GAME THEORY

László Gulyás
Associate Professor, ELTE, AI Department
lgulyas@inf.elte.hu

Tamás Takács
PhD student, ELTE, AI Department
tamastheactual@inf.elte.hu
tamastheactual.github.io

Game Theory

Lecture 4

Extensive Form Games and Welfare Concepts

1 Extensive Form Games

2 Subgame Perfect E.

3 Welfare Analysis

Best Response: Definition

For player i and opponents' mixed strategy π_{-i},

$$BR_i(\pi_{-i}) = \arg\max_{\pi_i \in \Delta(A_i)} R_i(\pi_i, \pi_{-i})$$

- May be multi-valued
- Always nonempty for finite games
- Contains all optimal mixtures against π_{-i}

Upper Hemicontinuity: Intuition

If $\pi_{-i}^{k} \to \pi_{-i}$ and $\pi_i^k \in BR_i(\pi_{-i}^{k})$ with $\pi_i^k \to \pi_i$, then $\pi_i \in BR_i(\pi_{-i})$.

Small changes in beliefs do not create discontinuous jumps in optimal responses.

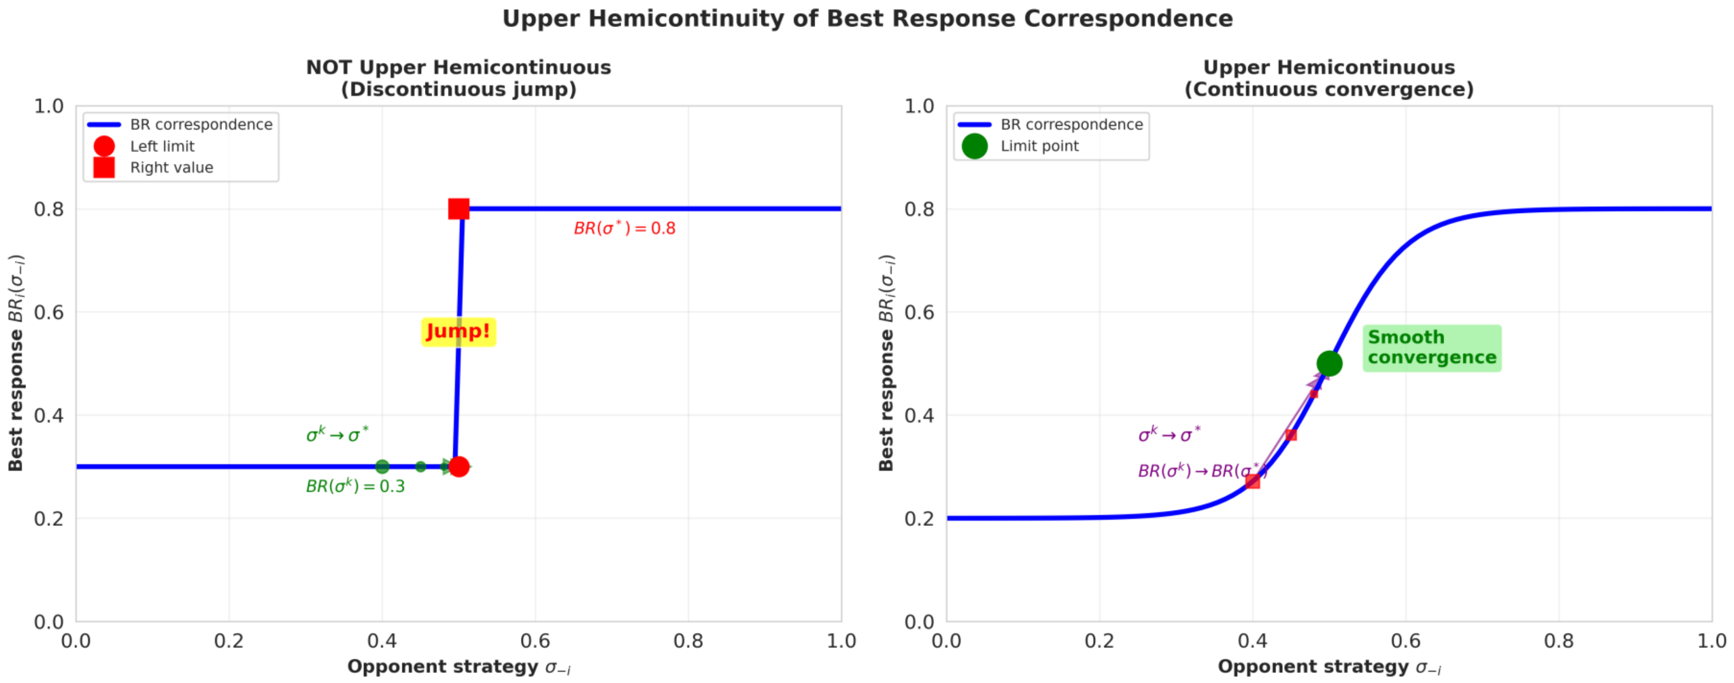

Nash Equilibrium: Formal Definition

In finite $(N, \{A_i\}, \{R_i\})$, π^* is a Nash equilibrium if for all i and all π_i,

$$R_i(\pi_i^*, \pi_{-i}^*) \geq R_i(\pi_i, \pi_{-i}^*)$$

- Pure NE if each π_i^* is a point mass
- Mixed NE if some π_i^* is a distribution

2x2 Mixed NE: Template

For

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}, \; B = \begin{pmatrix} e & f \\ g & h \end{pmatrix},$$

let the row player play U with probability p, and the column player play L with probability q.

- Row indifference: $aq + b(1-q) = cq + d(1-q) \Rightarrow$ solve for q
- Column indifference: $ep + g(1-p) = fp + h(1-p) \Rightarrow$ solve for p

Relationship & Hierarchy

- **Brouwer** is a special case of **Kakutani** (when the correspondence is single-valued, i.e., a function).
- **Banach** is fundamentally different, deals with iterative contractions. Not generally present in game theoretic contexts.
- In Nash's context, **Kakutani** is needed due to set-valuedness of the best response correspondence.

