Chapter 1

What Is Game Theory?

1.1 Strategic interaction everywhere

Consider the following situations:

- Two firms decide simultaneously how to price a similar product.
- Countries negotiate a climate treaty, each tempted to free-ride on the others' emission cuts.
- Bidders compete in an online advertising auction that runs billions of times per day.
- Two drivers approach an intersection; each must decide whether to yield.
- A poker player decides whether to bluff.

What do they have in common? In each case, the *best choice for one participant depends on what the others do.* This interdependence is the defining feature of a *game* in the technical sense, and game theory is the mathematical language for analyzing such situations.

Key Idea: Strategic interdependence

A decision problem becomes a *game* when the outcome of your choice depends on the choices of other purposeful decision-makers, and everyone is aware of this mutual dependence. Game theory asks: *what will (or should) rational players do?*

Contrast this with a single-person decision problem, like choosing a route with known traffic: there, you optimize against a fixed environment. In a game, the "environment" consists of other optimizers who are simultaneously reasoning about *you.* This circularity—I think about what you think about what I think—is what makes game theory subtle, and what its solution concepts are designed to cut through.

1.2 The ingredients of a game

Every game-theoretic model specifies three things:

1. **Players**: the decision-makers. They may be people, firms, governments, animals, or algorithms.
2. **Strategies**: a complete description of what each player can do. A strategy is a full contingent plan—it says what the player would do in *every* situation that could arise.

3. **Payoffs**: a numerical evaluation of every possible outcome for every player, representing their preferences (Chapter 2 makes this precise using utility theory).

Two further modeling choices matter enormously:

- **Timing**: do players move simultaneously (*static* games, Chapters 3–5) or in sequence (*dynamic* games, Chapter 6)?
- **Information**: does everyone know everyone else's payoffs (*complete information*) or not (*incomplete information*, Chapter 10)? Can players observe past moves (*perfect monitoring*) or not?

The four combinations of timing and information organize the whole classical theory, and each has its own equilibrium concept, summarized in Table 1.1.

	Complete information	**Incomplete information**
Static	Nash equilibrium (Ch. 4–5)	Bayesian Nash equilibrium (Ch. 10)
Dynamic	Subgame perfect equilibrium (Ch. 6, 9)	Perfect Bayesian equilibrium (Ch. 11)

Table 1.1: The classical roadmap: solution concepts by timing and information.

1.3 A first game: the Prisoner's Dilemma

The most famous game of all illustrates the ingredients. Two suspects are interrogated in separate rooms. Each can **Cooperate** with the other suspect (stay silent) or **Defect** (betray the partner). If both stay silent, each serves 1 year. If one betrays while the other stays silent, the betrayer walks free and the silent one serves 10 years. If both betray, each serves 5 years. Writing payoffs as $-$(years served):

	Cooperate	**Defect**
Cooperate	$-1, -1$	$-10, 0$
Defect	$0, -10$	$-5, -5$

Table 1.2: The Prisoner's Dilemma. Row player's payoff listed first.

Look at the game from the row player's perspective. If the column player cooperates, defecting gives $0 > -1$. If the column player defects, defecting gives $-5 > -10$. *Defection is better no matter what the other does*—it is a *dominant strategy* (Chapter 3). Since the game is symmetric, both players defect, and both serve 5 years—even though both staying silent would have given each only 1 year.

Key Idea: Individual rationality vs. collective optimality

The Prisoner's Dilemma shows that the pursuit of individual self-interest can lead to outcomes that are worse *for everyone* than an available alternative. This tension drives much of game theory: arms races, price wars, overfishing, doping in sports, and climate change all share this structure. Later chapters study how repetition (Chapter 9), contracts, and mechanism design (Chapter 14) can restore cooperation.

1.4 Cooperative vs. non-cooperative theory

Game theory has two grand branches.

Non-cooperative theory models individual actions. Binding agreements are impossible; anything that is to be sustained must be *self-enforcing*. Most of this book (Parts I–II and much of Part III) is non-cooperative theory.

