choosing the same or complementary actions.
3. **Cournot Oligopoly with Linear Demand**: The potential function is:

$$\Phi(q_1, \ldots, q_n) = (a - c) \sum_{i=1}^{n} q_i - \frac{1}{2} \left(\sum_{i=1}^{n} q_i \right)^2 - \frac{1}{2} \sum_{i=1}^{n} q_i^2$$

Theorem 2.10 (Monderer and Shapley, 1996): A game is an exact potential game if and only if it is isomorphic to a congestion game.

3. Dynamic Games of Complete Information

3.1 Extensive Form Games

The extensive form represents a game as a tree, capturing the sequential nature of decision-making.

Definition 3.1 (Extensive Form Game): An extensive form game consists of:

1. A set of players $N = 1, 2, \ldots, n$
2. A game tree with nodes representing decision points and edges representing actions
3. A player function assigning each non-terminal node to a player
4. Information sets for each player
5. Payoffs for each player at each terminal node

Components of an Extensive Form Game:

- **Root**: The starting node of the game
- **Terminal nodes**: Nodes with no successors, where payoffs are assigned
- **Player function**: Maps each non-terminal node to the player who moves at that node
- **Action set**: The set of available actions at each decision node
- **Payoff function**: Maps each terminal node to a vector of payoffs

Definition 3.2 (Path): A path in the game tree is a sequence of nodes $(x^0, x^1, \ldots, x^k)$ such that x^0 is the root, x^k is a terminal node, and for each $j = 0, 1, \ldots, k-1$, x^{j+1} is a successor of x^j.

Definition 3.3 (History): A history is a sequence of actions $(a^1, a^2, \ldots, a^k)$ that lead from the root to a particular node.

3.2 Information Sets and Strategies

Definition 3.4 (Information Set): An information set h_i for player i is a set of decision nodes such that:

1. All nodes in h_i belong to player i
2. When the play reaches any node in h_i, player i does not know which specific node in h_i has been reached
3. The available actions at all nodes in h_i are the same

Definition 3.5 (Perfect Information): A game has perfect information if every information set is a singleton (i.e., at each decision point, the acting player knows exactly where they are in the game tree).

Strategies in Extensive Form Games:

Definition 3.6 (Pure Strategy): A pure strategy s_i for player i is a complete contingent plan that specifies an action at each of player i's information sets.

Let H_i be the set of all information sets for player i, and $A(h)$ be the set of available actions at information set h. Then:

$$s_i : H_i \to \cup_{h \in H_i} A(h) \text{ such that } s_i(h) \in A(h) \text{ for all } h \in H_i$$

Definition 3.7 (Behavioral Strategy): A behavioral strategy β_i for player i specifies a probability distribution over the available actions at each of player i's information sets:

$$\beta_i : H_i \to \cup_{h \in H_i} \Delta(A(h)) \text{ such that } \beta_i(h) \in \Delta(A(h)) \text{ for all } h \in H_i$$

Definition 3.8 (Perfect Recall): A game has perfect recall if no player forgets information they once knew. Formally, for every player i, if nodes x and y are in the same information set, then the sequences of player i's own actions leading to x and y must be the same.

Theorem 3.1 (Kuhn, 1953): In games with perfect recall, any mixed strategy can be replaced by an equivalent behavioral strategy that yields the same expected payoffs.

Definition 3.9 (Strategy Profile): A strategy profile $s = (s_1, s_2, \ldots, s_n)$ is a tuple of strategies, one for each player.

Definition 3.10 (Outcome): The outcome of a strategy profile s is the terminal node reached when all players follow their strategies in s.

Definition 3.11 (Equivalent Strategies): Two strategies s_i and s'_i for player i are equivalent if for every profile s_{-i} of the other players' strategies, the outcomes of (s_i, s_{-i}) and (s'_i, s_{-i}) are the same.

3.3 Backward Induction and Subgame Perfect Equilibrium

Definition 3.12 (Subgame): A subgame of an extensive form game is a subset of the game tree starting at a single node (the root of the subgame) and containing all successor nodes, such that if a node is in an information set, then all nodes in that information set are also in the subgame.

Backward Induction: A method for solving games of perfect information:

1. Start at the terminal nodes and work backward
2. For each decision node, determine the optimal choice for the player who moves at that node
3. Replace the subtree with the payoff resulting from optimal play
4. Continue until reaching the root

Definition 3.13 (Subgame Perfect Equilibrium, SPE): A strategy profile s^* is a subgame perfect equilibrium if it induces a Nash equilibrium in every subgame of the original game.

Theorem 3.2: Every finite extensive form game with perfect information has a subgame perfect equilibrium, which can be found by backward induction.

Proof (Sketch):

1. At each terminal node, the payoffs are fixed.
2. At each decision node, the player chooses the action that maximizes their payoff, given the choices at all successor nodes.
3. This process yields a strategy profile where no player can improve their payoff by deviating in any subgame.

Example 3.1 (Entry Deterrence): Firm 1 decides whether to enter a market (E) or stay out (O). If firm 1 enters, firm 2 can either accommodate (A) or fight (F).

Payoffs:

- (O): (0, 2)
- (E, A): (1, 1)
- (E, F): (-1, -1)

The game has two Nash equilibria: (O, F) and (E, A). However, only (E, A) is subgame perfect because (O, F) relies on a non-credible threat by firm 2 to fight if firm 1 enters.

Definition 3.14 (Credible Threat): A threat is credible if carrying it out is optimal for the player making the threat when called upon to do so.

Theorem 3.3 (One-Stage Deviation Principle): A strategy profile s^* is a subgame perfect equilibrium if and only if no player can improve their payoff by deviating from s^* at a single decision node and then following s^* at all other nodes.

Proof (Sketch):

- ($\Rightarrow$) If s^* is a SPE, then no player can improve their payoff by any deviation, including a one-stage deviation.

- (⇐) If no player can improve their payoff by a one-stage deviation, then no player can improve their payoff by any finite sequence of deviations, because any such sequence can be broken down into a series of one-stage deviations. Since the game is finite, any improvement plan must be finite, so s^* is a SPE.

3.4 Minimax and Zero-Sum Games

Definition 3.15 (Zero-Sum Game): A two-player game is zero-sum if for any strategy profile s:

$$u_1(s) + u_2(s) = 0$$

We can represent a zero-sum game with a single payoff matrix showing player 1's payoffs (player 2's payoffs are the negatives of these values).

Definition 3.16 (Maxmin Value): The maxmin value for player i is:

$$v_i = \max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i})$$

It's the maximum payoff player i can guarantee regardless of what the other players do.

Definition 3.17 (Minmax Value): The minmax value for player i is:

$$\bar{v}_i = \min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_i, s_{-i})$$

It's the minimum payoff the other players can force on player i if they coordinate against i.

Theorem 3.4: In any game, $v_i \leq \bar{v}_i$ for all players i.

Proof: Let s_i^* be a strategy that achieves the maxmin value for player i:

$$v_i = \min_{s_{-i} \in S_{-i}} u_i(s_i^*, s_{-i})$$

Then for any $s_{-i} \in S_{-i}$:

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

But by definition of the maximum:

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

So for any s_{-i}:

$$\max_{s_i \in S_i} u_i(s_i, s_{-i}) \geq v_i$$

Taking the minimum over s_{-i}:

$$\min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_i, s_{-i}) \geq v_i$$

But this is precisely $\bar{v}_i$, so $\bar{v}_i \geq v_i$.

Theorem 3.5 (Minimax Theorem, von Neumann, 1928): In a two-player zero-sum game, $v_1 = \bar{v}_1 = -v_2 = -\bar{v}_2$. This common value is called the value of the game.

Moreover, (s_1^*, s_2^*) is a Nash equilibrium if and only if:

$$s_1^* \in \arg\max_{s_1} \min_{s_2} u_1(s_1, s_2)$$

and

$$s_2^* \in \arg\max_{s_2} \min_{s_1} u_2(s_1, s_2)$$

Mixed Strategy Minimax: For mixed strategies, the maxmin value is:

$$v_i^m = \max_{\sigma_i \in \Delta(S_i)} \min_{\sigma_{-i} \in \Delta(S_{-i})} u_i(\sigma_i, \sigma_{-i})$$

And the minmax value is:

$$\bar{v}_i^m = \min_{\sigma_{-i} \in \Delta(S_{-i})} \max_{\sigma_i \in \Delta(S_i)} u_i(\sigma_i, \sigma_{-i})$$

Theorem 3.6: In any finite game, $v_i^m = \bar{v}_i^m$ for all players i.

