Introduction to Game Theory

Lecture 10

Div Bhagia

Introduction

Why Game Theory?

So far, we’ve mostly modeled agents as making decisions against an anonymous market: a consumer takes prices as given, a competitive firm takes the market price as given. Nobody’s individual action affects the environment that others face.

Game theory is the toolkit for situations where that assumption fails, where the payoff to your action depends on what others do, and they know the same about you.

- Two firms setting prices in a duopoly
- A worker deciding how hard to bargain with an employer
- Countries choosing whether to honor a trade agreement
- A central bank setting policy while private actors anticipate it

What Makes a Game?

Every game has three ingredients:

1. **Players**: who is making decisions
2. **Strategies**: what each player can do
3. **Payoffs**: what each player gets as a function of the strategies chosen

We'll also need to specify the **information structure** (who knows what when) and the **timing** (who moves when).

A game is fully specified once you know all of these.

Two Representations

Games come in two standard forms:

- **Normal (strategic) form:** a payoff matrix listing strategies and payoffs. Best for **simultaneous-move** games.
- **Extensive form:** a game tree showing the order of moves and information sets. Best for **sequential-move** games.

Both representations can describe the same game. The choice is about which aspect, strategies or timing, you want to emphasize.

Normal Form Games

The Prisoner’s Dilemma

Two prisoners are separated into individual rooms and cannot communicate with each other. Each has two options: **Cooperate** (stay silent) or **Defect** (testify). The payoffs in years they will have to serve are:

	B Cooperates	**B Defects**
A Cooperates	(-1, -1)	(-3, 0)
A Defects	(0, -3)	(-2, -2)

Each player does better by defecting regardless of what the other does:

- If B cooperates: A gets 0 years by defecting vs. 1 year by cooperating
- If B defects: A gets 2 years by defecting vs. 3 years by cooperating

Dominant Strategies

A strategy is **strictly dominant** for a player if it yields a strictly higher payoff than any alternative, regardless of what opponents do.

In the Prisoner's Dilemma, **Defect** is strictly dominant for both players. So the prediction is (Defect, Defect), giving both players 2 years each, even though (Cooperate, Cooperate) would give them 1 year each.

This is the central puzzle: individually rational behavior leads to a collectively bad outcome.

The Prisoner's Dilemma shows up everywhere: arms races, price wars, overfishing, tax evasion, doping in sports.

Nash Equilibrium

Not all games have dominant strategies. We need a more general solution concept.

Nash equilibrium (Nash, 1950): a profile of strategies $(s_1^*, \ldots, s_n^*)$ such that each player's strategy is a **best response** to the others'. No player has a profitable unilateral deviation.

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*) \quad \text{for all } s_i, \ \text{for all } i$$

Interpretation: a Nash equilibrium is a "no regrets" profile. Given what everyone else is doing, no one wishes they had chosen differently.

Finding Nash Equilibria

To find Nash equilibria in a normal-form game:

1. For each player, fix the opponent's strategy and find the best response
2. Underline or circle the payoffs corresponding to those best responses
3. A Nash equilibrium is any cell where **both players are playing best responses**

In the Prisoner's Dilemma, (Defect, Defect) is the unique Nash equilibrium: both cells where the other defects have Defect underlined, and both cells where the other cooperates also have Defect underlined.

Example: Coordination Game

Two friends want to meet but can't communicate. Each can go to a **Cafe** or a **Bar**.

	B: Cafe	**B: Bar**
A: Cafe	(2, 2)	(0, 0)
A: Bar	(0, 0)	(1, 1)

Two pure-strategy Nash equilibria: (Cafe, Cafe) and (Bar, Bar).

The game has **multiple equilibria**, raising the question of how players coordinate on one, and possibly whether some equilibria are more "focal" than others.

Example: Battle of the Sexes

Same structure, but the players prefer different outcomes. A prefers **Opera**, B prefers **Football**, but both prefer being together.

	B: Opera	**B: Football**
A: Opera	(2, 1)	(0, 0)
A: Football	(0, 0)	(1, 2)

Two pure-strategy Nash equilibria: (Opera, Opera) and (Football, Football). Coordination is good, but there's **conflict** about which equilibrium to land on.

Example: Matching Pennies

A zero-sum game. Player A wants the pennies to **match**, Player B wants them to **differ**.

