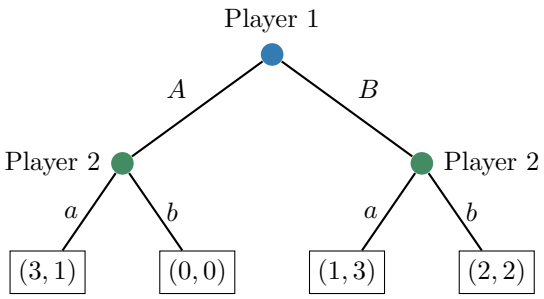
Game Theory

A Complete Learning Path

From Beginner to Advanced



Strategic Form Games · Nash Equilibrium · Dynamic Games
Bayesian Games · Mechanism Design · Cooperative Games
Evolutionary Dynamics · Algorithmic Game Theory

Compiled Learning Materials
August 3, 2026

Preface

Game theory is the mathematical study of strategic interaction: situations in which the outcome for each participant depends not only on their own decisions but also on the decisions of others. Born in the 1940s from the collaboration of John von Neumann and Oskar Morgenstern, and transformed by John Nash's equilibrium concept in 1950, game theory has become one of the most influential analytical frameworks of the modern era. It underpins large parts of economics, is central to political science and evolutionary biology, and—through auction design, market design, and algorithmic game theory—now shapes the infrastructure of the internet economy.

How this book is organized

This book is a self-contained learning path organized in three parts of increasing sophistication.

Part I: Beginner. No prior exposure is assumed. We build the vocabulary of games—players, strategies, payoffs—and develop the central solution concepts: dominance, Nash equilibrium (pure and mixed), and backward induction in dynamic games. Every concept is introduced through classic examples such as the Prisoner's Dilemma, the Battle of the Sexes, and Matching Pennies.

Part II: Intermediate. We apply the theory to economics (oligopoly, public goods), then extend it in three major directions: *time* (repeated games and the Folk Theorem, bargaining), *information* (Bayesian games, auctions, signaling, cheap talk), and *conflict* (zero-sum games and the Minimax Theorem).

Part III: Advanced. We treat the theory rigorously and broadly: existence proofs via fixed-point theorems, equilibrium refinements, mechanism design and auction theory (the Revelation Principle, VCG, Myerson's optimal auction), cooperative game theory (the core, the Shapley value, matching markets), evolutionary game theory, and algorithmic game theory (complexity of equilibria, price of anarchy, learning dynamics), closing with a tour of current research frontiers.

Appendix A reviews the mathematical background (sets, probability, optimization, fixed points). Appendix B contains full solutions to the end-of-chapter exercises.

How to use this book

- **If you are new to the subject**, read Part I linearly. Do the exercises: game theory is learned by solving games, not by reading about them.
- **If you know the basics**, skim Part I for notation and start with Part II.
- **If you are preparing for graduate study or research**, Part III together with the cited literature (Fudenberg–Tirole, Osborne–Rubinstein, Myerson, Maschler–Solan–Zamir) forms a solid core.

Throughout, boxed **Definitions**, **Theorems**, **Examples**, and **Key Ideas** mark the essential material, while *Historical Notes* place the ideas in context. A guide to further reading concludes the book.

Contents

Preface **iii**

I Beginner: Foundations of Strategic Thinking **1**

1 What Is Game Theory? **3**
1.1 Strategic interaction everywhere . . . 3
1.2 The ingredients of a game . . . 3
1.3 A first game: the Prisoner's Dilemma . . . 4
1.4 Cooperative vs. non-cooperative theory . . . 5
1.5 A brief history . . . 5
1.6 What game theory is—and is not . . . 5

2 Rationality, Preferences, and Expected Utility **7**
2.1 Preferences and utility . . . 7
2.2 Choice under risk: lotteries . . . 8
2.3 Risk attitudes . . . 9
2.4 Beliefs, best responses, and rationality in games . . . 9

3 Games in Normal Form and Dominance **11**
3.1 The normal (strategic) form . . . 11
3.2 Classic 2×2 games . . . 11
3.3 Dominant and dominated strategies . . . 12
3.4 Iterated elimination of dominated strategies . . . 13
3.5 Dominance by mixed strategies . . . 14
3.6 Rationalizability . . . 14

4 Nash Equilibrium **17**
4.1 Definition and interpretation . . . 17
4.2 Finding pure-strategy Nash equilibria . . . 18
4.3 Games with continuous action spaces . . . 18
4.4 Existence, multiplicity, and selection . . . 19
4.5 Correlated equilibrium: a first look . . . 20

5 Mixed Strategies **23**
5.1 Mixed strategies and expected payoffs . . . 23
5.2 The indifference principle . . . 23
5.3 Interpreting mixed strategies . . . 24
5.4 Computing equilibria in larger games . . . 25

6 Extensive-Form Games and Backward Induction **27**
6.1 Game trees . . . 27
6.2 Backward induction . . . 28
6.3 Subgames and subgame perfection . . . 28
6.4 The ultimatum game and the limits of backward induction . . . 29
6.5 Imperfect information: information sets . . . 30

II Intermediate: Dynamics, Information, and Applications **33**

7 Economic Applications: Markets and Strategy **35**
7.1 Cournot competition: quantities . . . 35
7.2 Bertrand competition: prices . . . 35
7.3 Public goods and free-riding . . . 36
7.4 Elections: the median voter . . . 36
7.5 Congestion: selfish routing . . . 37
7.6 Strategic complements vs. substitutes: a unifying lens . . . 37

8 Bargaining **39**
8.1 The axiomatic approach: Nash bargaining . . . 39
8.2 The strategic approach: alternating offers . . . 40
8.3 Bargaining with outside options . . . 41
8.4 Applications . . . 41

9 Repeated Games and the Emergence of Cooperation **43**
9.1 Finitely repeated games . . . 43
9.2 Infinitely repeated games . . . 43
9.3 The Folk Theorem . . . 44
9.4 Cartels, price wars, and imperfect monitoring . . . 45
9.5 Reputation . . . 45

10 Bayesian Games: Incomplete Information **47**
10.1 Harsanyi's framework: types and beliefs . . . 47
10.2 Auctions: private values . . . 48
10.3 Common values and the winner's curse . . . 49
10.4 Global games: equilibrium selection through uncertainty . . . 49

11 Dynamic Games with Incomplete Information: Signaling and Screening **51**
11.1 Perfect Bayesian equilibrium . . . 51
11.2 Signaling: Spence's education model . . . 51
11.3 Cheap talk . . . 52
11.4 Screening: the uninformed side designs the menu . . . 53
11.5 Reputation and dynamic signaling (a synthesis) . . . 53

12 Zero-Sum Games and the Minimax Theorem **55**
12.1 Maxmin reasoning . . . 55
12.2 Minimax and linear programming duality . . . 56
12.3 Solving games by fictitious play and no-regret learning . . . 56
12.4 Applications of pure conflict . . . 57

III Advanced: Theory, Design, and Computation **59**

13 Existence of Equilibrium and Refinements **61**
13.1 Fixed-point theorems . . . 61
13.2 Nash's theorem . . . 61
13.3 Trembling-hand perfection . . . 62
13.4 Sequential equilibrium . . . 63
13.5 Epistemic foundations . . . 63

14 Mechanism Design and Auction Theory **65**
14.1 The setting and the Revelation Principle . . . 65
14.2 Dominant-strategy implementation: VCG . . . 66
14.3 Bayesian mechanism design: Myerson's optimal auction . . . 66
14.4 Impossibility: bilateral trade . . . 67
14.5 Social choice: Arrow and Gibbard–Satterthwaite . . . 67

15 Cooperative Game Theory **71**
15.1 Coalitional games . . . 71
15.2 The core . . . 71
15.3 The Shapley value . . . 72
15.4 The nucleolus and other solutions . . . 73
15.5 Matching markets (NTU cooperation) . . . 73

16 Evolutionary Game Theory **75**
16.1 Evolutionarily stable strategies . . . 75
16.2 Replicator dynamics . . . 76
16.3 The evolution of cooperation . . . 76
16.4 Learning in games (the bridge to algorithms) . . . 77

17 Algorithmic Game Theory **79**
17.1 The complexity of computing Nash equilibria . . . 79
17.2 Potential games and best-response dynamics . . . 80
17.3 The price of anarchy . . . 80
17.4 Learning dynamics and correlated equilibria . . . 81
17.5 Computational mechanism design and Bitcoin's game theory . . . 81

18 Frontiers: Behavioral, Experimental, and Modern Directions **83**
18.1 Experimental game theory: the evidence . . . 83
18.2 Behavioral game theory: amended models . . . 83
18.3 Stochastic games and dynamic equilibria . . . 84
18.4 Games with a continuum of players and mean-field games . . . 84
18.5 Information design and Bayesian persuasion . . . 84
18.6 Game theory and AI . . . 84
18.7 A reading roadmap . . . 85

A Mathematical Background **87**
A.1 Sets, functions, correspondences . . . 87
A.2 Optimization facts . . . 87
A.3 Probability . . . 87

A.4 Fixed points 88
A.5 Linear programming and duality 88
A.6 Basic dynamical systems 88

B Solutions to Exercises **89**

Bibliography and Further Reading **99**

Part I

Beginner: Foundations of Strategic Thinking

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.)

Chapter 2

Rationality, Preferences, and Expected Utility

Before analyzing games we must be precise about what a "payoff" is and what "rational" means. This chapter builds the decision-theoretic foundation on which all of game theory rests. Readers eager for games may skim it and return later; but the ideas here—especially expected utility—are used constantly from Chapter 5 onward.

2.1 Preferences and utility

Decision theory starts not from numbers but from *preferences*. Let X be a set of outcomes. A decision-maker's preferences are captured by a binary relation $\succsim$ on X, where $x \succsim y$ reads "x is at least as good as y." We write $x \succ y$ (strict preference) when $x \succsim y$ but not $y \succsim x$, and $x \sim y$ (indifference) when both hold.

Definition 2.1: Rational preference relation

A preference relation $\succsim$ on X is *rational* if it is

- **complete**: for all $x, y \in X$, either $x \succsim y$ or $y \succsim x$ (or both); and
- **transitive**: for all $x, y, z \in X$, if $x \succsim y$ and $y \succsim z$ then $x \succsim z$.

Completeness says the decision-maker can always rank alternatives; transitivity rules out preference cycles. Neither is innocuous—real people violate both in experiments—but together they permit a dramatic simplification:

Theorem 2.1: Utility representation

If X is finite (or countable) and $\succsim$ is rational, then there exists a *utility function* $u : X \to \mathbb{R}$ such that

$$x \succsim y \iff u(x) \geq u(y).$$

Moreover, u is unique only up to strictly increasing transformations: if f is strictly increasing, $f \circ u$ represents the same preferences.

Proof. Enumerate $X = \{x_1, x_2, \dots\}$ and set $u(x) = \#\{y \in X : x \succsim y\}$ when X is finite (the countable case requires a small extra construction). Completeness and transitivity guarantee $x \succsim y \iff u(x) \geq u(y)$: transitivity implies that the set of outcomes weakly worse than y is

contained in the set weakly worse than x whenever $x \succsim y$, giving $u(x) \geq u(y)$; completeness gives the converse direction. □

Remark: Ordinal vs. cardinal utility

Theorem 2.1 delivers *ordinal* utility: only the ranking matters, not magnitudes. Saying "$u(x) = 10$ and $u(y) = 5$, so x is twice as good" is meaningless at this stage. Cardinal meaning arrives only with risk, below.

2.2 Choice under risk: lotteries

Games force players to evaluate *risky* prospects: if my opponent randomizes, my action induces a probability distribution over outcomes. Let $X = \{x_1, \dots, x_n\}$ be a finite outcome set. A *lottery* is a probability distribution $p = (p_1, \dots, p_n)$ with $p_i \geq 0$ and $\sum_i p_i = 1$. Let $\Delta(X)$ denote the set of all lotteries over X. The decision-maker now has preferences $\succsim$ over $\Delta(X)$, not just over X.

Definition 2.2: von Neumann–Morgenstern axioms

A preference relation $\succsim$ on $\Delta(X)$ satisfies the vNM axioms if it is rational (complete, transitive) and additionally:

- **Continuity**: for all p, q, r with $p \succ q \succ r$, there exist $\alpha, \beta \in (0, 1)$ such that $\alpha p + (1-\alpha) r \succ q \succ \beta p + (1-\beta) r$.
- **Independence**: for all p, q, r and $\alpha \in (0, 1]$,

$$p \succsim q \iff \alpha p + (1-\alpha) r \succsim \alpha q + (1-\alpha) r.$$

Independence is the crucial (and most debated) axiom: mixing both lotteries with a common third lottery should not reverse the ranking.

Theorem 2.2: Expected utility (von Neumann–Morgenstern, 1944)

$\succsim$ on $\Delta(X)$ satisfies the vNM axioms if and only if there exists $u : X \to \mathbb{R}$ (a *Bernoulli utility function*) such that for all lotteries p, q:

$$p \succsim q \iff \sum_{i=1}^{n} p_i \, u(x_i) \;\geq\; \sum_{i=1}^{n} q_i \, u(x_i).$$

Moreover u is unique up to positive affine transformations: $v = au + b$ with $a > 0$ represents the same preferences.

Key Idea: Why game payoffs are expected utilities

Theorem 2.2 is the license for everything we do later: when we write a payoff matrix and compute *expected* payoffs against a randomizing opponent (Chapter 5), we are implicitly assuming players are vNM expected-utility maximizers. The affine-uniqueness clause tells us which operations on payoffs are harmless: adding a constant to all of a player's payoffs, or rescaling them by a positive factor, changes nothing strategically.

2.3 Risk attitudes

Suppose outcomes are money amounts. A decision-maker with Bernoulli utility u facing a lottery with random prize W compares $\mathbb{E}[u(W)]$ with $u(\mathbb{E}[W])$.

Definition 2.3: Risk aversion

A decision-maker is *risk-averse* if she prefers the expected value of any lottery for sure to the lottery itself: $u(\mathbb{E}[W]) \geq \mathbb{E}[u(W)]$ for all W. By Jensen's inequality this holds iff u is concave. She is *risk-neutral* iff u is affine, and *risk-loving* iff u is convex.

Example 2.1: The St. Petersburg puzzle

A fair coin is flipped until the first heads; if this happens on flip k you win 2^k dollars. The expected prize is $\sum_{k=1}^{\infty} 2^{-k} \cdot 2^k = \sum_k 1 = \infty$, yet nobody would pay more than a modest amount to play. Daniel Bernoulli's 1738 resolution: people maximize expected *utility*, not expected money. With $u(x) = \log_2 x$, the certainty equivalent is finite and small: $\mathbb{E}[u] = \sum_{k \geq 1} 2^{-k} \cdot k = 2$, i.e. the gamble is worth $2^2 = 4$ dollars. This is the historical origin of expected utility theory.

2.4 Beliefs, best responses, and rationality in games

In a game, a player's optimal action depends on what the others do—which she does not observe in a simultaneous-move game. The modern view: a rational player forms a *belief* (a probability distribution over opponents' strategy choices) and chooses a *best response* to that belief.

Definition 2.4: Best response to a belief

Let player i have belief μ over the opponents' action profiles A_{-i}. Action a_i is a *best response to* μ if

$$\sum_{a_{-i}} \mu(a_{-i})\, u_i(a_i, a_{-i}) \;\geq\; \sum_{a_{-i}} \mu(a_{-i})\, u_i(a_i', a_{-i}) \quad \text{for every } a_i' \in A_i.$$

A player is *rational* if she plays a best response to *some* belief.

Rationality alone predicts little: almost any action is a best response to some belief. Predictive power comes from assumptions about what players believe about *each other*:

Definition 2.5: Common knowledge

An event E is *common knowledge* if everyone knows E, everyone knows that everyone knows E, everyone knows that everyone knows that everyone knows E, and so on ad infinitum.

The classical assumption of game theory is *common knowledge of rationality* (CKR): all players are rational, all know this, all know that all know it, etc. In Chapter 3 we will see exactly what CKR implies (iterated elimination of dominated strategies / rationalizability), and in Chapter 4 what additional assumptions Nash equilibrium requires.

Example 2.2: The email/coordinated-attack puzzle

Two generals must attack simultaneously to win. General 1 sees the enemy is weak and emails General 2: "attack at dawn." Messages fail with small probability, so General 2

replies to confirm, General 1 confirms the confirmation, *ad infinitum*. After any finite number of confirmations, "attack at dawn" is mutually known to high order but never *common* knowledge—and one can show coordinated attack cannot be guaranteed. The example (Rubinstein's email game) shows common knowledge is a genuinely stronger and more delicate notion than "everybody knows."

Exercises

Exercise 2.1: Transitivity and money pumps

A decision-maker has cyclic preferences: $x \succ y$, $y \succ z$, $z \succ x$. Show that if she is willing to pay a small amount $\varepsilon > 0$ to swap any held item for a strictly preferred one, a trader holding x, y, z can extract unbounded money from her. (This "money pump" is the classic pragmatic defense of transitivity.)

Exercise 2.2: Affine invariance

Player 1's payoffs in some game are given by u_1. Show that replacing u_1 by $v_1 = 3u_1 + 7$ leaves her set of best responses to every belief unchanged. Then give an example showing that a *nonlinear* increasing transformation (say $u \mapsto u^3$ with payoffs that can be negative) can change best responses against a *randomizing* opponent even though it preserves the ordinal ranking of pure outcomes.

Exercise 2.3: Certainty equivalents

An agent has Bernoulli utility $u(x) = \sqrt{x}$ and faces a lottery paying 100 with probability $1/2$ and 0 otherwise. Compute the expected value, the expected utility, the certainty equivalent, and the risk premium.

Exercise 2.4: Allais paradox

Consider lotteries over prizes $\{0, 1\text{M}, 5\text{M}\}$ (millions): $A = (0, 1, 0)$ (1M for sure); $B = (0.01, 0.89, 0.10)$; $C = (0.89, 0.11, 0)$; $D = (0.90, 0, 0.10)$. Most people report $A \succ B$ and $D \succ C$. Show these two rankings jointly violate the independence axiom.

Chapter 3

Games in Normal Form and Dominance

We now formalize simultaneous-move games and develop the most conservative solution concepts: those that rely only on rationality and common knowledge of rationality.

3.1 The normal (strategic) form

Definition 3.1: Normal-form game

A *normal-form* (or *strategic-form*) game is a triple $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ where

- $N = \{1, \dots, n\}$ is a finite set of *players*;
- A_i is the set of *actions* (pure strategies) available to player i; an *action profile* is $a = (a_1, \dots, a_n) \in A = A_1 \times \cdots \times A_n$;
- $u_i : A \to \mathbb{R}$ is player i's *payoff function* (a vNM utility).

We write a_{-i} for the profile of actions of everyone except i, and $a = (a_i, a_{-i})$.

Everything about the game—the sets N, A_i, and functions u_i—is assumed common knowledge among the players (this is what *complete information* means).

Two-player finite games are displayed as *payoff matrices*: rows are player 1's actions, columns player 2's, and each cell shows (u_1, u_2).

3.2 Classic 2×2 games

Four canonical games recur throughout the book. Internalize them—they are the fruit flies of game theory.

Example 3.1: Prisoner's Dilemma

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

(Payoffs rescaled from Chapter 1; only the ordering matters.) Defection (D) strictly dominates cooperation; the unique prediction (D, D) is Pareto-dominated by (C, C). *Theme:*

conflict between individual and group rationality.

Example 3.2: Coordination / Stag Hunt

Two hunters can jointly hunt a stag (needs both) or individually hunt hares:

	Stag	Hare
Stag	$4, 4$	$0, 2$
Hare	$2, 0$	$2, 2$

Both (Stag, Stag) and (Hare, Hare) are stable (each is a Nash equilibrium, as we'll see): the first is *payoff-dominant*, the second *risk-dominant* (safer if you doubt your partner). *Theme: trust and equilibrium selection.*

Example 3.3: Battle of the Sexes

A couple wants to spend the evening together but one prefers the Ballet, the other the Football match:

	B	F
B	$3, 1$	$0, 0$
F	$0, 0$	$1, 3$

Coordination is good but the players disagree about *where* to coordinate. *Theme: coordination with distributional conflict.*

Example 3.4: Matching Pennies

Each player shows Heads or Tails; player 1 wins if they match, player 2 if they differ:

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

Pure conflict (zero-sum): whatever one wins the other loses. No pair of pure actions is stable—each player wants to outguess the other. *Theme: the need for randomization (Chapter 5).*

3.3 Dominant and dominated strategies

Definition 3.2: Dominance

Action a_i *strictly dominates* a_i' if

$$u_i(a_i, a_{-i}) > u_i(a_i', a_{-i}) \quad \text{for every } a_{-i} \in A_{-i}.$$

It *weakly dominates* a_i' if the inequality is weak for all a_{-i} and strict for at least one. An action is *strictly (weakly) dominant* if it strictly (weakly) dominates every other action; it is *strictly dominated* if some other action (pure or mixed—see Section 3.5) strictly dominates it.

A rational player never plays a strictly dominated action, whatever her beliefs. When every

player has a strictly dominant action, the game is solved by rationality alone—no strategic reasoning about others is even needed. The Prisoner's Dilemma is the leading case.

Example 3.5: Second-price auction (a preview)

A seller auctions one object. Each bidder i values it at v_i and submits a sealed bid b_i; the highest bidder wins and pays the *second*-highest bid. Then bidding truthfully, $b_i = v_i$, *weakly dominates* every other bid: overbidding only matters when it wins at a price above v_i (a loss), underbidding only matters when it loses an object priced below v_i (a forgone gain). This remarkable dominance argument (Vickrey, 1961) needs no assumptions about opponents at all, and is the seed of mechanism design (Chapter 14).

3.4 Iterated elimination of dominated strategies

Rationality removes dominated actions. *Common knowledge* of rationality lets us iterate: I know you won't play dominated actions, so I can delete them from my analysis of *your* play, which may make some of *my* actions dominated in the reduced game, and so on.

Definition 3.3: IESDS

Iterated elimination of strictly dominated strategies (IESDS) is the process of repeatedly deleting, for every player, all strictly dominated actions in the current reduced game, until no more deletions are possible. A game is *dominance solvable* if IESDS leaves exactly one action profile.

Theorem 3.1: Order independence

In a finite game, the set of action profiles surviving IESDS does not depend on the order or speed of elimination.

(For *weak* dominance, order matters—one reason strict dominance is the preferred notion for iteration.)

Example 3.6: A dominance-solvable game

	L	M	R
T	$4, 3$	$2, 7$	$0, 4$
B	$5, 5$	$5, -1$	$-4, -2$

Step 1: for player 2, R is strictly dominated by M ($7 > 4$ against T; $-1 > -2$ against B). Delete R. Step 2: in the reduced game, no pure action of player 1 dominates the other ($4 < 5$ against L but $2 < 5$ against M: in fact B dominates T here since $5 > 4$ and $5 > 2$). Delete T. Step 3: against B, player 2 prefers L ($5 > -1$). Delete M. The unique survivor is (B, L) with payoffs $(5, 5)$.

Example 3.7: The beauty contest (guessing game)

n students each pick a number in $[0, 100]$; the winner is whoever is closest to $\frac{2}{3}$ of the class average. Any guess above $\frac{2}{3} \cdot 100 = 66.7$ is dominated (the target can never exceed it). Once those are removed, any guess above $\frac{2}{3} \cdot 66.7 = 44.4$ is dominated, and so on: IESDS spirals

down to **0**, the unique rationalizable guess. In experiments the winning guess is typically around 20–35: real people iterate only a few steps. Keynes used this game as a metaphor for stock markets ("anticipating what average opinion expects average opinion to be").

3.5 Dominance by mixed strategies

An action can be strictly dominated by a *randomization* even when no pure action dominates it.

Example 3.8: Mixed domination

	L	R
T	$4, 1$	$0, 2$
M	$1, 1$	$1, 0$
B	$0, 3$	$4, 1$

Neither T nor B dominates M. But the mixture $\frac{1}{2}T + \frac{1}{2}B$ yields expected payoff 2 against both L and R, strictly beating M's payoff of 1. So M is strictly dominated and a rational player 1 never plays it.

From now on, "strictly dominated" always means dominated by some pure *or mixed* strategy.

3.6 Rationalizability

What exactly does common knowledge of rationality imply? An action is justifiable if it is a best response to *some* belief about opponents' play; it survives deeper scrutiny only if that belief itself assigns positive probability only to justifiable actions; and so on.