Theorem 3.7: In a two-player zero-sum game, if (s_1^*, s_2^*) and (s_1', s_2') are both Nash equilibria, then (s_1^*, s_2') and (s_1', s_2^*) are also Nash equilibria.

Computing Minimax: For a two-player zero-sum game, the minimax strategy can be computed using linear programming.

For a payoff matrix A, player 1's problem is:

Maximize v Subject to:

- $\sum_{i=1}^{m} p_i A_{ij} \geq v$ for all $j = 1, 2, \ldots, n$
- $\sum_{i=1}^{m} p_i = 1$
- $p_i \geq 0$ for all $i = 1, 2, \ldots, m$

Player 2's problem is:

Minimize w Subject to:

- $\sum_{j=1}^{n} A_{ij} q_j \leq w$ for all $i = 1, 2, \ldots, m$
- $\sum_{j=1}^{n} q_j = 1$
- $q_j \geq 0$ for all $j = 1, 2, \ldots, n$

By the minimax theorem, $v = w$ at the optimum.

3.5 Stackelberg Equilibrium

In a Stackelberg game, players move sequentially rather than simultaneously.

Definition 3.18 (Stackelberg Equilibrium): Consider a two-player game where player 1 (the leader) moves first, and player 2 (the follower) moves second. A strategy profile (s_1^*, s_2^*) is a Stackelberg equilibrium if:

1. $s_2^*(s_1) \in \arg\max_{s_2 \in S_2} u_2(s_1, s_2)$ for all $s_1 \in S_1$ (best response of follower)
2. $s_1^* \in \arg\max_{s_1 \in S_1} u_1(s_1, s_2^*(s_1))$ (leader's optimal choice)

Theorem 3.8: The Stackelberg equilibrium can be found by backward induction:

1. Determine the follower's best response function $s_2^*(s_1)$
2. Substitute this into the leader's payoff function
3. Find the leader's optimal strategy s_1^*
4. The Stackelberg equilibrium is $(s_1^*, s_2^*(s_1^*))$

Example 3.2 (Stackelberg Duopoly): In a Stackelberg duopoly, firm 1 (leader) chooses output q_1, and then firm 2 (follower) chooses output q_2 after observing q_1.

The best response function for firm 2 is:

$$BR_2(q_1) = \frac{a - c - q_1}{2}$$

Anticipating this, firm 1 chooses:

$$q_1^* = \frac{a - c}{2}$$

Which leads to:

$$q_2^* = \frac{a - c}{4}$$

With profits:

$$\pi_1^* = \frac{(a - c)^2}{8} \quad \text{and} \quad \pi_2^* = \frac{(a - c)^2}{16}$$

First-Mover Advantage: In Stackelberg games, the leader often has an advantage. In the Stackelberg duopoly:

$$\pi_1^* = \frac{(a - c)^2}{8} > \frac{(a - c)^2}{9} = \pi_1^{Cournot}$$

Theorem 3.9: In a Stackelberg game with continuous action spaces and payoff functions that are continuous and quasiconcave in own actions, a Stackelberg equilibrium always exists.

3.6 Time Consistency

Time consistency relates to whether a player's optimal plan remains optimal as time passes.

Definition 3.19 (Time Consistency): A plan is time-consistent if at each point in time, the continuation of the original plan remains optimal.

Example 3.3 (Resource Allocation): A player has a fixed resource budget $K = 1$ to allocate over three periods. The total payoff is a (possibly discounted) sum of stage payoffs:

$$v(x_1, x_2, x_3) = u(x_1) + \delta u(x_2) + \delta^2 u(x_3)$$

where $u(x) = \log(1 + x)$ and $\delta \in (0, 1]$ is the discount factor.

The problem is:

$$\max v(x_1, x_2, x_3) \quad \text{subject to } x_1 + x_2 + x_3 = 1$$

If $\delta = 1$ (no discounting), the optimal allocation is $x_1^* = x_2^* = x_3^* = \frac{1}{3}$.

If $\delta = 0.8$, the optimal allocation is:

$$x_1^* = 0.6393, x_2^* = 0.3115, x_3^* = 0.0492$$

To check time consistency, we see if the planned allocations for periods 2 and 3 remain optimal at the beginning of period 2.

At period 2, the player has $1 - x_1^*$ resources left and maximizes:

$$w = u(x_2) + \delta u(x_3) \quad \text{subject to } x_2 + x_3 = 1 - x_1^*$$

The first-order condition gives:

$$\frac{u'(x_2)}{u'(x_3)} = \delta$$

For $u(x) = \log(1 + x)$, this becomes:

$$\frac{1 + x_3}{1 + x_2} = \delta$$

With $\delta = 0.8$ and $x_1^* = 0.6393$, we get:

$$x_2 = 0.3115, x_3 = 0.0492$$

which matches the original plan.

Theorem 3.10: With exponential discounting (δ^t), optimal plans are time-consistent.

Theorem 3.11: With hyperbolic discounting ($\beta\delta^t$ for $t \geq 1$, where $\beta < 1$), optimal plans are generally time-inconsistent.

3.7 Multistage Games

Definition 3.20 (Multistage Game): A multistage game is a sequence of T stage games, where:

1. The same set of players participate in each stage
2. Players move simultaneously within each stage
3. Players observe the outcome of previous stages before making decisions in the current stage
4. The overall payoff is a (possibly discounted) sum of stage payoffs

For a finite horizon multistage game with T stages, the payoff for player i is:

$$U_i(s) = \sum_{t=1}^{T} \delta^{t-1} u_i^t(s^t)$$

where $\delta \in (0, 1]$ is the discount factor, s^t is the strategy profile in stage t, and u_i^t is player i's stage utility in period t.

Definition 3.21 (History): A history h^t at stage t is the sequence of realized action profiles in previous stages:

$$h^t = (a^1, a^2, \ldots, a^{t-1})$$

where a^k is the action profile at stage k.

Definition 3.22 (Strategy in Multistage Games): A strategy s_i for player i in a multistage game specifies an action at each stage as a function of the history:

$$s_i^t : H^t \to A_i^t$$

where H^t is the set of possible histories at stage t and A_i^t is the set of actions available to player i at stage t.

Theorem 3.12: If the stage game has a unique Nash equilibrium s^*, then the unique subgame perfect equilibrium of the finite horizon multistage game is to play s^* in every stage.

Proof: By backward induction, in the last stage T, the only equilibrium is to play s^*. Given that, in stage $T - 1$, the only equilibrium is to play s^*, and so on.

Strategic Connection: If the stage game has multiple Nash equilibria, then strategies in earlier stages can influence which equilibrium is played in later stages, potentially leading to cooperation.

Example 3.4 (Prisoner's Dilemma with Reward/Punishment): Consider a two-stage game where the first stage is a Prisoner's Dilemma and the second stage is a coordination game.

First stage:

	C	**D**
C	(3,3)	(0,4)

	C	D
D	(4,0)	(1,1)

Second stage (if both cooperated in the first stage):

	A	B
A	(5,5)	(0,0)
B	(0,0)	(0,0)

Second stage (otherwise):

	A	B
A	(0,0)	(0,0)
B	(0,0)	(1,1)

In the SPE, both players cooperate in the first stage and play A in the second stage, yielding payoffs of (3+5, 3+5) = (8, 8).

3.8 Repeated Games

Definition 3.23 (Repeated Game): A repeated game is a multistage game where the same stage game is played in every period.

For an infinitely repeated game with discount factor $\delta \in (0, 1)$, the payoff for player i is:

$$U_i(s) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i(s^t)$$

where the factor $(1 - \delta)$ normalizes the payoffs to be comparable to single-stage payoffs.

Definition 3.24 (Trigger Strategy): A simple class of strategies in repeated games:

- Start by playing a cooperative action
- Continue cooperating as long as others cooperate
- If anyone deviates, switch to a punishment action forever after

Definition 3.25 (Grim Trigger Strategy (GrT)): A specific trigger strategy:

- Start by playing the cooperative action at $t = 1$
- At stage $t > 1$, play the cooperative action if all players have played the cooperative action in all previous stages; otherwise, play the Nash equilibrium of the stage game

Example 3.5 (Repeated Prisoner's Dilemma): In the one-shot Prisoner's Dilemma (where the stage game payoffs are as in Example 3.4), the unique Nash equilibrium is (D, D). In the infinitely repeated version, mutual cooperation (C, C) can be sustained as a subgame perfect equilibrium using grim trigger strategies if δ is sufficiently high.