Cooperative (coalitional) theory takes binding agreements as possible and asks what coalitions can achieve and how the gains from cooperation should be divided. Its central concepts—the core and the Shapley value—are developed in Chapter 15.

1.5 A brief history

Historical Note

- **1713–1913 (prehistory):** early probabilistic analyses of card games (Waldegrave's minimax solution of *le Her*); Cournot's 1838 duopoly model anticipates Nash equilibrium; Zermelo's 1913 theorem on chess anticipates backward induction.
- **1928:** John von Neumann proves the *Minimax Theorem* for two-player zero-sum games—the founding theorem of the field (Chapter 12).
- **1944:** von Neumann and Oskar Morgenstern publish *Theory of Games and Economic Behavior*, creating the discipline and expected utility theory.
- **1950–1953:** John Nash defines the equilibrium concept for general games and proves existence; he also founds axiomatic bargaining theory. Nash, John Harsanyi (incomplete information), and Reinhard Selten (subgame perfection) share the 1994 Nobel Prize.
- **1960s–70s:** Harsanyi shows how to model incomplete information (1967–68); Selten refines Nash equilibrium for dynamic games (1965, 1975); Aumann develops repeated games and correlated equilibrium; Maynard Smith exports the theory to biology (evolutionarily stable strategies, 1973).
- **1980s–90s:** mechanism design and auction theory mature (Myerson's optimal auctions, 1981); refinements flourish; experimental and behavioral game theory test the theory in the lab.
- **2000s–present:** algorithmic game theory emerges at the interface with computer science (complexity of computing Nash equilibria, price of anarchy, online ad auctions); market design (matching, spectrum auctions, kidney exchange) earns Nobel Prizes for Vickrey–Mirrlees (1996), Hurwicz–Maskin–Myerson (2007), Roth–Shapley (2012), Tirole (2014), Hart–Holmström (2016), and Milgrom–Wilson (2020).

1.6 What game theory is—and is not

Game theory is sometimes caricatured as either a cynical manual for selfishness or an ivory-tower abstraction. Neither is accurate. Three clarifications:

Remark: Payoffs already include everything players care about

A "payoff" is not money; it is a utility index that can encode altruism, fairness, spite, or reputation. If a player genuinely cares about her opponent's welfare, that belongs *in her payoff function*. Game theory takes preferences as inputs; it does not prescribe selfishness.

Remark: Descriptive vs. normative use

Game theory can be used *descriptively* (predicting behavior—how do bidders actually behave in auctions?), *normatively* (advising a player—how should I bid?), and *prescriptively for designers* (how should the auction itself be designed?). The third use, mechanism design, is often called "reverse game theory" and has been spectacularly successful in practice.

Remark: Rationality is a modeling benchmark

Real people are not perfectly rational, and behavioral game theory documents systematic deviations. But the rational benchmark is indispensable: it identifies the forces at work, often predicts well in high-stakes settings with experienced players (markets, auctions), and is the reference point against which "behavioral" deviations are defined and measured. See Chapter 18.

Exercises

Exercise 1.1: Classifying interactions

For each situation, identify the players, their strategies, whether moves are simultaneous or sequential, and whether information is complete or incomplete: (a) rock–paper–scissors; (b) a used-car sale where only the seller knows the car's quality; (c) chess; (d) a sealed-bid auction for a painting.

Exercise 1.2: A real-world dilemma

Give an example from your own experience (work, sports, family, online life) with the structure of a Prisoner's Dilemma. Specify the two actions, and explain why defection is individually tempting but collectively harmful. What mechanisms (if any) sustain cooperation in your example?

Exercise 1.3: Payoffs are not money

Two siblings split a cake. The older sibling proposes a split; the younger accepts or rejects (if rejected, the cake is thrown away). If both care only about their own cake, what split do you expect? Now suppose the younger sibling feels anger worth -10 units when offered less than a third. How does the prediction change? (This previews bargaining, Chapter 8.)