Definition 3.4: Rationalizability (informal)

An action is *rationalizable* if it is a best response to a belief that puts weight only on opponents' actions that are themselves best responses to such beliefs, recursively ad infinitum.

Theorem 3.2: Rationalizability and IESDS (Bernheim 1984; Pearce 1984)

In finite two-player games, the set of rationalizable actions coincides exactly with the set of actions surviving iterated elimination of strictly dominated strategies. (With three or more players the equivalence holds if beliefs may be correlated across opponents.)

Key Idea: The epistemic ladder

Rationality $\Rightarrow$ no strictly dominated actions. *Rationality + knowledge of others' rationality* $\Rightarrow$ survives two rounds of IESDS. *Common knowledge of rationality* $\Rightarrow$ rationalizable actions. Nash equilibrium (next chapter) requires something more: *correct* beliefs, not merely coherent ones.

Exercises

Exercise 3.1: Dominance hunting

	L	C	R
T	$2,0$	$1,1$	$4,2$
M	$3,4$	$1,2$	$2,3$
B	$1,3$	$0,2$	$3,0$

Solve by IESDS, stating at each step which action is eliminated and by what it is dominated.

Exercise 3.2: Weak dominance and order dependence

	L	R
T	$1,1$	$0,0$
M	$1,1$	$2,1$
B	$0,0$	$2,1$

Show that different orders of eliminating *weakly* dominated strategies can lead to different final predictions.

Exercise 3.3: Cournot by dominance

Two firms choose quantities $q_1, q_2 \geq 0$; market price is $P = \max\{12 - q_1 - q_2, 0\}$ and production is costless, so $u_i = q_i P$. Show that any $q_i > 6$ is strictly dominated. In the reduced game with $q_i \in [0, 6]$, show any $q_i < 3$ is strictly dominated. Iterate and show the process converges to $q_1 = q_2 = 4$. (Compare Chapter 7.)

Exercise 3.4: Traveler's dilemma

Two travelers independently claim an integer amount in $\{2, \dots, 100\}$ for lost luggage. Both are paid the *lower* claim; additionally the lower claimant gets a bonus of 2 and the higher pays a penalty of 2. Show that IESDS (using weak dominance at the top, then strict) leads to both claiming 2, and comment on whether you find the prediction plausible.

Chapter 4

Nash Equilibrium

Dominance arguments are compelling but often silent: in the Battle of the Sexes or the Stag Hunt, nothing is dominated, and rationalizability permits everything. The central solution concept of game theory adds one more requirement: beliefs must be *correct*.

4.1 Definition and interpretation

Definition 4.1: Best response correspondence

Player i's *best response correspondence* maps each profile of opponents' actions to the set of i's optimal actions:

$$\mathrm{BR}_i(a_{-i}) \;=\; \arg\max_{a_i \in A_i} u_i(a_i, a_{-i}).$$

Definition 4.2: Nash equilibrium

An action profile $a^* = (a_1^*, \ldots, a_n^*)$ is a *Nash equilibrium* if for every player i,

$$u_i(a_i^*, a_{-i}^*) \;\geq\; u_i(a_i, a_{-i}^*) \qquad \text{for all } a_i \in A_i,$$

equivalently $a_i^* \in \mathrm{BR}_i(a_{-i}^*)$ for every i: *every player's action is a best response to the others' equilibrium actions.*

Key Idea: No profitable unilateral deviation

A Nash equilibrium is a profile of mutually consistent plans: if everyone expects a^* to be played, no individual can gain by deviating *alone*. It is a *stability* or *self-enforcement* condition—not a claim that the outcome is good (the Prisoner's Dilemma equilibrium is bad for both players) nor that players somehow negotiate it.

Why should we expect equilibrium play? Several complementary interpretations:

- **Self-enforcing agreement**: if players discuss beforehand and reach a non-binding agreement, only Nash profiles are viable—anything else, someone would break.
- **Stable social convention**: in a population repeatedly playing the game (drivers at intersections, market participants), a convention persists only if no one benefits from unilaterally deviating—i.e., only if it is Nash.
- **Rational expectations**: if a theory of the game's play were known to all players, rational

players would follow it only if it prescribed a Nash equilibrium; any other prediction refutes itself.

- **Limit of learning/evolution**: many adaptive dynamics, when they converge, converge to Nash equilibria (Chapters 16, 17).

4.2 Finding pure-strategy Nash equilibria

For matrix games, the mechanical method is *best-response marking*: for each column, underline the row player's best payoff(s); for each row, underline the column player's best payoff(s); cells where both payoffs are underlined are Nash equilibria.

Example 4.1: Marking best responses

	L	C	R
T	$\underline{7},\underline{6}$	$5,5$	$0,4$
M	$5,1$	$\underline{6},\underline{8}$	$1,2$
B	$3,3$	$4,4$	$\underline{2},\underline{5}$

Check: against L, row's best is T (7); against C, M (6); against R, B (2). Against T, column's best is L (6); against M, C (8); against B, R (5). All three diagonal cells are Nash equilibria—a reminder that equilibria need not be unique and can be Pareto-ranked.

Revisiting the classics of Chapter 3:

- **Prisoner's Dilemma**: unique equilibrium (D, D).
- **Stag Hunt**: two pure equilibria, (S, S) and (H, H).
- **Battle of the Sexes**: two pure equilibria, (B, B) and (F, F).
- **Matching Pennies**: *no* pure equilibrium—the motivation for mixed strategies (Chapter 5).

4.3 Games with continuous action spaces

When actions are continuous, we find equilibria by solving the players' first-order conditions simultaneously.

Example 4.2: Cournot duopoly

Firms 1 and 2 simultaneously choose outputs $q_1, q_2 \geq 0$. The market price is $P = a - q_1 - q_2$ (when positive) and both firms have unit cost $c < a$, so

$$u_i(q_i, q_j) = q_i\,(a - q_i - q_j - c).$$

Best response: maximizing over q_i gives $\mathrm{BR}_i(q_j) = \frac{a-c-q_j}{2}$ (when nonnegative). Solving $q_1 = \mathrm{BR}_1(q_2)$ and $q_2 = \mathrm{BR}_2(q_1)$:

$$q_1^* = q_2^* = \frac{a-c}{3}, \qquad P^* = \frac{a+2c}{3}, \qquad u_i^* = \frac{(a-c)^2}{9}.$$

Each firm produces more than half the monopoly quantity $\frac{a-c}{2}$ in total; joint profit is below the monopoly level—a continuous cousin of the Prisoner's Dilemma (each firm's output exerts a negative externality on the other). Full treatment in Chapter 7.

Example 4.3: Tragedy of the commons

n herders each choose how many goats $g_i \geq 0$ to graze on a common pasture. The value per goat is $v(G) = 100 - G$ where $G = \sum_j g_j$, so $u_i = g_i(100 - G)$. The first-order condition for herder i is $100 - G - g_i = 0$. By symmetry $g_i = g$ and $(n+1)g = 100$:

$$g^* = \frac{100}{n+1}, \qquad G^* = \frac{100n}{n+1} \xrightarrow{n\to\infty} 100.$$

The socially optimal total (maximizing $G(100 - G)$) is $G^{\text{opt}} = 50$; the commons is overgrazed, the more so the more herders there are. Each herder ignores the cost his goats impose on others—an *externality*, the structural cause of inefficient equilibria.

4.4 Existence, multiplicity, and selection

Three basic facts frame how we use Nash equilibrium.

Theorem 4.1: Nash's existence theorem (1950) — statement

Every finite game has at least one Nash equilibrium, possibly in mixed strategies.

The proof, via fixed-point theorems, is given in Chapter 13. Existence makes Nash equilibrium a universally applicable prediction tool.

Multiplicity. Many games have several equilibria (Stag Hunt, Battle of the Sexes). Game theory then needs a theory of *selection*:

- **Focal points** (Schelling 1960): cultural or contextual salience can pick an equilibrium. Asked to meet a stranger in New York with no communication, many people say "noon at Grand Central"—nothing in the mathematics, but a powerful coordinating force.
- **Payoff dominance vs. risk dominance**: in the Stag Hunt, (S, S) gives higher payoffs, but (H, H) is safer against strategic uncertainty. Evolutionary and learning models often select the risk-dominant equilibrium.
- **Communication and cheap talk** can help select equilibria (though not always credibly—Chapter 11).
- **Refinements** sharpen predictions in dynamic games (Chapters 6, 13).

Efficiency. Nash equilibria may be socially inefficient (Prisoner's Dilemma, commons). Quantifying this inefficiency—the *price of anarchy*—is a central theme of algorithmic game theory (Chapter 17).

Historical Note

John Nash's 1950 PhD thesis at Princeton was 27 pages long. Its central concept was initially dismissed by von Neumann as "just a fixed point theorem," yet it became the cornerstone of modern economics. Nash shared the 1994 Nobel Prize with Harsanyi and Selten. His life—including his struggle with schizophrenia and late-life recognition—is

portrayed in *A Beautiful Mind*.

4.5 Correlated equilibrium: a first look

Nash equilibrium assumes independent randomization. Aumann (1974) observed that a *correlation device* can enlarge the set of stable outcomes.

Example 4.4: Traffic light

In the Chicken/Hawk–Dove game

	Yield	Dare
Yield	$4,4$	$2,5$
Dare	$5,2$	$0,0$

the pure equilibria are (Dare, Yield) and (Yield, Dare); the symmetric mixed equilibrium gives expected payoff $10/3$ each. Now install a fair traffic light: with probability $\frac{1}{2}$ it shows (Green, Red), else (Red, Green); each driver sees only their own light, and the convention is Green $\to$ Dare, Red $\to$ Yield. Given a red light you know the other is daring, so yielding is optimal ($2 > 0$); given green, the other yields, so daring is optimal ($5 > 4$). Nobody deviates, and the average payoff is $\frac{5+2}{2} = 3.5 > 10/3$. This is a *correlated equilibrium* that is not a Nash equilibrium. The general theory—including why correlated equilibria are computationally far more tractable than Nash—appears in Chapter 17.

Exercises

Exercise 4.1: Find all pure equilibria

	X	Y	Z
A	$3,2$	$0,0$	$1,1$
B	$0,0$	$2,3$	$1,0$
C	$1,1$	$0,1$	$2,2$

Use best-response marking to find all pure-strategy Nash equilibria.

Exercise 4.2: Partnership game

Two partners each choose effort $e_i \in [0,4]$. Output is $e_1 + e_2 + \frac{1}{2}e_1e_2$, shared equally; effort costs e_i^2. So $u_i = \frac{1}{2}(e_1 + e_2 + \frac{1}{2}e_1e_2) - e_i^2$. Find the Nash equilibrium efforts and compare with the efforts that maximize the sum of payoffs. Explain the direction of the inefficiency.

Exercise 4.3: Hotelling's beach

Two ice-cream vendors choose locations on a beach $[0,1]$; customers are uniformly distributed and walk to the nearest vendor (splitting equally if equidistant). Each vendor maximizes market share. Show that both locating at $\frac{1}{2}$ is the unique Nash equilibrium. Relate this to the observed similarity of political platforms (the median voter theorem).

Exercise 4.4: Equilibrium is not optimality

Construct a 2×2 game in which the unique Nash equilibrium gives both players their *worst* payoff among all four cells, or prove that this is impossible. What is the strongest statement of this kind that *is* possible?

Chapter 5

Mixed Strategies

Matching Pennies has no pure Nash equilibrium: any deterministic play can be exploited. The way out, foreshadowed by the practice of every poker player and penalty-kick taker, is to *randomize*.

5.1 Mixed strategies and expected payoffs

Definition 5.1: Mixed strategy

A *mixed strategy* for player i is a probability distribution $\sigma_i \in \Delta(A_i)$ over her pure actions. The *support* of σ_i, written $\operatorname{supp}(\sigma_i)$, is the set of actions played with positive probability. A profile $\sigma = (\sigma_1, \dots, \sigma_n)$ induces, via independent randomization, expected payoffs

$$u_i(\sigma) \;=\; \sum_{a \in A} \Big(\prod_{j \in N} \sigma_j(a_j) \Big)\, u_i(a).$$

By Theorem 2.2 (expected utility), players evaluate mixed profiles by these expectations. A *mixed-strategy Nash equilibrium* is defined exactly as before: σ^* such that $u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*)$ for every player i and every mixed strategy σ_i.

5.2 The indifference principle

The single most useful tool for computing mixed equilibria:

Proposition 5.1: Indifference principle

σ^* is a Nash equilibrium if and only if, for every player i:

1. every action in $\operatorname{supp}(\sigma_i^*)$ yields the same expected payoff against σ_{-i}^*; and
2. no action outside the support yields a higher expected payoff against σ_{-i}^*.

Proof. If some action in the support were worse than another, shifting its probability to the better action would raise u_i—contradicting equilibrium. Conversely, if all support actions are equally good and no outside action is better, every mixture over the support is a best response, including σ_i^* itself. □

Key Idea: You mix to make your *opponent* indifferent

In equilibrium, a player's mixing probabilities are determined by the requirement of keeping the *other* player indifferent—not herself well-off. This is counterintuitive and powerful: your equilibrium mixture depends on your opponent's payoffs, not your own.

Example 5.1: Matching Pennies solved

Let player 1 play H with probability p and player 2 play H with probability q. Player 1 is willing to mix only if H and T give her equal expected payoffs:

$$\underbrace{q(1)+(1-q)(-1)}_{u_1(H)} = \underbrace{q(-1)+(1-q)(1)}_{u_1(T)} \implies 2q-1 = 1-2q \implies q = \tfrac{1}{2}.$$

Symmetrically $p = \frac{1}{2}$. The unique equilibrium is both mixing 50–50, with value 0 each.

Example 5.2: Battle of the Sexes: the mixed equilibrium

In Example 3.3, besides the two pure equilibria there is a mixed one. Let row play B with probability p; column mixes only if indifferent: $u_2(B) = p \cdot 1$, $u_2(F) = (1-p)\cdot 3$, so $p = \frac{3}{4}$. Symmetrically column plays B with probability $q = \frac{1}{4}$. Expected payoffs are $\frac{3}{4}$ each—*worse for both* than either pure equilibrium (3 and 1): the players miscoordinate with probability $1 - \left(\frac{3}{4}\cdot\frac{1}{4} + \frac{1}{4}\cdot\frac{3}{4}\right) = \frac{5}{8}$.

Example 5.3: Hawk–Dove and mixing in populations

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

Two animals contest a resource of value V; fighting costs $C > V$. The symmetric mixed equilibrium has each playing Hawk with probability $p^* = V/C$. In biology this is reinterpreted: a stable *population* contains a fraction V/C of aggressive types. This reinterpretation of mixed equilibrium as population frequencies is the gateway to evolutionary game theory (Chapter 16).

5.3 Interpreting mixed strategies

Do people literally roll dice? Four standard interpretations:

1. **Deliberate randomization**: in strictly competitive settings (poker bluffs, soccer penalties, tax audits, military tactics), unpredictability is genuinely valuable and players consciously randomize.
2. **Population frequencies**: the mixture describes proportions of a population playing each pure action (Example 5.3).
3. **Beliefs**: σ_i is not what i does but what *others believe* i will do; equilibrium is a consistency condition on beliefs.
4. **Purification (Harsanyi 1973)**: each player has tiny private payoff perturbations and

plays a *pure* best response to her realized type; to outsiders this looks exactly like the mixed equilibrium. The mixed equilibrium is the limit of pure equilibria of nearby incomplete-information games.

Remark: Empirical evidence

Professional soccer penalty kicks are a celebrated test: kickers and goalkeepers choose sides with frequencies closely matching the mixed equilibrium computed from scoring rates, and their choices are serially independent, as theory requires (Palacios-Huerta 2003). Novices in labs, by contrast, alternate too much—true randomness is hard for humans.

5.4 Computing equilibria in larger games

For $2 \times n$ and small games, a practical algorithm:

1. Find all pure equilibria by best-response marking.

2. For each candidate *support profile* (subset of actions for each player), impose indifference within supports, solve the linear equations, and check: probabilities in $[0, 1]$, and no profitable deviation outside the support.

The number of support profiles grows exponentially—a first hint that computing Nash equilibria is genuinely hard (formalized in Chapter 17).

Example 5.4: A 2×3 game

	L	M	R
T	$2, 2$	$0, 3$	$1, 0$
B	$1, 0$	$2, 1$	$0, 2$

No pure equilibrium (check by marking). Suppose player 1 mixes $pT + (1-p)B$ with $p \in (0, 1)$. Player 2's expected payoffs: $u_2(L) = 2p$, $u_2(M) = 3p + (1-p) = 2p + 1$, $u_2(R) = 2(1-p)$. Since $u_2(M) > u_2(L)$ always, L is never in player 2's best responses when 1 mixes; try support $\{M, R\}$: indifference $2p + 1 = 2 - 2p$ gives $p = \frac{1}{4}$. Player 1's indifference between T and B against $qM + (1-q)R$: $0 \cdot q + 1(1-q) = 2q + 0(1-q)$, so $q = \frac{1}{3}$. Check deviations: player 2's payoff from L is $2p = \frac{1}{2} < u_2(M) = \frac{3}{2}$. Equilibrium: $\sigma_1 = (\frac{1}{4}, \frac{3}{4})$, $\sigma_2 = (0, \frac{1}{3}, \frac{2}{3})$.

Theorem 5.2: Oddness (Wilson 1971)

Almost every finite game (all games outside a closed set of measure zero in payoff space) has a finite and *odd* number of Nash equilibria.

This explains a pattern you may have noticed: Battle of the Sexes has 2 pure +1 mixed = 3; coordination games likewise; Matching Pennies has exactly 1.

Exercises

Exercise 5.1: Penalty kicks

A kicker chooses Left or Right; so does the goalkeeper. Scoring probabilities for the kicker: 0.6 if both choose the same side (keeper guesses right), 0.95 if kicker goes Left and keeper Right, 0.9 if kicker Right and keeper Left. Model as a zero-sum game (kicker maximizes scoring probability), find the mixed equilibrium, and compute the equilibrium scoring probability.

Exercise 5.2: Rock–Paper–Scissors

Show that uniform randomization by both players is the unique Nash equilibrium of Rock–Paper–Scissors (win $= 1$, loss $= -1$, tie $= 0$). Then modify the game so a win with Rock pays 2 instead of 1 (for both players' Rock wins) and recompute the equilibrium. Which action's probability *falls*? Explain the intuition.

Exercise 5.3: Volunteer's dilemma

Each of n neighbors can call the police (cost c) or not. If at least one calls, every neighbor gets benefit $v > c$; if nobody calls, all get 0. Find the symmetric mixed equilibrium probability of calling, and show that the probability that *nobody* calls *increases* with n. Discuss the connection to the bystander effect.

Exercise 5.4: Support checking

	L	R
T	$6, 2$	$0, 6$
M	$2, 4$	$4, 2$
B	$3, 3$	$3, 3$

Find *all* Nash equilibria (pure and mixed), checking every possible support profile for player 1.

Chapter 6

Extensive-Form Games and Backward Induction

So far, players moved simultaneously. Most real interactions unfold over *time*—entry then retaliation, offer then counteroffer—and timing changes everything: threats must be *credible* to matter.

6.1 Game trees

Definition 6.1: Extensive-form game with perfect information

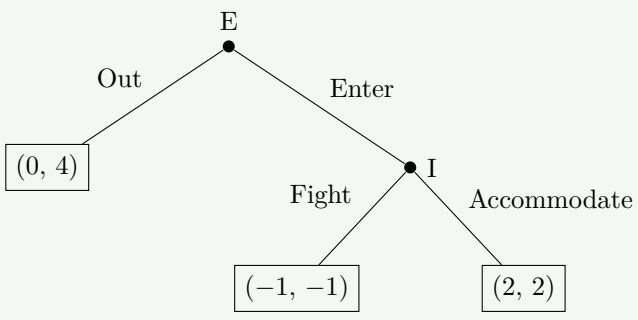
An *extensive-form game with perfect information* consists of:

- a finite *game tree*: a set of nodes with a unique root, where each non-terminal node is assigned to a player (or to *chance*, with fixed move probabilities) and each edge is an available *move*;
- *payoffs* u_i assigned to every terminal node;
- *perfect information*: each player, when moving, observes the entire history of play so far.

A *strategy* for player i is a function assigning a move to *every* node where i moves—a complete contingent plan, even for nodes that the strategy itself makes unreachable.

Example 6.1: The Entry Game

A potential entrant (E) decides whether to enter a market or stay out; if entry occurs, the incumbent (I) decides whether to fight a price war or accommodate.



Payoffs are (Entrant, Incumbent). The incumbent would love to deter entry by threatening to fight. Is the threat credible?

6.2 Backward induction

Definition 6.2: Backward induction

Solve the game from the end: at each final decision node, the mover chooses her best move; replace that node by the resulting payoff; repeat up the tree until the root. The resulting strategy profile is the *backward induction solution*.

In the Entry Game: at I's node, Accommodate (2) beats Fight (−1), so I accommodates. Anticipating this, E compares Out (0) with Enter (2) and *enters*. Backward induction outcome: (Enter, Accommodate).

Key Idea: Credibility

"Fight if you enter" is an *incredible threat*: once entry has occurred, carrying it out hurts the incumbent. Backward induction disciplines predictions by insisting that plans be optimal *at every point where they might be executed*—not just along the anticipated path. This idea, generalized as subgame perfection, is due to Selten.

Note the contrast with the normal form: if we write the Entry Game as a matrix, the profile (Out, Fight) is also a Nash equilibrium—E's "Out" is optimal given I's plan to fight, and I's plan is costless because it is never tested. Backward induction eliminates precisely such equilibria supported by empty threats.

Theorem 6.1: Zermelo's theorem (1913); Kuhn's theorem for perfect information

Every finite extensive-form game with perfect information has a pure-strategy Nash equilibrium, obtainable by backward induction. If no player is indifferent between any two terminal nodes, the backward induction solution is unique.

Corollary 6.2: Determinacy of chess

In chess (a finite two-player win/lose/draw game with perfect information), exactly one holds: White can force a win, or Black can force a win, or both can force at least a draw. (Which one is true remains computationally unknown.)

6.3 Subgames and subgame perfection

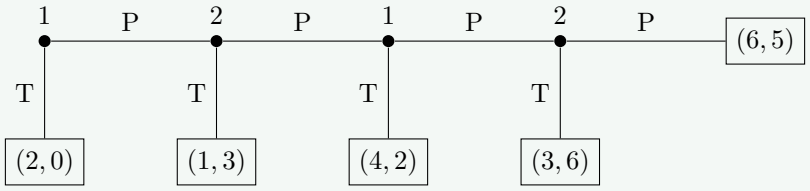
Definition 6.3: Subgame; subgame perfect equilibrium

A *subgame* is the part of the game tree starting at any single node and including all its successors. A strategy profile is a *subgame perfect equilibrium* (SPE) if it induces a Nash equilibrium *in every subgame*—including subgames off the anticipated path of play.

For perfect-information games, SPE coincides with backward induction. The concept's real power appears in games with simultaneous stages or imperfect information (Chapters 9, 11), where "backward induction" alone is not well-defined but subgame perfection still bites.

Example 6.2: The Centipede Game

Players alternate deciding whether to Take a growing pot or Pass. The pot starts at $(2, 0)$; each Pass multiplies total value but shifts advantage:



Backward induction: at the last node, 2 Takes ($6 > 5$); so at the third node, 1 Takes ($4 > 3$); so 2 Takes at the second ($3 > 2$); so 1 Takes *immediately*, yielding $(2, 0)$—even though passing throughout would give $(6, 5)$. In experiments, people pass for several rounds. The game sharpens the question of what common knowledge of rationality really implies, and whether a rational player should "believe" backward induction after observing a deviation from it.

Example 6.3: Stackelberg duopoly (first-mover advantage)

Take the Cournot setting of Example 4.2 but let firm 1 choose q_1 first, observed by firm 2. Backward induction: firm 2 best-responds, $q_2 = \frac{a-c-q_1}{2}$. Firm 1 anticipates this and maximizes

$$u_1 = q_1\Big(a - q_1 - \tfrac{a-c-q_1}{2} - c\Big) = \tfrac{q_1}{2}\,(a - c - q_1),$$

giving $q_1^* = \frac{a-c}{2} > \frac{a-c}{3}$ and $q_2^* = \frac{a-c}{4} < \frac{a-c}{3}$. Firm 1 earns $\frac{(a-c)^2}{8} > \frac{(a-c)^2}{9}$: *commitment is valuable*. Note the paradox of power: firm 2 knows more when moving (it observes q_1) yet earns less—in games, information and options can hurt you if others anticipate you have them.