For player 1, the payoff from cooperation is:

$$U_1(\text{cooperate}) = (1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\cdot 3 = 3$$

The payoff from deviation is:

$$U_1(\text{deviate}) = (1-\delta)\cdot 4 + (1-\delta)\sum_{t=2}^{\infty}\delta^{t-1}\cdot 1 = (1-\delta)\cdot 4 + \delta\cdot 1 = 4 - 3\delta$$

For cooperation to be a SPE, we need $U_1(\text{cooperate}) \geq U_1(\text{deviate})$, which gives:

$$3 \geq 4 - 3\delta \implies \delta \geq \frac{1}{3}$$

Definition 3.26 (Tit-for-Tat Strategy): A strategy in repeated games where a player begins by cooperating and then mimics the opponent's previous move.

Theorem 3.13: In infinitely repeated games with sufficiently patient players, Tit-for-Tat can sustain cooperation as a subgame perfect equilibrium if the discount factor is sufficiently high.

Theorem 3.14: In the infinitely repeated Prisoner's Dilemma, Tit-for-Tat is a subgame perfect equilibrium if $\delta \geq \frac{1}{2}$.

3.9 The Folk Theorem

The Folk Theorem states that in infinitely repeated games with patient players, any individually rational and feasible payoff can be sustained as a subgame perfect equilibrium.

Definition 3.27 (Feasible Payoff): A payoff vector $v = (v_1, v_2, \ldots, v_n)$ is feasible if it's a convex combination of payoffs from pure strategy profiles in the stage game.

Formally, v is feasible if there exist weights $\alpha(s) \geq 0$ for each strategy profile $s \in S$ such that $\sum_{s\in S}\alpha(s) = 1$ and $v_i = \sum_{s\in S}\alpha(s)u_i(s)$ for all $i \in N$.

Definition 3.28 (Individually Rational Payoff): A payoff vector $v = (v_1, v_2, \ldots, v_n)$ is individually rational if $v_i \geq \bar{v}_i$ for all players i, where $\bar{v}_i$ is player i's minmax value.

Theorem 3.15 (Folk Theorem): For any feasible and individually rational payoff vector v, there exists $\delta^* \in (0,1)$ such that for all $\delta \in (\delta^*, 1)$, there is a subgame perfect equilibrium of the infinitely repeated game with discount factor δ that yields average payoff v.

Proof (Sketch):

1. For each player i, define a punishment strategy profile p^i such that $u_i(p^i) = \bar{v}_i$.
2. Construct a strategy profile where:
   - Players follow a sequence of pure strategy profiles that yield average payoff v.
   - If any player i deviates, all players switch to p^i for a finite number of periods, then return to the original sequence.
3. For sufficiently high δ, the cost of punishment outweighs the gain from deviation.

Corollary 3.1: In the infinitely repeated Prisoner's Dilemma, any feasible payoff vector that gives each player more than their minmax value (1) can be sustained as a subgame perfect equilibrium for sufficiently high δ.

Extended Folk Theorem (Fudenberg and Maskin, 1986): For any feasible payoff vector v with $v_i > \bar{v}_i$ for all i, there exists $\delta^* \in (0, 1)$ such that for all $\delta \in (\delta^*, 1)$, there is a subgame perfect equilibrium of the infinitely repeated game with discount factor δ that yields average payoff v.

3.10 Applications of Dynamic Games

Stackelberg Duopoly

As seen in Example 3.2, in a Stackelberg duopoly, the leader produces more and earns higher profits than in a Cournot duopoly.

Dynamic Cournot with Collusion

In an infinitely repeated Cournot duopoly, firms can sustain collusion at the monopoly output level:

$$q_1 = q_2 = \frac{a - c}{4}$$

yielding profits:

$$\pi_1 = \pi_2 = \frac{(a - c)^2}{8}$$

This is sustainable if:

$$\delta \geq \frac{9}{17}$$

If $\delta < \frac{9}{17}$, partial collusion may still be possible, with:

$$q^* = \frac{(a - c)(9 - 5\delta)}{3(9 - \delta)}$$

which increases from the Cournot quantity $\frac{a-c}{3}$ to the collusive quantity $\frac{a-c}{4}$ as δ increases from 0 to $\frac{9}{17}$.

Dynamic Bargaining

Definition 3.29 (Rubinstein Bargaining Model): Two players take turns making offers to divide a pie of size 1. If an offer is accepted, the game ends. If an offer is rejected, the game continues to the next period with discounting.

For discount factors $\delta_1, \delta_2 \in (0,1)$, the unique subgame perfect equilibrium division is:

$$x_1 = \frac{1-\delta_2}{1-\delta_1\delta_2}, \quad x_2 = \frac{\delta_2(1-\delta_1)}{1-\delta_1\delta_2}$$

where x_i is player i's share.

In the special case where $\delta_1 = \delta_2 = \delta$, this simplifies to:

$$x_1 = \frac{1}{1+\delta}, \quad x_2 = \frac{\delta}{1+\delta}$$

As $\delta \to 1$, both players get equal shares: $x_1 = x_2 = \frac{1}{2}$.

Theorem 3.16 (Rubinstein, 1982): In the alternating-offers bargaining game, there is a unique subgame perfect equilibrium where agreement is reached immediately.

4. Bayesian Games

4.1 Games with Incomplete Information

In games with incomplete information, some players have private information that others do not have.

Definition 4.1 (Bayesian Game): A Bayesian game consists of:

1. A set of players $N = 1, 2, \ldots, n$
2. A set of possible types T_i for each player i
3. A common prior probability distribution p over types $T = T_1 \times T_2 \times \ldots \times T_n$
4. A set of actions A_i for each player i
5. A utility function $u_i : A \times T \to \mathbb{R}$ for each player i, where $A = A_1 \times A_2 \times \ldots \times A_n$

Definition 4.2 (Common Prior Assumption): All players share the same prior beliefs about the distribution of types. This is common knowledge among the players.

Harsanyi Transformation: Harsanyi proposed transforming games of incomplete information into games of imperfect information by introducing Nature as a player who moves first and chooses the types of all players according to the prior probability distribution.

Private Values vs. Common Values:

- **Private Values**: Each player's utility depends only on their own type: $u_i(a,t) = u_i(a,t_i)$

- **Common Values**: A player's utility may depend on other players' types

Example 4.1 (First-Price Auction with Private Values): Two bidders have private valuations v_1 and v_2 drawn independently from a uniform distribution on $[0,1]$. Each bidder submits a sealed bid, and the highest bidder gets the object and pays their bid.

This is a Bayesian game where:

- $N = 1, 2$
- $T_i = [0,1]$ for $i \in 1, 2$
- $p(v_1, v_2) = 1$ for $(v_1, v_2) \in [0,1]^2$
- $A_i = [0, \infty)$ for $i \in 1, 2$
- $u_i(b_1, b_2, v_1, v_2) = \{v_i - b_i \quad \text{if } b_i > b_j \ \frac{v_i - b_i}{2} \quad \text{if } b_i = b_j \ 0 \quad \text{if } b_i < b_j$

4.2 Types and Beliefs

Types: A player's type $t_i \in T_i$ represents their private information. This could be:

- Their payoff function
- Their available actions
- Their beliefs about other players' types

Beliefs: Given a prior distribution p over type profiles, player i with type t_i has beliefs about other players' types given by the conditional distribution:

$$p(t_{-i}|t_i) = \frac{p(t_i, t_{-i})}{p(t_i)}$$

where $p(t_i) = \sum_{t_{-i} \in T_{-i}} p(t_i, t_{-i})$.

Independent Types: Types are independent if $p(t) = \prod_{i=1}^{n} p(t_i)$ for all $t \in T$.

Example 4.2 (Independent Types): If types are independent, then:

$$p(t_{-i}|t_i) = p(t_{-i}) = \prod_{j \neq i} p(t_j)$$

Example 4.3 (Correlated Types): Consider a two-player game where each player's type is either H or L, with joint distribution:

$$p(H,H) = 0.4, \quad p(H,L) = 0.1, \quad p(L,H) = 0.1, \quad p(L,L) = 0.4$$

Player 1's beliefs about player 2's type are:

$$p(t_2 = H|t_1 = H) = \frac{p(H,H)}{p(t_1 = H)} = \frac{0.4}{0.5} = 0.8$$

$$p(t_2 = L|t_1 = H) = \frac{p(H,L)}{p(t_1 = H)} = \frac{0.1}{0.5} = 0.2$$

$$p(t_2 = H|t_1 = L) = \frac{p(L, H)}{p(t_1 = L)} = \frac{0.1}{0.5} = 0.2$$

$$p(t_2 = L|t_1 = L) = \frac{p(L, L)}{p(t_1 = L)} = \frac{0.4}{0.5} = 0.8$$

4.3 Bayesian Nash Equilibrium

Definition 4.3 (Strategy in Bayesian Games): A pure strategy for player i is a function $s_i : T_i \to A_i$ that specifies an action for each possible type.