Epsilon-Nash: Definition

A profile π is an ε-Nash equilibrium if for all i,

$$R_i(\pi) \geq \max_{\pi_i'} R_i(\pi_i', \pi_{-i}) - \varepsilon$$

- $\varepsilon = 0$ gives an exact Nash equilibrium
- Useful when using numerical solvers or rounding

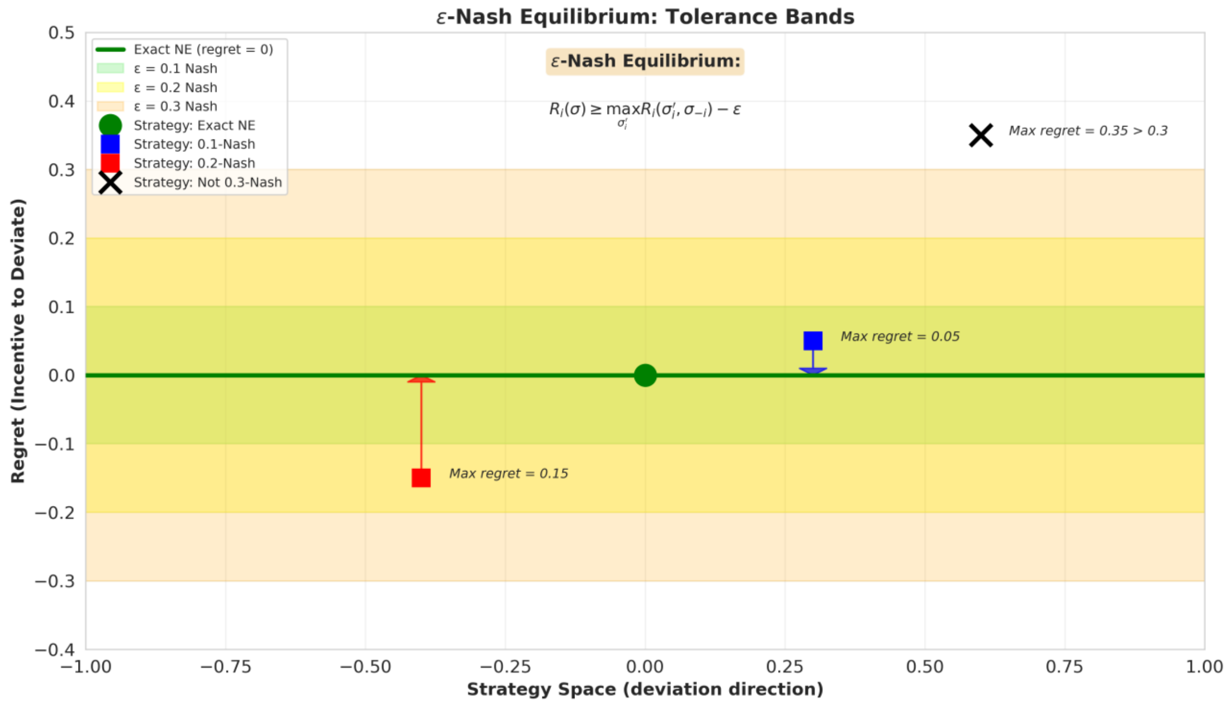

Measuring Epsilon in Finite Games

For each player i:

1. Compute u_i at (π_1, π_2)
2. Compute the best pure-response payoff u_i^{BR}
3. Set $\varepsilon_i = u_i^{BR} - u_i$

Report $\max_i \varepsilon_i$

Game Theory

Lecture 4

Extensive Form Games and Welfare Concepts

1 Extensive Form Games

2 Subgame Perfect E.

3 Welfare Analysis

Why Sequential Games?

Real-world strategic situations often involve:

- **Timing**: Actions happen in sequence, not simultaneously
- **Information**: Players observe some (but not all) past actions
- **Commitment**: First-movers can credibly commit to strategies
- **Threats**: Some threats are credible, others are not

Examples:

- Entry deterrence (incumbent vs entrant)
- Bargaining and negotiation
- Poker and card games
- Job market signaling

From Simultaneous to Sequential

	Extensive Form
Players choose simultaneously	Sequential moves
Complete plans at once	Adaptive strategies
No information revelation	Learning from observations
Nash equilibrium focus	Subgame perfection

Sequential structure allows us to identify and eliminate incredible threats

Example: Entry Deterrence

Setting: New firm considers entering a market. Incumbent can fight or accommodate.

Payoffs:

- Out, -: (0, 10) — monopoly for incumbent
- In, Accommodate: (2, 3) — peaceful duopoly
- In, Fight: (-1, 1) — costly price war

Incredible threat: "If you enter, I'll fight!" — but once entry occurs, fight yields payoff 1 while accommodate yields 3, so the threat is not credible: the incumbent prefers to accommodate.

Question: Should the entrant believe this threat?

Extensive Form Games

An extensive form game with **perfect information** consists of:

- **Nodes X**: decision points in the game tree
- **Player function P**: who moves at each node (or chance)
- **Actions A(h)**: available actions at history h
- **Successor function**: next node after action
- **Payoffs u_i**: terminal payoffs for each player

A **history** is a sequence of actions from the root to some node.

Complete Formal Framework

For imperfect information, an extensive form game is:

$$\Gamma = \langle N, H, P, \{I\_i\}\_\{i \in N\}, A, \rho, \{u\_i\}\_\{i \in N\} \rangle$$

Where:

- $N = \{1, \ldots, n\}$: finite set of players
- H: set of all histories (sequences of actions)
- $Z \subseteq H$: terminal histories (leaves of the game tree, where no further action is taken)
- $P: H \to N \cup \{c\}$: player function (c = chance)
- I_i: collection of information sets for player i
- $A(h)$ or $A(I)$: actions available at history / information set
- ρ: probability distribution for chance moves
- $u\_i: Z \to \mathbb{R}$: payoff function for player i ($\mathbb{R}$ = real numbers)

Histories in Detail

A **history** $h = (a_1, \dots, a_k)$ where:

- Empty history $\emptyset$ is the root
- Each $a_i \in A(h_{i-1})$ where $h_{i-1} = (a_1, \dots, a_{i-1})$
- $Z \subseteq H$ is the set of **terminal histories**

Precedence: $h' \sqsubseteq h$ if h' is a prefix of h

Immediate successors: (h, a) denotes history obtained by adding action a to h

Tree property: For any two histories h, h', either $h \sqsubseteq h'$, $h' \sqsubseteq h$, or they are incomparable

Game Trees and Information Sets

Perfect information: At each node, the player knows the complete history
Imperfect information: Some histories are indistinguishable (grouped into information sets)
Information set I: A set of nodes a player cannot distinguish between

I denotes a single information set; $\mathcal{I}$_i (script I) denotes the collection of all information sets for player i; each I ∈ $\mathcal{I}$_i