Key Idea: Commitment

The ability to bind yourself—burn bridges, sign contracts, delegate to a tough agent—can be strategically valuable precisely because it removes your future flexibility and thereby changes others' behavior. Thomas Schelling received the 2005 Nobel Prize largely for developing this insight (deterrence theory in the Cold War).

6.4 The ultimatum game and the limits of backward induction

Example 6.4: Ultimatum game

A proposer offers a split $(10 - x, x)$ of \$10, $x \in \{0, 1, \ldots, 10\}$; the responder accepts (split implemented) or rejects (both get \$0). Backward induction with purely self-interested players: the responder accepts any $x \geq 1$ (and is indifferent at 0), so the proposer offers $x = 1$ (or 0). In experiments across many cultures, proposers typically offer 40–50%, and offers below 20–30% are rejected about half the time. The resolution is not that game theory is wrong, but that *payoffs are not money* (Remark in Chapter 1): responders derive negative utility from being treated unfairly, and proposers know it. With fairness-adjusted payoffs, generous offers *are* the backward induction solution.

6.5 Imperfect information: information sets

To connect the extensive and normal forms fully we need games where a mover does *not* know exactly where in the tree she is.

Definition 6.4: Information set

An *information set* for player i is a collection of decision nodes of i such that (i) i has the same available moves at each node, and (ii) when play reaches the set, i knows only that she is at *some* node of the set, not which one. A game has *perfect information* iff every information set is a singleton.

A simultaneous-move game is the special case where player 2's single information set contains all nodes following player 1's moves. Conversely, every extensive-form game has a normal-form representation (list each player's full contingent strategies and the resulting payoffs). Mixed strategies over full plans are called *mixed*; move-by-move randomizations are *behavior strategies*.

Theorem 6.3: Kuhn's theorem (1953)

In every extensive-form game with *perfect recall* (no player forgets what she once knew or did), mixed and behavior strategies are equivalent: each induces the same distributions over outcomes as some strategy of the other kind.

Subgame perfection extends verbatim: a subgame starts at any *singleton* information set whose successors don't cut across other information sets. When information sets prevent the tree from decomposing, subgame perfection loses power, and we will need sequential rationality with explicit beliefs—perfect Bayesian equilibrium (Chapter 11) and sequential equilibrium (Chapter 13).

Exercises

Exercise 6.1: Backward induction practice

Player 1 chooses L or R; after L, player 2 chooses ℓ (payoffs $(3,1)$) or r ($(1,2)$); after R, player 2 chooses ℓ' ($(2,1)$) or r' ($(0,0)$). Draw the tree, solve by backward induction, then write the 2×4 normal form (player 2 has four strategies) and find all pure Nash equilibria. Which are subgame perfect?

Exercise 6.2: Burning money / commitment

In the Battle of the Sexes (Example 3.3), suppose before play the row player can publicly burn 1 unit of her own utility (or not), after which the game is played with all past moves observed. Show by backward induction on the normal form (using iterated weak dominance) that the row player achieves her favorite equilibrium *without actually burning anything*. Comment on how commitment opportunities can matter even when unused.

Exercise 6.3: Sequential Battle of the Sexes

Let the row player move first in the Battle of the Sexes, observed by column. Solve by backward induction. Who benefits from moving first? Compare with Matching Pennies,

where moving first is fatal. What feature of the payoffs determines whether there is a first-mover or second-mover advantage?

Exercise 6.4: Pirates

Five perfectly rational pirates (seniority 1–5) divide 100 gold coins. The most senior proposes a division; all pirates (including the proposer) vote; a strict majority (or exactly half including the proposer's vote, if you prefer that variant—state your convention) passes the proposal; otherwise the proposer is thrown overboard and the next most senior proposes. Pirates value, in order: survival, gold, seeing others thrown overboard. Find the backward induction outcome with the strict-majority convention.

Part II

Intermediate: Dynamics, Information, and Applications

Chapter 7

Economic Applications: Markets and Strategy

The concepts of Part I earn their keep in applications. This chapter analyzes the canonical models of imperfect competition and collective action—workhorses of industrial organization and public economics.

7.1 Cournot competition: quantities

Recall the setting (Examples 4.2, 6.3): n firms simultaneously choose outputs $q_i \geq 0$; inverse demand is $P(Q) = a - Q$ with $Q = \sum_i q_i$; each firm has constant marginal cost $c < a$.

Proposition 7.1: Cournot oligopoly with n firms

The unique Nash equilibrium is symmetric with

$$q_i^* = \frac{a-c}{n+1}, \qquad Q^* = \frac{n}{n+1}(a-c), \qquad P^* = \frac{a+nc}{n+1}, \qquad u_i^* = \left(\frac{a-c}{n+1}\right)^2.$$

As $n \to \infty$: $P^* \to c$ (perfect competition), and industry profit vanishes.

Proof. Firm i maximizes $q_i(a - q_i - Q_{-i} - c)$ where $Q_{-i} = \sum_{j \neq i} q_j$. The first-order condition $a - 2q_i - Q_{-i} - c = 0$ gives the best response $q_i = \frac{a-c-Q_{-i}}{2}$. Summing FOCs over i: $n(a - c) - (n+1)Q^* = 0$, so $Q^* = \frac{n}{n+1}(a-c)$; symmetry and substitution yield the rest. Uniqueness follows because best responses are linear with slope $-\frac{1}{2}$ in aggregate opponents' output, so the system has a unique solution. □

Key Idea: Competition as a continuum

Cournot interpolates smoothly between monopoly ($n = 1$: $Q^* = \frac{a-c}{2}$, the monopoly output) and perfect competition ($n \to \infty$: price equals marginal cost). Market structure maps into market power via a strategic model—the founding insight of industrial organization.

7.2 Bertrand competition: prices

Now let two identical firms set *prices*; consumers buy from the cheaper firm (splitting demand $D(p)$ at ties), and each firm has marginal cost c.

Proposition 7.2: Bertrand paradox

The unique Nash equilibrium has both firms pricing at marginal cost: $p_1 = p_2 = c$, with zero profits—even with only *two* firms.

Proof sketch. No equilibrium can have $\min(p_1, p_2) > c$: the higher-priced firm (or either firm, if tied) can undercut by a penny and capture the whole market at nearly the same margin. Prices below c yield losses, and asymmetric prices above cost invite undercutting. Checking $p_1 = p_2 = c$: any deviation upward sells nothing; downward loses money. □

The stark contrast with Cournot ($n = 2$ gives substantial profits) teaches an important modeling lesson: *the choice of strategic variable matters.* Real markets escape the Bertrand trap through capacity constraints (Kreps–Scheinkman: capacity choice then price competition reproduces Cournot outcomes), product differentiation, search costs, or repeated interaction (Chapter 9).

Example 7.1: Bertrand with differentiated products

Let demands be $q_1 = 1 - p_1 + bp_2$ and $q_2 = 1 - p_2 + bp_1$ with $0 < b < 1$ (imperfect substitutes), costs zero. Firm i maximizes $p_i q_i$; the FOC gives $\text{BR}_i(p_j) = \frac{1+bp_j}{2}$. Note best responses *increase* in the rival's price: prices are *strategic complements* (while Cournot quantities are *strategic substitutes*). Equilibrium: $p^* = \frac{1}{2-b}$, rising in b. Differentiation softens competition.

7.3 Public goods and free-riding

Example 7.2: Voluntary contribution to a public good

Each of n citizens has wealth w and chooses a contribution $g_i \in [0, w]$ to a public good. Payoffs: $u_i = w - g_i + m\sum_{j=1}^{n} g_j$ where $\frac{1}{n} < m < 1$ (the *marginal per capita return*). Since $\partial u_i/\partial g_i = m - 1 < 0$, contributing nothing is a *strictly dominant strategy*; the unique equilibrium is $g_i = 0$ for all. Yet full contribution would give everyone $w(nm) > w$: an n-player Prisoner's Dilemma. Laboratory subjects typically contribute 40–60% at first, decaying with repetition—and contributions are dramatically sustained if players can *punish* free-riders at a cost (Fehr–Gächter 2000), previewing the role of sanctions in Chapter 9.

7.4 Elections: the median voter

Proposition 7.3: Median voter theorem (Hotelling–Downs)

Two office-motivated candidates simultaneously choose platforms on a policy line; voters vote for the closer platform, and voter ideal points have median x_m. The unique Nash equilibrium has both candidates locating at x_m.

Proof sketch. If a candidate stands at $x \neq x_m$, the rival can locate between x and x_m (slightly on the median side) and capture more than half the electorate. Only (x_m, x_m) resists all such deviations, since deviating from the median concedes a majority to the opponent. □

The model explains the centrist pull in two-party systems—and its failures (primaries, turnout, multi-dimensional policy, entry of third candidates) map exactly onto its assumptions,

illustrating how a sharp game-theoretic benchmark organizes empirical debate.

7.5 Congestion: selfish routing

Example 7.3: Pigou's example

A unit mass of drivers travels from A to B via two roads: road 1 has fixed travel time 1; road 2 has travel time x equal to the fraction x of drivers using it. Equilibrium: drivers switch to road 2 until it is as slow as road 1, i.e. $x = 1$ (everyone on road 2), with average travel time 1. The social optimum splits traffic ($x = \frac{1}{2}$), for average time $\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \frac{1}{2} = \frac{3}{4}$. The ratio $\frac{1}{3/4} = \frac{4}{3}$ is the *price of anarchy*, and it is the worst possible for affine congestion costs (Chapter 17).

Example 7.4: Braess's paradox

Adding a zero-cost expressway to a congested network can make *every* driver worse off in equilibrium: the new link attracts traffic that congests links everyone must use. Removing roads has, on occasion, genuinely improved real urban traffic (documented cases in Seoul and New York). The paradox is impossible for a central planner—only selfish equilibrium behavior produces it—and is a vivid caution against "more options are always better" reasoning in strategic settings.

7.6 Strategic complements vs. substitutes: a unifying lens

Definition 7.1: Strategic complements and substitutes

In a two-player game with ordered action sets and smooth payoffs, actions are *strategic complements* if best responses are increasing in the opponent's action ($\partial^2 u_i / \partial a_i \partial a_j > 0$), and *strategic substitutes* if decreasing.

- **Substitutes** (Cournot quantities, public-good contributions, fishing effort): rivals' aggression makes you retreat. Comparative statics are moderate; free-riding is the typical pathology.
- **Complements** (Bertrand prices, bank runs, currency attacks, adoption of platforms, arms races): rivals' aggression makes you more aggressive. Multiplicity of equilibria, tipping, and self-fulfilling crises are the typical phenomena.

This classification (formalized in supermodular game theory—Topkis, Milgrom–Roberts, Vives) often predicts the *qualitative* behavior of a model before any computation.

Exercises

Exercise 7.1: Cournot with asymmetric costs

Two firms have costs $c_1 < c_2$, demand $P = a - Q$. Find the Nash equilibrium quantities and profits (assume both produce). Show the low-cost firm produces more, and find the cost gap at which firm 2 exits ($q_2^* = 0$).

Exercise 7.2: Bertrand with asymmetric costs

Two firms with costs $c_1 < c_2$ compete in prices for a market of unit demand with reservation price $v > c_2$. Characterize equilibrium outcomes (be careful: with continuous prices, use the convention that the low-cost firm wins ties, or work with discrete price grids and take limits). Who serves the market and at what price?

Exercise 7.3: Tullock rent-seeking contest

Two lobbyists spend $x_1, x_2 \geq 0$ competing for a prize V; lobbyist i wins with probability $x_i/(x_1 + x_2)$ (and $\frac{1}{2}$ if both spend 0). Payoffs are $\frac{x_i}{x_1+x_2}V - x_i$. Find the symmetric equilibrium and show total dissipation $x_1^* + x_2^* = V/2$: half the prize's value is burned in the contest.

Exercise 7.4: Bank run as coordination

Two depositors each placed D in a bank that invested in an illiquid asset. Each may Withdraw now or Wait. If both wait, the investment matures and each gets $R > D$. If one withdraws and the other waits, the withdrawer recovers D in full while the waiter receives less than D (liquidation loss). If both withdraw, each gets some amount strictly between those two (partial recovery). Set up a payoff matrix with concrete numbers satisfying these orderings, find both pure equilibria, and explain why deposit insurance eliminates the bad one (Diamond–Dybvig 1983, Nobel 2022).

Chapter 8

Bargaining

How is a surplus divided when two parties must agree to realize it? Bargaining theory has two complementary traditions: Nash's *axiomatic* approach (what properties should a solution have?) and Rubinstein's *strategic* approach (model the back-and-forth explicitly and solve for equilibrium). Their eventual convergence is one of game theory's most satisfying results.

8.1 The axiomatic approach: Nash bargaining

Definition 8.1: Bargaining problem

A (two-player) *bargaining problem* is a pair (S, d) where $S \subset \mathbb{R}^2$ is a compact convex set of feasible utility pairs, and $d = (d_1, d_2) \in S$ is the *disagreement point* (what each gets if talks fail), with some $s \in S$ satisfying $s > d$ componentwise. A *bargaining solution* is a map $f(S, d) \in S$.

Theorem 8.1: Nash bargaining solution (1950)

There is exactly one bargaining solution satisfying:

1. **Pareto efficiency**: no feasible point dominates $f(S, d)$;
2. **Symmetry**: if S is symmetric and $d_1 = d_2$, then $f_1 = f_2$;
3. **Invariance to affine transformations** of either player's utility;
4. **Independence of irrelevant alternatives (IIA)**: if $T \subseteq S$ and $f(S, d) \in T$, then $f(T, d) = f(S, d)$.

It is the point maximizing the *Nash product*:

$$f(S, d) \;=\; \arg\max_{s \in S,\; s \geq d} (s_1 - d_1)(s_2 - d_2).$$

Example 8.1: Splitting a dollar with risk aversion

Two players split \$1; utilities are $u_1(x) = x$ and $u_2(x) = x^\alpha$ with $\alpha \in (0, 1)$ (player 2 is risk-averse), $d = (0, 0)$. Maximizing $x \cdot (1 - x)^\alpha$ gives $x^* = \frac{1}{1+\alpha}$. As α falls (more risk aversion), player 2's share $\frac{\alpha}{1+\alpha}$ *shrinks*: risk aversion is a bargaining weakness. Impatience will play the identical role in the strategic model.

The generalized (asymmetric) Nash solution maximizes $(s_1 - d_1)^{\beta}(s_2 - d_2)^{1-\beta}$, where β captures "bargaining power"—a black box that the strategic approach opens.

8.2 The strategic approach: alternating offers

Definition 8.2: Rubinstein's alternating-offers game (1982)

Players alternate proposing divisions $(x, 1-x)$ of a pie of size 1: player 1 proposes in periods $0, 2, 4, \ldots$, player 2 in periods $1, 3, 5, \ldots$. A proposal accepted in period t gives discounted payoffs $(\delta_1^t x,\ \delta_2^t(1-x))$, where $\delta_i \in (0,1)$ are discount factors. Rejection moves to the next period; perpetual disagreement gives $(0, 0)$.

Theorem 8.2: Rubinstein (1982)

The alternating-offers game has a *unique* subgame perfect equilibrium. Agreement is immediate; player 1 proposes and player 2 accepts

$$x^* = \frac{1-\delta_2}{1-\delta_1\delta_2}, \qquad 1 - x^* = \frac{\delta_2(1-\delta_1)}{1-\delta_1\delta_2}.$$

Proof idea (stationary equilibrium + uniqueness sketch). Look for values (x, y): x is what the current proposer gets when player 1 proposes, y what player 2 gets as proposer. A responder accepts anything at least as good as her continuation value from proposing next period:

$$1 - x = \delta_2\, y, \qquad 1 - y = \delta_1\, x.$$

Solving: $x = \frac{1-\delta_2}{1-\delta_1\delta_2}$, $y = \frac{1-\delta_1}{1-\delta_1\delta_2}$. Uniqueness (Shaked–Sutton argument): let $\bar{x}, \underline{x}$ be the supremum and infimum of player 1's SPE payoffs as proposer. Optimal play bounds $\bar{x} \le 1 - \delta_2(1 - \delta_1\bar{x})$ and $\underline{x} \ge 1 - \delta_2(1 - \delta_1\underline{x})$, which force $\bar{x} = \underline{x} = x^*$. □

Key Idea: What determines bargaining power

In Rubinstein's model, shares are pinned down by *relative patience*: as $\delta_2 \to 0$ (player 2 desperate to settle), player 1 takes almost everything; with equal discounting $\delta_1 = \delta_2 = \delta$, player 1 gets $\frac{1}{1+\delta} > \frac{1}{2}$—a vanishing first-mover advantage as $\delta \to 1$. Outside options, risk of breakdown, and commitment tactics all map into the same calculus. "Bargaining power" is no longer a free parameter but an implication of the environment.

Theorem 8.3: Convergence (Binmore–Rubinstein–Wolinsky 1986)

As the time between offers shrinks to zero, the unique SPE outcome of the alternating-offers game converges to the (generalized) Nash bargaining solution, with disagreement point given by the payoffs during delay/breakdown and bargaining weights determined by relative discount rates.

This "Nash program"—grounding axiomatic solutions in explicit non-cooperative games—is a template used throughout cooperative theory (Chapter 15).

8.3 Bargaining with outside options

Suppose player 2 can leave and take an outside option worth w at any response node.

Proposition 8.4: Outside option principle

If $w \leq \delta_2 y^*$ (the option is worse than 2's equilibrium continuation), the equilibrium is unchanged—the option is an *empty threat*. If $w > \delta_2 y^*$, player 2 receives exactly w: the option binds but confers no extra rent.

The lesson—threats affect outcomes only when credible, and then only at their exact value—recurs in wage negotiation, litigation-vs-settlement, and M&A practice. It also warns against the naïve "split the difference plus my outside option" intuition.

8.4 Applications

Example 8.2: Union–firm wage bargaining

A firm and a union bargain over the wage w; production yields revenue R per period, the union's strike payoff is s per period. With alternating offers and a common discount factor near 1, the outcome approximates the Nash solution with disagreement $(0, s)$: $w^* \approx \frac{R+s}{2}$. The union's strike fund s raises the wage with weight $\frac{1}{2}$ even though no strike ever occurs in equilibrium. *Strength that is never used still shapes the deal.*

Example 8.3: Pre-trial settlement

Litigation will cost each party c and transfer J from defendant to plaintiff with probability p (common knowledge). The settlement surplus relative to trial is $2c$; Nash bargaining predicts settlement at pJ plus/minus the split of costs saved. Settlement fails in practice mainly through *asymmetric information* (each side optimistic about p)—the subject of Chapters 10 and 11, and a leading explanation for costly disputes, strikes, and wars.

Exercises

Exercise 8.1: Nash product computation

S is the triangle with vertices $(0,0)$, $(4,0)$, $(0,2)$; $d = (1, \frac{1}{2})$. Compute the Nash bargaining solution. Verify IIA by re-solving on the smaller set $S' = S \cap \{s_1 \leq 3\}$ and checking the solution is unchanged (or explain if it moved).

Exercise 8.2: Three rounds of alternating offers

Solve by backward induction the three-period alternating-offers game (player 1 proposes at $t = 0$ and $t = 2$, player 2 at $t = 1$; if the $t = 2$ proposal is rejected both get 0), with common discount factor δ. Verify that as the horizon grows, proposer shares approach $\frac{1}{1+\delta}$.

Exercise 8.3: Impatience vs. risk of breakdown

In the alternating-offers game, replace discounting with an exogenous breakdown risk: after each rejection, with probability $1 - \pi$ negotiations collapse and players receive (b_1, b_2). Derive the analogue of the Rubinstein shares and interpret the limit $\pi \to 1$.

Exercise 8.4: Deadline effect

Modify the alternating-offers game so the pie disappears after period T (no discounting). Show the last proposer takes (almost) everything, and explain why real negotiations so often settle "on the courthouse steps."

Chapter 9

Repeated Games and the Emergence of Cooperation

The Prisoner's Dilemma predicts defection—yet cooperation among self-interested parties is everywhere: cartels hold, merchants extend credit, countries honor treaties, strangers leave tips at restaurants they'll revisit. The unifying explanation: *the interaction is repeated*, so today's betrayal can be punished tomorrow.

9.1 Finitely repeated games

Let G be a *stage game*, repeated T times with observed history and payoffs summed (or discounted).

Proposition 9.1: Unraveling

If the stage game G has a *unique* Nash equilibrium, then the finitely repeated game has a unique subgame perfect equilibrium: the stage equilibrium is played in every period regardless of history.

Proof. In the last period T, no future exists, so play must be the stage Nash equilibrium in every subgame. Then period $T-1$ actions cannot influence period-T outcomes, so period $T-1$ must also be stage-Nash; induct backwards. □

So the finitely repeated Prisoner's Dilemma yields defection throughout—backward induction's cold logic again (compare the Centipede, Example 6.2). Experiments show cooperation until near the end, collapsing in the final rounds. When the stage game has *multiple* equilibria, however, even finite repetition can support non-equilibrium stage outcomes early on (rewarding good behavior by selecting the better continuation equilibrium).

9.2 Infinitely repeated games

Let G be played in periods $t = 0, 1, 2, \ldots$ with common discount factor $\delta \in (0, 1)$; player i's repeated-game payoff is the normalized sum $(1-\delta)\sum_t \delta^t u_i(a^t)$. Interpret δ either as time preference or as the probability the relationship continues another period—so "infinite" really means *indefinite*: no commonly known final date from which unraveling can start.

Definition 9.1: Trigger strategies

In the repeated Prisoner's Dilemma, *grim trigger* plays C initially and after any history of mutual cooperation, and plays D forever after any defection. *Tit-for-tat* plays C first and thereafter copies the opponent's previous move.

Proposition 9.2: Cooperation via grim trigger

In the Prisoner's Dilemma of Example 3.1 (payoffs $3,3$ / $0,4$ / $4,0$ / $1,1$), grim trigger by both players is a subgame perfect equilibrium iff

$$\underbrace{(1-\delta)\cdot 4+\delta\cdot 1}_{\text{defect today, punished forever}} \;\le\; \underbrace{3}_{\text{cooperate forever}} \quad\Longleftrightarrow\quad \delta \ge \tfrac{1}{3}.$$

Proof. On-path one-period deviation: gain $4-3$ today, lose $3-1$ every period after, discounted—the displayed inequality. Off path (after a defection), both play D forever, a stage Nash equilibrium repeated, hence immune to deviations. By the *one-shot deviation principle* (it suffices to check single-period deviations in discounted repeated games), the profile is SPE. □

Key Idea: The shadow of the future

Cooperation is sustainable when players are patient enough that the long-run value of the relationship outweighs the short-run gain from betrayal. The condition $\delta \ge \delta^{\min}$ formalizes folk wisdom: relationships sustain honesty where one-shot encounters do not; cartels are more stable in frequent-interaction markets; online platforms manufacture a "future" via reputation scores precisely to enable trade among strangers.

9.3 The Folk Theorem

Which payoffs can patient players sustain? Essentially all reasonable ones.

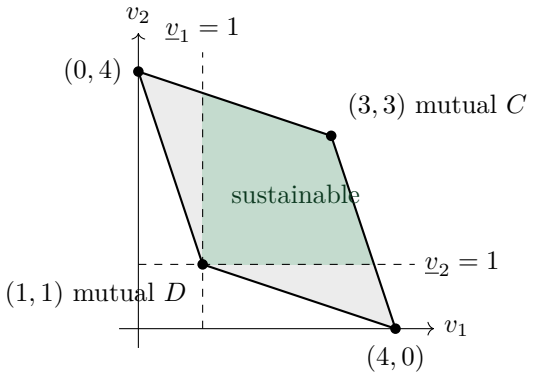
Definition 9.2: Minmax and feasible payoffs

Player i's *minmax payoff* is $\underline{v}_i = \min_{\alpha_{-i}} \max_{a_i} u_i(a_i, \alpha_{-i})$: the worst the others can inflict on a best-responding i. A payoff vector is *feasible* if it is a convex combination of stage-game payoff vectors, and *individually rational* if it gives each player strictly more than $\underline{v}_i$.

Theorem 9.3: Folk Theorem (Fudenberg–Maskin 1986)

Let v be any feasible and strictly individually rational payoff vector of the stage game. Then (under a mild dimensionality condition) for all δ sufficiently close to 1, the infinitely repeated game has a *subgame perfect* equilibrium with payoffs v.