Definition 4.4 (Bayesian Nash Equilibrium): A strategy profile $s^* = (s_1^*, s_2^*, \ldots, s_n^*)$ is a Bayesian Nash Equilibrium (BNE) if for each player i, each type $t_i \in T_i$, and each alternative action $a_i \in A_i$:

$$\sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) u_i(s_i^*(t_i), s_{-i}^*(t_{-i}), t_i, t_{-i}) \geq \sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) u_i(a_i, s_{-i}^*(t_{-i}), t_i, t_{-i})$$

This means that for each type t_i, player i's strategy maximizes their expected utility given their beliefs about other players' types and assuming other players follow their equilibrium strategies.

Theorem 4.1: Every finite Bayesian game has at least one Bayesian Nash equilibrium (possibly in mixed strategies).

Proof (Sketch): A Bayesian game can be converted to a normal form game where each player's strategy space is the set of functions from types to actions. The Nash theorem then guarantees the existence of a Nash equilibrium.

Example 4.4 (Double Auction): A seller and a buyer have private valuations v_s and v_b for an object, drawn independently from uniform distributions on $[0, 1]$. The seller submits an ask price a, and the buyer submits a bid price b. If $b \geq a$, trade occurs at price $p = \frac{a+b}{2}$.

In a Bayesian Nash equilibrium, the seller's strategy is $a(v_s) = \frac{2v_s+1}{3}$ and the buyer's strategy is $b(v_b) = \frac{2v_b}{3}$.

4.4 Applications of Bayesian Nash Equilibrium

First-Price Sealed-Bid Auction

In a first-price sealed-bid auction with independent private values drawn from a uniform distribution on $[0, 1]$, the Bayesian Nash equilibrium bidding strategy for player i with value v_i is:

$$b_i(v_i) = \frac{n-1}{n} v_i$$

where n is the number of bidders.

Proof (Sketch): Assume all other bidders follow the strategy $b_j(v_j) = \alpha v_j$ for some $\alpha \in (0, 1)$. The probability that bidder i wins with bid b_i is:

$$\Pr(\text{win with bid } b_i) = \Pr(b_i > \max_{j \neq i} b_j) = \Pr\left(\frac{b_i}{\alpha} > \max_{j \neq i} v_j\right) = \left(\frac{b_i}{\alpha}\right)^{n-1}$$

The expected utility is:

$$\mathbb{E}[u_i] = (v_i - b_i) \Pr(\text{win with bid } b_i) = (v_i - b_i)\left(\frac{b_i}{\alpha}\right)^{n-1}$$

Taking the derivative with respect to b_i and setting it to zero:

$$\frac{d\mathbb{E}[u_i]}{db_i} = -\left(\frac{b_i}{\alpha}\right)^{n-1} + (v_i - b_i)\frac{n-1}{\alpha}\left(\frac{b_i}{\alpha}\right)^{n-2} = 0$$

Solving for b_i:

$$b_i = \frac{n-1}{n} v_i$$

So $\alpha = \frac{n-1}{n}$, which confirms our assumption.

Cournot Duopoly with Incomplete Information

Consider a Cournot duopoly where firm i has marginal cost c_i which is private information. The costs are independently drawn from a uniform distribution on $[0, 1]$.

The inverse demand function is $P(Q) = a - Q$ where $Q = q_1 + q_2$.

In a Bayesian Nash equilibrium, firm i with cost c_i produces:

$$q_i(c_i) = \frac{2a - 3c_i + 1}{6}$$

4.5 Signaling Games

Definition 4.5 (Signaling Game): A signaling game is a two-player Bayesian game where:

1. Player 1 (the sender) has private information represented by their type $t_1 \in T_1$
2. Player 1 sends a message $m \in M$ to player 2
3. Player 2 (the receiver) observes the message (but not the type) and takes an action $a \in A$
4. Payoffs depend on the type, message, and action: $u_1(t_1, m, a)$ and $u_2(t_1, m, a)$

Definition 4.6 (Belief System): A belief system μ specifies, for each message $m \in M$, a probability distribution $\mu(\cdot|m)$ over the sender's types.

Types of Equilibria in Signaling Games:

1. **Separating Equilibrium**: Different types of sender choose different messages, fully revealing their types
2. **Pooling Equilibrium**: All types of sender choose the same message, revealing no information
3. **Semi-Separating/Hybrid Equilibrium**: Some types mix between different messages, partially revealing information

Example 4.5 (Job Market Signaling): A worker knows their productivity (high or low) but the firm does not. The worker can choose to get education (which is costly but less costly for high-productivity workers) before applying for a job. The firm observes the worker's education level and sets a wage.

Types: $T_1 = H, L$ with prior $p(H) = p(L) = 0.5$. Messages: $M = E, N$ (education or no education). Actions: $A = \mathbb{R}_+$ (wage). Payoffs:

- $u_1(t_1, m, a) = a - c(t_1, m)$ where $c(H, E) = 1$, $c(L, E) = 2$, $c(H, N) = c(L, N) = 0$.
- $u_2(t_1, m, a) = t_1 - a$ where $H = 2$ and $L = 1$.

In a separating equilibrium:

- High-type worker chooses education: $s_1(H) = E$
- Low-type worker chooses no education: $s_1(L) = N$
- Firm offers wage 2 to educated workers: $s_2(E) = 2$
- Firm offers wage 1 to uneducated workers: $s_2(N) = 1$
- Beliefs: $\mu(H|E) = 1$, $\mu(H|N) = 0$

Theorem 4.2: In a signaling game with finite types and messages, if the single-crossing property holds, then there exists a separating equilibrium.

4.6 Perfect Bayesian Equilibrium

Definition 4.7 (Perfect Bayesian Equilibrium, PBE): A PBE consists of:

1. A strategy profile s^*
2. A belief system μ that specifies, for each information set, a probability distribution over the nodes in that information set

such that:

1. The strategy profile s^* is sequentially rational given beliefs μ
2. The beliefs μ are derived from s^* using Bayes' rule whenever possible

Sequential Rationality: For each player i and each information set h_i, the strategy s_i^* maximizes player i's expected payoff given the beliefs $\mu(h_i)$ and the strategies s_{-i}^* of the other players.

Bayes' Rule Requirement: For information sets that are reached with positive probability under s^*, the beliefs μ must be consistent with s^* and the prior using Bayes' rule.

Example 4.6 (Entry Deterrence with Incomplete Information): Consider a market entry game where an entrant (player 1) decides whether to enter a market. The incumbent (player 2) can be either strong (with low costs) or weak (with high costs). The incumbent's type is private information, with prior probability p of being strong. If the entrant enters, the incumbent decides whether to fight or accommodate.

Types: $T_2 = S, W$ with prior $p(S) = p$, $p(W) = 1 - p$. Actions: $A_1 = E, O$ (enter or out), $A_2 = F, A$ (fight or accommodate). Payoffs:

- $u_1(O, t_2, \cdot) = 0$ for all t_2
- $u_1(E, S, F) = -1$, $u_1(E, S, A) = 2$, $u_1(E, W, F) = 1$, $u_1(E, W, A) = 2$
- $u_2(O, S, \cdot) = u_2(O, W, \cdot) = 2$
- $u_2(E, S, F) = 1$, $u_2(E, S, A) = 0$, $u_2(E, W, F) = -1$, $u_2(E, W, A) = 0$

In a PBE, the strong incumbent always fights, and the weak incumbent always accommodates. The entrant enters if $p < \frac{2}{3}$ and stays out if $p > \frac{2}{3}$.

Theorem 4.3 (One-Shot Deviation Principle for PBE): A strategy profile s^* and belief system μ constitute a PBE if and only if no player can profit by deviating from s^* at a single information set and then following s^* at all other information sets, given the beliefs μ.