	B: Heads	**B: Tails**
A: Heads	$(1, -1)$	$(-1, 1)$
A: Tails	$(-1, 1)$	$(1, -1)$

No pure-strategy Nash equilibrium. For any pure profile, one player wants to switch.

We need **mixed strategies**.

Mixed Strategies

A **mixed strategy** is a probability distribution over pure strategies.

A mixed-strategy Nash equilibrium is a profile of randomizations such that each player is **indifferent** between the pure strategies in the support of their mix, given the opponents’ mix.

Why indifference? If a player were strictly better off playing one pure strategy, they would play it with probability 1, not mix. Mixing is only rational when you are indifferent.

Matching Pennies: Solving for the Mix

Let p = probability A plays Heads, q = probability B plays Heads.

A is indifferent between Heads and Tails when B makes both equally attractive:

$$q \cdot 1 + (1 - q)(-1) = q(-1) + (1 - q)(1)$$

$$2q - 1 = 1 - 2q \implies q = 1/2$$

By symmetry, $p = 1/2$. The unique Nash equilibrium is $(p, q) = (1/2, 1/2)$: each player randomizes 50-50.

Expected payoff: 0 for each player. Exactly what you'd guess by symmetry.

Existence

Nash's theorem (1950): Every finite game (finitely many players, finitely many pure strategies) has at least one Nash equilibrium, possibly in mixed strategies.

This is reassuring: we can always find *some* equilibrium. But the theorem doesn't tell us which one is selected in games with multiple equilibria, or whether players will actually play it.

Sequential Games

Extensive Form

So far, our games have been **simultaneous**. But many real situations are **sequential**: one player moves first, and the second responds after observing the first move.

The extensive form uses a **game tree**:

- Nodes represent decision points (who moves)
- Branches represent actions
- Payoffs sit at the end of each path
- Information sets capture what a player knows when deciding

Entry Game

An entrant (E) decides whether to **Enter** or **Stay Out** of a market. If E enters, the incumbent (I) decides to **Fight** or **Accommodate**.

Path	Payoffs (E, I)
Stay Out	(0, 10)
Enter, Accommodate	(2, 5)
Enter, Fight	(−3, 0)

What happens? It depends on what E believes I will do if entry occurs.