For the Prisoner's Dilemma the theorem's set is the shaded region below:



Remark: Blessing and curse

The Folk Theorem is good news (cooperation is possible) and an embarrassment of riches (so is almost everything else, including unequal and inefficient outcomes). Repetition converts the theory from a prediction machine into a *possibility* theory; which equilibrium prevails is then a matter of history, norms, focality, and renegotiation—active research areas.

9.4 Cartels, price wars, and imperfect monitoring

Example 9.1: Collusion in Cournot competition

Two Cournot firms (Example 4.2) can sustain the monopoly output (each producing $\frac{a-c}{4}$, earning $\frac{(a-c)^2}{8}$ each) with grim trigger reversion to Cournot ($\frac{(a-c)^2}{9}$ each) iff $\delta \geq \delta^{\min} = \frac{9}{17}$. (The optimal single-period cheat against a colluding rival earns $\frac{9(a-c)^2}{64}$; plug into the deviation inequality.) Antitrust implications: collusion is easier with fewer firms, frequent interaction, transparency of rivals' prices—which is why publishing detailed transaction prices can be *anti*-competitive.

When rivals' actions are observed only through noisy signals (market price moves with demand shocks), punishments trigger *on the equilibrium path*: occasional price wars are not collusion breaking down but the very machinery that sustains it (Green–Porter 1984). The general theory of repeated games with imperfect monitoring (Abreu–Pearce–Stacchetti; Fudenberg–Levine–Maskin) is a highlight of modern economic theory.

9.5 Reputation

Example 9.2: Chain store paradox and reputation

An incumbent faces entrants in T markets sequentially (Example 6.1 repeated). Backward induction says: accommodate everywhere, so all enter—Selten's *chain store paradox*. Kreps–Milgrom–Roberts–Wilson (1982) resolution: if entrants assign even tiny probability ε to the incumbent being a "tough" type who enjoys fighting, the rational incumbent *mimics* the tough type by fighting early entrants, deterring entry for most of a long horizon. A drop of incomplete information plus repetition generates the reputation effects that plain backward induction destroys. (Tools for this analysis: Chapter 11.)

Historical Note: Axelrod's tournaments

In Robert Axelrod's famous computer tournaments (1980), strategies for the repeated Prisoner's Dilemma were submitted by game theorists and competed round-robin. The winner—both times—was the simplest entry: Anatol Rapoport's TIT-FOR-TAT. Axelrod distilled its virtues: it is *nice* (never defects first), *retaliatory* (punishes immediately), *forgiving* (returns to cooperation), and *clear*. Caveat: TFT is not subgame perfect, and in noisy environments two TFT players fall into echoing vendettas—more forgiving variants ("generous TFT", "win-stay lose-shift") do better under noise.

Exercises

Exercise 9.1: Threshold discount factors

For the Prisoner's Dilemma with payoffs $C, C \to (2, 2)$; $D, C \to (5, 0)$; $C, D \to (0, 5)$; $D, D \to (1, 1)$: find the minimum δ sustaining mutual cooperation via (a) grim trigger; (b) one-period punishment (defect once, then return to cooperation if the opponent also defected during the punishment). Explain why (b) requires a higher δ.

Exercise 9.2: Tit-for-tat is not subgame perfect

In the Prisoner's Dilemma of Example 3.1, show that mutual tit-for-tat fails the one-shot deviation test after the history in which player 1 defected once: alternating punishment cycles give lower payoffs than an immediate return to cooperation for some δ. Find the range of δ (if any) for which TFT survives the relevant one-shot deviations.

Exercise 9.3: Folk theorem geometry

For the Battle of the Sexes (Example 3.3), draw the feasible set, compute both players' minmax payoffs, and shade the payoffs sustainable by patient players. Can the players sustain the alternating scheme (Ballet in even periods, Football in odd)? Compute its payoff vector for $\delta \to 1$.

Exercise 9.4: Community enforcement

A large population is randomly matched each period to play the Prisoner's Dilemma with a new partner. Personal punishment is impossible (you never meet again), but each player carries a public label "good" or "bad": defecting against a good partner makes you bad. Design a strategy rule under which cooperation with good partners is an equilibrium for large δ, and discuss what happens if labels are sometimes wrong (interpret: gossip, credit scores, platform ratings).

Chapter 10

Bayesian Games: Incomplete Information

Until now, everyone knew everyone's payoffs. In reality, a bidder doesn't know rivals' valuations; an insurer doesn't know the applicant's risk; a poker player doesn't see opponents' cards. Modeling such *incomplete information* coherently stumped the field until Harsanyi's breakthrough.

10.1 Harsanyi's framework: types and beliefs

Harsanyi's idea (1967–68): convert incomplete information into *imperfect* information by introducing a chance move. "Nature" first draws each player's private characteristics—her *type*—from a commonly known prior distribution; each player observes her own type only, then actions are chosen. Infinite hierarchies of beliefs (what I believe about your payoff, what I believe you believe about mine, ...) collapse into a tractable object.

Definition 10.1: Bayesian game

A *Bayesian game* is $\langle N, (A_i), (\Theta_i), p, (u_i) \rangle$ where Θ_i is player i's *type space*, $p \in \Delta(\Theta_1 \times \cdots \times \Theta_n)$ is the common prior, and payoffs $u_i(a; \theta)$ may depend on everyone's types. A *(pure) strategy* is a function $s_i : \Theta_i \to A_i$—a full plan of what to do for each type you might be.

Definition 10.2: Bayesian Nash equilibrium

A strategy profile $(s_1, \ldots, s_n)$ is a *Bayesian Nash equilibrium* (BNE) if for every player i and every type θ_i, the action $s_i(\theta_i)$ maximizes i's expected payoff given her belief $p(\theta_{-i} \mid \theta_i)$ about others' types and their strategies:

$$s_i(\theta_i) \in \arg\max_{a_i \in A_i} \; \mathbb{E}_{\theta_{-i}}\big[\, u_i(a_i, s_{-i}(\theta_{-i}); \, \theta_i, \theta_{-i}) \,\big|\, \theta_i \big].$$

Operationally: treat each *type* as a separate player and find an ordinary Nash equilibrium of the expanded game.

Example 10.1: Battle of the Sexes with uncertain interest

Row loves Column, but is unsure whether Column wants to meet (type θ^m, probability $\frac{1}{2}$) or avoid her (type θ^a): meeting types have BoS payoffs; avoiding types get the BoS payoffs with Column's coordination payoffs flipped to miscoordination. Solving type-by-type: the avoiding type has a dominant tendency to mismatch, and Row optimally goes to her favorite event regardless; the meeting type follows Row. Working through the best responses gives a strict BNE: Row plays B; Column plays B if θ^m, F if θ^a. The point: equilibrium play of *every* type must be specified, even types that do not occur in the realized play.

Example 10.2: Cournot with unknown costs

Firm 1's cost is c (common knowledge); firm 2's cost is c_L w.p. μ or c_H w.p. $1-\mu$, privately known. Demand $P = a - q_1 - q_2$. Each firm-2 type sets $q_2(c_2) = \frac{a-c_2-q_1}{2}$; firm 1 best-responds to the *expected* output $\bar{q}_2 = \mu q_2(c_L) + (1-\mu)q_2(c_H)$. Solving the three linear equations gives, with $\bar{c}_2 = \mu c_L + (1-\mu)c_H$:

$$q_1^* = \frac{a - 2c + \bar{c}_2}{3}, \qquad q_2^*(c_2) = \frac{a - 2c_2 + c}{3} + \frac{c_2 - \bar{c}_2}{6}.$$

The high-cost type produces more than it would under complete information (its rival, not knowing costs are high, produces less aggressively than it otherwise would). Information asymmetries change strategic behavior even for the *uninformed* party.

10.2 Auctions: private values

Auctions are the flagship application of Bayesian games—and of game theory in practice, running trillions of dollars of ad, spectrum, treasury, and electricity markets.

Setup (*independent private values*, IPV): one object; n bidders; valuations v_i drawn i.i.d. from a distribution F on $[0, \bar{v}]$, each privately known.

Example 10.3: Second-price (Vickrey) auction

Highest bidder wins, pays the second-highest bid. As shown in Example 3.5, bidding $b_i = v_i$ is weakly dominant. Truthful bidding requires no knowledge of F, of n, or of others' rationality—an ideal of *robust* mechanism design. Revenue: the expected second-highest valuation.

Example 10.4: First-price auction

Highest bidder wins, pays *her own* bid. Now bidding v_i guarantees zero profit; optimal bids *shade* below the valuation, trading off win probability against margin. Look for a symmetric increasing equilibrium $b(v)$. Bidder i with value v choosing bid $b(z)$ (imitating type z) wins iff all rivals' values are below z, probability $F(z)^{n-1}$; expected profit $\pi(z) = F(z)^{n-1}(v - b(z))$. The FOC at $z = v$ plus the boundary condition $b(0) = 0$ yields

$$b(v) \;=\; v - \frac{\int_0^v F(x)^{n-1}\,dx}{F(v)^{n-1}} \;=\; \mathbb{E}\big[\max_{j\neq i} v_j \,\big|\, \max_{j\neq i} v_j < v\big].$$

With uniform values on $[0,1]$: $b(v) = \frac{n-1}{n}v$. You bid the *expected highest rival value*

conditional on winning—shading exactly to the point where you'd tie with your strongest rival.

Theorem 10.1: Revenue Equivalence (Vickrey 1961; Myerson 1981)

In the IPV setting with risk-neutral bidders, *any* auction in which (i) the object always goes to the highest-value bidder in equilibrium and (ii) a bidder with the lowest possible value gets zero expected payoff, yields the *same expected revenue* to the seller. In particular, first-price, second-price, English (ascending), and Dutch (descending) auctions are all revenue-equivalent.

The proof is a corollary of the mechanism-design machinery in Chapter 14. The theorem's force is in its failure modes: risk aversion, correlated values, asymmetric bidders, and collusion break the equivalence in predictable directions, and auction *design* becomes consequential (Chapter 14).

10.3 Common values and the winner's curse

Example 10.5: The winner's curse

An oil tract has the same unknown value V to all bidders; each observes a noisy private signal $s_i = V + \varepsilon_i$. If everyone bids their unbiased estimate, the *winner* is the bidder with the largest positive error: winning is bad news. Rational bidders shade for this selection effect, bidding as if their signal were the highest—i.e., conditioning on the event of winning. Inexperienced bidders systematically overpay (documented in offshore oil leases, book publishing advances, free agency in sports, and corporate takeovers).

Key Idea: Condition on winning

In any common-value interaction—auctions, trading against possibly better-informed counterparties, accepting a deal others declined—evaluate the object *conditional on your action succeeding*. "If my bid wins / my offer is accepted, what does that tell me about what others know?" Ignoring this is among the most costly strategic errors in practice.

A related classic: *trade with adverse selection.*

Example 10.6: Akerlof's lemons market

A seller knows her car's quality $q \sim U[0,1]$; the buyer values quality q at $\frac{3}{2}q$, the seller at q. At price p, only sellers with $q \le p$ sell; average quality traded is $p/2$, worth $\frac{3}{4}p < p$ to the buyer. The market unravels: only $q = 0$ trades in equilibrium, though *every* car has positive gains from trade. Asymmetric information alone can destroy a market—the founding result of information economics (Akerlof 1970, Nobel 2001), and the reason for warranties, inspections, and certification (signaling: next chapter).

10.4 Global games: equilibrium selection through uncertainty

Incomplete information can *reduce* the multiplicity of equilibria.

Example 10.7: Currency attack (Morris–Shin 1998, sketch)

Speculators decide whether to attack a currency peg; attacking costs t and pays 1 if the peg breaks, which happens if enough speculators attack relative to fundamentals θ. With θ common knowledge, there is a range of θ with two equilibria (attack/no attack)—self-fulfilling crises. Now let each speculator observe θ with tiny idiosyncratic noise. Iterated dominance from the extreme regions (very bad θ: attacking dominant; very good: abstaining dominant) propagates inward and selects a *unique* equilibrium: attack iff the signal is below a cutoff θ^*. Vanishingly small uncertainty destroys the multiplicity, and the selected cutoff corresponds to the *risk-dominant* behavior. Global games are now a standard tool for bank runs, debt crises, and regime change.

Exercises

Exercise 10.1: First-price auction with two uniform bidders

Two bidders, $v_i \sim U[0,1]$ i.i.d. Verify directly that $b(v) = v/2$ is a symmetric BNE of the first-price auction. Compute the seller's expected revenue and check it equals the expected second-highest value $\mathbb{E}[\min(v_1, v_2)] = \frac{1}{3}$.

Exercise 10.2: All-pay auction

Two bidders with $v_i \sim U[0,1]$; both pay their bids, highest wins. Show $b(v) = v^2/2$ is a symmetric BNE, and verify revenue equivalence with the first- and second-price auctions. Name two real-world contests that are effectively all-pay.

Exercise 10.3: Public good with private values

Two neighbors decide simultaneously whether to contribute to a streetlight costing $c = 1$ to each contributor. The light is built if at least one contributes; enjoying it is worth $\theta_i \sim U[0,2]$, privately known. Find the symmetric cutoff BNE (contribute iff $\theta_i \geq \theta^*$) and compare total surplus with the first-best rule.

Exercise 10.4: Trading against the informed

You hold an asset worth $\theta \sim U[0,100]$ to you; a counterparty who *knows* θ offers to buy at price p of her choosing, which you accept or reject; the asset is worth 1.5θ to her. Characterize which offers get made and accepted in equilibrium, and compute your expected payoff from a policy of accepting any $p \geq 50$. Relate to the lemons example: why should an attractive-looking offer from an informed counterparty make you suspicious?

Chapter 11

Dynamic Games with Incomplete Information: Signaling and Screening

When informed players move first, their *actions speak*: choices reveal, conceal, or distort private information. This chapter develops perfect Bayesian equilibrium and its two flagship applications—signaling and cheap talk—plus the screening problem faced by uninformed movers.

11.1 Perfect Bayesian equilibrium

Subgame perfection has no grip when information sets prevent the tree from decomposing into subgames. The fix: require optimality at *every information set*, given explicitly specified *beliefs* that are consistent with Bayes' rule.

Definition 11.1: Perfect Bayesian equilibrium

A *perfect Bayesian equilibrium* (PBE) consists of a strategy profile and a *belief system* (a probability distribution over nodes at each information set) such that:

1. **Sequential rationality**: at every information set, the mover's strategy is optimal from that point on, given her beliefs and others' strategies;
2. **Belief consistency**: on the equilibrium path, beliefs are derived from the strategies by Bayes' rule; off the path (probability-zero events), beliefs are unrestricted (but must still support sequentially rational behavior).

The freedom in off-path beliefs generates multiplicity; *refinements* such as the intuitive criterion (below) and sequential equilibrium (Chapter 13) discipline it.

11.2 Signaling: Spence's education model

Example 11.1: Job-market signaling (Spence 1973)

A worker is of high ability $\theta_H = 2$ (probability μ) or low ability $\theta_L = 1$. She chooses education level $e \geq 0$ *before* entering the job market; education is pure signal (does not raise productivity) and costs e/θ —*cheaper for the able*. Competitive firms then pay a wage equal to expected productivity given e. Worker's payoff: $w - e/\theta$.
Separating equilibrium: types choose different education levels, revealing themselves;

$e_L^* = 0$, $w(e_L^*) = 1$ and $e_H^* = e^s$, $w(e^s) = 2$. Incentive compatibility requires the low type not to imitate: $2 - e^s/1 \leq 1 \Rightarrow e^s \geq 1$; and the high type to prefer signaling: $2 - e^s/2 \geq 1 \Rightarrow e^s \leq 2$. Any $e^s \in [1, 2]$ works, supported by beliefs that any $e < e^s$ means low ability.

Pooling equilibria also exist: both types choose e^p and are paid $1 + \mu$, supported by pessimistic off-path beliefs.

Definition 11.2: Single-crossing property

Signaling works when the signal is *differentially costly*: the marginal cost of the signal is lower for higher types, so indifference curves of different types cross exactly once in signal–transfer space. This *single-crossing* (Spence–Mirrlees) condition is the engine of all signaling and screening models.

Theorem 11.1: The intuitive criterion (Cho–Kreps 1987)

Fix a PBE. Suppose an off-path signal is *equilibrium-dominated* for some type (that type could not gain by sending it under *any* belief-response), but some other type *could* gain if it convinced the receiver. Then beliefs after that signal must put zero probability on the dominated type. In the Spence model, this refinement eliminates all pooling equilibria and all separating equilibria except the least-cost one ($e^s = 1$: the *Riley outcome*).

Key Idea: Costly signals must be differentially costly

Burning resources credibly communicates only what your *willingness to burn* reveals. Education requirements, extravagant advertising, warranties, engagement rings, lavish courtship displays (the peacock's tail—Zahavi's handicap principle), IPO underpricing: all are equilibria in which the cost of the signal is the message. Note the social ambivalence: in pure-signaling models, the signaling expenditure is *wasteful*—the private return to education exceeds its social return.

11.3 Cheap talk

What if talk is free—no signal cost at all? Credible communication is still possible, but only partially, and only when interests align enough.

Example 11.2: Crawford–Sobel (1982) uniform-quadratic model

A sender observes a state $\theta \sim U[0, 1]$ and sends a costless message to a receiver, who then chooses action $a \in \mathbb{R}$. Payoffs: receiver loses $(a - \theta)^2$; sender loses $(a - \theta - b)^2$, where $b > 0$ measures the conflict of interest (the sender wants action biased upward).

Full revelation is never an equilibrium for $b > 0$ (the sender would inflate). All equilibria are *partitional*: the sender reveals only which interval $[\theta_{k-1}, \theta_k]$ the state lies in. Indifference at the cut points forces interval lengths to grow by $4b$ each step, so the maximal number of intervals is the largest K with $2K(K-1)b < 1$. As $b \to 0$, $K \to \infty$ (communication becomes perfect); for $b \geq 1/4$, only *babbling* (uninformative talk) survives.

Remark: Babbling is always an equilibrium

In any cheap-talk game there is an equilibrium in which messages are ignored and hence carry no information. Cheap talk can create meaning only as one equilibrium among several—language itself is an equilibrium convention.

11.4 Screening: the uninformed side designs the menu

Reverse the timing: the *uninformed* party moves first, offering a menu of contracts to induce types to self-select. (This is a bridge to mechanism design, Chapter 14.)

Example 11.3: Price discrimination via menus

A seller with cost zero faces a buyer whose taste is $\theta_H = 2$ (probability μ) or $\theta_L = 1$; a buyer of type θ enjoying quality q and paying t gets $\theta q - t$; quality costs the seller $q^2/2$. The seller offers a menu $\{(q_L, t_L), (q_H, t_H)\}$ maximizing profit subject to *individual rationality* (each type accepts) and *incentive compatibility* (each type picks its own contract). The solution:

$$q_H = \theta_H \text{ (efficient—“no distortion at the top”)}, \qquad q_L = \theta_L - \frac{\mu}{1-\mu}(\theta_H - \theta_L) \;<\; \theta_L,$$

with the low type’s quality distorted *downward* (possibly excluded entirely) and the high type earning an *information rent*. Every airline (economy vs. business), software vendor (feature tiers), and insurer (deductible menus) solves this problem; the deliberately degraded low-end product is not a design failure but an incentive device.

11.5 Reputation and dynamic signaling (a synthesis)

The chain-store analysis of Example 9.2 can now be made precise: in each period the incumbent’s fighting decision is a *signal* about its type; entrants update by Bayes’ rule; and the rational incumbent randomizes between mimicking the tough type and revealing itself. Key general results:

- **Reputation bounds** (Fudenberg–Levine 1989): a long-run player facing a sequence of short-run opponents, with any positive prior on a “commitment type,” secures nearly the payoff of committing to that type’s strategy. Small doubts have large, permanent effects.
- **Reputation dynamics**: incentives to invest in reputation are strongest when uncertainty about type is greatest; established reputations may be milked (temporary quality cuts) and rebuilt—matching observed brand dynamics.

Exercises

Exercise 11.1: Beer–Quiche and the intuitive criterion

A defendant is strong (prob. 0.9) or weak; he chooses breakfast: Beer or Quiche. Strong types prefer Beer, weak prefer Quiche (worth 1 to eat your preferred breakfast); avoiding a duel is worth 2 to either type. An observer sees breakfast, then duels iff he believes the defendant is more likely weak. Find both pooling PBEs (all-Beer and all-Quiche), then show the all-Quiche equilibrium fails the intuitive criterion.

Exercise 11.2: Separating wages

In the Spence model with $\theta_L = 1$, $\theta_H = 3$ and education cost e^2/θ, find the set of separating education levels e^s and the Riley outcome. Then let productivity itself rise with education (wage $\theta + \frac{1}{2}e$ for a worker of known type θ with education e) and recompute the least-cost separating level. Is education still (partly) wasteful?

Exercise 11.3: Two-interval cheap talk

In the Crawford–Sobel model with $b = 1/10$, compute the maximal number of intervals K, the equilibrium cutpoints for the 2-interval equilibrium, and the receiver's expected loss under babbling vs. 2-interval communication.

Exercise 11.4: Insurance screening

High-risk and low-risk drivers have accident probabilities $p_H > p_L$, common wealth w, loss L, and identical concave utility. Show that in a competitive (Rothschild–Stiglitz) screening equilibrium—if it exists—the high-risk type gets full insurance at its fair premium and the low-risk type gets *partial* coverage, and explain why a pooling contract cannot survive competition.

Chapter 12

Zero-Sum Games and the Minimax Theorem

Pure conflict—what one wins, the other loses—is where game theory began, and where its theory is most complete. Zero-sum games also connect game theory to optimization (linear programming) and to machine learning (adversarial training, GANs), making this chapter a bridge between Parts II and III.

12.1 Maxmin reasoning

Definition 12.1: Zero-sum game; value

A two-player game is *zero-sum* if $u_2 = -u_1$; write $u = u_1$ and think of player 1 (the *maximizer*) and player 2 (the *minimizer*). Player 1's *security level* (maxmin) and player 2's (minmax) are

$$\underline{v} = \max_{\sigma_1} \min_{\sigma_2} u(\sigma_1, \sigma_2), \qquad \overline{v} = \min_{\sigma_2} \max_{\sigma_1} u(\sigma_1, \sigma_2).$$

Always $\underline{v} \le \overline{v}$. When they coincide, the common number is the *value* of the game, and the optimizers are the players' *optimal strategies*.

In pure strategies the gap can be strict (Matching Pennies: $\underline{v} = -1 < 1 = \overline{v}$). Von Neumann's theorem says randomization closes it:

Theorem 12.1: Minimax Theorem (von Neumann 1928)

Every finite two-player zero-sum game has a value in mixed strategies:

$$\max_{\sigma_1 \in \Delta(A_1)} \min_{\sigma_2 \in \Delta(A_2)} u(\sigma_1, \sigma_2) \;=\; \min_{\sigma_2 \in \Delta(A_2)} \max_{\sigma_1 \in \Delta(A_1)} u(\sigma_1, \sigma_2) \;=\; v.$$

Moreover (σ_1^*, σ_2^*) is a Nash equilibrium if and only if each σ_i^* is an optimal (security-maximizing) strategy; all equilibria yield the same payoff v; and equilibrium strategies are *interchangeable*.

Key Idea: Why zero-sum games are "solved"

In zero-sum games the troubles of general games vanish: equilibrium payoffs are unique (no selection problem), equilibrium strategies are interchangeable (no coordination problem),

and playing your maxmin strategy is unconditionally safe (it guarantees v regardless of the opponent). Game-theoretic advice is unambiguous: play optimally and you cannot be exploited. This is precisely the property modern superhuman poker programs exploit.

Example 12.1: Solving a 2×2 game

u	L	R
T	3	-1
B	0	2

(Entries are payments from 2 to 1.) No saddle point in pure strategies ($\underline{v}_{\text{pure}} = 0$, $\overline{v}_{\text{pure}} = 2$). Mix: player 1 plays T with p making 2 indifferent: $3p = -p + 2(1-p) \Rightarrow p = \frac{1}{3}$; player 2 plays L with q making 1 indifferent: $3q - (1-q) = 2(1-q) \Rightarrow q = \frac{1}{2}$. Value $v = 3 \cdot \frac{1}{3} \cdot \frac{1}{2} \cdot 2 = 1$ (compute: $u = 3pq - p(1-q) + 2(1-p)(1-q) = 1$).