4.7 Reputation Building

Reputation Model: Consider a finitely repeated game where:

1. One player may be of different types (e.g., rational or "committed" to a specific strategy)
2. The other player has uncertainty about this player's type
3. The player with private information can build a reputation by mimicking a committed type

Example 4.7 (Chain Store Paradox): A monopolist faces potential entrants in a sequence of markets. The monopolist could be "tough" (always fights entry) or "strategic" (chooses the best response). By fighting early entrants, a strategic monopolist can build a reputation for being tough, deterring future entrants.

Theorem 4.4 (Reputation Effects, Kreps and Wilson, 1982): In finitely repeated games with one-sided incomplete information and sufficiently patient players, if there is a small probability that one player is committed to a specific strategy, that player can guarantee themselves a payoff close to what they would get if they were known to be committed to that strategy.

Example 4.8 (Reputation in the Prisoner's Dilemma): Consider a finitely repeated Prisoner's Dilemma where player 1 may be either a rational player or a "grim trigger" player who starts with cooperation and then cooperates if and only if the opponent has never defected.

If player 2 assigns probability $p > 0$ to player 1 being a grim trigger, then in a game with T stages where T is large, both players cooperate until near the end of the game.

Theorem 4.5 (Fudenberg and Levine, 1989): In an infinitely repeated game with incomplete information where one long-run player faces a sequence of short-run players, if the long-run player is sufficiently patient, they can guarantee themselves a payoff close to their Stackelberg payoff (the payoff they would receive if they could commit to a strategy and the short-run players best-responded).

4.8 Bayesian Cooperative Games

Bayesian cooperative games extend the concept of coalitional games to settings with incomplete information.

Definition 4.8 (Bayesian Cooperative Game): A Bayesian cooperative game consists of:

1. A set of players $N = 1, 2, \ldots, n$
2. A set of possible types T_i for each player i
3. A common prior probability distribution p over types $T = T_1 \times T_2 \times \ldots \times T_n$
4. A characteristic function $v : 2^N \times T \to \mathbb{R}$ that assigns a value to each coalition and type profile

Example 4.9 (Study Group Problem): Two students with abilities $t_1, t_2 \in [0, 1]$ can work together on a project. The project succeeds if at least one student puts in effort. The value of success to student i is t_i^2. The cost of effort is a constant c.

If student i puts in effort, their payoff is $t_i^2 - c$ regardless of what the other student does. If student i doesn't put in effort, their payoff is t_i^2 if the other student puts in effort, and 0 otherwise.

In a Bayesian Nash equilibrium, student i puts in effort if and only if $t_i \geq \sqrt{c}$ when working alone. When working together, student i puts in effort if and only if $t_i \geq \sqrt{\frac{c}{1-\Pr[s_j(t_j)=E]}}$, where $\Pr[s_j(t_j) = E]$ is the probability that the other student puts in effort.

Under uniform priors and with symmetric strategies, the equilibrium threshold is $t_i \geq \sqrt[3]{c}$.

Definition 4.9 (Ex-Ante/Interim/Ex-Post Core): In a Bayesian cooperative game:

1. **Ex-Ante Core**: Allocations that cannot be blocked by any coalition before types are realized
2. **Interim Core**: Allocations that cannot be blocked by any coalition after each player learns their own type

3. **Ex-Post Core**: Allocations that cannot be blocked by any coalition after all types are realized

Theorem 4.6: The ex-post core is contained in the interim core, which is contained in the ex-ante core.

Theorem 4.7: If the characteristic function v is superadditive and convex for each type profile, then the ex-ante core is non-empty.

5. Advanced Topics

5.1 Auctions

Definition 5.1 (Auction): A mechanism for selling an item to one of several bidders.

Common Auction Formats:

1. **First-Price Sealed-Bid Auction**: Bidders submit sealed bids; highest bidder wins and pays their bid
2. **Second-Price Sealed-Bid Auction (Vickrey)**: Bidders submit sealed bids; highest bidder wins and pays the second-highest bid
3. **English Auction (Ascending-Price)**: Price increases until only one bidder remains
4. **Dutch Auction (Descending-Price)**: Price decreases until a bidder accepts the current price

Valuation Models:

1. **Independent Private Values**: Each bidder's valuation is drawn independently from a distribution, and bidders know their own valuations but not others'.
2. **Common Value**: The item has the same value to all bidders, but bidders have different information about this value.
3. **Interdependent Values**: Bidders' valuations depend on both their own information and others' information.

Theorem 5.1 (Revenue Equivalence): Under the independent private values model with risk-neutral bidders drawn from the same distribution, any auction mechanism that allocates the item to the highest bidder and gives zero expected payoff to a bidder with the lowest possible valuation yields the same expected revenue.

Proof (Sketch): Consider any auction mechanism that satisfies the conditions. Let $P(v)$ be the expected payment by a bidder with valuation v, and let $W(v)$ be the probability that the bidder wins the auction. By incentive compatibility:

$$v \in \arg\max_{v'} vW(v') - P(v')$$

Taking the first-order condition:

$$vW'(v') - P'(v') = 0 \quad \text{at } v' = v$$

So:

$$P'(v) = vW'(v)$$

Integrating:

$$P(v) = P(0) + \int_0^v tW'(t)dt = P(0) + vW(v) - \int_0^v W(t)dt$$

Since the lowest-value bidder pays zero in expectation, $P(0) = 0$. Thus:

$$P(v) = vW(v) - \int_0^v W(t)dt$$

This formula holds for any auction mechanism satisfying the conditions, so all such mechanisms yield the same expected payment for each bidder type, and hence the same expected revenue.

Example 5.1 (Second-Price Auction): In a second-price auction, bidding truthfully (i.e., bidding one's true valuation) is a weakly dominant strategy.

Proof: Let v_i be bidder i's valuation, and let b_i be their bid. Let $b^*_{-i} = \max_{j \neq i} b_j$ be the highest bid among other bidders.

Case 1: $v_i > b^*_{-i}$. If $b_i > b^*_{-i}$, bidder i wins and pays b^*_{-i}, for a payoff of $v_i - b^*_{-i}$. If $b_i < b^*_{-i}$, bidder i loses and gets a payoff of 0. Since $v_i - b^*_{-i} > 0$, bidding $b_i = v_i$ is better than bidding $b_i < b^*_{-i}$.

Case 2: $v_i < b^*_{-i}$. If $b_i > b^*_{-i}$, bidder i wins and pays b^*_{-i}, for a payoff of $v_i - b^*_{-i}$. If $b_i < b^*_{-i}$, bidder i loses and gets a payoff of 0. Since $v_i - b^*_{-i} < 0$, bidding $b_i = v_i$ is better than bidding $b_i > b^*_{-i}$.

Case 3: $v_i = b^*_{-i}$. Bidder i is indifferent between winning and losing, so any bid is optimal.

Thus, bidding $b_i = v_i$ is a weakly dominant strategy.

Example 5.2 (First-Price Auction): In a first-price auction with n bidders whose valuations are drawn independently from a uniform distribution on $[0, 1]$, the Bayesian Nash Equilibrium bidding strategy is $b(v) = \frac{n-1}{n}v$.

Winner's Curse: In common value auctions, the winner tends to be the bidder who most overestimates the value, leading to negative expected profits.

Definition 5.2 (All-Pay Auction): An auction where all bidders pay their bids, but only the highest bidder receives the item.

5.2 Mechanism Design

Definition 5.3 (Mechanism Design): The study of how to design games or institutions that achieve desired outcomes when players act strategically.

Components of a Mechanism:

1. A set of possible outcomes X
2. A set of possible types T_i for each player i
3. A utility function $u_i(x, t_i)$ for each player i, where $x \in X$ and $t_i \in T_i$
4. A mechanism (M, g) where $M = M_1 \times M_2 \times \ldots \times M_n$ is a message space and $g : M \to X$ is an outcome function

Definition 5.4 (Direct Mechanism): A mechanism where $M_i = T_i$ for all i (players report their types directly).

Definition 5.5 (Incentive Compatibility): A direct mechanism is incentive compatible if truthfully reporting one's type is a Bayesian Nash Equilibrium.

Definition 5.6 (Dominant Strategy Incentive Compatible/Strategy-Proof): A direct mechanism is strategy-proof if truthfully reporting one's type is a weakly dominant strategy.

Definition 5.7 (Efficiency): A mechanism is efficient if it maximizes the sum of players' utilities.

Definition 5.8 (Individual Rationality): A mechanism is individually rational if each player's expected utility from participating is at least as high as their expected utility from not participating.

5.3 Revelation Principle

Theorem 5.2 (Revelation Principle for Dominant Strategies): For any mechanism and dominant strategy equilibrium of that mechanism, there exists a strategy-proof direct mechanism that yields the same outcome.

