Module 1: Foundations of Game Theory

CSCE 631 — Algorithmic Game Theory meets LLMs

Week 1: May 26 (Tue) – May 30 (Fri)
(Short week — Memorial Day Monday)

Contents

Learning Objectives

By the end of this module, you should be able to:

1. Define a normal-form game formally: players, strategy sets, and utility functions.
2. Compute best responses and verify whether a given strategy profile is a Nash equilibrium.
3. Compute maxmin values and maxmin strategies in two-player games.
4. Explain the difference between Nash equilibrium, maxmin, and correlated equilibrium as solution concepts.
5. Identify and eliminate strictly dominated strategies using iterated elimination (IESDS).

1 Lecture 1: Introduction to Game Theory — Games in Normal Form (49 min)

Game theory provides the mathematical language for reasoning about strategic interaction. When two or more decision-makers have interdependent outcomes—where what is best for one depends on what the others do—we leave the realm of single-agent optimization and enter the world of games. This lecture introduces the fundamental representation we will use throughout Weeks 1–2: the *normal form*.

1.1 The Normal-Form Game

Definition: Normal-Form Game

A **normal-form game** (also called a *strategic-form game*) is a tuple $\Gamma = (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ where:

- $N = \{1, 2, \ldots, n\}$ is a finite set of **players**.
- S_i is the finite set of **pure strategies** available to player i.
- $u_i : S_1 \times S_2 \times \cdots \times S_n \to \mathbb{R}$ is the **utility function** (or payoff function) for player i.

A **strategy profile** $s = (s_1, s_2, \ldots, s_n)$ specifies one strategy per player. We write s_{-i} for the strategies of all players except i, so that $s = (s_i, s_{-i})$.

1.2 Mixed Strategies

A **mixed strategy** for player i is a probability distribution $\sigma_i \in \Delta(S_i)$ over the pure strategies. If player i plays mixed strategy σ_i and the other players play σ_{-i}, the expected utility is:

$$u_i(\sigma) = \sum_{s \in S} \left(\prod_{j \in N} \sigma_j(s_j) \right) u_i(s).$$

This expression is *multilinear*: linear in each σ_j when the others are fixed. Multilinearity is what makes best-response analysis tractable—it means that a best response to a fixed opponent strategy always includes a pure strategy.

1.3 Canonical Examples

Example: Prisoner's Dilemma

Two suspects are interrogated separately. Each can *cooperate* (stay silent) or *defect* (testify). The payoff matrix is:

	C	D
C	$-1, -1$	$-4,\ 0$
D	$0, -4$	$-3, -3$

Both players have a *dominant strategy* to defect, leading to the outcome $(-3, -3)$, which is Pareto-dominated by $(-1, -1)$ (the mutual cooperation outcome). This tension between individual rationality and collective welfare is the central insight.

Example: Matching Pennies

A zero-sum game with no pure-strategy Nash equilibrium:

	H	T
H	$1, -1$	$-1,\ 1$
T	$-1,\ 1$	$1, -1$

The unique Nash equilibrium is the mixed strategy $(1/2,\ 1/2)$ for both players. This game motivates why mixed strategies are essential.

Common Pitfalls

- Confusing the *game* (the structure—players, strategies, payoffs) with a *solution* (a Nash equilibrium). The payoff matrix describes the game; the equilibrium is a prediction about play.
- Assuming mixed strategies require physical randomization. A mixed strategy can represent uncertainty about an opponent's belief.
- Forgetting that expected utility is *multilinear* in mixed strategies (linear in each player's mixture given the others fixed)—this is what makes best-response analysis tractable.

Connection: PA1

PA1 Task 1 asks you to compute expected utilities under mixed strategies. Make sure you can set up the bilinear form $u_i(\sigma_i, \sigma_{-i})$ for a two-player game.

2 Lecture 2: Chess (41 min)

Chess serves as the canonical example of a two-player, zero-sum, perfect-information, sequential, deterministic game. It illustrates concepts that generalize far beyond chess: game trees, backward induction, and the computational limits of exact solutions.

2.1 Game Trees and Backward Induction

Definition: Game Tree

A **game tree** is a rooted tree where internal nodes are labeled with the acting player, edges represent actions, and leaves carry payoff vectors. The tree represents the sequential structure of the game.

Backward induction solves a finite game tree from the leaves upward. At each decision node, the acting player selects the action that maximizes their payoff, assuming all future players also play optimally. For two-player zero-sum games with perfect information, this computes the *game value*—the payoff that both players can guarantee with optimal play.