12.2 Minimax and linear programming duality

The Minimax Theorem is linear programming duality in disguise. Player 1 seeks a mixture $x \in \Delta(A_1)$ maximizing the worst-case column payoff:

$$\text{(P)} \quad \max_{x, v} \; v \quad \text{s.t.} \quad \sum_i x_i \, u(i,j) \geq v \;\; \forall j, \quad \textstyle\sum_i x_i = 1, \; x \geq 0,$$

and player 2's problem is exactly the LP dual; strong duality delivers $\underline{v} = \overline{v}$. Conversely, LP duality can be derived from the Minimax Theorem: the two results are equivalent. Historically, Dantzig and von Neumann discovered this connection in a famous 1947 meeting; computationally, it means zero-sum games are solvable in *polynomial time*—in sharp contrast to general-sum games (Chapter 17).

12.3 Solving games by fictitious play and no-regret learning

Optimal strategies can be *learned* by iterative self-play, an idea with modern resonance.

Definition 12.2: Fictitious play (Brown 1951)

Each round, every player best-responds to the *empirical frequency* of the opponent's past actions.

Theorem 12.2: Robinson (1951)

In finite two-player zero-sum games, the empirical frequencies of fictitious play converge to the set of optimal (maxmin) strategies, and the average payoff converges to the value.

Modern algorithms strengthen this: if both players use any *no-regret* learning rule (e.g., multiplicative weights), time-averaged play converges to optimal strategies at rate $O(1/\sqrt{T})$—the engine behind poker solvers (counterfactual regret minimization) and adversarial training of neural networks (Chapter 17).

12.4 Applications of pure conflict

Example 12.2: Colonel Blotto

Two colonels allocate armies across k battlefields; each field is won by the larger force, and the player winning more fields wins. Blotto games model resource allocation in elections (campaign spending across swing states), R&D races, cybersecurity, and sports scheduling. Optimal strategies are highly random even though nothing in the environment is random—unpredictability across fronts is the essence of the solution. (Closed forms are known only for special cases; the continuous symmetric case involves uniform marginals on $[0, 2/k]$.)

Example 12.3: Inspection games

A tax authority audits with probability p (cost c); a taxpayer evades or complies. Evasion caught costs a fine F; uncaught evasion saves tax T. The equilibrium is mixed: the authority audits just often enough to make the taxpayer indifferent ($p^* = T/F$), and evasion occurs just often enough to justify the audit budget. Comparative static worth remembering: raising the fine F reduces *auditing* in equilibrium, not evasion incentives at the margin—the indifference logic of Chapter 5 at work in law enforcement design.

Historical Note: Poker, solved

Heads-up limit Texas hold'em was *essentially solved* in 2015 (Bowling et al., *Science*): a strategy computed by counterfactual regret minimization is so close to the game's maxmin optimum that a human lifetime of play cannot distinguish it from perfection. No-limit variants (Libratus 2017, Pluribus 2019) followed via approximation. The underlying mathematics is exactly this chapter's: zero-sum extensive-form games, behavior strategies, and convergence of no-regret dynamics to minimax equilibrium.

Exercises

Exercise 12.1: Saddle points

u	L	M	R
T	4	2	3
C	1	0	5
B	3	1	2

Find $\underline{v}$ and $\overline{v}$ in pure strategies. Does the game have a saddle point? If so, identify all of them; if not, solve for the mixed-strategy value.

Exercise 12.2: Guess-the-larger-number

Each player secretly writes 1, 2, or 3. If the numbers are equal, the payoff is 0. If they differ by 1, the player with the *larger* number pays 2 to the smaller. If they differ by 2, the player with the larger number receives 1 from the smaller. Write the payoff matrix, verify the game is symmetric with value 0, and find an optimal strategy.

Exercise 12.3: LP formulation

Write down explicitly the linear program (P) for Matching Pennies and its dual, solve both, and confirm the value 0 and optimal strategies $(\frac{1}{2}, \frac{1}{2})$.

Exercise 12.4: Fictitious play by hand

Run six rounds of fictitious play on Matching Pennies starting from both players playing H in round 1 (break best-response ties toward H). Record empirical frequencies and observe the cycling-but-converging pattern.

Part III

Advanced: Theory, Design, and Computation

Chapter 13

Existence of Equilibrium and Refinements

This chapter supplies the rigorous core of non-cooperative theory: why equilibria exist, and how to discipline the multiplicity of equilibria in dynamic games.

13.1 Fixed-point theorems

Theorem 13.1: Brouwer's fixed-point theorem

Every continuous function $f : K \to K$ from a nonempty compact convex set $K \subset \mathbb{R}^m$ to itself has a fixed point: $f(x^*) = x^*$.

Theorem 13.2: Kakutani's fixed-point theorem (1941)

Let $K \subset \mathbb{R}^m$ be nonempty, compact, convex, and let $F : K \rightrightarrows K$ be a correspondence such that (i) $F(x)$ is nonempty and convex for every x, and (ii) F has a closed graph. Then F has a fixed point: $x^* \in F(x^*)$.

13.2 Nash's theorem

Theorem 13.3: Nash (1950)

Every finite game $G = \langle N, (A_i), (u_i) \rangle$ has at least one Nash equilibrium in mixed strategies.

Proof. Let $\Sigma = \prod_i \Delta(A_i)$, a nonempty compact convex subset of a Euclidean space. Define the best-response correspondence $B(\sigma) = \prod_i \mathrm{BR}_i(\sigma_{-i})$, where $\mathrm{BR}_i(\sigma_{-i}) = \arg\max_{\tau_i \in \Delta(A_i)} u_i(\tau_i, \sigma_{-i})$. Each $\mathrm{BR}_i(\sigma_{-i})$ is nonempty (a continuous function is maximized on a compact set) and convex (the maximizers of a linear function on a simplex form a face, which is convex). The graph of B is closed because u_i is continuous in σ (Berge's maximum theorem, or a direct limit argument). By Kakutani's theorem, B has a fixed point σ^*, which by construction satisfies $\sigma_i^* \in \mathrm{BR}_i(\sigma_{-i}^*)$ for all i: a Nash equilibrium. □

Remark: Nash's second proof via Brouwer

Nash (1951) gave a self-contained proof using Brouwer directly: define the "gain" $g_{ia}(\sigma) = \max\{0,\, u_i(a, \sigma_{-i}) - u_i(\sigma)\}$ and the continuous map

$$T_{ia}(\sigma) = \frac{\sigma_i(a) + g_{ia}(\sigma)}{1 + \sum_{a'} g_{ia'}(\sigma)}.$$

Fixed points of T are exactly the Nash equilibria (if σ is not an equilibrium, probability flows toward better replies and σ moves). This map is also the starting point for computational approaches—and for understanding why computing the fixed point is hard (Chapter 17).

Theorem 13.4: Games with continuous action spaces (Debreu–Glicksberg–Fan 1952)

If each A_i is a nonempty compact convex subset of $\mathbb{R}^m$, and each u_i is continuous in a and *quasi-concave* in own action a_i, then a *pure-strategy* Nash equilibrium exists.

This covers Cournot, Bertrand with differentiated goods, public-good games, etc. Glicksberg (1952) further extends Nash's theorem to continuous games (compact metric action spaces, continuous payoffs) with mixed strategies. Discontinuous games (all-pay auctions, Bertrand) need more delicate arguments (Reny's better-reply security, 1999).

13.3 Trembling-hand perfection

Nash equilibrium can rely on weakly dominated actions and incredible behavior at unreached information sets. Selten's idea: robustness to small mistakes.

Definition 13.1: Trembling-hand perfect equilibrium (Selten 1975)

A mixed profile σ^* of a finite normal-form game is *(trembling-hand) perfect* if there is a sequence of *totally mixed* profiles $\sigma^k \to \sigma^*$ (every action played with positive probability) such that each σ_i^* is a best response to σ_{-i}^k for every k.

Proposition 13.5: Properties of perfection

Every finite game has a perfect equilibrium. Every perfect equilibrium is Nash, and assigns zero probability to weakly dominated actions. In two-player games, an equilibrium is perfect iff both strategies are undominated.

Example 13.1: Ruling out dominated equilibria

	L	R
T	$1, 1$	$0, 0$
B	$0, 0$	$0, 0$

Both (T, L) and (B, R) are Nash equilibria, but (B, R) uses weakly dominated actions on both sides: any tremble by the opponent makes T (resp. L) strictly better. Only (T, L) is perfect.

13.4 Sequential equilibrium

For extensive-form games, the workhorse refinement pairs strategies with beliefs, like PBE, but disciplines off-path beliefs by a limiting consistency requirement.

Definition 13.2: Sequential equilibrium (Kreps–Wilson 1982)

An *assessment* (σ, μ)—a behavior strategy profile plus a belief system —is a *sequential equilibrium* if:

1. **Sequential rationality**: at every information set, the mover's continuation strategy is optimal given μ and σ;
2. **Consistency**: $(\sigma, \mu) = \lim_k (\sigma^k, \mu^k)$ for some sequence of totally mixed σ^k with μ^k derived from σ^k by Bayes' rule everywhere.

Consistency says off-path beliefs must be explicable as limits of small correlated trembles—ruling out arbitrary belief threats that plain PBE permits. Every finite extensive-form game with perfect recall has a sequential equilibrium; every sequential equilibrium is subgame perfect; trembling-hand perfection in the *agent normal form* implies sequentiality.

Remark: The refinements hierarchy

For finite extensive-form games with perfect recall:

$$\text{THP (agent form)} \subseteq \text{sequential} \subseteq \text{PBE} \subseteq \text{subgame perfect} \subseteq \text{Nash}.$$

Further refinements—properness (Myerson), stability (Kohlberg–Mertens), forward induction, the intuitive criterion (Theorem 11.1)—address signaling and forward-looking inferences. The refinements program of the 1980s never produced a single universally accepted concept, but its ideas (credibility, consistency of beliefs, forward induction) permanently shaped applied modeling.

13.5 Epistemic foundations

What assumptions on knowledge and belief justify each solution concept? A sample of landmark results from *epistemic game theory*:

- **Rationalizability**: common knowledge of rationality characterizes (correlated) rationalizability, i.e. IESDS survivors (Theorem 3.2).
- **Nash equilibrium** (Aumann–Brandenburger 1995): in two-player games, mutual knowledge of rationality and of the players' *conjectures* about each other implies the conjectures form a Nash equilibrium—common knowledge is not needed, but knowing the opponent's belief is a strong requirement.
- **Correlated equilibrium** (Aumann 1987): if players share a common prior and are "Bayes rational" given their information, their behavior is a correlated equilibrium. This gives correlated equilibrium a cleaner epistemic pedigree than Nash.
- **Agreeing to disagree** (Aumann 1976): with a common prior, if two players' posteriors about an event are common knowledge, they must be *equal*—rational agents cannot agree to disagree. Corollary (no-trade theorems, Milgrom–Stokey 1982): rational risk-

averse traders with a common prior cannot trade on differences of information alone—if someone wants to bet with you, that itself is information (compare the winner's curse, Example 10.5).

Exercises

Exercise 13.1: Existence via Kakutani, hands-on

Consider a two-player game where each player picks $x_i \in [0,1]$ and $u_i = -(x_i - \frac{1}{2}x_j)^2$. Compute best responses, verify the conditions of Theorem 13.4, and find the unique pure equilibrium.

Exercise 13.2: Perfection in the normal form

	L	C	R
T	$2,2$	$1,1$	$0,0$
M	$1,1$	$1,1$	$1,1$
B	$0,0$	$1,1$	$2,2$

Find all pure Nash equilibria and determine which are trembling-hand perfect.

Exercise 13.3: Sequential vs. PBE

Construct (or recall from the literature) a three-player extensive-form game with a PBE whose off-path beliefs cannot arise as limits of trembles, so the assessment is not sequential. Explain exactly where consistency bites.

Exercise 13.4: No-trade

Two risk-neutral agents with a common prior each receive a private signal about an asset's value and are offered the chance to bet against each other; betting is voluntary. Prove that in any equilibrium, the probability of a bet actually occurring is zero. Where exactly does the common prior enter, and what happens without it?

Chapter 14

Mechanism Design and Auction Theory

Mechanism design inverts game theory: instead of analyzing a given game, *design* the game so that self-interested equilibrium play produces the outcome you want. It is the theory behind spectrum auctions, ad exchanges, school choice, kidney exchange, and much of market design—"economic engineering."

14.1 The setting and the Revelation Principle

An *environment*: agents $i \in N$ with private types $\theta_i \in \Theta_i$, an outcome set X, and quasilinear utilities $u_i = v_i(x, \theta_i) + t_i$ (outcome x plus a monetary transfer t_i). A *social choice function* (SCF) $f : \Theta \to X$ specifies the designer's goal (e.g., allocate the object to whoever values it most). A *mechanism* $(M_1, \dots, M_n, g)$ endows each agent with a message space M_i and maps messages to outcomes; together with types it defines a Bayesian game. The mechanism *implements* f if equilibrium play yields $f(\theta)$ for every type profile.

Theorem 14.1: Revelation Principle (Gibbard 1973; Myerson 1979)

If some mechanism implements f in dominant strategies (resp. Bayesian Nash equilibrium), then f is implementable by a *direct* mechanism in which each agent simply reports her type and *truthful reporting* is a dominant strategy (resp. a Bayesian equilibrium).

Proof idea. Build the direct mechanism that asks for θ_i and then plays each agent's equilibrium strategy of the original mechanism on her behalf. If lying profitably in the direct mechanism were possible, the corresponding deviation would have been profitable in the original mechanism—contradiction. □

Key Idea: Search space reduction

The Revelation Principle collapses the unimaginably large space of all possible institutions to the tractable family of *truthful direct mechanisms*, characterized by incentive-compatibility (IC) and participation (IR) constraints. Virtually every result in this chapter is proved in this reduced space. (Caveats: it presumes full commitment by the designer and breaks down under limited commitment, renegotiation, or collusion.)

14.2 Dominant-strategy implementation: VCG

Definition 14.1: Vickrey–Clarke–Groves mechanism

For allocation problems where the designer chooses $x \in X$ to maximize total value: each agent reports $\hat{v}_i(\cdot)$; the mechanism selects $x^* = \arg\max_x \sum_i \hat{v}_i(x)$ and charges agent i the *externality* she imposes on others:

$$t_i \;=\; \underbrace{\max_x \sum_{j\neq i} \hat{v}_j(x)}_{\text{others' best without } i} \;-\; \underbrace{\sum_{j\neq i} \hat{v}_j(x^*)}_{\text{others' value at } x^*} .$$

Theorem 14.2: VCG properties

In the VCG mechanism, truthful reporting is a dominant strategy for every agent, the chosen allocation maximizes true total surplus, payments are nonnegative (in standard allocation settings), and losing a single-object auction costs nothing. For a single object, VCG *is* the second-price auction.

Proof of dominant-strategy truthfulness. Agent i's payoff given reports is $v_i(x^*) + \sum_{j\neq i} \hat{v}_j(x^*) - \max_x \sum_{j\neq i} \hat{v}_j(x)$. The last term is outside her control; the first two are maximized when x^* maximizes $v_i(x) + \sum_{j\neq i} \hat{v}_j(x)$—which is exactly what the mechanism does when she reports truthfully. Misreporting can only move x^* away from her true optimum. □

VCG aligns each agent's objective with society's by making her *internalize her externality*. Its practical weaknesses are as instructive as its elegance: revenue can be very low or non-monotone, it is vulnerable to collusion and false-name bidding, and computing x^* in combinatorial settings is NP-hard—one reason real combinatorial auctions (spectrum) use carefully engineered variants.

Theorem 14.3: Green–Laffont / Holmström

In rich (smoothly connected) type spaces, any efficient, dominant-strategy incentive-compatible mechanism with quasilinear utilities is a VCG (Groves) mechanism. Moreover, no such mechanism is also exactly budget-balanced: efficiency, dominant strategies, and budget balance are jointly impossible in general.

14.3 Bayesian mechanism design: Myerson's optimal auction

Now target *revenue*. One object, bidders with independent private values $v_i \sim F_i$ (density f_i), and the designer may use any selling procedure whatsoever.

Definition 14.2: Virtual value

Bidder i's *virtual valuation* is

$$\psi_i(v_i) \;=\; v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}.$$

The environment is *regular* if each ψ_i is increasing.

Theorem 14.4: Myerson (1981)

In the regular IPV setting, the revenue-maximizing mechanism awards the object to the bidder with the *highest nonnegative virtual value* (to no one if all are negative), with payments pinned down by incentive compatibility. Consequences:

1. **Symmetric bidders**: the optimum is a second-price auction with reserve price r^* solving $\psi(r^*) = 0$—e.g. $r^* = \frac{1}{2}$ for $U[0,1]$ values, *independent of the number of bidders.*
2. **Revenue equivalence** (Theorem 10.1) falls out: expected revenue of any mechanism equals the expectation of the winner's virtual value, so mechanisms with the same allocation rule earn the same revenue.
3. **Asymmetric bidders**: the optimal auction is biased—it sometimes awards the object to a weaker bidder with a lower value (handicapping strong bidders), sacrificing efficiency for revenue.

The key lemma behind all of this (and much of contract theory):

Lemma 14.5: Payoff/envelope characterization (Myerson)

In any IC mechanism, bidder i's interim expected utility satisfies $U_i'(v_i) = Q_i(v_i)$ (her win probability), hence $U_i(v_i) = U_i(\underline{v}) + \int_{\underline{v}}^{v_i} Q_i(s)\,ds$: *the allocation rule determines all payoffs up to a constant.* IC holds iff Q_i is nondecreasing.

Example 14.1: Reserve prices in practice

One $U[0,1]$ bidder: posting price $\frac{1}{2}$ earns $\frac{1}{4}$, versus 0 for any auction without reserve. Two bidders: optimal reserve $\frac{1}{2}$ raises expected revenue from $\frac{1}{3}$ to $\frac{5}{12}$. But note **Bulow–Klemperer (1996)**: one *extra bidder* in a plain second-price auction beats the optimal reserve with the original bidders. Practical moral: *promoting competition usually beats optimizing the fine print.*

14.4 Impossibility: bilateral trade

Theorem 14.6: Myerson–Satterthwaite (1983)

A seller values an object at $c \sim F_S$, a buyer at $v \sim F_B$, with overlapping supports. *No* mechanism is simultaneously (i) Bayesian incentive-compatible, (ii) interim individually rational, (iii) ex-post efficient (trade occurs iff $v > c$), and (iv) budget-balanced.

Private information creates an irreducible friction: some mutually beneficial trades must fail (or a deficit must be subsidized). This impossibility rationalizes observed bargaining failures—strikes, litigation, wars—as consequences of information rents rather than irrationality, and it sharpens the Coase theorem: with private information, *how* property rights and trading institutions are designed genuinely matters.

14.5 Social choice: Arrow and Gibbard–Satterthwaite

Beyond quasilinear settings, mechanism design meets voting.

Theorem 14.7: Arrow's impossibility theorem (1951)

With three or more alternatives, any social welfare function (mapping profiles of individual rankings to a social ranking) satisfying *unanimity* and *independence of irrelevant alternatives* is a *dictatorship*.

Theorem 14.8: Gibbard–Satterthwaite (1973, 1975)

With three or more alternatives, any onto, strategy-proof social choice function on the full domain of preferences is a dictatorship.

Escape routes define entire fields: money and quasilinearity (VCG), restricted domains (single-peaked preferences $\Rightarrow$ median voting is strategy-proof), and matching without transfers:

Example 14.2: Stable matching and deferred acceptance (Gale–Shapley 1962)

Doctors and hospitals rank each other. The *deferred acceptance* algorithm (doctors propose; hospitals tentatively hold their best proposal, rejecting others; repeat) always terminates in a *stable* matching—no doctor–hospital pair prefers each other to their assignments. The doctor-proposing version is optimal for doctors and strategy-proof for them (though not for hospitals; no stable mechanism is strategy-proof for both sides). This mechanism runs the U.S. medical residency match (NRMP), school choice in New York and Boston, and inspired kidney exchange—the work recognized by the 2012 Nobel to Roth and Shapley. (More in Chapter 15.)

Exercises

Exercise 14.1: VCG in a combinatorial auction

Two items $\{A, B\}$; bidder 1 values only the pair: $v_1(AB) = 10$; bidder 2 values $v_2(A) = 6$; bidder 3 values $v_3(B) = 6$. Compute the VCG allocation and payments. Observe the revenue and compare with selling the items separately by second-price auctions. What does this teach about VCG revenue?

Exercise 14.2: Deriving a reserve price

Values are exponential: $F(v) = 1 - e^{-v}$ on $[0, \infty)$. Compute the virtual value, verify regularity, and find the optimal reserve price. Then compute it for the Pareto distribution $F(v) = 1 - v^{-2}$ on $[1, \infty)$ and interpret what goes wrong with revenue maximization for heavy tails.

Exercise 14.3: Myerson–Satterthwaite by hand

Seller cost and buyer value both $U[0, 1]$. Show that the *efficient* trading rule requires expected subsidies, by computing the maximal expected gains-from-trade extractable under IC + IR (the second-best achieves trade iff $v \geq c + \frac{1}{4}$). Compare the first-best and second-best trade regions in a diagram.

Exercise 14.4: Single-peaked voting

Voters have single-peaked preferences over a line and report their peaks. Show that choosing the *median* report is strategy-proof, and that adding fixed "phantom" peaks preserves strategy-proofness (Moulin 1980). Why does Gibbard–Satterthwaite not apply?

Chapter 15

Cooperative Game Theory

Cooperative (coalitional) theory abstracts from individual moves and asks two questions: *which coalitions form*, and *how are the gains divided*? Its solution concepts—the core and the Shapley value—are now everyday tools far beyond economics, from cost allocation to explainable machine learning.

15.1 Coalitional games

Definition 15.1: TU coalitional game

A *transferable-utility (TU) coalitional game* is a pair (N, v) where N is a finite player set and $v : 2^N \to \mathbb{R}$ with $v(\emptyset) = 0$ assigns each *coalition* $S \subseteq N$ its *worth*—the total payoff its members can guarantee by cooperating among themselves. The game is *superadditive* if $v(S \cup T) \geq v(S) + v(T)$ for disjoint S, T. An *allocation* $x \in \mathbb{R}^N$ is *efficient* if $\sum_i x_i = v(N)$.

Example 15.1: A gloves market

Players $L = \{1\}$ owns a left glove; $R = \{2, 3\}$ each own a right glove. A pair is worth \$1: $v(S) = \min(|S \cap L|, |S \cap R|)$. Which divisions of the \$1 generated by a pair are "stable"?

15.2 The core

Definition 15.2: Core

The *core* of (N, v) is the set of efficient allocations no coalition can improve upon:

$$C(v) \;=\; \Big\{ x \in \mathbb{R}^N : \textstyle\sum_{i \in N} x_i = v(N) \text{ and } \sum_{i \in S} x_i \geq v(S) \;\; \forall S \subseteq N \Big\}.$$

Example 15.2: Gloves, continued

In Example 15.1, $v(N) = 1$. Core constraints: $x_1 + x_2 \geq 1$, $x_1 + x_3 \geq 1$, $x_i \geq 0$, $\sum x_i = 1$. Adding the first two: $2x_1 + x_2 + x_3 \geq 2$, so $x_1 \geq 1$: the core is the *single point* $(1, 0, 0)$. The long side of the market gets nothing—competition between the two right-glove owners drives their "price" to zero. The core captures market discipline with brutal clarity.

Example 15.3: An empty core

Three players; any two-player coalition is worth 1; the grand coalition is worth 1; singletons worth 0. Core: $x_i + x_j \geq 1$ for all pairs sums to $2(x_1 + x_2 + x_3) \geq 3 > 2$—impossible. *No* division is stable: any agreement between two players can be undercut by the third. Empty cores model genuinely unstable situations (odd-player majority bargaining, some cost-sharing problems).

Theorem 15.1: Bondareva–Shapley (1963, 1967)

A TU game has a nonempty core if and only if it is *balanced*: for every collection of weights $\lambda_S \in [0,1]$ with $\sum_{S \ni i} \lambda_S = 1$ for each player i, we have $\sum_S \lambda_S v(S) \leq v(N)$. (Proof: LP duality applied to the core's linear program—the Minimax/duality toolkit of Chapter 12 again.)

15.3 The Shapley value

The core may be empty or huge; Shapley (1953) asked instead for a unique, fair one-point solution.