Proof (Sketch): Let (M, g) be a mechanism and let s^* be a dominant strategy equilibrium of that mechanism. Define a direct mechanism (T, h) where $h(t) = g(s^*(t))$ for all $t \in T$. Then truthful reporting is a dominant strategy in (T, h).

Theorem 5.3 (Revelation Principle for BNE): For any mechanism and Bayesian Nash Equilibrium of that mechanism, there exists an incentive-compatible direct mechanism that yields the same outcome.

Proof (Sketch): Let (M, g) be a mechanism and let s^* be a BNE of that mechanism. Define a direct mechanism (T, h) where $h(t) = g(s^*(t))$ for all $t \in T$. Then truthful reporting is a BNE in (T, h).

Theorem 5.4 (Gibbard-Satterthwaite): If there are at least three possible outcomes, any strategy-proof social choice function that is onto (i.e., every outcome is chosen for some preference profile) must be dictatorial.

Definition 5.9 (Dictatorial Mechanism): A mechanism is dictatorial if there is a player whose preferred outcome is always chosen.

Example 5.3 (Vickrey-Clarke-Groves (VCG) Mechanism): The VCG mechanism is a direct mechanism where:

1. Players report their valuations $v_i(x)$ for each outcome $x \in X$
2. The mechanism chooses the outcome x^* that maximizes the sum of reported valuations: $x^* \in \arg\max_{x \in X} \sum_{i=1}^n v_i(x)$
3. Each player i pays a "pivotal payment" equal to the reduction in others' welfare caused by their presence:

$$p_i = \max_{x \in X} \sum_{j \neq i} v_j(x) - \sum_{j \neq i} v_j(x^*)$$

Theorem 5.5: The VCG mechanism is strategy-proof and efficient.

Proof (Sketch): For efficiency, note that the mechanism explicitly chooses the outcome that maximizes the sum of reported valuations.

For strategy-proofness, consider player i's utility when reporting truthfully versus reporting v_i':

When reporting truthfully, player i's utility is:

$$u_i = v_i(x^*) - p_i = v_i(x^*) - \left(\max_{x \in X} \sum_{j \neq i} v_j(x) - \sum_{j \neq i} v_j(x^*) \right)$$

$$= v_i(x^*) + \sum_{j \neq i} v_j(x^*) - \max_{x \in X} \sum_{j \neq i} v_j(x)$$

$$= \sum_{j=1}^n v_j(x^*) - \max_{x \in X} \sum_{j \neq i} v_j(x)$$

When reporting v_i', the mechanism chooses x' to maximize $v_i'(x) + \sum_{j \neq i} v_j(x)$. Player i's utility is:

$$u_i' = v_i(x') - p_i' = v_i(x') - \left(\max_{x \in X} \sum_{j \neq i} v_j(x) - \sum_{j \neq i} v_j(x') \right)$$

$$= v_i(x') + \sum_{j \neq i} v_j(x') - \max_{x \in X} \sum_{j \neq i} v_j(x)$$

$$= \sum_{j=1}^n v_j(x') - \max_{x \in X} \sum_{j \neq i} v_j(x)$$

Since x^* maximizes $\sum_{j=1}^n v_j(x)$, we have $\sum_{j=1}^n v_j(x^*) \geq \sum_{j=1}^n v_j(x')$, so $u_i \geq u_i'$.

Definition 5.10 (Budget Balance): A mechanism is budget-balanced if the sum of payments is zero: $\sum_{i=1}^n p_i = 0$.

Theorem 5.6 (Myerson-Satterthwaite): In a bilateral trading problem with private values and continuous type spaces, no mechanism can be simultaneously efficient, incentive compatible, individually rational, and budget-balanced.

5.4 Coalitional Games

Definition 5.11 (Coalitional Game with Transferable Utility): A pair (N, v) where:

1. $N = 1, 2, \ldots, n$ is a set of players
2. $v : 2^N \to \mathbb{R}$ is a characteristic function with $v(\emptyset) = 0$

The value $v(S)$ represents the total payoff that coalition S can secure for itself.

Definition 5.12 (Superadditive Game): A game is superadditive if for all disjoint coalitions $S, T \subseteq N$:

$$v(S \cup T) \geq v(S) + v(T)$$

Definition 5.13 (Convex Game): A game is convex if for all coalitions $S, T \subseteq N$:

$$v(S \cup T) + v(S \cap T) \geq v(S) + v(T)$$

Definition 5.14 (Imputation): An imputation is a payoff vector $x = (x_1, x_2, \ldots, x_n)$ satisfying:

1. Efficiency: $\sum_{i=1}^n x_i = v(N)$
2. Individual rationality: $x_i \geq v(i)$ for all $i \in N$

Definition 5.15 (Core): The core of a coalitional game (N, v) is the set of imputations x such that $\sum_{i \in S} x_i \geq v(S)$ for all coalitions $S \subseteq N$.

Theorem 5.7: A superadditive and convex game has a non-empty core.

Theorem 5.8 (Bondareva-Shapley): A coalitional game has a non-empty core if and only if it is balanced.

Definition 5.16 (Balanced Game): A game is balanced if for any balanced collection of weights $a_{S S \in 2^N \setminus \emptyset}$ (i.e., for all $i \in N$, $\sum_{S: i \in S} a_S = 1$), we have $\sum_{S \in 2^N \setminus \emptyset} a_S v(S) \leq v(N)$.

Definition 5.17 (Shapley Value): The Shapley value ϕ allocates to each player i the amount:

$$\phi_i(v) = \sum_{S \subseteq N \setminus i} \frac{|S|!, (n - |S| - 1)!}{n!} [v(S \cup i) - v(S)]$$

Properties of the Shapley Value:

1. Efficiency: $\sum_{i=1}^n \phi_i(v) = v(N)$
2. Symmetry: If $v(S \cup i) = v(S \cup j)$ for all $S \subseteq N \setminus i, j$, then $\phi_i(v) = \phi_j(v)$
3. Dummy player: If $v(S \cup i) = v(S)$ for all $S \subseteq N \setminus i$, then $\phi_i(v) = 0$
4. Additivity: $\phi_i(v + w) = \phi_i(v) + \phi_i(w)$ for all games v, w

Theorem 5.9: The Shapley value is the unique payoff vector satisfying the properties of efficiency, symmetry, dummy player, and additivity.

Example 5.4 (Voting Game): In a weighted voting game, player i has weight w_i, and a coalition S can pass a proposal if the sum of weights exceeds a threshold q:

$$v(S) = \{1 \quad \text{if } \sum_{i \in S} w_i \geq q \, 0 \quad \text{otherwise}$$

The Shapley value of player i represents their power in this voting system, also known as the Shapley-Shubik power index.

Definition 5.18 (Non-Transferable Utility (NTU) Game): A pair (N, V) where:

1. $N = 1, 2, \ldots, n$ is a set of players
2. V is a function that assigns to each coalition $S \subseteq N$ a set $V(S) \subseteq \mathbb{R}^{|S|}$ of feasible payoff vectors

Definition 5.19 (Core of an NTU Game): The core of an NTU game (N, V) is the set of payoff vectors $x \in V(N)$ such that there is no coalition S and payoff vector $y \in V(S)$ with $y_i > x_i$ for all $i \in S$.

5.5 Nash Bargaining

Definition 5.20 (Nash Bargaining Problem): A pair (S, d) where:

1. $S \subseteq \mathbb{R}^2$ is a compact, convex set of feasible utility pairs
2. $d \in S$ is the disagreement point

Definition 5.21 (Nash Bargaining Solution): The Nash bargaining solution is the point $s^* \in S$ that maximizes the Nash product:

$$(s_1 - d_1)(s_2 - d_2)$$

subject to $s_1 \geq d_1$ and $s_2 \geq d_2$.

Nash's Axioms:

1. **Pareto Efficiency**: If $s \in S$ and there exists $t \in S$ with $t \geq s$ and $t \neq s$, then t is the solution.
2. **Symmetry**: If S is symmetric around the line $s_1 = s_2$ and $d_1 = d_2$, then $s_1^* = s_2^*$.
3. **Independence of Irrelevant Alternatives**: If $S \subseteq T$ and the solution to (T, d) is in S, then it is also the solution to (S, d).

4. **Scale Invariance**: For positive affine transformations, the solution transforms accordingly.

Theorem 5.10 (Nash, 1950): The Nash bargaining solution is the unique solution satisfying the four axioms.