For each player $i \in N$, information sets $\mathcal{I}\_i$ partition {h ∈ H : P(h) = i}

Requirements for each I ∈ $\mathcal{I}$_i:
- Same actions: A(h) = A(h') for all h, h' ∈ I
- No self-revelation: if h, h' ∈ I and h ⊑ h', then h = h'

Write A(I) = A(h) for any h ∈ I

Example: In poker, after opponent bets, you don't know their cards

Perfect vs Imperfect Information

Perfect information: Every information set is a singleton

- Player always knows exact position in tree
- Examples: Chess, Go, tic-tac-toe
- Every node induces a subgame

Imperfect information: Some information sets contain multiple nodes

- Player uncertain about past moves
- Examples: Poker, simultaneous-move subgames
- Fewer subgames (only at start of info sets)

Perfect Recall

Player i has **perfect recall** if:

1. They remember their own past actions
2. They remember which information sets they have visited

Formal definition: If h, h' belong to the same information set I of player i, and (h, a) is a continuation where i moved, then there exists a' such that (h', a') is also a continuation and the information sets containing the nodes after (h, a) and (h', a') are the same. In other words, player i cannot be in the same information set at two nodes that differ in i's own prior moves.

Importance:

- Ensures consistency of beliefs and strategies
- Kuhn's theorem (1953): With perfect recall, mixed and behavioral strategies are outcome-equivalent — a player can randomize locally at each information set rather than committing to a global lottery over pure strategies
- Most real games satisfy perfect recall

Counter-example: "Absent-minded driver"

Strategies in Extensive Form

Pure strategy s_i: complete contingent plan

- Specifies an action for *every* information set
- Even off the equilibrium path
- Strategy space: $S_i = \prod_{I \in \mathcal{I}_i} A(I)$

Behavioral strategy b_i: independent randomization at each information set

- $b_i(I)$: probability distribution over $A(I)$
- Local randomization at each decision point

Mixed strategy: distribution over pure strategies

- Global correlation across information sets
- $\sigma_i \in \Delta(S_i)$

Kuhn's Theorem

Theorem (Kuhn, 1953): In games with perfect recall, for any mixed strategy there exists an outcome-equivalent behavioral strategy, and vice versa.

Proof sketch:

- Mixed strategy: probability distribution over pure strategies
- Behavioral strategy: local randomization at each information set
- Perfect recall ensures no conflicts in conditioning

Implications:

- Can work with simpler behavioral strategies
- Probabilities assigned locally at each information set
- No need to track complex correlations

Without perfect recall: Equivalence fails!

Example: Sequential Game Tree

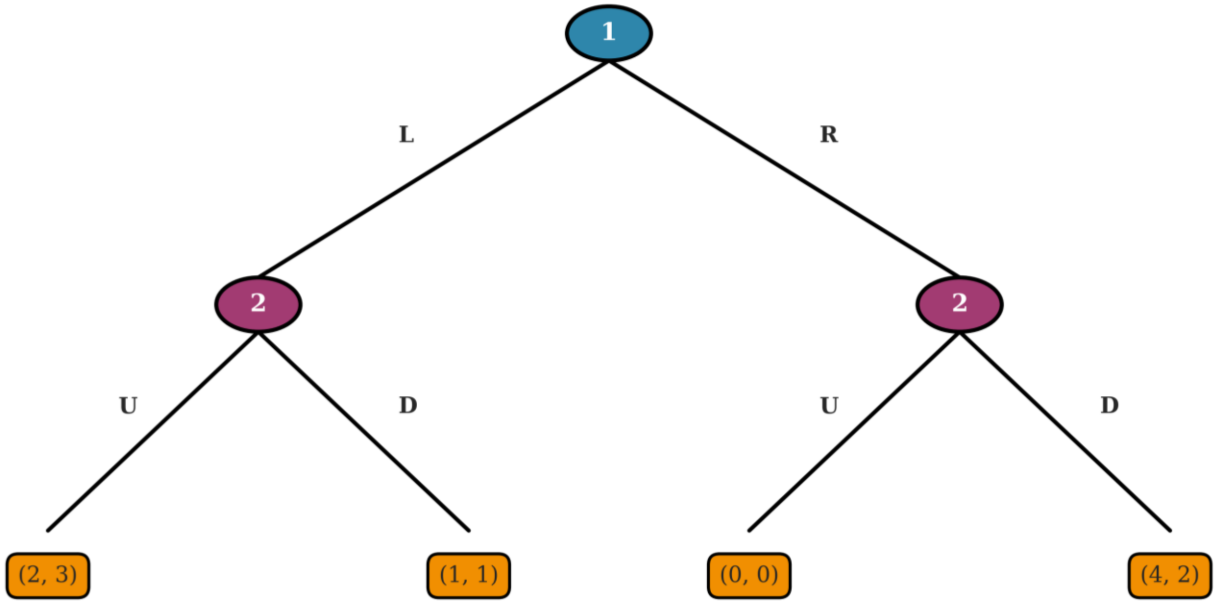

Backward induction solution:

- At left node, Player 2 chooses U (payoff 3 vs 1)
- At right node, Player 2 chooses D (payoff 2 vs 0)
- At root, Player 1 anticipates and chooses R (payoff 4 vs 2)
- **Outcome**: (R, D) with payoffs (4, 2)

Backward Induction

For finite games with **perfect information**:

1. **Initialize**: Mark all terminal nodes with payoffs
2. **Iterate**: For each non-terminal node h where all successors are marked:
   - Let $i = P(h)$ be player to move
   - Find $a^* \in \arg\max_{a \in A(h)} v_i(h, a)$ where $v_i(h, a)$ is marked value at successor
   - Mark h with value $\left(v_1(h, a^*), \ldots, v_n(h, a^*)\right)$
   - Record a^* as BI action at h
3. **Terminate**: When root is marked

Output: Backward induction strategy profile and outcome

Properties of Backward Induction

Theorem: In finite perfect information games, backward induction produces:

1. A well-defined strategy profile
2. The unique subgame perfect equilibrium (in generic games)
3. An outcome path

Genericity: Ties at any node can create multiple BI solutions

Computation: Linear in size of game tree (single pass)

The Centipede Game

Game structure:

- Two players alternate T times
- Each can Take (T) or Pass (P)
- Payoffs grow along the path
- Taking ends game with slight advantage to taker

Payoff structure (simplified):

$$\text{If P1 takes at round } k\text{:} \quad (k+1, k)$$

The Centipede Game (Cont.)