Definition 15.3: Shapley value

The *Shapley value* gives player i her average marginal contribution over all $|N|!$ orders of arrival:

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

Theorem 15.2: Shapley's axiomatization (1953)

The Shapley value is the unique allocation rule satisfying: **efficiency** ($\sum_i \phi_i = v(N)$); **symmetry** (interchangeable players get the same); **null player** (a player adding nothing to any coalition gets 0); and **additivity** ($\phi(v + w) = \phi(v) + \phi(w)$).

Example 15.4: Airport cost game

Three airlines need runways of length small ($c = 1$), medium ($c = 2$), large ($c = 3$); one runway serving the largest need serves all. Fair division of the \$3 cost? Model the *cost* game and compute Shapley: the first kilometer is shared by all three ($\frac{1}{3}$ each), the second by the two larger airlines ($\frac{1}{2}$ each), the third only by the largest. Shapley cost shares: $(\frac{1}{3},\ \frac{1}{3} + \frac{1}{2},\ \frac{1}{3} + \frac{1}{2} + 1) = (0.33,\ 0.83,\ 1.83)$. This "divide each cost increment among its users" logic is used for actual airport landing fees, water infrastructure, and network cost sharing.

Example 15.5: Shapley values in machine learning (SHAP)

To explain a model's prediction, treat features as players and the prediction (relative to baseline) as $v(S)$ for each feature subset S; each feature's Shapley value is its fair credit for the prediction. The axioms translate to desirable explanation properties (local accuracy, missingness, consistency)—which is why SHAP became the standard attribution method.

Exact computation is exponential in the number of features; practical algorithms sample permutations—the same computational tension as everywhere in game theory (Chapter 17).

Remark: Core vs. Shapley; power indices

The Shapley value need not lie in the core (it does for *convex* games, where it is the core's center of gravity). For voting games ($v(S) \in \{0,1\}$: winning or losing coalitions), the Shapley value becomes the *Shapley–Shubik power index*; a close cousin is the Banzhaf index. Applied to the UN Security Council, the five permanent members hold over 98% of the power by Shapley–Shubik—formalizing the intuition that vetoes, not vote counts, are what matter.

15.4 The nucleolus and other solutions

The *nucleolus* (Schmeidler 1969) picks the allocation that lexicographically minimizes the largest coalition dissatisfaction $e(S,x) = v(S) - \sum_{i\in S} x_i$. It always exists, is unique, and lies in the core whenever the core is nonempty. Its pedigree is ancient: the nucleolus reproduces the disputed-garment divisions of the Babylonian Talmud (Aumann–Maschler 1985)—a 2000-year-old text anticipating a 1969 solution concept.

15.5 Matching markets (NTU cooperation)

Without transferable utility, stability remains the guiding idea. Recall deferred acceptance (Example 14.2); the coalitional view adds:

Theorem 15.3: Structure of stable matchings

In the marriage market, the set of stable matchings is a nonempty *lattice*: there is a man-optimal stable matching (simultaneously best for every man and worst for every woman among stable matchings) and dually a woman-optimal one. The core of the matching game equals the set of stable matchings.

Example 15.6: Kidney exchange

Patient–donor pairs that are immunologically incompatible can swap donors. With no money allowed (illegal almost everywhere), this is a pure matching-market design problem: top trading cycles and chain algorithms (Roth–Sönmez–Ünver) now run national kidney exchange programs, saving thousands of lives—cooperative game theory at its most concrete.

Exercises

Exercise 15.1: Core computation

$N = \{1,2,3\}$; $v(\{i\}) = 0$; $v(\{1,2\}) = 90$, $v(\{1,3\}) = 80$, $v(\{2,3\}) = 70$; $v(N) = 135$. Find the core (describe it as a polytope and give a concrete member), and compute the Shapley value. Is the Shapley value in the core?

Exercise 15.2: Production economy

A landowner (L) and w identical workers produce only together: $v(S) = 0$ if $L \notin S$, and $v(S) = g(k)$ if $L \in S$ with k workers, where g is increasing concave. Show core allocations pay each worker at most the marginal product $g(w) - g(w-1)$, and compute the Shapley value for $g(k) = \sqrt{k}$, $w = 3$.

Exercise 15.3: Shapley–Shubik index

A committee has voters with weights $(4, 3, 2, 1)$ and quota 6. List the pivotal positions and compute the Shapley–Shubik power index of each voter. Compare power with weight; comment.

Exercise 15.4: Instability in roommates

Show by a four-player example that the *roommate problem* (one-sided matching: everyone can be matched with anyone) can have an empty set of stable matchings, in contrast to the two-sided marriage market. Explain what structural feature of two-sided markets the example destroys.

Chapter 16

Evolutionary Game Theory

Maynard Smith's export of game theory to biology replaced rationality with *selection*: strategies that earn higher payoffs (fitness) reproduce faster. The resulting theory illuminates animal conflict, the evolution of cooperation, social norms—and, read as a model of imitation and learning, the behavior of decidedly non-Darwinian agents like firms and traders.

16.1 Evolutionarily stable strategies

Consider a large population whose members are randomly matched in pairs to play a symmetric two-player game with payoff $u(a, b)$ to a player using a against b.

Definition 16.1: Evolutionarily stable strategy (Maynard Smith–Price 1973)

A (possibly mixed) strategy σ^* is an *evolutionarily stable strategy* (ESS) if for every mutant strategy $\sigma \neq \sigma^*$ there is an invasion barrier $\bar{\varepsilon} > 0$ such that for all $\varepsilon \in (0, \bar{\varepsilon})$:

$$u(\sigma^*,\, (1-\varepsilon)\sigma^* + \varepsilon\sigma) \;>\; u(\sigma,\, (1-\varepsilon)\sigma^* + \varepsilon\sigma).$$

Equivalently: (i) (σ^*, σ^*) is a symmetric Nash equilibrium, and (ii) if $u(\sigma, \sigma^*) = u(\sigma^*, \sigma^*)$ for a mutant $\sigma \neq \sigma^*$, then $u(\sigma^*, \sigma) > u(\sigma, \sigma)$.

Key Idea: Stability without rationality

ESS refines symmetric Nash equilibrium using no rationality at all: condition (i) says no mutant does strictly better against the incumbent population; condition (ii) says that neutral mutants (equally good against incumbents) must do strictly worse against *themselves*, so they cannot drift in. Equilibrium emerges from differential survival, not deliberation—answering the oldest objection to Nash equilibrium ("real agents don't compute fixed points").

Example 16.1: Hawk–Dove revisited

In Example 5.3 with $C > V$, the mixed strategy $\sigma^* =$ Hawk with probability V/C is the unique ESS: a population of pure Doves is invaded by Hawks (condition (i) fails), pure Hawks by Doves, and the mixed population resists all mutants. Prediction: ritualized limited conflict in nature, with escalation frequency governed by the value/cost ratio—one of theoretical biology's most celebrated explanations.

Example 16.2: ESS may fail to exist

Rock–Paper–Scissors has a unique symmetric equilibrium (uniform mixing), but it is *not* an ESS: any pure mutant earns exactly the equilibrium payoff against the incumbent and ties against itself, so condition (ii) fails with equality. Nature concurs: the side-blotched lizard *Uta stansburiana* has three male morphs (orange beats blue beats yellow beats orange) whose frequencies *cycle* over years instead of settling—exactly the dynamic behavior the missing ESS predicts.

16.2 Replicator dynamics

Definition 16.2: Replicator dynamics (Taylor–Jonker 1978)

Let $x_a(t)$ be the population share playing pure strategy a, and let $u(a, x)$ be its expected payoff against the population mix x. The *replicator equation* is

$$\dot{x}_a \;=\; x_a\big[\,u(a,x) - \bar{u}(x)\,\big], \qquad \bar{u}(x) = \textstyle\sum_b x_b\, u(b,x)\;:$$

strategies grow in proportion to their payoff advantage over the population average.

Theorem 16.1: Folk theorem of evolutionary dynamics

For the replicator dynamics: (i) every Nash equilibrium is a rest point; (ii) every stable rest point is a Nash equilibrium; (iii) every ESS is asymptotically stable (Hofbauer–Schuster–Sigmund); (iv) strictly dominated strategies are eliminated along interior trajectories. The converses generally fail: rest points need not be Nash (unstable vertices), and asymptotically stable states need not be ESS.

Example 16.3: Stag Hunt dynamics and basins of attraction

For the Stag Hunt (Example 3.2), the replicator dynamics on the share x of Stag players has three rest points: $x = 0$, $x = 1$, and the mixed equilibrium $x^* = \frac{1}{2}$ (with those payoffs, indifference at $4x = 2$, i.e. $x^* = 1/2$). Both pure equilibria are stable; the mixed one is the unstable watershed between their basins. Which convention evolves depends on initial conditions—and the *risk-dominant* equilibrium has the larger basin. Stochastic evolutionary models (Kandori–Mailath–Rob 1993; Young 1993) go further: with rare mutations, the population spends almost all its long-run time at the risk-dominant equilibrium. Evolution provides the equilibrium *selection* theory that rationality alone could not (Chapter 4).

16.3 The evolution of cooperation

How can cooperation evolve when defection is dominant in the one-shot Prisoner's Dilemma? The main mechanisms (Nowak's "five rules"), each translatable into a precise payoff condition:

1. **Direct reciprocity**: repeated interaction with the same partner (Chapter 9); cooperation viable when the continuation probability exceeds the cost–benefit ratio c/b.
2. **Indirect reciprocity**: reputation—help those who help others; viable when the probability your action is observed exceeds c/b. ("For direct reciprocity you need a face; for indirect reciprocity, a name.")

3. **Kin selection**: Hamilton's rule $rb > c$ (relatedness r).
4. **Network/spatial reciprocity**: cooperators forming clusters on graphs survive when b/c exceeds the average degree (Ohtsuki–Nowak).
5. **Group selection**: competition between groups can favor within-group cooperators.

Historical Note: From lizards to language

Evolutionary game theory now spans: animal signaling (the handicap principle as costly signaling, cf. Chapter 11); the evolution of fairness norms (ultimatum-game behavior, Example 6.4, as an evolved commitment device); microbial games (bacteriocin-producing *E. coli* strains playing rock–paper–scissors in petri dishes); cancer modeling (tumor subclones as competing strategies, with "evolutionary therapy" designed as a game against the tumor); and the evolution of language conventions (signaling games where meaning itself is the equilibrium, cf. the babbling discussion in Chapter 11).

16.4 Learning in games (the bridge to algorithms)

Replicator dynamics also describes *imitation* in human populations, and its discrete-time cousin is precisely the *multiplicative weights* algorithm: agents increase the probability of strategies in proportion (exponentially) to accumulated payoffs. This is no coincidence—evolution, learning, and optimization dynamics form one mathematical family, taken up in Chapter 17: no-regret learning converges (in time average) to coarse correlated equilibria in general games and to minimax play in zero-sum games, while last-iterate behavior may cycle or be chaotic (as the lizards already demonstrated).

Exercises

Exercise 16.1: ESS verification

For the symmetric game with payoffs $u(A,A) = 2$, $u(A,B) = 5$, $u(B,A) = 1$, $u(B,B) = 3$: find all symmetric Nash equilibria and determine which are ESS.

Exercise 16.2: Replicator phase line

Write the replicator equation for the Hawk–Dove game with $V = 2$, $C = 4$ as a single ODE for the Hawk share x. Find all rest points, determine their stability by the sign of $\dot{x}$, and sketch the phase line.

Exercise 16.3: Dominated strategies die

Prove claim (iv) of Theorem 16.1 for a strategy a strictly dominated by pure strategy b: show $\frac{d}{dt}\log(x_a/x_b) \le -\varepsilon < 0$ along interior replicator trajectories, so $x_a \to 0$.

Exercise 16.4: Invasion barriers

In the Stag Hunt of Example 3.2, a fraction ε of Stag-playing mutants enters a population of Hare players. Compute the payoff of each strategy as a function of ε and find the minimal mutant mass needed for Stag to spread. Interpret: why are "better" conventions hard to

ignite, and what real-world interventions (assurance contracts, coordinated switching dates) lower the barrier?

Chapter 17

Algorithmic Game Theory

Computer science met game theory when the internet became a system of interacting selfish agents—routers, bidders, miners, platforms. Three founding questions organize the field (Nisan–Roughgarden–Tardos–Vazirani): *Can equilibria be computed efficiently? How bad is selfish behavior? How should we design mechanisms under computational constraints?*

17.1 The complexity of computing Nash equilibria

Zero-sum games are solvable in polynomial time via linear programming (Chapter 12). General games appear fundamentally harder.

Definition 17.1: The class PPAD (Papadimitriou 1994)

PPAD ("Polynomial Parity Argument, Directed") is the class of total search problems reducible to END-OF-LINE: given a concisely described directed graph where every node has in- and out-degree at most one, and given one unbalanced node, find another. Existence is guaranteed by a parity argument—just as Nash equilibria are guaranteed by a fixed-point theorem—but no efficient algorithm for actually *finding* the object is known.

Theorem 17.1: Daskalakis–Goldberg–Papadimitriou (2006); Chen–Deng (2006)

Computing a Nash equilibrium of a finite game is PPAD-complete, even for two-player games and even for any polynomial approximation precision. Under the widely believed conjecture PPAD $\neq$ P, there is no polynomial-time algorithm for Nash equilibrium.

Key Idea: A critique of the concept itself

"If your laptop can't find it, neither can the market" (Kamal Jain). The hardness result is not merely about algorithms: it questions Nash equilibrium as a *prediction* for large systems, since no plausible adaptive process can reach it quickly in the worst case. It redirects attention to tractable alternatives—correlated equilibria and the outcomes of learning dynamics—and to structured subclasses (zero-sum, potential games) where computation is easy.

Related results: finding Nash equilibria with specific properties (e.g., maximizing welfare) is NP-hard; the Lemke–Howson pivoting algorithm (the field's simplex-method analogue) has exponential worst cases yet works well in practice.

17.2 Potential games and best-response dynamics

Definition 17.2: Potential game (Monderer–Shapley 1996)

G is an *exact potential game* if there exists $\Phi : A \to \mathbb{R}$ with

$$u_i(a_i', a_{-i}) - u_i(a_i, a_{-i}) \;=\; \Phi(a_i', a_{-i}) - \Phi(a_i, a_{-i})$$

for every player i and all actions: every unilateral deviation changes a single global function by exactly the deviator's gain.

Theorem 17.2: Properties of potential games

Every finite potential game has a pure-strategy Nash equilibrium (any maximizer of Φ), and best-response dynamics converges to one in finite time (each step strictly increases Φ). Congestion games (Rosenthal 1973)—players choose resource subsets, each resource's cost depends on its load—are exactly the finite potential games, with $\Phi(a) = \sum_r \sum_{k=1}^{\ell_r(a)} c_r(k)$.

Traffic routing, network formation, and load balancing are congestion games: this is why decentralized traffic reliably *settles* (an equilibrium selection dynamic that provably works), even though finding the equilibrium a system will settle on can still be hard (PLS-complete).

17.3 The price of anarchy

Definition 17.3: Price of anarchy (Koutsoupias–Papadimitriou 1999)

For a game with social cost function C,

$$\mathrm{PoA} \;=\; \frac{\max_{\text{equilibria } a} C(a)}{\min_{a'} C(a')} \,,$$

the worst-case ratio of equilibrium cost to optimal cost.

Theorem 17.3: Selfish routing (Roughgarden–Tardos 2002)

In nonatomic selfish routing with affine cost functions, the price of anarchy is at most $\frac{4}{3}$—on *every* network topology, with the Pigou example (Example 7.3) attaining the bound. With polynomial costs of degree d, the bound grows as $\Theta(d/\ln d)$.

Remark: Smoothness and robustness

Roughgarden's *smoothness* framework shows that most known PoA bounds automatically extend from Nash equilibria to coarse correlated equilibria—hence to the long-run outcomes of arbitrary no-regret learners. The pessimistic worst case over a hard-to-reach concept becomes a robust guarantee about plausible adaptive behavior: one of the field's deepest ideas.

17.4 Learning dynamics and correlated equilibria

Definition 17.4: External regret; coarse correlated equilibrium

An online algorithm choosing mixed actions σ^t against payoff sequences has *external regret* $R_T = \max_a \sum_{t \le T} u(a, a^t_{-i}) - \sum_{t \le T} u(\sigma^t, a^t_{-i})$; it is *no-regret* if $R_T/T \to 0$ against any opponent. A distribution μ over action profiles is a *coarse correlated equilibrium* (CCE) if no player gains by a constant unilateral deviation from a profile drawn from μ; a *correlated equilibrium* (CE, Aumann 1974) if no player gains by any *swap function* applied to her recommended action. Nash $\subseteq$ CE $\subseteq$ CCE.

Theorem 17.4: Learning converges to equilibrium sets

(i) The multiplicative weights algorithm achieves regret $O(\sqrt{T \log |A|})$. (ii) If all players use no-external-regret algorithms, the empirical distribution of play converges to the set of CCE. (iii) If all use no-*swap*-regret algorithms (Foster–Vohra; Hart–Mas-Colell's regret matching; Blum–Mansour), it converges to the set of *correlated* equilibria. (iv) In two-player zero-sum games, the marginal empirical frequencies converge to minimax-optimal strategies.

Correlated equilibria are also computable in polynomial time by linear programming, even in succinct multiplayer games—the tractable face of equilibrium theory. This cluster of results powers practice: counterfactual regret minimization solved poker (Chapter 12); self-play learning underlies AlphaZero-style systems; and ad platforms analyze bidder learning dynamics rather than assuming instant equilibrium.

17.5 Computational mechanism design and Bitcoin's game theory

- **Combinatorial auctions**: welfare maximization is NP-hard, and VCG with approximate optimization loses truthfulness; designing mechanisms that are simultaneously truthful and computationally feasible is a rich frontier (maximal-in-range mechanisms, posted prices).

- **Simple vs. optimal**: real ad markets abandoned exact optimality for robust simplicity. Landmark bounds: a single second-price auction with an extra bidder beats Myerson (Bulow–Klemperer, Example 14.1); simultaneous first-price auctions for many items achieve a constant fraction of optimal welfare at equilibrium (smoothness again); posted prices with the right thresholds achieve half of optimal (prophet inequalities)—the theory behind "buy-it-now" pricing.

- **Incentives in protocols**: Bitcoin's security is a game-theoretic, not purely cryptographic, claim—honest mining is (approximately) an equilibrium only under bounds on miner size; *selfish mining* (Eyal–Sirer 2014) is profitable for pools above about a quarter of total hash power. Transaction fee mechanisms (EIP-1559), maximal extractable value, and validator incentives are live mechanism-design problems.

Exercises

Exercise 17.1: A potential for Cournot

Show the two-firm Cournot game with linear demand (Example 4.2) is an (ordinal) potential game with $\Phi(q_1, q_2) = q_1 q_2 \cdot (-1) + \dots$ — construct an exact potential explicitly, verify the defining identity, and use it to re-derive the equilibrium.

Exercise 17.2: Price of anarchy in Pigou networks

Generalize Example 7.3: road 2 has cost x^d. Compute equilibrium and optimal flows and show PoA $\to \infty$ as $d \to \infty$, at rate $\Theta(d/\ln d)$.

Exercise 17.3: CCE that is not CE

Exhibit a small game and a distribution over its action profiles that is a coarse correlated equilibrium but not a correlated equilibrium, and identify the profitable swap deviation.

Exercise 17.4: Load balancing

n jobs (players) each choose one of 2 identical machines; a job's cost is the number of jobs on its machine. Show this is a congestion game, compute its exact potential, find all pure equilibria, and bound the price of anarchy for the makespan (maximum load) objective.

Chapter 18

Frontiers: Behavioral, Experimental, and Modern Directions

The final chapter surveys where game theory is being tested, revised, and extended. Each section states the core departure from the classical theory, the leading models, and pointers into the literature.

18.1 Experimental game theory: the evidence

Fifty years of laboratory and field experiments yield robust regularities:

- **Ultimatum/dictator games**: proposers offer 40–50%; low offers are rejected (Example 6.4); even *dictators* (no rejection possible) give around 20%—evidence for fairness preferences.
- **Public goods**: initial contributions of 40–60% decay with repetition; costly *punishment* opportunities restore high cooperation (Example 7.2); anti-social punishment varies across societies.
- **Beauty contests**: winning guesses far above the rationalizable 0 (Example 3.7)—people iterate reasoning only 1–3 steps.
- **Coordination**: subjects find focal points reliably; risk-dominance predicts behavior in Stag Hunts better than payoff-dominance; conventions once established are sticky.
- **Auctions**: overbidding relative to risk-neutral theory in first-price auctions; the winner's curse afflicts inexperienced common-value bidders exactly as Example 10.5 suggests.
- **Mixed strategies**: professionals (soccer, tennis) approximate equilibrium mixing; novices alternate excessively.

The pattern: theory performs best with experienced players, high stakes, competitive market-like settings, and strategic simplicity; it needs behavioral amendment where fairness, cognition limits, or novelty dominate.

18.2 Behavioral game theory: amended models

Social preferences. Fehr–Schmidt inequity aversion: $u_i = x_i - \alpha_i \max(x_j - x_i, 0) - \beta_i \max(x_i - x_j, 0)$; Bolton–Ockenfels ERC; Rabin's intentions-based reciprocity (kindness begets kindness) built on *psychological game theory*, where payoffs depend on beliefs themselves.

Level-k and cognitive hierarchy. Level-0 acts nonstrategically (random or focal); level-k best-responds to level-$(k-1)$; estimated levels concentrate on 1–2, organizing beauty-contest and auction data (Nagel; Camerer–Ho–Chong).

Quantal response equilibrium (McKelvey–Palfrey 1995). Players best-respond *noisily* (logit choice), in equilibrium with each other's noise. QRE rationalizes systematic deviations (e.g. own-payoff effects in Matching Pennies) while nesting Nash as the zero-noise limit.

Limited foresight and complexity. Finite automata playing repeated games, sampling equilibria, and models of costly reasoning connect to the computational themes of Chapter 17.

18.3 Stochastic games and dynamic equilibria

Shapley's *stochastic games* (1953)—state variables evolving with play—underlie modern dynamic oligopoly (Markov perfect equilibrium, Maskin–Tirole), resource economics, and multi-agent reinforcement learning. Deep results: existence of stationary equilibria, the *uniform value* of zero-sum stochastic games (Mertens–Neyman 1981), and the algorithmic theory now central to multi-agent AI.

18.4 Games with a continuum of players and mean-field games

Large anonymous populations justify *nonatomic games* (Schmeidler) and *mean-field games* (Lasry–Lions 2007; Huang–Caines–Malhamé): each agent optimizes against the population *distribution*, whose evolution is generated by the agents' own optimal behavior—a fixed point in distribution space, described by coupled Hamilton–Jacobi–Bellman and Kolmogorov equations. Applications: crowd dynamics, energy markets, systemic risk, and large-scale multi-agent learning.

18.5 Information design and Bayesian persuasion

Mechanism design fixes the game and designs incentives; *information design* fixes incentives and designs *what players learn*.

Example 18.1: Bayesian persuasion (Kamenica–Gentzkow 2011)

A prosecutor (sender) commits to an investigation policy (a signal about guilt) to persuade a judge (receiver) who convicts iff the posterior probability of guilt is at least $\frac{1}{2}$; the prior is 0.3. Full disclosure convicts 30% of defendants. Optimal persuasion: a signal that says "guilty" always when guilty and with probability $\frac{3}{7}$ when innocent—making the posterior upon "guilty" exactly $\frac{1}{2}$. Conviction rate: $0.3 + 0.7 \cdot \frac{3}{7} = 0.6$, double full disclosure. The general solution method—*concavification* of the sender's value over posterior beliefs—has made information design a flourishing field (ratings, stress tests, recommendation systems, traffic apps).

18.6 Game theory and AI

- **Superhuman play**: checkers solved (2007); heads-up limit poker essentially solved (2015); Go, chess, shogi via self-play reinforcement learning (AlphaZero 2017); no-limit poker

(Libratus, Pluribus); Diplomacy with natural-language negotiation (Cicero 2022)—each a milestone in computing (approximate) equilibria of ever richer games.