Proof (Sketch): First, show that the Nash product maximizer satisfies all four axioms. Then, show that any solution satisfying the four axioms must be the Nash product maximizer.

For the first part, Pareto efficiency follows from the fact that increasing one player's utility while holding the other's constant increases the Nash product. Symmetry follows from the symmetry of the Nash product. IIA follows because if the Nash product maximizer in T is in S , then it must also maximize the Nash product in S. Scale invariance follows because the Nash product transforms monotonically under positive affine transformations.

For the second part, consider any solution f satisfying the four axioms. By scale invariance, we can normalize the problem so that $d = (0, 0)$ and $f(S, d) = (1, 1)$. By symmetry and Pareto efficiency, f must select $(1, 1)$ in the symmetric case. By IIA, f must select the Nash product maximizer in general.

Example 5.5 (Nash Bargaining for Resource Allocation): Two players need to divide a resource of size 1. The utility of player i from receiving x_i is $u_i(x_i)$. The disagreement point is $(0, 0)$.

The Nash bargaining solution maximizes $u_1(x_1) \cdot u_2(1 - x_1)$. The first-order condition is:

$$u_1'(x_1) \cdot u_2(1 - x_1) = u_1(x_1) \cdot u_2'(1 - x_1)$$

If $u_i(x_i) = x_i^{\alpha_i}$, then the solution is:

$$x_1 = \frac{\alpha_1}{\alpha_1 + \alpha_2}$$

Kalai-Smorodinsky Solution: An alternative to the Nash bargaining solution that replaces the IIA axiom with monotonicity. It lies at the intersection of the Pareto frontier with the line connecting the disagreement point to the "ideal point" (where each player gets their maximum possible utility).

Definition 5.22 (Kalai-Smorodinsky Solution): Given a bargaining problem (S, d) and the ideal point $i = (\bar{s}_1, \bar{s}_2)$ where $\bar{s}_j = \max s_j : (s_1, s_2) \in S, s_1 \geq d_1, s_2 \geq d_2$ for $j = 1, 2$, the Kalai-Smorodinsky solution is the point $s^* \in S$ on the Pareto frontier such that:

$$\frac{s_1^* - d_1}{\bar{s}_1 - d_1} = \frac{s_2^* - d_2}{\bar{s}_2 - d_2}$$

5.6 Cheap Talk

Definition 5.23 (Cheap Talk): Communication between players that does not directly affect payoffs.

Crawford-Sobel Model:

1. A sender has private information about a state $\theta \in [0, 1]$
2. The sender sends a message m to a receiver
3. The receiver takes an action $a \in \mathbb{R}$
4. Payoffs depend on the state and action but not on the message

The sender's utility is $U_S(a, \theta, b) = -(a - (\theta + b))^2$, and the receiver's utility is $U_R(a, \theta) = -(a - \theta)^2$, where $b > 0$ is the sender's bias.

Results:

1. If the sender and receiver have aligned preferences ($b = 0$), full information revelation is possible.
2. If their preferences differ by a small amount ($b > 0$ but small), partial information revelation is possible (partition equilibria).
3. If their preferences differ significantly (large b), no information revelation is possible (babbling equilibrium).

Theorem 5.11 (Crawford-Sobel, 1982): In the cheap talk game, all equilibria are partition equilibria, where the state space is divided into intervals, and the sender reveals only which interval the state belongs to.

Example 5.6 (Legislative Committees): A legislator (receiver) relies on a committee (sender) with specialized knowledge to propose policies. The committee has a bias b relative to the legislator. The state space $\theta \in [-w, w]$ is uniformly distributed.

If $b \leq w/2$, there is an informative equilibrium where the committee reveals whether θ is in $[-w, 0]$ or $[0, w]$.

Applications of Cheap Talk:

1. Political communication
2. Expert advice
3. Organizational communication
4. Negotiation and mediation

5.7 Algorithmic Game Theory

Definition 5.24 (Algorithmic Game Theory): The study of computational aspects of game theory, including the complexity of finding equilibria and the design of algorithms for strategic settings.

Complexity Classes:

1. **P**: Problems solvable in polynomial time

2. **NP**: Problems where a solution can be verified in polynomial time
3. **PPAD**: Polynomial Parity Arguments on Directed graphs, a complexity class containing the problem of finding a Nash equilibrium

Theorem 5.12 (Daskalakis, Goldberg, Papadimitriou, 2006): Finding a Nash equilibrium in a normal-form game with three or more players is PPAD-complete.

Theorem 5.13 (Chen and Deng, 2006): Finding a Nash equilibrium in a two-player normal-form game is PPAD-complete.

Definition 5.25 (ϵ-Nash Equilibrium): A strategy profile s is an ϵ-Nash equilibrium if no player can improve their payoff by more than ϵ by unilaterally deviating:

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}) - \epsilon \quad \forall i \in N, \forall s'_i \in S_i$$

Theorem 5.14 (Lipton, Markakis, Mehta, 2003): For any game and any $\epsilon > 0$, there exists an ϵ-Nash equilibrium with support size $O(\log n/\epsilon^2)$, where n is the number of players.

Algorithmic Mechanism Design: The study of how to design mechanisms that are both economically sound and computationally efficient.

Definition 5.26 (Communication Complexity): The minimum number of bits that need to be exchanged to compute a function whose inputs are distributed among multiple parties.

Theorem 5.15 (Nisan and Segal, 2006): Implementing an efficient allocation in a combinatorial auction requires exponential communication in the number of items.

5.8 Fictitious Play and Learning in Games

Definition 5.27 (Fictitious Play): A learning process where each player best responds to the empirical distribution of others' past actions.

Formally, let $f_i^t(s_j)$ be the frequency with which player j has played strategy s_j up to time t. At time $t+1$, player i chooses:

$$s_i^{t+1} \in \arg\max_{s_i \in S_i} \sum_{s_{-i} \in S_{-i}} \left(\prod_{j \neq i} f_j^t(s_j) \right) u_i(s_i, s_{-i})$$

Theorem 5.16 (Robinson, 1951): In a zero-sum game, the empirical distributions of play in fictitious play converge to the set of Nash equilibria.

Theorem 5.17 (Monderer and Shapley, 1996): In a potential game, the empirical distributions of play in fictitious play converge to the set of Nash equilibria.

Definition 5.28 (Best-Response Dynamics): A learning process where each player best responds to the current strategy profile of others.

At time $t+1$, player i chooses:

$$s_i^{t+1} \in \arg\max_{s_i \in S_i} u_i(s_i, s_{-i}^t)$$

Theorem 5.18: In a potential game, best-response dynamics converge to a Nash equilibrium.

Definition 5.29 (No-Regret Learning): A learning algorithm has no regret if, in the limit, the average payoff is at least as good as the payoff from always playing any fixed strategy.

Formally, a sequence of strategies $s_i^t{}_{t=1}^T$ has no regret if for all $s_i \in S_i$:

$$\lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \left[u_i(s_i, s_{-i}^t) - u_i(s_i^t, s_{-i}^t)\right] \leq 0$$

Theorem 5.19 (Freund and Schapire, 1999): If all players use a no-regret learning algorithm, the empirical distributions of play converge to the set of coarse correlated equilibria.

Definition 5.30 (Coarse Correlated Equilibrium): A distribution σ over strategy profiles is a coarse correlated equilibrium if for each player i and each strategy $s_i \in S_i$:

$$\sum_{s\in S} \sigma(s) u_i(s) \geq \sum_{s\in S} \sigma(s) u_i(s_i, s_{-i})$$

5.9 Distributed Optimization

Definition 5.31 (Distributed Optimization): The problem of optimizing a global objective function by multiple agents, each with access to only partial information.

Game-Theoretic Approach: Formulate the distributed optimization problem as a game where each agent's objective is aligned with the global objective.

Definition 5.32 (Potential Game Formulation): Design a game with a potential function that corresponds to the global objective function. The Nash equilibria of the game then correspond to local optima of the objective function.

Example 5.7 (Distributed Resource Allocation): Consider n agents sharing m resources. Each agent i chooses an allocation $x_i \in \mathbb{R}_+^m$ subject to constraints. The cost of resource j is a function c_j of the total allocation $\sum_{i=1}^n x_{ij}$.