Backward induction:
- At final node, P2 takes (gets $T + 1$ vs T)
- At penultimate node, P1 takes (gets T vs $T - 1$)
- Unraveling continues to root
- Prediction: P1 takes immediately!

Experimental evidence:
- Most players pass several times
- Cooperation emerges
- Payoffs much higher than BI prediction

Explanations:
- Bounded rationality (limited lookahead)
- Social preferences (altruism)
- Beliefs about opponent's rationality
- Trembling hand (fear of mistakes)

The Centipede Game (Cont. 2)

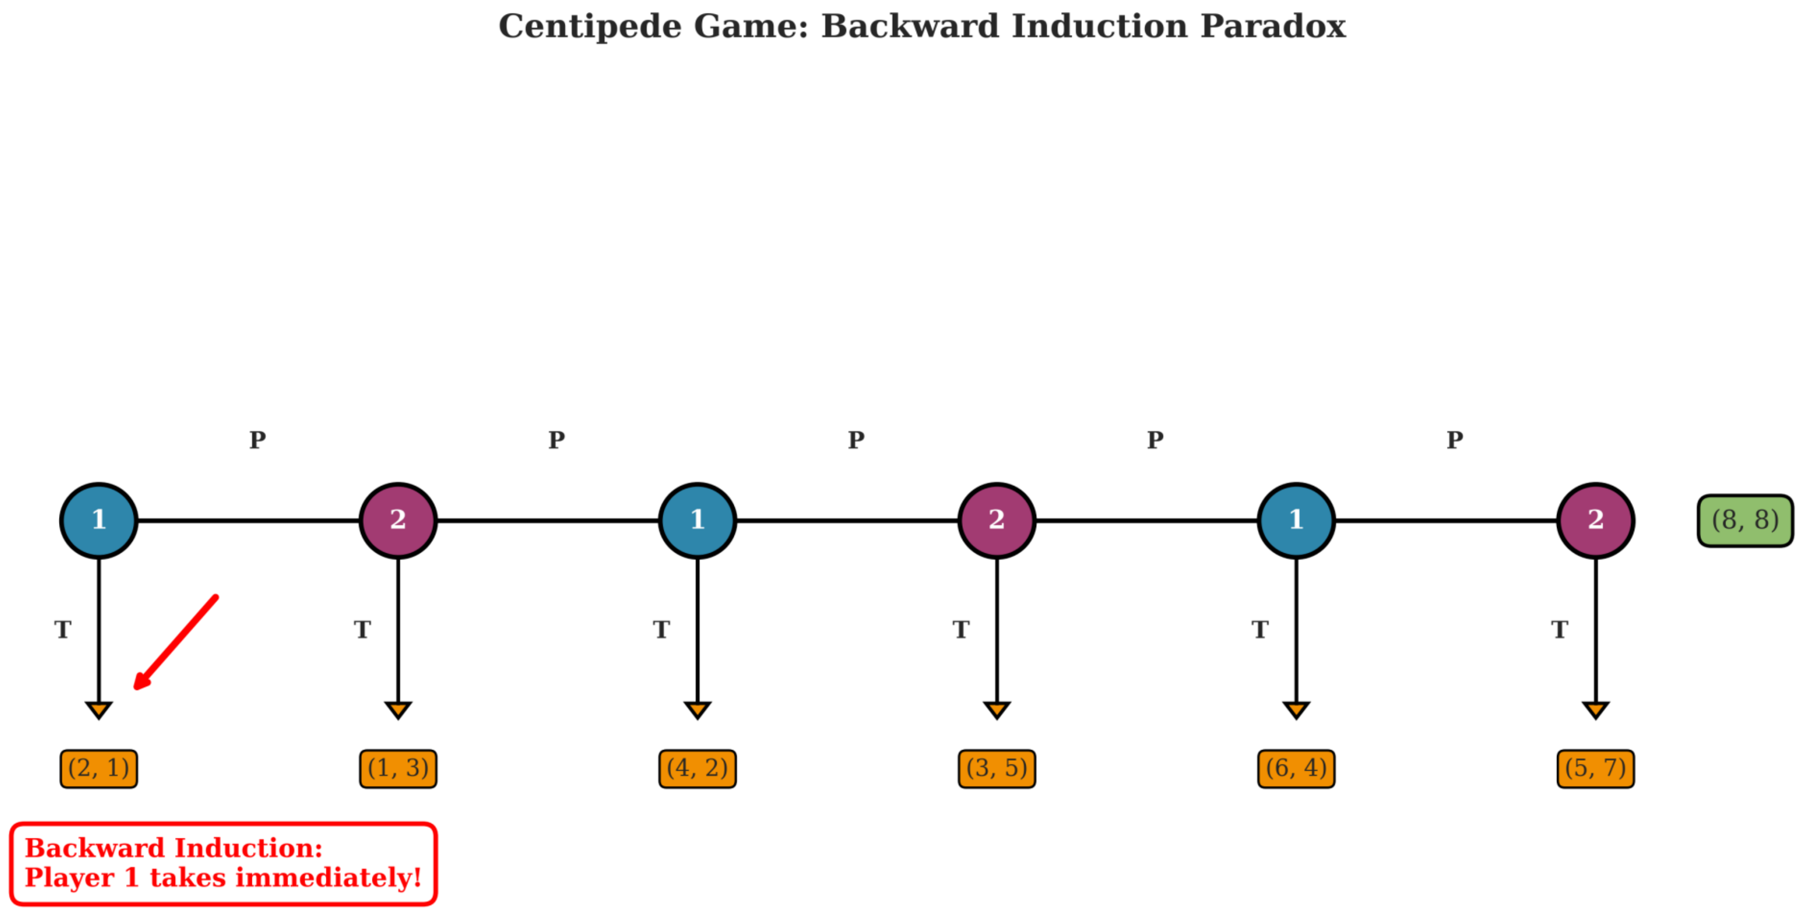

Game Theory

Lecture 4

Extensive Form Games and Welfare Concepts

1 Extensive Form Games

2 Subgame Perfect E.

3 Welfare Analysis

Subgames: Rigorous Definition

A **subgame** of Γ starting at $h \in H$ consists of:

- All histories $h' \sqsupseteq h$
- For each information set I:
- If $I \cap \{h' : h' \sqsupseteq h\} \neq \emptyset$
- Then $I \subseteq \{h' : h' \sqsupseteq h\}$