Theorem: Zermelo's Theorem (1913)

Every finite two-player zero-sum game with perfect information has a determined value. Both players have optimal pure strategies computable by backward induction.

2.2 Computational Limits

The **Shannon number** ($\approx 10^{120}$) is the estimated game-tree complexity of chess. This makes exact backward induction infeasible. **Alpha-beta pruning** reduces the effective search by eliminating branches that cannot affect the minimax value at the root, but the worst case remains exponential.

Common Pitfalls

- Backward induction only applies in finite, perfect-information games. It breaks down with imperfect information (Week 3).
- Zermelo's theorem (perfect-information games) is distinct from von Neumann's minimax theorem (zero-sum normal-form games with mixed strategies).

3 Lecture 3: Solution Concepts — Nash Equilibria, Maxmin (68 min)

This lecture develops the two most important solution concepts for normal-form games: Nash equilibrium and maxmin strategies. Understanding when they coincide (zero-sum games) and when they diverge (general-sum games) is fundamental.

3.1 Best Responses and Nash Equilibrium

Definition: Best Response

Strategy σ_i^* is a **best response** to σ_{-i} if

$$u_i(\sigma_i^*, \sigma_{-i}) \geq u_i(\sigma_i, \sigma_{-i}) \quad \text{for all } \sigma_i \in \Delta(S_i).$$

We write $\sigma_i^* \in \mathrm{BR}_i(\sigma_{-i})$.

Definition: Nash Equilibrium

A strategy profile $\sigma^* = (\sigma_1^*, \ldots, \sigma_n^*)$ is a **Nash equilibrium** if every player is simultaneously playing a best response:

$$\sigma_i^* \in \mathrm{BR}_i(\sigma_{-i}^*) \quad \text{for all } i \in N.$$

Theorem: Nash's Theorem (1951)

Every finite game (with a finite number of players and action profiles) has at least one Nash equilibrium (possibly in mixed strategies). The proof uses Brouwer's fixed-point theorem applied to a continuous function constructed from the best-response correspondence.

3.2 The Support Condition

In a mixed-strategy NE, every pure strategy in the *support* (played with positive probability) must yield the same expected payoff. This **indifference principle** is the workhorse for solving 2×2 games by hand: set up the equations requiring the opponent to be indifferent, then solve.

3.3 Maxmin Strategies

Definition: Maxmin Value

The **maxmin value** of player i is:

$$\overline{v}_i = \max_{\sigma_i} \min_{\sigma_{-i}} u_i(\sigma_i, \sigma_{-i}).$$

This is the highest payoff player i can *guarantee*, regardless of what the opponents do. The maximizing strategy is the **maxmin strategy**.

Theorem: Von Neumann's Minimax Theorem (1928)

In two-player zero-sum games:

$$\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2) = \min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2).$$

The maxmin value equals the minimax value. Every NE achieves this value, and every NE strategy is a maxmin strategy.

3.4 NE vs. Maxmin in General-Sum Games

In general-sum games, NE and maxmin can differ substantially:

- NE payoffs are always individually rational: $u_i(\sigma^*) \geq \overline{v}_i$.
- Multiple NE can exist with different payoffs (unlike zero-sum).
- Maxmin strategies are pessimistic; NE strategies assume mutual rationality.

Common Pitfalls

- A Nash equilibrium is not necessarily efficient—the Prisoner's Dilemma NE is Pareto-dominated.
- The indifference principle only applies in non-degenerate games.

- Maxmin is *not* the same as Nash in general-sum games.

Connection: PA1

PA1 Tasks 2–3 ask you to check pure best responses and verify Nash equilibria (both weak and strict). The indifference principle is essential for solving 2×2 games by hand.

4 Lecture 4: Dominated Strategies, Correlated Equilibria (66 min)

This lecture extends the solution-concept landscape with several important ideas: dominated strategies and rationalizability (which simplify analysis), minimax regret (an alternative to maxmin), and correlated equilibrium (which relaxes Nash's independence assumption and can achieve better outcomes).

4.1 Dominated Strategies

Definition: Strict Dominance

Strategy s_i **strictly dominates** s_i' if $u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$ for *all* opponent profiles s_{-i}. A rational player never plays a strictly dominated strategy.

Iterated elimination of strictly dominated strategies (IESDS) removes dominated strategies iteratively. The order of elimination does not matter, and the result is unique. IESDS preserves all Nash equilibria.