- **Multi-agent learning**: convergence of independent learners, the emergence of cooperation and communication among agents, and the game dynamics of GAN training (a zero-sum game whose last-iterate behavior can cycle—Chapter 12's theory of averages vs. iterates, live in ML).
- **Mechanism design by/for AI**: automated mechanism design; auction design via deep learning; incentive alignment questions for LLM agents negotiating, bidding, and delegating on behalf of humans—game theory's next large application domain.
- **AI safety as game design**: principal–agent framings of alignment, debate and verification games, and commitment/credibility problems between humans and capable systems.

18.7 A reading roadmap

Consolidate (advanced undergraduate). Osborne, *An Introduction to Game Theory*; Gibbons, *Game Theory for Applied Economists*; Dixit–Nalebuff, *The Art of Strategy* (popular but sound).

Graduate core. Fudenberg–Tirole, *Game Theory*; Osborne–Rubinstein, *A Course in Game Theory* (free online); Myerson, *Game Theory: Analysis of Conflict*; Maschler–Solan–Zamir, *Game Theory* (most comprehensive and mathematical).

Fields. Krishna, *Auction Theory*; Börgers, *An Introduction to Mechanism Design*; Roth–Sotomayor, *Two-Sided Matching*; Mailath–Samuelson, *Repeated Games and Reputations*; Weibull, *Evolutionary Game Theory*; Sandholm, *Population Games and Evolutionary Dynamics*; Nisan–Roughgarden–Tardos–Vazirani (eds.), *Algorithmic Game Theory* (free online); Roughgarden, *Twenty Lectures on Algorithmic Game Theory*; Camerer, *Behavioral Game Theory*; Perea, *Epistemic Game Theory*.

Primary sources worth reading in the original. Nash (1950); Schelling, *The Strategy of Conflict* (1960); Aumann (1974, 1976); Myerson (1981); Rubinstein (1982); Kreps–Wilson (1982); Milgrom–Weber (1982); Kamenica–Gentzkow (2011).

Key Idea: The unity of the subject

Look back at the arc of this book: a handful of ideas—best response, equilibrium, credibility, incentive compatibility, conditioning on others' information, the shadow of the future—generate the entire edifice, from penalty kicks to spectrum auctions to blockchain protocols. Mastering game theory is not memorizing models; it is learning to see every interactive situation through these few lenses. That skill, once acquired, is permanent.

Appendix A

Mathematical Background

This appendix collects the mathematical prerequisites, stated for reference rather than developed from scratch.

A.1 Sets, functions, correspondences

A *correspondence* $F : X \rightrightarrows Y$ assigns to each $x \in X$ a subset $F(x) \subseteq Y$ (e.g., best-response sets). F has a *closed graph* if $x^k \to x$, $y^k \to y$, $y^k \in F(x^k)$ imply $y \in F(x)$. A set $K \subseteq \mathbb{R}^m$ is *compact* iff it is closed and bounded (Heine–Borel), and *convex* if it contains the segment between any two of its points. The *simplex* $\Delta(A) = \{p \in \mathbb{R}^{|A|}_{\geq 0} : \sum_a p_a = 1\}$—the home of mixed strategies—is compact and convex.

A.2 Optimization facts

- **Weierstrass**: a continuous function on a nonempty compact set attains a maximum.
- **First-order condition**: at an interior maximum of a differentiable function, the gradient vanishes; with concavity, the FOC is also sufficient.
- **Berge's maximum theorem**: if $u(x, \theta)$ is continuous and the constraint set $C(\theta)$ is a continuous, compact-valued correspondence, then the value function is continuous and the maximizer correspondence is upper hemicontinuous with closed graph—the workhorse behind continuity of best responses.
- **Envelope theorem**: the derivative of a value function with respect to a parameter equals the partial derivative of the objective at the optimum; integral form (Milgrom–Segal) underlies Lemma 14.5.
- **Jensen's inequality**: for concave u, $\mathbb{E}[u(W)] \leq u(\mathbb{E}[W])$—the definition of risk aversion.

A.3 Probability

Finite probability spaces suffice for most of this book. Key tools: expectation and its linearity; conditional probability and **Bayes' rule**

$$\mathbb{P}(\theta \mid s) \;=\; \frac{\mathbb{P}(s \mid \theta)\,\mathbb{P}(\theta)}{\sum_{\theta'} \mathbb{P}(s \mid \theta')\,\mathbb{P}(\theta')},$$

which drives every updating argument in Chapters 10–11; independence; and for auction theory, *order statistics*: with n i.i.d. draws from F, the maximum has CDF $F(v)^n$, and for $U[0,1]$ the expected k-th highest of n is $\frac{n+1-k}{n+1}$.

A.4 Fixed points

Stated as Theorems 13.1 and 13.2. Intuition for Brouwer in one dimension: a continuous $f : [0,1] \to [0,1]$ must cross the diagonal (intermediate value theorem applied to $f(x) - x$). Kakutani generalizes to correspondences; convex-valuedness cannot be dropped (rotating the circle has no fixed point when values are allowed to be disconnected). Sperner's lemma provides the standard combinatorial proof route and also explains the computational hardness of finding fixed points (Chapter 17).

A.5 Linear programming and duality

The primal–dual pair

$$\max\{c^\top x : Ax \le b,\, x \ge 0\} \qquad \longleftrightarrow \qquad \min\{b^\top y : A^\top y \ge c,\, y \ge 0\}$$

satisfies *weak duality* (any feasible pair has $c^\top x \le b^\top y$) and *strong duality* (optimal values coincide when either problem is feasible and bounded). Complementary slackness links optimal solutions. Applications in this book: the Minimax Theorem (Chapter 12), the Bondareva–Shapley theorem (Chapter 15), and polynomial-time computation of correlated equilibria (Chapter 17).

A.6 Basic dynamical systems

For an ODE $\dot{x} = f(x)$ on the simplex: *rest points* satisfy $f(x^*) = 0$; a rest point is *Lyapunov stable* if trajectories starting nearby stay nearby, and *asymptotically stable* if they also converge to it. In one dimension, stability is read off the sign of f around the rest point; in higher dimensions, from the eigenvalues of the Jacobian or via a Lyapunov function. These suffice for the replicator analyses of Chapter 16.

Appendix B

Solutions to Exercises

Solutions are deliberately compact: they give the key steps and answers, leaving routine algebra to the reader. Attempt each exercise seriously before reading.

Chapter 1

Exercise 1.1. (a) Two players, three actions each, *simultaneous*, complete information. (b) Seller and buyer; seller moves first (set price / describe car), buyer accepts or not; *sequential* with *incomplete* information (quality is the seller's private type)—a signaling/lemons situation (Chapters 10–11). (c) Two players, *sequential*, *perfect* information, finite (with repetition rules); by Zermelo's theorem it is determined (Chapter 6). (d) n bidders move *simultaneously* (sealed bids); *incomplete* information (private valuations): a Bayesian game (Chapter 10).

Exercise 1.2. Any answer with: two parties, mutual cooperation better than mutual defection, but defection individually profitable at each opponent action. Typical sustaining mechanisms: repetition and retaliation (Chapter 9), reputation, contracts, norms with sanctions.

Exercise 1.3. With money-only preferences, backward induction gives the younger sibling accepting any positive amount, so the older offers the smallest possible slice. With the anger term, accepting less than a third yields (slice -10) < 0 for small slices, so rejection is credible; the older sibling optimally offers exactly one third. Preferences, not the game tree, changed—and so did the prediction.

Chapter 2

Exercise 2.1. Start her with z. She pays ε to swap $z \to y$ (since $y \succ z$), ε to swap $y \to x$, then—by the cycle $z \succ x$—ε to swap $x \to z$. She is back at z having paid 3ε; repeat forever. Cyclic preferences are exploitable, hence "irrational."

Exercise 2.2. Expected utility of action a_i under belief μ is linear in payoffs, so under $v_1 = 3u_1 + 7$ every expected payoff is transformed by the same increasing affine map; rankings, hence arg max sets, are unchanged. For the nonlinear part: let action S (safe) pay 0 always and action R pay $+1$ with probability $2/3$ and -2 with probability $1/3$ (probabilities induced by the opponent's mixture). Under u: $\mathbb{E}[R] = \frac{2}{3} - \frac{2}{3} = 0 = \mathbb{E}[S]$ (indifferent). Under u^3:

$\mathbb{E}[R] = \frac{2}{3}(1) + \frac{1}{3}(-8) = -2 < 0 = \mathbb{E}[S]$. The cube preserves the ordinal ranking of sure outcomes but changes the best response.

Exercise 2.3. $\mathbb{E}[W] = 50$. $\mathbb{E}[u] = \frac{1}{2}\sqrt{100} = 5$. Certainty equivalent: $u(CE) = 5 \Rightarrow CE = 25$. Risk premium $= 50 - 25 = 25$.

Exercise 2.4. Write A and B as mixtures involving the common component: with $P =$ (1M for sure) and $Q =$ ($\frac{1}{11}$ chance 0, $\frac{10}{11}$ chance $5M$): $A = 0.11P + 0.89P$ and $B = 0.11Q + 0.89P$, while $C = 0.11P + 0.89\,\delta_0$ and $D = 0.11Q + 0.89\,\delta_0$ ($\delta_0 =$ zero for sure). Independence applied to the common 0.89-components forces $A \succsim B \iff P \succsim Q \iff C \succsim D$. The modal pattern $A \succ B$, $D \succ C$ contradicts this.

Chapter 3

Exercise 3.1. Step 1: for player 1, $T = (2, 1, 4)$ strictly dominates $B = (1, 0, 3)$; delete B. Step 2: in the reduced game, for player 2, R (payoffs 2, 3 vs T, M) strictly dominates C $(1, 2)$; delete C. No further strict domination: the survivors are $\{T, M\} \times \{L, R\}$. Within the reduced game, best-response marking gives two pure Nash equilibria, (M, L) and (T, R) (plus one mixed). Moral: IESDS need not reach a single profile; it brackets the predictions.

Exercise 3.2. Order 1: delete T (weakly dominated by M); then L (weakly dominated by R: $1 = 1$ vs M, $0 < 1$ vs B); outcome payoffs $(2, 1)$. Order 2: delete B (weakly dominated by M); then R (now weakly dominated by L: $1 = 1$ vs M, $0 < 1$ vs T); remaining outcomes pay $(1, 1)$. Different orders, different predictions—why iteration uses *strict* dominance.

Exercise 3.3. $u_i = q_i(12 - q_i - q_j)$. Best response $\mathrm{BR}(q_j) = (12 - q_j)/2 = 6 - q_j/2$. Any $q_i > 6$ is strictly dominated by $q_i = 6$ (marginal profit $12 - 2q_i - q_j < 0$ beyond 6 for all $q_j \geq 0$). Given $q_j \leq 6$, BR ≥ 3: outputs below 3 are dominated. Iterating: the surviving interval $[a_k, b_k]$ satisfies $a_{k+1} = 6 - b_k/2$, $b_{k+1} = 6 - a_k/2$, contracting to the fixed point $q = 6 - q/2 \Rightarrow q = 4$. The game is dominance solvable at the Cournot–Nash output $q_1 = q_2 = 4$.

Exercise 3.4. Claiming 100 is weakly dominated by 99 (if the rival claims 100 you get $99 + 2$ instead of 100; other cases tie or win). Once 100 is removed, 99 is dominated by 98, and so on down to the unique survivor $(2, 2)$. Most people (and experimental subjects) claim 95–100: with a small penalty/bonus parameter, the IESDS logic feels implausibly aggressive—a benchmark case for level-k thinking (Chapter 18).

Chapter 4

Exercise 4.1. Best responses — row: vs X: A; vs Y: B; vs Z: C. Column: vs A: X; vs B: Y; vs C: Z. Pure equilibria: (A, X) with $(3, 2)$, (B, Y) with $(2, 3)$, (C, Z) with $(2, 2)$.

Exercise 4.2. FOC for i: $\frac{1}{2}(1 + \frac{1}{2}e_j) - 2e_i = 0 \Rightarrow e_i = \frac{1}{4} + e_j/8$. Symmetric solution: $e^* = \frac{1}{4} + e^*/8 \Rightarrow e^* = \frac{2}{7}$. Joint surplus maximization: $1 + \frac{1}{2}e_j - 2e_i = 0$ symmetric gives $e^{\text{opt}} = \frac{2}{3} > \frac{2}{7}$. Each partner bears the full cost of effort but captures only half the extra output—classic team free-riding (Example 7.2 in miniature).

Exercise 4.3. At $(\frac{1}{2}, \frac{1}{2})$ each gets $\frac{1}{2}$; deviating to $x \neq \frac{1}{2}$ yields market share $\frac{x+1/2}{2} < \frac{1}{2}$ (for $x < \frac{1}{2}$)—no gain. If the profile has some $x_i \neq \frac{1}{2}$, the rival can move adjacent to x_i on the longer side and capture more than half, so no other profile is an equilibrium. Political reading: with single-peaked voters, two vote-maximizing candidates converge to the median voter's position.

Exercise 4.4. Impossible. If (a_1^*, a_2^*) is Nash, then $u_1(a^*) \geq u_1(a_1, a_2^*)$ for the other row(s): player 1's equilibrium payoff is (weakly) the best *in its column*, so it cannot be the strict minimum of all four cells. Strongest possible statement: the unique equilibrium can be strictly Pareto-dominated by another cell, and can give each player their worst payoff *outside their own deviation set*—the Prisoner's Dilemma achieves both players' second-worst payoffs, which is the worst possible.

Chapter 5

Exercise 5.1. Kicker mixes p on Left to equalize keeper's sides: $0.6p + 0.9(1-p) = 0.95p + 0.6(1-p) \Rightarrow p = \frac{6}{13}$. Keeper mixes q on Left to equalize kicker: $0.6q + 0.95(1-q) = 0.9q + 0.6(1-q) \Rightarrow q = \frac{7}{13}$. Equilibrium scoring probability $= 0.6 + 0.3 \cdot \frac{7}{13} = \frac{9.9}{13} \approx 0.762$.

Exercise 5.2. Standard RPS: by symmetry the value is 0; if any action had probability $> \frac{1}{3}$, its counter would earn strictly positive payoff, contradiction—so uniform is the unique equilibrium. Modified game (Rock beats Scissors for ± 2): indifference conditions $-p+2s = r-s = -2r+p = 0$ with $r+p+s = 1$ give $(r, p, s) = (\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$. *Rock's* probability falls (and Paper's rises): making Rock's win bigger makes opponents avoid Scissors and play Paper, which punishes Rock. You exploit an action's strength by playing its *counter's counter*, not the action itself.

Exercise 5.3. Symmetric equilibrium: calling gives $v - c$; not calling gives $v\,[1-(1-p)^{n-1}]$. Indifference: $(1-p)^{n-1} = c/v$, so $p^* = 1 - (c/v)^{1/(n-1)}$, decreasing in n. Probability nobody calls: $(1-p^*)^n = (c/v)^{n/(n-1)}$, *increasing* in n (exponent falls toward 1 and $c/v < 1$). More bystanders, more diffusion of responsibility, less help—the equilibrium logic of the bystander effect.

Exercise 5.4. Let q = probability of L. Row payoffs: $u(T) = 6q$, $u(M) = 4 - 2q$, $u(B) = 3$. There is no pure equilibrium (marking fails everywhere). All three row payoffs are equal only at $q = \frac{1}{2}$ (all give 3). Column is indifferent when $2t + 4m + 3b = 6t + 2m + 3b \Rightarrow m = 2t$. Hence a *component* of equilibria: $\sigma_2 = (\frac{1}{2}, \frac{1}{2})$ and $\sigma_1 = (t,\ 2t,\ 1-3t)$ for any $t \in [0, \frac{1}{3}]$ (all give payoffs $(3, 3)$). This degenerate game illustrates why "generic" games (Theorem 5.2) have finitely many equilibria but special payoff ties create continua.

Chapter 6

Exercise 6.1. Backward induction: after L, player 2 picks r $(2 > 1)$; after R, ℓ' $(1 > 0)$. Player 1 compares $L \to (1, 2)$ with $R \to (2, 1)$ and plays R: outcome $(2, 1)$. Normal form (player 2's strategies $\ell\ell', \ell r', r\ell', rr'$): pure Nash equilibria are $(R, r\ell')$—the SPE—and (L, rr'), which is sustained by the incredible plan to play r' after R (at that unreached node r' pays $0 < 1$).

Exercise 6.2. After burning, row's payoffs are shifted down by 1. Key steps of iterated weak dominance on the 4×4 normal form: (i) for row, "burn then play F" is dominated (best case 0); (ii) given (i), column infers from observed burning that row will play B, so column plays B after burning; (iii) then "burn, B" guarantees row 2, which dominates "don't burn, F" (at most 1); (iv) column, anticipating (iii), plays B even when row does *not* burn; (v) so row plays "don't burn, B" and gets 3. The unused option to burn selects row's favorite equilibrium—forward induction reasoning in action.

Exercise 6.3. Sequential BoS: row commits to B; column's best response is B; outcome $(3, 1)$ —a *first-mover advantage*, because the game rewards being predictable (coordination). In Matching Pennies the first mover loses for the same reason reversed: pure conflict rewards being *un*predictable, creating a second-mover advantage. The dividing line: does your opponent want to match you (move first) or mismatch you (move second)?

Exercise 6.4. Convention: strict majority; indifferent pirates vote against. With 2 pirates (4 proposes to 5): 5 rejects everything (he then gets all 100 *and* sees 4 overboard), so 4 dies: with 2 left, proposer 4 cannot win. Hence with 3 pirates, 4 votes for *anything* 3 proposes (voting no means the 2-pirate subgame where he dies): 3 proposes $(100, 0, 0)$ and wins with votes $\{3, 4\}$. With 4 pirates, proposer 2 needs 3 of 4 votes: buy the two pirates who get 0 in the 3-pirate continuation—4 and 5—with 1 coin each: $(98, 0, 1, 1)$. With 5 pirates, proposer 1 needs 3 of 5: continuation (from the 4-pirate case) gives $(-, 98, 0, 1, 1)$ to pirates $2, 3, 4, 5$; the two cheapest votes are pirate 3 (offer 1) and pirate 4 or 5 (offer 2): **proposal** $(97, 0, 1, 2, 0)$ (or $(97, 0, 1, 0, 2)$). Seniority plus foresight is worth 97 coins.

Chapter 7

Exercise 7.1. FOCs: $q_i = \frac{a - c_i - q_j}{2}$. Solving: $q_i^* = \frac{a - 2c_i + c_j}{3}$, $u_i^* = (q_i^*)^2$. Lower cost $\Rightarrow$ higher output and profit. Firm 2 is driven to zero when $a - 2c_2 + c_1 \leq 0$, i.e. $c_2 \geq \frac{a + c_1}{2}$.

Exercise 7.2. With the convention that the efficient (low-cost) firm wins ties: equilibrium $p_1 = p_2 = c_2$; firm 1 serves the market earning margin $c_2 - c_1$; firm 2 earns 0. (Without the tie-break convention, no equilibrium exists in continuous prices: firm 1 wants to price "just below" c_2; on a discrete grid it prices one tick below.) Prices are pinned by the *second-lowest* cost—the Bertrand analogue of the second-price auction logic.

Exercise 7.3. FOC for i: $\frac{x_j}{(x_1 + x_2)^2} V = 1$. Symmetry ($x_1 = x_2 = x$): $\frac{x}{4x^2} V = 1 \Rightarrow x^* = V/4$ each. Total rent dissipation $x_1^* + x_2^* = V/2$: half the prize is burned; each contestant nets $V/4$. (With n contestants, dissipation $\to V$: rent-seeking can consume the entire rent.)

Exercise 7.4. Consistent numbers: $D = 100$, both wait $\to$ 150 each; one withdraws $\to$ $(100, 60)$; both withdraw $\to$ 80 each.

	Wait	Withdraw
Wait	$150, 150$	$60, 100$
Withdraw	$100, 60$	$80, 80$

Both (Wait, Wait) and (Withdraw, Withdraw) are strict Nash equilibria: a run is a self-fulfilling coordination failure, not irrationality. Deposit insurance guarantees the waiter $100 \geq$ any

withdrawal payoff, making Wait a (weakly) dominant strategy and deleting the run equilibrium—which is why insurance works largely *without being used*.

Chapter 8

Exercise 8.1. Pareto frontier: $s_2 = 2 - s_1/2$. Maximize $(s_1 - 1)(\frac{3}{2} - \frac{s_1}{2})$: derivative $2 - s_1 = 0 \Rightarrow s^* = (2, 1)$. On $S' = S \cap \{s_1 \leq 3\}$ the point $(2, 1)$ is still feasible and the same maximization applies: unchanged, confirming IIA.

Exercise 8.2. $t = 2$: proposer 1 takes 1 (offers 0). $t = 1$: player 2 must offer player 1 at least $\delta \cdot 1$; keeps $1 - \delta$. $t = 0$: player 1 must offer player 2 $\delta(1 - \delta)$; keeps $x_3 = 1 - \delta + \delta^2$. Generally $x_T = 1 - \delta + \delta^2 - \cdots \to \frac{1}{1+\delta}$, the Rubinstein share.

Exercise 8.3. Stationary conditions: $1 - x = (1 - \pi)b_2 + \pi y$ and $1 - y = (1 - \pi)b_1 + \pi x$. Solving: $x = \frac{1-b_2+\pi b_1}{1+\pi}$ (and symmetrically for y). As $\pi \to 1$: $x \to b_1 + \frac{1-b_1-b_2}{2}$—the Nash bargaining solution with the *breakdown payoffs* as the disagreement point. What matters for threat-point purposes is what happens if bargaining *fails*, not impatience per se.

Exercise 8.4. With a hard deadline at T and no discounting, backward induction gives the period-T proposer the whole pie (the responder accepts anything over 0); at $T - 1$ the responder can secure that by waiting, so the $T - 1$ proposer concedes everything, and so on: the equilibrium division is determined entirely by who proposes *last*, and agreement (in the frictionless model) is immediate at that division. Adding any small per-period cost or risk makes early agreement strict—but the *terms* still reflect deadline power, which is why real negotiations cluster near deadlines: delay is costly, and bargaining power shifts discontinuously at T.

Chapter 9

Exercise 9.1. (a) Grim: $5(1 - \delta) + \delta \leq 2 \Rightarrow \delta \geq \frac{3}{4}$. (b) One-period punishment: the deviation gains 3 today and costs only $(2 - 1)\delta$ next period: $3(1 - \delta) \leq \delta(1 - \delta) \Rightarrow \delta \geq 3$—impossible. With these payoffs a single punishment period can never deter: the punishment must last several periods (roughly k with $3 \leq \sum_{j=1}^{k} \delta^j$). Shorter, milder punishments demand more patience; sometimes no feasible patience suffices.

Exercise 9.2. After 1's unilateral defection, mutual TFT prescribes the alternating path $(C, D), (D, C), (C, D), \ldots$ Check player 2's one-shot deviation at the start of the punishment phase: following TFT (play D) yields the alternation worth $\frac{4}{1+\delta}$; deviating to C restores (C, C) forever, worth 3. TFT requires $\frac{4}{1+\delta} \geq 3 \Rightarrow \delta \leq \frac{1}{3}$. But sustaining cooperation on-path requires $\delta \geq \frac{1}{3}$. Both hold only at the knife-edge $\delta = \frac{1}{3}$: mutual TFT is essentially never subgame perfect, despite being a Nash equilibrium for all $\delta \geq \frac{1}{3}$.

Exercise 9.3. Feasible set: convex hull of $(3, 1)$, $(1, 3)$, $(0, 0)$. Minmax for player 1: player 2 chooses q on B to minimize $\max(3q, 1 - q)$: at $3q = 1 - q$, $q = \frac{1}{4}$, value $\frac{3}{4}$; symmetrically $\underline{v}_2 = \frac{3}{4}$. Sustainable set: feasible payoffs dominating $(\frac{3}{4}, \frac{3}{4})$. The alternating scheme's average payoff $\to (2, 2)$ as $\delta \to 1$: feasible, individually rational, hence sustainable —and indeed the natural "fair" convention.

Exercise 9.4. Rule: "cooperate iff your partner is labeled good; defecting on a good partner makes you bad (forever, or for k periods)." A bad-labeled player is defected on by all future partners, so the one-shot gain g is weighed against a permanent per-period loss, exactly as in grim trigger: cooperation is an equilibrium for δ above a threshold. With noisy labels, permanent exclusion destroys too much surplus by accident; finite punishments, forgiveness, and label-repair (apologies, paying fines, rebuilding scores) become optimal—which is what credit systems and platform reputation mechanisms implement.