Intuition: Information sets cannot "straddle" the boundary

Consequence: Under perfect information, every h starts a subgame

Restriction: Strategies restricted to subgame form valid strategies for subgame

SPE: Complete Definition

Definition: Strategy profile s* is a **subgame perfect equilibrium** if for every subgame $\Gamma(h)$:

$$u_i(s^*|_h) \geq u_i(s_i, s^*_{-i}|_h) \quad \forall i, \forall s_i$$

where $s|_h$ denotes the restriction of strategy profile s to the subgame starting at node h (the "|_h" subscript notation means "conditional on being in subgame h")

Equivalently: s* restricted to any subgame is a Nash equilibrium of that subgame

SPE vs Nash

Every SPE is a Nash equilibrium (taking $h = \emptyset$ as the whole game)

Not every Nash is SPE: Non-credible threats can support Nash

One-Shot Deviation Principle

Theorem: In finite games, s^* is SPE if no player can profitably deviate at any single information set, holding all other decisions fixed

Formally: For all i, all $I \in \mathcal{I}_i$, all $a \in A(I)$:

$$u_i(s^*) \geq u_i\left(s_i^{a,I}, s_{-i}^*\right)$$

where $s_i^{a,I}$ plays a at I and follows s_i^* elsewhere

Useful for verification: Only need to check local deviations

Computing SPE in Perfect Information

Relation to Backward Induction:

Theorem (Kuhn 1953): In finite perfect information games:

1. Backward induction outcome is SPE
2. If game is generic (no ties), BI gives unique SPE

Proof sketch:

- Every node is a subgame
- BI ensures optimal play in every subgame
- Therefore SPE

Stackelberg: Full Analysis

Stackelberg duopoly: two-stage quantity competition. Firm 1 (leader) commits to a quantity first; Firm 2 (follower) observes q1 and responds. Solved by backward induction (SPE)

Model:

- Leader chooses $q1 \geq 0$
- Follower observes q1, chooses $q2 \geq 0$
- Price: $P(Q) = a - b(q1 + q2)$
- Cost: c per unit

Follower's problem: max_{q2} [(a – b(q1+q2))q2 – c·q2]

FOC (first-order condition, set $\partial\pi/\partial q2 = 0$): $a - bq1 - 2bq2 - c = 0$

Best response: $q2^*(q1) = (a - c - bq1) / 2b$

Stackelberg: Leader's Problem

Leader solves:

$$\max_{q_1}\left(a - b\left(q_1 + q_2^*(q_1)\right) - c\right) q_1$$

Substituting $q_2^*(q_1)$:

$$= \max_{q_1}\left(\frac{a-c-bq_1}{2}\right) q_1 \max_{q_1}\left(a - b\left(q_1 + \frac{a-c-bq_1}{2b}\right) - c\right) q_1$$

FOC: $\frac{a-c-2bq_1}{2} = 0$

Solution: $q_1^* = \frac{a-c}{2b}, q_2^* = \frac{a-c}{4b}$

Stackelberg vs Cournot

Metric	Cournot (simultaneous)	Stackelberg (sequential)
Leader quantity	$\frac{a-c}{3b}$	$\frac{a-c}{2b}$
Follower quantity	$\frac{a-c}{3b}$	$\frac{a-c}{4b}$
Total output	$\frac{2(a-c)}{3b}$	$\frac{3(a-c)}{4b}$
Leader profit	$\frac{(a-c)^2}{9b}$	$\frac{(a-c)^2}{8b}$
Follower profit	$\frac{(a-c)^2}{9b}$	$\frac{(a-c)^2}{16b}$

First-mover advantage: Leader earns $\frac{(a-c)^2}{8b} > \frac{(a-c)^2}{9b}$

Signaling Games

Timing:

1. Nature draws type $t \in T$ for sender with probability $p(t)$
2. Sender observes t, sends message $m \in M$
3. Receiver observes m (not t), chooses action $a \in A$
4. Payoffs: $u_S(t, m, a)$ and $u_R(t, m, a)$

Information structure:

- Sender knows t
- Receiver forms posterior $\mu(t|m)$ from observing m
- Asymmetric information is key

Spence's Job Market Model

Setup:

- Worker types: $t \in \{H, L\}$ (high / low ability) — privately known to the worker
- Prior: $Pr(t = H) = p$
- Education choice: $e \geq 0$ (observable signal; does not directly raise productivity)
- Employer sets wage w after observing e

Payoffs:

- Worker: $u_W(t, e, w) = w - c(t, e)$, where $c(t, e)$ is the cost of education for type t
- Employer: $u_E(t, w) = v(t) - w$, where $v(t)$ is worker productivity ($v(H) > v(L)$)

Single-crossing condition: $c(H, e) < c(L, e)$ for all $e > 0$ — education is less costly for the high-ability type. This is the key assumption that makes signalling possible: acquiring the signal is relatively cheap for the type it is meant to identify

Separating Equilibrium Construction

Strategy:

- High type: $e_H > 0$
- Low type: $e_L = 0$
- Employer: $w(e_H) = v(H), w(0) = v(L)$

Beliefs: $\mu(H|e_H) = 1, \mu(L|0) = 1$

Incentive constraints:

1. High type prefers e_H: $v(H) - c(H, e_H) \geq v(L)$
2. Low type prefers 0: $v(L) \geq v(H) - c(L, e_H)$

Existence: Requires $\frac{v(H)-v(L)}{c(L,e_H)-c(H,e_H)} \in (0,1)$

Separating vs Pooling Equilibria

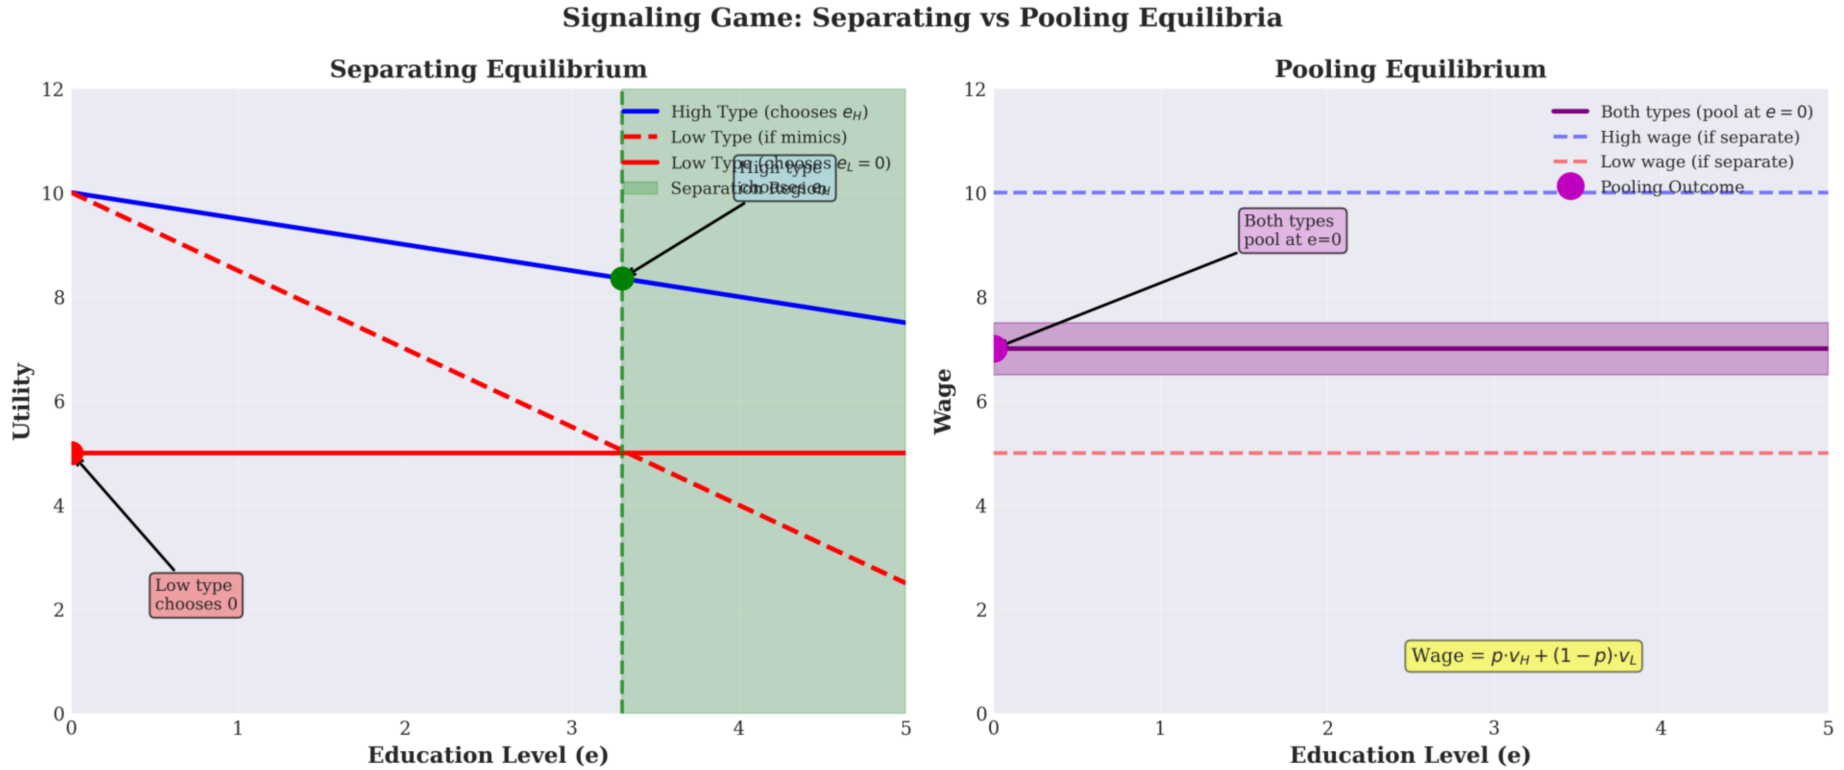

Cost differential between types enables separation. In separating equilibrium, high types signal credibly through costly education that low types won't mimic.

Pooling Equilibrium

Strategy: Both types choose $e^* = 0$

Beliefs: $\mu(H|0) = p$ (prior)

Wage: $w(0) = p \cdot v(H) + (1 - p) \cdot v(L)$

Off-equilibrium beliefs: Must specify $\mu(\cdot\,|e)$ for $e > 0$

Stability: Typically fails refinements (Cho-Kreps intuitive criterion)

Perfect Bayesian Equilibrium

Components: (s, μ) where $s = strategies, \mu = beliefs$

Requirements:

- **Sequential rationality:** At each information set, action maximizes expected payoff given beliefs
- **Consistency:** Beliefs derived via Bayes' rule on equilibrium path
- **Off-path beliefs:** Must be specified (various refinements restrict these)

For sender at t:

$$m^*(t) \in \arg\max_m E_{a \sim a^*(m)}[u_S(t, m, a)]$$

For receiver observing m:

$$a^*(m) \in \arg\max_a E_{t \sim \mu(\cdot|m)}[u_R(t, m, a)]$$

Refinements of PBE

Cho-Kreps Intuitive Criterion: Eliminate equilibria where off-path beliefs are "unreasonable"

Idea: If only one type benefits from deviation, receiver should believe deviation came from that type

Universal Divinity: Stronger refinement based on divinity ordering

D1 Refinement: Even stronger, often selects unique equilibrium

Application: Typically eliminates pooling equilibria in Spence model

Game Theory

Lecture 4

Extensive Form Games and Welfare Concepts

1 Extensive Form Games

2 Subgame Perfect E.

3 Welfare Analysis

Pareto Efficiency: Formal Definition

Outcome z is **Pareto dominated** by z' if:

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

with strict inequality for at least one i

Outcome z is **Pareto efficient** (or Pareto optimal) if no z' Pareto dominates it

Set of Pareto efficient outcomes: Pareto frontier

Computing Pareto Frontier

For finite games:

1. List all possible outcomes $z_1, \ldots, z_m$
2. For each outcome z_k:
   - Check if $\exists z_j$ such that $u_i(z_j) \geq u_i(z_k)$ for all i with strict inequality for some i
   - If yes, z_k is dominated; if no, z_k is Pareto efficient
3. Pareto frontier $=$ all non-dominated outcomes

Computational complexity: $O(m^2 \cdot n)$ for m outcomes, n players

Example: Prisoner's Dilemma Revisited

	C	D
C	$(3, 3)$	$(0, 4)$
D	$(4, 0)$	$(1, 1)$

Nash equilibrium: (D,D) with payoffs (1,1)

Pareto frontier: Only (C,C) with payoffs (3,3)

Welfare loss: Nash equilibrium is Pareto dominated!