Weak dominance: s_i weakly dominates s_i' if $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$ for all s_{-i}, with strict inequality for at least one s_{-i}. Unlike strict dominance, the order of weak-dominance elimination *does* matter, and eliminating weakly dominated strategies may remove Nash equilibria.

Very weak dominance: s_i very weakly dominates s_i' if $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$ for all s_{-i} (no strict inequality required).

4.2 Rationalizability

A strategy is **rationalizable** if it is a best response to some belief about the other players' strategies, where those beliefs are themselves consistent with rationality. Formally, define an iterative process: start with all strategies, and at each step keep only those that are best responses to beliefs supported on the surviving strategies of the opponents.

In two-player games, the set of rationalizable strategies coincides exactly with the set of strategies surviving iterated removal of strictly dominated strategies.

4.3 Minimax Regret

The **regret** of playing action a_i when opponents play a_{-i} is $\max_{a_i' \in A_i} u_i(a_i', a_{-i}) - u_i(a_i, a_{-i})$: the gap between the best possible payoff and the actual payoff given the opponent's actions. The **maximum regret** of action a_i is the worst-case regret over all opponent profiles. The **minimax-regret** actions minimize this worst-case regret.

4.4 Correlated Equilibrium

Definition: Correlated Equilibrium

A **correlated equilibrium** (CE) is a probability distribution p over strategy profiles $S = S_1 \times \cdots \times S_n$ such that, for each player i and each pair of strategies $s_i, s_i' \in S_i$:

$$\sum_{s_{-i}} p(s_i, s_{-i}) \left[u_i(s_i, s_{-i}) - u_i(s_i', s_{-i}) \right] \geq 0.$$

Intuitively: given the signal "play s_i," no player wants to deviate.

A **coarse correlated equilibrium** (CCE) is weaker: no player can improve by deviating *before* seeing their signal. The inclusion hierarchy is:

$$\text{NE} \subseteq \text{CE} \subseteq \text{CCE}.$$

The set of correlated equilibria is a *convex polytope* defined by linear inequalities, so any linear objective can be optimized over it via LP in polynomial time. This computational advantage over NE is a major reason CE is practically important.

Example: Coin Flip in Battle of the Sexes

In the Battle of the Sexes, a trusted coin flip (correlation device) can coordinate the two players to alternate between (LW, LW) and (WL, WL) with equal probability, achieving a correlated equilibrium with expected payoff 1.5 per player—better than the mixed-strategy Nash equilibrium.

Common Pitfalls

- Don't eliminate weakly dominated strategies freely—doing so can remove Nash equilibria.
- In two-player games, rationalizable strategies = strategies surviving IESDS, but this equivalence does not hold for $n > 2$ players.
- A correlated equilibrium requires a trusted mediator. Without one, players are limited to NE.
- CE $\neq$ "players randomize together." The signal structure matters: each player sees only their own recommendation.

Connection: PA1

PA1 Task 4 covers strictly dominated strategy detection (pairwise pure-strategy dominance). Task 5 asks you to verify correlated equilibrium constraints. Understanding the distinction between strict and weak dominance is crucial.

Key Definitions Summary

Term	Definition
Normal-form game	$(N, \{S_i\}, \{u_i\})$
Mixed strategy	$\sigma_i \in \Delta(S_i)$
Best response	$\sigma_i^* = \arg\max_{\sigma_i} u_i(\sigma_i, \sigma_{-i})$
Nash equilibrium	All players simultaneously play best responses
Maxmin value	$\max_{\sigma_i} \min_{\sigma_{-i}} u_i(\sigma_i, \sigma_{-i})$
Dominated strategy	s_i such that some other s_i' always yields higher payoff
Correlated equilibrium	Distribution over profiles with no profitable deviations given signals

Key Takeaways

1. Normal-form games capture simultaneous strategic interaction with the tuple $(N, \{S_i\}, \{u_i\})$.
2. Nash equilibrium is a fixed point of the best-response correspondence; Nash's theorem guarantees existence in finite games.
3. In zero-sum games, NE, maxmin, and minimax all coincide—the game has a unique value. In general-sum games, they can differ.
4. Correlated equilibrium relaxes NE's independence requirement and is computable in polynomial time via LP.
5. IESDS is a safe simplification tool; weak dominance elimination is not.

Suggested Reading

- Nisan, Roughgarden, Tardos, Vazirani: *Algorithmic Game Theory*, Chapters 1–2.
- Shoham & Leyton-Brown: *Multiagent Systems*, Chapters 3–4.