Chapter 10

Exercise 10.1. If the rival bids $v_j/2$, bidding b wins with probability $\mathbb{P}(v_j/2 < b) = 2b$ (for $b \leq \frac{1}{2}$); expected profit $(v-b)\,2b$ is maximized at $b = v/2$. Revenue: $\mathbb{E}[\max(v_1, v_2)/2] = \frac{1}{2} \cdot \frac{2}{3} = \frac{1}{3} = \mathbb{E}[\min(v_1, v_2)]$, the second-price revenue—revenue equivalence verified.

Exercise 10.2. Imitating type z against a rival using $b(v) = v^2/2$: profit $= zv - z^2/2$ (win probability z, pay always). FOC: $v - z = 0$ at $z = v$: equilibrium confirmed. Revenue $= \mathbb{E}[b(v_1)] + \mathbb{E}[b(v_2)] = \mathbb{E}[v^2] = \frac{1}{3}$: equivalent again. Real-world all-pay contests: lobbying, patent races, sports training, wars of attrition, crowdfunded prize competitions.

Exercise 10.3. Cutoff θ^*: contributing yields $\theta - 1$; abstaining yields $\theta \cdot \mathbb{P}(\text{other contributes}) = \theta\frac{2-\theta^*}{2}$. Indifference at θ^*: $\theta^* - 1 = \theta^*\frac{2-\theta^*}{2} \Rightarrow (\theta^*)^2 = 2 \Rightarrow \theta^* = \sqrt{2} \approx 1.41$. First-best: build iff $\max_i \theta_i \geq 1$ with one contributor (or $\theta_1 + \theta_2 \geq 1$ with cost sharing)—the equilibrium underprovides: types in $(1, \sqrt{2})$ free-ride hoping the other contributes.

Exercise 10.4. The informed buyer offers p only when the asset is worth more than p to her *and* the offer is profitable: she buys exactly when $1.5\theta > p$, i.e. your acceptance selects against you. Accepting all $p \geq 50$: offers arrive (at $p = 50$) only when $\theta > 100/3$; you receive 50 for assets worth on average $\mathbb{E}[\theta \mid \theta > 33.3] = 66.7$. Expected payoff $\approx \frac{1}{3}(16.7) + \frac{2}{3}(50) = 38.9 < 50 =$ never trading. Like the lemons market: the willingness of an informed counterparty to trade at p is itself bad news about p.

Chapter 11

Exercise 11.1. Pooling on Beer: on-path belief = prior (0.9 strong) $\Rightarrow$ no duel; off-path (Quiche) belief "weak" $\Rightarrow$ duel. Strong: 3 vs 0; weak: 2 vs 1: both stay—PBE. Pooling on Quiche: symmetric construction works (2 vs 1 for strong, 3 vs 0 for weak)—also a PBE. Intuitive criterion on the Quiche equilibrium: deviating to Beer is *equilibrium-dominated* for the weak type (best conceivable payoff $0 + 2 = 2 < 3$, his equilibrium payoff) but not for the strong type $(1 + 2 = 3 > 2)$. Beliefs after Beer must then assign probability 1 to strong $\Rightarrow$ no duel $\Rightarrow$ the strong type profitably deviates: the all-Quiche equilibrium fails the criterion; all-Beer survives.

Exercise 11.2. Pure signaling, cost e^2/θ: low type: $e_L = 0$, $w_L = 1$. Separation requires the low type not to imitate: $3 - (e^s)^2 \leq 1 \Rightarrow e^s \geq \sqrt{2}$; and the high type to comply: $3 - (e^s)^2/3 \geq 1 \Rightarrow e^s \leq \sqrt{6}$. Riley outcome: $e^s = \sqrt{2}$. Productive education (wage $\theta + e/2$): full-information optima $e_\theta = \theta/4$ (from $\frac{1}{2} = 2e/\theta$). Separation now requires $3 + e/2 - e^2 \leq$ (low type's honest payoff ≈ 1.06), giving $e^s \approx 1.66$, far above the high type's efficient 0.75: the

high type *over*-invests to separate. Education is productive but still partly wasteful at the margin—the signaling distortion survives.

Exercise 11.3. $b = \frac{1}{10}$: condition $2K(K-1)b < 1 \iff K(K-1) < 5 \Rightarrow K = 2$. Two intervals $[0, a], [a, 1]$ with actions $a/2$ and $(1+a)/2$; the boundary type a must be indifferent between the two actions, which gives $a/2 + b = (1-a)/2 - b$, hence $a = \frac{1}{2} - 2b = 0.3$. Receiver's expected loss: babbling $= \mathrm{Var}(U[0,1]) = \frac{1}{12} \approx 0.083$; two-interval $= \frac{0.3^3 + 0.7^3}{12} \approx 0.031$. Communication is coarse but valuable.

Exercise 11.4. Pooling fails: a contract slightly better for low-risks only (lower premium, lower coverage) attracts them and leaves the pool high-risk, making the pooling contract unprofitable—cream-skimming. Separating equilibrium: high-risks get full insurance at their fair premium (their IC binds, no distortion for them); low-risks get partial coverage at their own fair rate—coverage limited exactly enough that high-risks prefer their full-coverage contract. If high-risks are too rare, even this fails and no equilibrium exists (Rothschild–Stiglitz)—an early sign that competitive markets with adverse selection are fragile.

Chapter 12

Exercise 12.1. Row minima: $T : 2$, $C : 0$, $B : 1 \Rightarrow \underline{v} = 2$. Column maxima: $L : 4$, $M : 2$, $R : 5 \Rightarrow \overline{v} = 2$. Saddle point at (T, M), value 2 (unique: no other cell is simultaneously a row-max of its column and column-min of its row).

Exercise 12.2. Payoff matrix (row's payoff): $u(1,\cdot) = (0, 2, -1)$, $u(2,\cdot) = (-2, 0, 2)$, $u(3,\cdot) = (1, -2, 0)$. Antisymmetry ($u(i,j) = -u(j,i)$) confirms value 0. Equalizing: $2x_2 - x_3 = 0$, $-2x_1 + 2x_3 = 0$, $x_1 - 2x_2 = 0$ give $x^* = (\frac{2}{5}, \frac{1}{5}, \frac{2}{5})$. The "middle" action is played least—it never wins big.

Exercise 12.3. Primal: $\max v$ s.t. $x_H - x_T \geq v$ (vs H), $-x_H + x_T \geq v$ (vs T), $x_H + x_T = 1$, $x \geq 0$. Adding the constraints: $0 \geq 2v$, so $v \leq 0$; $x = (\frac{1}{2}, \frac{1}{2})$ achieves $v = 0$. The dual is the column player's identical program; both optima are 0 with uniform strategies.

Exercise 12.4. With the stated tie-breaking, one obtains (one legitimate run): $(H,H), (H,T), (H,T), (T,T), (T$ each player chases the other's empirical frequency, overshoots, and reverses. Frequencies after 6 rounds are already near $(\frac{1}{2}, \frac{1}{2})$; play cycles forever but the *averages* converge to optimal play, exactly as Robinson's theorem promises (and the value estimate converges to 0).

Chapter 13

Exercise 13.1. $\mathrm{BR}_i(x_j) = x_j/2$ (unique, continuous). $A_i = [0,1]$ compact convex; u_i continuous and strictly concave in own action: Theorem 13.4 applies. Fixed point: $x_1 = x_2/2$, $x_2 = x_1/2 \Rightarrow (0,0)$, the unique equilibrium.

Exercise 13.2. Pure equilibria: (T,L), (B,R), and (M,C). The strict equilibria (T,L) and (B,R) are trivially perfect. (M,C): against the *symmetric* totally mixed tremble $\sigma_2 = (\varepsilon, 1-2\varepsilon, \varepsilon)$, all three of row's actions pay exactly 1, so M is a best response, and likewise for C: (M,C) is also perfect (though not strict). Perfection kills weakly dominated behavior—but

M is not dominated, merely payoff-equivalent to $\frac{1}{2}T + \frac{1}{2}B$. The example shows perfection is weaker than one might hope in non-generic games.

Exercise 13.3. Standard example structure: players 1 and 2 move in sequence, player 3 moves at an information set reached only if *both* deviate. PBE permits player 3 to hold *any* belief there—including beliefs that treat 1's and 2's deviations as perfectly correlated in an arbitrary way. Consistency (limits of independent trembles) forces 3's belief to be the *product* of the two small deviation probabilities, pinning the conditional to specific ratios. Sequential rationality against those pinned-down beliefs can then eliminate equilibria PBE allows. Moral: consistency disciplines off-path beliefs by requiring a single coherent "theory of mistakes."

Exercise 13.4. Suppose a bet occurs with positive probability at some pair of signals. Willingness to bet is common knowledge at that point (betting is voluntary and observed). Player 1 bets only if his conditional expected gain is positive; player 2 likewise—but the bet is zero-sum, so the two conditional expectations, taken with respect to the *common prior* conditional on the same public event, cannot both be positive (iterated expectations). Aumann's agreement logic turns "I'm happy to bet" into information that destroys the other's willingness. Without the common prior, players can knowingly disagree, and betting on differences of *opinion* (not information) becomes possible.

Chapter 14

Exercise 14.1. Efficient allocation: $A \to 2$, $B \to 3$ (total $12 > 10$). Payments: without bidder 2, the best the others can do is 10 (bidder 1 takes the pair); others' value at the actual allocation is 6, so bidder 2 pays 4; symmetrically bidder 3 pays 4. Revenue $= 8$. VCG revenue can be fragile: raise $v_2(A) = v_3(B) = 10$ and each pays $10 - 10 = 0$—efficient allocation, *zero* revenue. Adding bidders can *reduce* VCG revenue; this non-monotonicity is why high-stakes combinatorial auctions modify VCG (core-selecting designs).

Exercise 14.2. Exponential: $\frac{1-F}{f} = 1$, so $\psi(v) = v - 1$ (increasing: regular); $r^* = 1$. Pareto $F = 1 - v^{-2}$ on $[1, \infty)$: $\frac{1-F}{f} = \frac{v^{-2}}{2v^{-3}} = \frac{v}{2}$, so $\psi(v) = v/2 > 0$ everywhere: the "reserve" is the support minimum (always sell). Indeed, for one bidder a posted price p earns $p\,(1 - F(p)) = p \cdot p^{-2} = p^{-1}$, maximized at the corner $p = 1$. Heavy tails make revenue maximization degenerate—and for tail index 1, expected values are infinite and every price earns the same revenue. Optimal auction theory quietly assumes tails are thin enough.

Exercise 14.3. With $c, v \sim U[0, 1]$, efficient trade (iff $v > c$) requires expected transfers that violate budget balance under IC and IR; running the Myerson–Satterthwaite calculation, the *second-best* mechanism trades iff $v \geq c + \frac{1}{4}$: one maximizes expected gains from trade subject to IC, IR, and budget balance; the binding budget constraint enters as a Lagrange multiplier on the virtual surplus, and solving for the multiplier that exactly balances the budget yields the $\frac{1}{4}$ wedge (this is also the outcome of the linear equilibrium of the Chatterjee–Samuelson double auction). First-best trade region: the triangle above the diagonal; second-best: the smaller triangle shifted up by $\frac{1}{4}$—trades with gains below $\frac{1}{4}$ are sacrificed to control information rents.

Exercise 14.4. Median voting: a voter left of the median can change the outcome only by reporting on the other side of the median, moving the chosen point *further* from her peak;

misreporting on her own side changes nothing. Single-peakedness makes "further in the wrong direction" unambiguously worse: strategy-proof. Phantoms preserve the argument verbatim (the median of true peaks plus constants). Gibbard–Satterthwaite does not apply because single-peaked preferences are a *restricted domain*—the theorem needs all preference profiles to be possible.

Chapter 15

Exercise 15.1. Core: $\sum x_i = 135$, $x_1 + x_2 \geq 90$, $x_1 + x_3 \geq 80$, $x_2 + x_3 \geq 70$ (equivalently $x_3 \leq 45$, $x_2 \leq 55$, $x_1 \leq 65$), $x_i \geq 0$: a nonempty polytope; e.g. $(65, 45, 25)$. Shapley value: averaging marginal contributions over the 6 orders gives $\phi = (50, 45, 40)$. Check core membership: pairwise sums $95, 90, 85$ all clear their thresholds: yes, $\phi \in C(v)$.

Exercise 15.2. Core: if a worker received more than $g(w) - g(w-1)$, the coalition of L with the other $w-1$ workers could object; so worker pay is capped by the marginal product. Shapley with $g = \sqrt{k}$, $w = 3$: in a uniformly random order, the number of workers preceding L is uniform on $\{0, 1, 2, 3\}$, so $\phi_L = \frac{1}{4}(0+1+\sqrt{2}+\sqrt{3}) \approx 1.04$; each worker gets $(\sqrt{3} - \phi_L)/3 \approx 0.23$ (below the marginal product $\sqrt{3} - \sqrt{2} \approx 0.32$, consistent with the core cap).

Exercise 15.3. Enumerating pivots over the 24 orders of $(A{:}4, B{:}3, C{:}2, D{:}1)$ with quota 6: A pivots 10 times, B 6, C 6, D 2. Index: $(\frac{5}{12}, \frac{1}{4}, \frac{1}{4}, \frac{1}{12})$. Note B and C have *equal power* despite different weights—power is about pivotal position in coalitions, not vote count (the deep lesson of power indices).

Exercise 15.4. Players $\{1, 2, 3, 4\}$; preferences: $1 : 2 \succ 3 \succ 4$; $2 : 3 \succ 1 \succ 4$; $3 : 1 \succ 2 \succ 4$; 4: anyone. In any matching, someone among $\{1, 2, 3\}$ is paired with 4; that person is the *favorite* of one of the other two, who is currently matched to their least-preferred partner among $\{1, 2, 3\}$—the two form a blocking pair. All three matchings are blocked: no stable matching exists. Two-sidedness matters because it prevents exactly these odd "cyclic envy" configurations within one side.

Chapter 16

Exercise 16.1. $u(A, B) = 5 > 3 = u(B, B)$ and $u(A, A) = 2 > 1 = u(B, A)$: A strictly dominates B. The only symmetric equilibrium is (A, A), which is strict, hence the unique ESS. (The indifference equation $5 - 3p = 3 - 2p$ has no solution in $[0, 1]$, confirming no mixed equilibrium.)

Exercise 16.2. Payoffs: $u_H = -x + 2(1 - x) = 2 - 3x$, $u_D = 1 - x$. Replicator: $\dot{x} = x(1 - x)\,(u_H - u_D) = x(1 - x)(1 - 2x)$. Rest points $0, \frac{1}{2}, 1$. Sign: positive on $(0, \frac{1}{2})$, negative on $(\frac{1}{2}, 1)$: interior point $x^* = \frac{1}{2} = V/C$ is asymptotically stable; both monomorphic states are unstable. The population converges to the ESS mix from any interior start.

Exercise 16.3. $\frac{d}{dt} \log \frac{x_a}{x_b} = \frac{\dot{x}_a}{x_a} - \frac{\dot{x}_b}{x_b} = u(a, x) - u(b, x) \leq -\varepsilon$, where $\varepsilon = \min_x [u(b, x) - u(a, x)] > 0$ by strict domination and compactness. Hence $x_a / x_b \leq e^{-\varepsilon t} \cdot x_a(0)/x_b(0) \to 0$, and since $x_b \leq 1$, $x_a \to 0$.

Exercise 16.4. With mutant share ε: $u_{\text{Stag}} = 4\varepsilon + 0(1-\varepsilon) = 4\varepsilon$; $u_{\text{Hare}} = 2$. Stag grows iff $4\varepsilon > 2$, i.e. $\varepsilon > \frac{1}{2}$: the good convention needs the *majority* to switch at once—an enormous invasion barrier (the mixed equilibrium share). This is why inferior conventions persist and why coordinated "big push" interventions (assurance contracts, standards bodies announcing a switch date, critical-mass subsidies) are the practical remedy.

Chapter 17

Exercise 17.1. Take $\Phi(q_1, q_2) = (a-c)(q_1+q_2) - q_1^2 - q_2^2 - q_1 q_2$. Then $\partial\Phi/\partial q_i = a - c - 2q_i - q_j = \partial u_i/\partial q_i$: an exact (continuous) potential. Maximizing Φ: FOCs give $q_1 = q_2 = \frac{a-c}{3}$, the Cournot equilibrium—and concavity of Φ proves uniqueness and gives convergence of continuous best-response dynamics.

Exercise 17.2. Equilibrium: all traffic on road 2 (cost $x^d \le 1$ with equality at $x = 1$): social cost 1. Optimum: minimize $1 - x + x^{d+1}$: $x^{\text{opt}} = (d+1)^{-1/d}$, cost $1 - \frac{d}{d+1}(d+1)^{-1/d} \to 0$ like $\frac{\ln d}{d}$. Hence PoA $\sim \frac{d}{\ln d} \to \infty$: with very convex congestion, selfish routing can be arbitrarily bad—and the affine bound $\frac{4}{3}$ of Theorem 17.3 is seen to be a genuinely structural fact about *linear* congestion.

Exercise 17.3. Rock–Paper–Scissors; let μ be uniform on the three *diagonal* profiles $(R,R), (P,P), (S,S)$, giving each player payoff 0. CCE: any constant deviation faces a uniform opponent marginal, earning 0—no gain. Not CE: conditional on the recommendation "Rock," the opponent plays Rock for sure, so swapping to Paper earns $+1 > 0$. The extra power of swap deviations (using the recommendation as information) is exactly what separates the two concepts.

Exercise 17.4. Congestion game with $c(k) = k$; potential $\Phi = \sum_r \frac{\ell_r(\ell_r+1)}{2}$. Pure equilibria: profiles where no job can move to a strictly less loaded machine, i.e. $|\ell_1 - \ell_2| \le 1$: loads are balanced. Every equilibrium makespan is $\lceil n/2 \rceil$, which is also optimal: PoA $= 1$ here. (With m *unrelated* machines or weighted jobs, equilibria can be unbalanced and the PoA grows—the classic Koutsoupias–Papadimitriou setting.)

Bibliography and Further Reading

Textbooks

Osborne, M. J. (2004). *An Introduction to Game Theory*. Oxford University Press. *The most complete rigorous undergraduate text; the natural companion to Parts I–II of this book.*

Gibbons, R. (1992). *Game Theory for Applied Economists*. Princeton University Press. *Compact, applications-first; superb on the four-cell roadmap of Table 1.1.*

Dixit, A., Skeath, S., & Reiley, D. (2015). *Games of Strategy* (4th ed.). Norton. *Gentlest entry point.*

Osborne, M. J., & Rubinstein, A. (1994). *A Course in Game Theory*. MIT Press. *Elegant graduate core; freely available on Rubinstein's website.*

Fudenberg, D., & Tirole, J. (1991). *Game Theory*. MIT Press. *The encyclopedic graduate reference, strongest on dynamic games and refinements.*

Myerson, R. B. (1991). *Game Theory: Analysis of Conflict*. Harvard University Press. *Systematic, deep; unified treatment of noncooperative and cooperative theory.*

Maschler, M., Solan, E., & Zamir, S. (2020). *Game Theory* (2nd ed.). Cambridge University Press. *The most comprehensive and mathematically complete modern text.*

Krishna, V. (2010). *Auction Theory* (2nd ed.). Academic Press.

Börgers, T. (2015). *An Introduction to the Theory of Mechanism Design*. Oxford University Press.

Roth, A. E., & Sotomayor, M. (1990). *Two-Sided Matching*. Cambridge University Press.

Mailath, G. J., & Samuelson, L. (2006). *Repeated Games and Reputations*. Oxford University Press.

Weibull, J. W. (1995). *Evolutionary Game Theory*. MIT Press.

Sandholm, W. H. (2010). *Population Games and Evolutionary Dynamics*. MIT Press.

Nisan, N., Roughgarden, T., Tardos, É., & Vazirani, V. (eds.) (2007). *Algorithmic Game Theory*. Cambridge University Press. *Free online.*

Roughgarden, T. (2016). *Twenty Lectures on Algorithmic Game Theory*. Cambridge University Press.

Camerer, C. F. (2003). *Behavioral Game Theory*. Princeton University Press.

Perea, A. (2012). *Epistemic Game Theory: Reasoning and Choice*. Cambridge University Press.

Schelling, T. C. (1960). *The Strategy of Conflict.* Harvard University Press. *The classic on focal points, commitment, and credibility—still unmatched as strategic prose.*

von Neumann, J., & Morgenstern, O. (1944). *Theory of Games and Economic Behavior.* Princeton University Press. *Where it all began.*

Landmark articles

von Neumann, J. (1928). Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen* 100, 295–320. (Minimax Theorem.)

Nash, J. (1950). Equilibrium points in n-person games. *PNAS* 36, 48–49; and (1951) Non-cooperative games. *Annals of Mathematics* 54, 286–295.

Nash, J. (1950). The bargaining problem. *Econometrica* 18, 155–162.

Shapley, L. S. (1953). A value for n-person games. In *Contributions to the Theory of Games II*; and (1953) Stochastic games. *PNAS* 39, 1095–1100.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16, 8–37.

Gale, D., & Shapley, L. S. (1962). College admissions and the stability of marriage. *American Mathematical Monthly* 69, 9–15.

Selten, R. (1965, 1975). Subgame perfection and trembling-hand perfection. *Zeitschrift für die gesamte Staatswissenschaft* 121; *IJGT* 4, 25–55.

Harsanyi, J. C. (1967–68). Games with incomplete information played by "Bayesian" players, I–III. *Management Science* 14.

Akerlof, G. A. (1970). The market for "lemons." *QJE* 84, 488–500.

Maynard Smith, J., & Price, G. R. (1973). The logic of animal conflict. *Nature* 246, 15–18.

Spence, M. (1973). Job market signaling. *QJE* 87, 355–374.

Aumann, R. J. (1974). Subjectivity and correlation in randomized strategies. *JME* 1, 67–96; and (1976) Agreeing to disagree. *Annals of Statistics* 4, 1236–1239.

Kreps, D., & Wilson, R. (1982). Sequential equilibria. *Econometrica* 50, 863–894; and, with Milgrom and Roberts, the "gang of four" reputation papers, *JET* 27.

Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica* 50, 97–109.

Crawford, V., & Sobel, J. (1982). Strategic information transmission. *Econometrica* 50, 1431–1451.

Milgrom, P., & Weber, R. (1982). A theory of auctions and competitive bidding. *Econometrica* 50, 1089–1122.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research* 6, 58–73; and, with Satterthwaite (1983), Efficient mechanisms for bilateral trading. *JET* 29, 265–281.

Fudenberg, D., & Maskin, E. (1986). The folk theorem in repeated games with discounting or with incomplete information. *Econometrica* 54, 533–554.

Cho, I.-K., & Kreps, D. (1987). Signaling games and stable equilibria. *QJE* 102, 179–221.

McKelvey, R., & Palfrey, T. (1995). Quantal response equilibria for normal form games. *GEB* 10, 6–38.

Monderer, D., & Shapley, L. S. (1996). Potential games. *GEB* 14, 124–143.

Morris, S., & Shin, H. S. (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *AER* 88, 587–597.

Roughgarden, T., & Tardos, É. (2002). How bad is selfish routing? *JACM* 49, 236–259.

Daskalakis, C., Goldberg, P., & Papadimitriou, C. (2009). The complexity of computing a Nash equilibrium. *SIAM J. Computing* 39, 195–259; with Chen & Deng (2009), Settling the complexity of two-player Nash. *JACM* 56(3).

Kamenica, E., & Gentzkow, M. (2011). Bayesian persuasion. *AER* 101, 2590–2615.

Bowling, M., Burch, N., Johanson, M., & Tammelin, O. (2015). Heads-up limit hold'em poker is solved. *Science* 347, 145–149.

Free online resources

- Osborne & Rubinstein, *A Course in Game Theory* and Rubinstein's other books: `https://arielrubinstein.tau.ac.il/`
- *Algorithmic Game Theory* (Nisan et al.): free PDF from Cambridge/authors' pages; Roughgarden's lecture videos and notes: `https://timroughgarden.org/`
- Stanford/UBC Coursera course "Game Theory" (Jackson, Leyton-Brown, Shoham); Yale Open Courses ECON 159 "Game Theory" (Ben Polak)—excellent video lectures matching Parts I–II.
- *Stanford Encyclopedia of Philosophy*: entries "Game Theory," "Evolutionary Game Theory," "Epistemic Foundations of Game Theory."
- Game Theory Society: `https://gametheorysociety.org/` (surveys, world congress materials, book lists).