Intuition: Individual rationality conflicts with collective welfare

Visualizing the Pareto Frontier

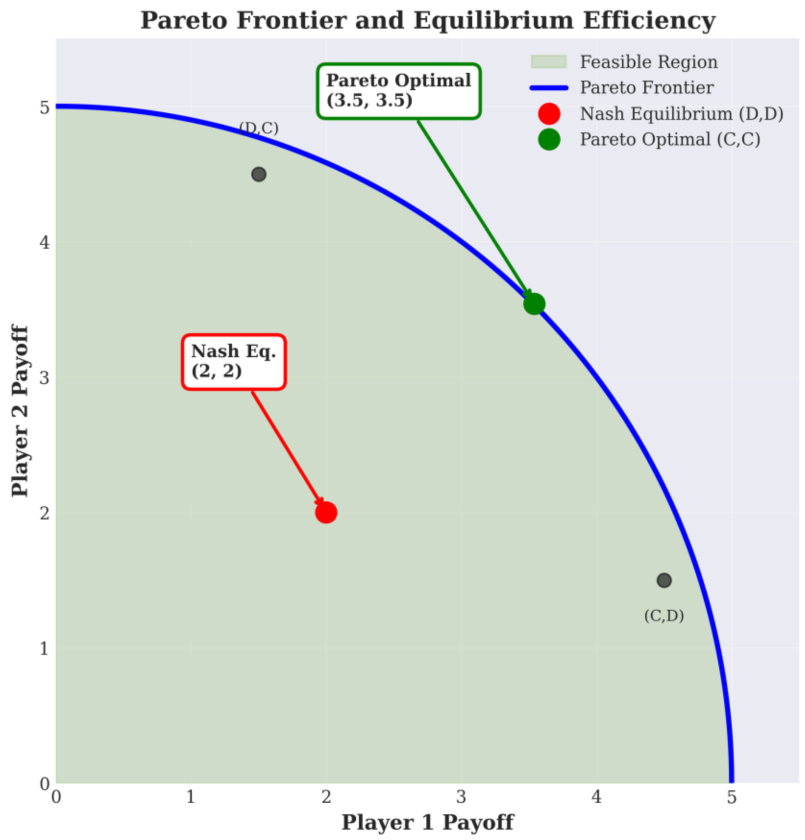

Nash equilibrium (red) is strictly dominated by Pareto optimal outcomes (green). The feasible region shows all achievable payoffs, and the Pareto frontier represents the efficient boundary.

Price of Anarchy (PoA)

Definition: Ratio of welfare at worst Nash equilibrium to optimal welfare

For maximization problems:

$$\text{PoA} = \frac{\max_{z \in Z} W(z)}{\min_{z \in \text{NE}} W(z)}$$

For minimization problems (e.g., congestion):

$$\text{PoA} = \frac{\max_{z \in \text{NE}} C(z)}{\min_{z \in Z} C(z)}$$

where W is welfare and C is cost

PoA: Interpretations

High PoA $(\gg 1)$**:**

- Large inefficiency of equilibrium
- Strong need for coordination or regulation
- Example: Braess's paradox in traffic

PoA = 1:

- All Nash equilibria are optimal
- No inefficiency from selfish behavior
- Rare but important class

Typical values:

- Prisoner's Dilemma: $\text{PoA} = 3$ (welfare $3\ vs\ 1$)
- Congestion games: Can be unbounded
- Auctions: Often modest $(< \ 2)$

Price of Anarchy in Various Games

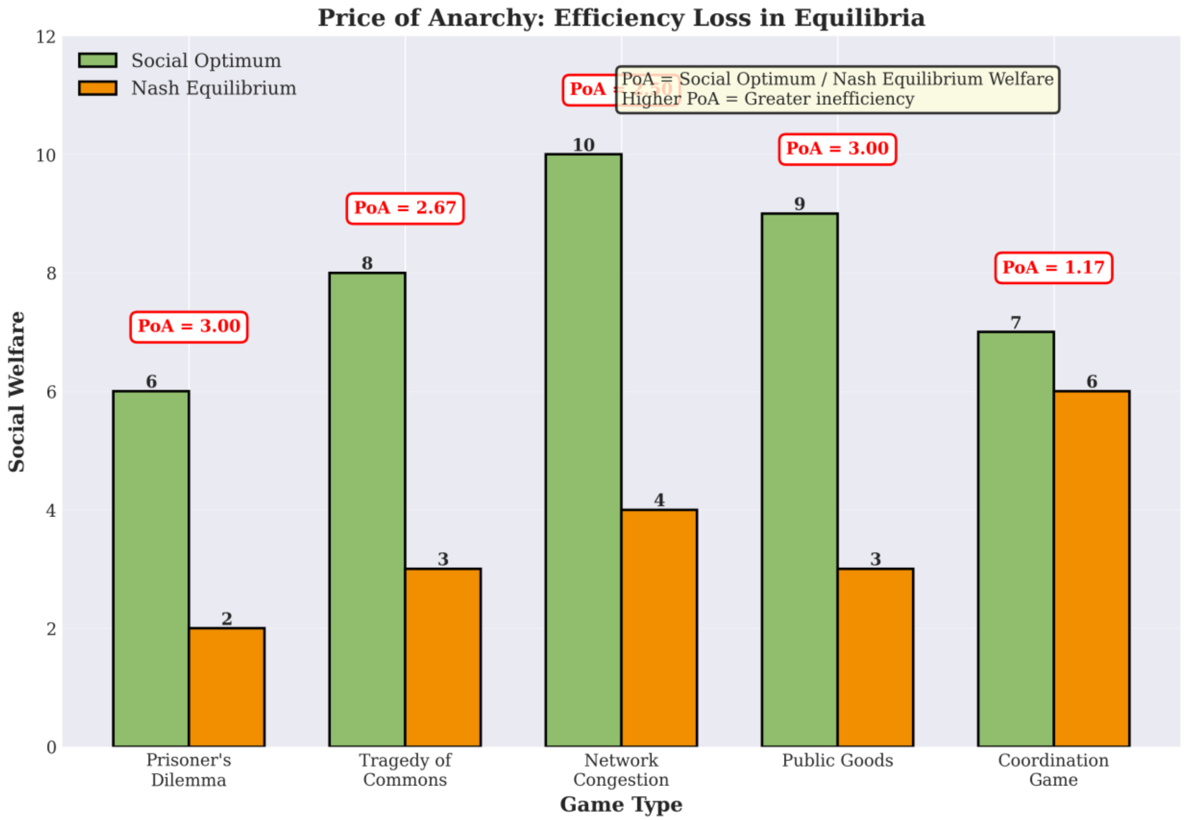

Higher PoA indicates greater inefficiency. Games like Prisoner's Dilemma show substantial welfare loss at equilibrium, motivating the need for cooperation mechanisms or regulation.