The global objective is to minimize the total cost:

$$\min_x \sum_{j=1}^{m} c_j \left(\sum_{i=1}^{n} x_{ij}\right)$$

This can be formulated as a potential game where agent i's cost is:

$$C_i(x_i, x_{-i}) = \sum_{j=1}^{m} x_{ij} \cdot c_j \left(\sum_{k=1}^{n} x_{kj}\right)$$

Theorem 5.20: The above game is a potential game with potential function:

$$\Phi(x) = \sum_{j=1}^{m} \int_{0}^{\sum_{i=1}^{n} x_{ij}} c_j(t)dt$$

Consensus Algorithms: Distributed algorithms where agents aim to agree on a common value.

Definition 5.33 (Consensus Protocol): Each agent i updates their value x_i based on their neighbors' values:

$$x_i^{t+1} = \sum_{j \in N_i} w_{ij} x_j^t$$

where N_i is the set of agent i's neighbors and w_{ij} are weights satisfying $\sum_{j \in N_i} w_{ij} = 1$.

Theorem 5.21: Under mild conditions on the weight matrix and the communication graph, the consensus protocol converges to a common value for all agents.

Price of Anarchy in Distributed Systems: The ratio between the social welfare at the worst Nash equilibrium and the optimal social welfare.

Theorem 5.22: In a network congestion game with linear latency functions, the price of anarchy is at most $\frac{4}{3}$.

5.10 Selfish Routing

Definition 5.34 (Selfish Routing Game): A game where players choose paths in a network to minimize their own latency, which depends on the congestion of each edge.

Formally, a selfish routing game consists of:

1. A directed graph $G = (V, E)$
2. A set of source-destination pairs (s_i, t_i) for $i \in 1, 2, \ldots, k$
3. A flow rate $r_i > 0$ for each pair
4. A latency function ℓ_e for each edge $e \in E$

Each player chooses a path from s_i to t_i, and the cost of a path is the sum of the latencies of its edges, where the latency of an edge depends on the total flow on that edge.

Definition 5.35 (Wardrop Equilibrium): A flow is a Wardrop equilibrium if for each source-destination pair, all used paths have the same latency, and this latency is minimal.

Theorem 5.23 (Existence and Uniqueness of Wardrop Equilibria): If the latency functions are continuous and non-decreasing, a Wardrop equilibrium always exists. If the latency functions are strictly increasing, the edge flows in a Wardrop equilibrium are unique.

Example 5.8 (Pigou's Example): Consider a routing game with a single source-destination pair and two parallel edges. The upper edge has constant latency 1, and the lower edge has latency x, where x is the flow on that edge.

The socially optimal solution is to route half the traffic on each edge, giving an average latency of $\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \frac{1}{2} = \frac{3}{4}$.

However, in the Wardrop equilibrium, all traffic uses the lower edge, giving a latency of 1. The price of anarchy is:

$$PoA = \frac{1}{\frac{3}{4}} = \frac{4}{3}$$

Theorem 5.24 (Roughgarden and Tardos, 2002): In a selfish routing game with latency functions in class $\mathcal{L}$, the price of anarchy is at most:

$$PoA \leq \sup_{f \in \mathcal{L}} \frac{\beta \cdot f(\beta)}{f(1) + (\beta - 1) \cdot f(\beta)}$$

where β is the unique solution to $f(1) = f(\beta) + \beta \cdot f'(\beta)$.

For linear latency functions ($f(x) = ax + b$), this gives $PoA \leq \frac{4}{3}$.

Braess's Paradox: Adding a new edge to a network can sometimes increase the latency of all players at equilibrium.

Example 5.9 (Braess's Paradox): Consider a network with four nodes: source s, destination t, and intermediate nodes v and w. There are edges from s to v, from s to w, from v to t, and from w to t. The edges (s, v) and (w, t) have latency x (the flow on the edge), while the edges (s, w) and (v, t) have constant latency 1.

In the Wardrop equilibrium, half the flow goes through v and half through w, for a total latency of $\frac{1}{2} + 1 = \frac{3}{2}$ per unit of flow.

Now add a zero-latency edge from v to w. In the new Wardrop equilibrium, all flow goes from s to v to w to t, for a total latency of $1 + 1 = 2$ per unit of flow.

The addition of the edge has increased the latency for all players.

5.11 Evolutionary Games

Definition 5.36 (Evolutionary Game): A game where a large population of agents repeatedly play a symmetric normal-form game. The strategies that yield higher payoffs tend to spread in the population.

Definition 5.37 (Replicator Dynamics): The differential equation describing how the proportions of strategies change over time:

$$\dot{x}_i = x_i \left[u_i(e_i, x) - \sum_{j=1}^{n} x_j u_j(e_j, x) \right]$$

where x_i is the proportion of the population playing strategy i, e_i is the strategy that always plays i, and $u_i(e_i, x)$ is the expected payoff to strategy i against the mixed strategy x.

Definition 5.38 (Evolutionarily Stable Strategy, ESS): A strategy x is an ESS if for all strategies $y \neq x$, there exists $\epsilon_y > 0$ such that for all $\epsilon \in (0, \epsilon_y)$:

$$u(x, (1-\epsilon)x + \epsilon y) > u(y, (1-\epsilon)x + \epsilon y)$$

Theorem 5.25: If x is an ESS, then (x, x) is a symmetric Nash equilibrium.

Theorem 5.26: If x is an ESS, then x is a locally asymptotically stable state of the replicator dynamics.

Example 5.10 (Hawk-Dove Game): Consider a population where individuals can be either "Hawks" (aggressive) or "Doves" (peaceful) when competing for resources.

	Hawk	**Dove**
Hawk	$\frac{V-C}{2}, \frac{V-C}{2}$	$V, 0$
Dove	$0, V$	$\frac{V}{2}, \frac{V}{2}$

where V is the value of the resource and C is the cost of fighting.

If $V < C$, the unique ESS is the mixed strategy $x = (\frac{V}{C}, 1 - \frac{V}{C})$.

If $V > C$, the unique ESS is pure strategy "Hawk".

Applications of Evolutionary Game Theory:

1. Biology: Evolution of behavior, altruism, cooperation
2. Economics: Evolution of preferences, norms, conventions
3. Computer Science: Evolution of protocols, algorithms
4. Sociology: Evolution of social norms, customs, institutions

5.12 Applications in Networks

Age of Information

Definition 5.39 (Age of Information, AoI): A metric that captures the freshness of information at a receiver. It is defined as the time elapsed since the generation of the most recent update that has been delivered to the receiver.

Game-Theoretic Model: Multiple sources compete for a shared channel to transmit updates to a common receiver. Each source aims to minimize its own AoI.

Nash Equilibrium: A strategy profile where each source optimally trades off the freshness of its own updates with the congestion caused by other sources.

Example 5.11 (AoI Minimization Game): Consider n sources, each with a rate constraint λ_i for generating updates. The AoI of source i is:

$$A_i = \frac{1}{2\lambda_i} + \frac{1}{\mu - \sum_{j=1}^{n} \lambda_j}$$

where μ is the service rate of the channel.

The Nash equilibrium rates are:

$$\lambda_i^* = \frac{\mu - (n-1)\lambda_j^*}{2}$$

If sources are symmetric, the unique symmetric Nash equilibrium is:

$$\lambda_i^* = \frac{\mu}{n+1}$$

Game-Theory Enabled MIMO Systems

Definition 5.40 (MIMO System): A multiple-input multiple-output system where multiple antennas are used at both the transmitter and receiver to improve communication performance.

Game-Theoretic Model: Multiple transmitters share a wireless channel, each with multiple antennas. Each transmitter aims to maximize its own data rate by choosing its transmit covariance matrix.

Definition 5.41 (MIMO Rate Maximization Game): Each player (transmitter) i chooses a transmit covariance matrix Q_i to maximize its rate:

$$R_i(Q_i, Q_{-i}) = \log\det\left(I + H_i Q_i H_i^H \left(I + \sum_{j \neq i} H_i Q_j H_i^H\right)^{-1}\right)$$

subject to a power constraint $\mathrm{tr}(Q_i) \leq P_i$.

Theorem 5.27: The MIMO rate maximization game is a potential game with potential function:

$$\Phi(Q) = \log\det\left(I + \sum_{i=1}^{n} H_i Q_i H_i^H\right)$$

Theorem 5.28: The MIMO rate maximization game has a unique Nash equilibrium.

Cognitive Radio Networks: Networks where devices dynamically adjust their transmission parameters based on the observed spectrum environment.

Game-Theoretic Model: Multiple cognitive radios compete for a shared spectrum, each choosing its transmission parameters (power, frequency, etc.) to maximize its own utility.

Theorem 5.29: Under certain conditions, the cognitive radio power control game has a unique Nash equilibrium, which can be reached through best-response dynamics.