Extensive Form for Entry Game

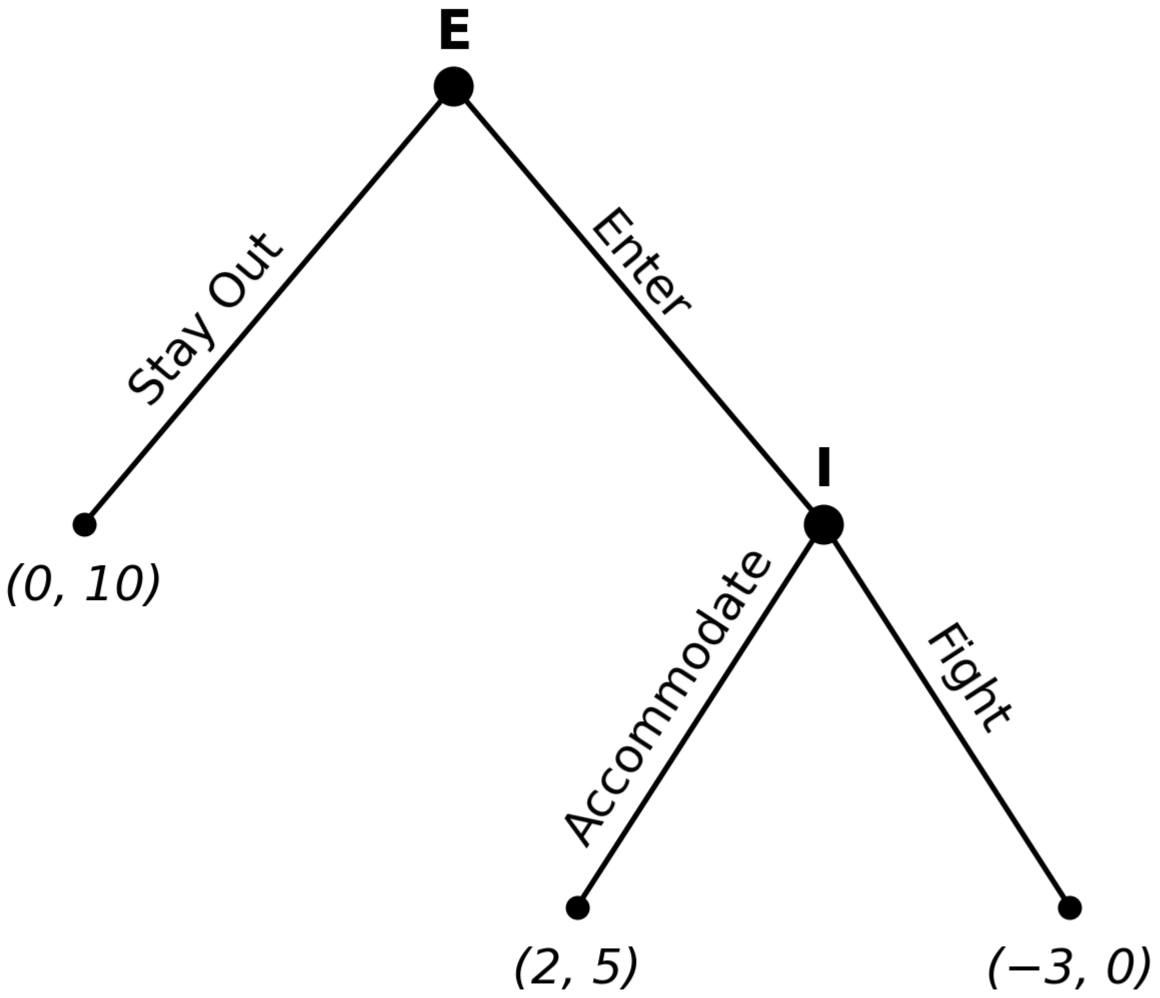

Backward Induction

Solve **from the end** of the tree back to the beginning.

Step 1: If E enters, what does I do? Accommodate (5) beats Fight (0). So I accommodates.

Step 2: Given that I will accommodate, what does E do? Enter (2) beats Stay Out (0). So E enters.

Prediction: (Enter, Accommodate), payoffs (2, 5).

Non-Credible Threats

In the entry game, the normal form actually has *two* Nash equilibria: (Enter, Accommodate) and (Stay Out, Fight).

At (Stay Out, Fight), neither player wants to deviate: given that E stays out, I’s strategy of “Fight if entry occurs” is never tested; and given I’s threat to fight, E prefers to stay out.

But the Fight threat is **not credible**: if entry actually occurred, I would not carry it out. Backward induction rules out such non-credible threats.

Subgame Perfect Equilibrium

Subgame perfect equilibrium (SPE): a strategy profile that induces a Nash equilibrium in *every* subgame, not just the whole game.

In finite games of perfect information, the SPE is exactly what **backward induction** gives you. It rules out Nash equilibria that rely on non-credible threats off the equilibrium path.

In the entry game, (Enter, Accommodate) is the unique SPE; (Stay Out, Fight) is Nash but not subgame perfect.

Commitment and Credibility

Schelling (1960): sometimes you are better off **limiting your own options**. A credible commitment device can change the equilibrium.

- A firm that **builds excess capacity** makes “Fight” credible and can deter entry
- A country that **publicly burns bridges** commits to not retreating
- A central bank with an **inflation targeting mandate** commits to tight policy

Commitment is valuable not despite restricting your future choices but *because* it does.

Repeated Games

Repeated Prisoner's Dilemma

In the one-shot Prisoner's Dilemma, the only Nash equilibrium is (Defect, Defect). What happens if the same game is played repeatedly?

Finitely repeated: backward induction still unravels cooperation. In the last round, both defect. Knowing that, they defect in the second-to-last round. And so on. Defection in every period is the unique SPE.

Infinitely repeated (or unknown horizon): cooperation can be sustained as an equilibrium if players are patient enough, using strategies that punish defection (e.g., **tit-for-tat** or **grim trigger**).

The Folk Theorem (Sketch)

In an infinitely repeated game with a sufficiently patient discount factor, essentially *any* individually rational payoff can be supported as a subgame perfect equilibrium.

Intuition: if defection today triggers enough future punishment, players prefer to cooperate.

The flip side: repeated interaction gives us cooperation, but at the cost of a huge **multiplicity of equilibria**. Predicting what will actually happen is hard.