Computing PoA

Algorithm:

1. Find all Nash equilibria
2. Compute welfare for each NE
3. Find worst-case NE welfare: $W_{\min}^{\text{NE}}$
4. Solve optimization problem for maximum welfare: W^*
5. Compute ratio: $\text{PoA} = W^*/W_{\min}^{\text{NE}}$

Challenges:

- Finding all NE can be hard (PPAD-complete)
- Computing optimal welfare may be NP-hard
- Bounds often derived analytically

Example: Coordination Game

	A	B
A	$(4, 4)$	$(0, 0)$
B	$(0, 0)$	$(1, 1)$

Nash equilibria: (A,A) with welfare 8, and (B,B) with welfare 2

Optimal welfare: 8 (at (A,A))

PoA: $8/2 \;=\; 4$

Interpretation: Coordination failure leads to $4\times$ welfare loss

Social Welfare Functions

Utilitarian (sum):

$$W_{\text{util}}(z) = \sum_{i=1}^{n} u_i(z)$$

Egalitarian (min):

$$W_{\text{egal}}(z) = \min_{i=1,\ldots,n} u_i(z)$$

Nash bargaining:

$$W_{\text{Nash}}(z) = \prod_{i=1}^{n} (u_i(z) - d_i)$$

where d_i is disagreement payoff

Mechanism Design Preview

Central question: Can we design game rules so that equilibrium achieves desired outcome?

Revelation principle: Any outcome achievable by any mechanism can be achieved by truthful direct mechanism

Applications:

- Auctions (maximize revenue or efficiency)
- Voting (aggregate preferences)
- Matching markets (stable assignments)
- Public goods provision

Application: Auction Design

English auction (ascending):

- Sequential: bidders observe and respond
- SPE: bid up to true value
- Efficient outcome

Sealed-bid first-price:

- Simultaneous: single-stage game
- Nash: shade bids below value
- Revenue equivalence (under symmetry)

Sequential structure can extract more information

Application: Bargaining

Rubinstein bargaining:

- Infinite horizon alternating offers
- Players discount future at rate $\delta < 1$
- SPE exists and is unique!

Outcome: Split pie as $\left(\frac{1}{1+\delta}, \frac{\delta}{1+\delta}\right)$

Interpretation: First-mover advantage diminishes with patience

Application: Reputation

Reputation as commitment:

- Long-run player faces sequence of short-run opponents
- Even if accommodating is optimal in one-shot, fighting can be optimal to build reputation
- Reputation effects require sequential structure

Chain store paradox (Selten):

- Backward induction says accommodate
- But with many potential entrants, building tough reputation valuable
- Resolution: Incomplete information about incumbent type

Finite vs Infinite Horizon

Finite horizon:

- Backward induction applies
- Unraveling from final period
- Cooperation difficult to sustain

Infinite horizon:

- No final period to unravel from
- Folk theorems: many outcomes sustainable

Exercise 1

Consider the following 3-player sequential game:

- Player 1 chooses L or R
- If L: Player 2 chooses U or D
- If R: Player 3 chooses U or D
- Payoffs: (L,U): $(3,2,1)$, (L,D): $(1,4,2)$, (R,U): $(2,3,5)$, (R,D): $(0,1,3)$

Find all SPE

Exercise 2

In a signaling game with three types $\{L, M, H\}$ with probabilities$(0.2,0.3,0.5)$ and two signals $\{e_0, e_1\}$:

- Costs: $c(L, e_1) = 4, c(M, e_1) = 2, c(H, e_1) = 1, c(\cdot, e_0) = 0$
- Values: $v(L) = 2, v(M) = 4, v(H) = 6$

Find a separating PBE where types separate into two groups.

Exercise 3

Consider a coordination game where both players choose A or B:

- Payoffs: (A, A): (x, x), (B, B): (y, y), miscoordination: $(0,0)$
- Assume $x > y > 0$

1. Characterize all Nash equilibria
2. Compute the Price of Anarchy
3. cHow does PoA change with x/y?

Exercise 4

In a traffic routing game:

- Two routes: fast road (capacity 100) and slow road (unlimited capacity)
- Cost on fast road: $c_f(n) = n$ if $n \leq 100$, ∞ otherwise
- Cost on slow road: $c_s(n) = 50$ for any n
- 150 drivers

Find:

1. Nash equilibrium flow
2. Social optimum flow
3. Price of Anarchy

Advanced Topic: Trembling Hand Perfection

Idea: Players may make mistakes with small probability

Trembling hand perfect equilibrium:

- Limit of equilibria in perturbed games where each action played with probability $\geq \varepsilon$
- As $\varepsilon \to 0$, converges to trembling hand perfect equilibrium

Property: Eliminates weakly dominated strategies

Relation to SPE: More restrictive refinement

Summary

- **Extensive form**: Trees, histories, information sets, perfect recall
- **Backward induction**: Algorithm for perfect information games
- **SPE**: Eliminates non-credible threats, refines Nash
- **Signaling games**: Asymmetric information, PBE, separating vs pooling
- **Welfare**: Pareto efficiency, Price of Anarchy, mechanism design

Course Textbooks

- Bonanno, G. (2024). ***Game Theory (3rd ed.)***. University of California, Davis. Received from: GT Book
- Axelrod, R. (1984). ***The Evolution of Cooperation***. Basic Books. Received from: Axelrod Article
- Nisan, N., Roughgarden, T., Tardos, É., & Vazirani, V. V. (2007). ***Algorithmic Game Theory***. Cambridge University Press. Received from: AGT Book
- Myerson, R. B. (1991). ***Game Theory: Analysis of Conflict.*** Harvard University Press. Received from: GT Book 2F.
- Christianos et al., ***Multi-Agent Reinforcement Learning: Foundations and Modern Approaches***, 2023. Received from: MARL Book.pdf
- Shoham, Y., & Leyton-Brown, K. (2008). ***Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations***. Cambridge University Press. Received from: MARL Book.pdf
- `***nashpy***` documentation (readthedocs). Link: NashPy Docs

That's All for Today!