

Game-Theoretic Design, Learning, and Engineering

A Course in the Theory of Games



Taka Khoo

Thayer School of Engineering
Dartmouth College

Game-Theoretic Design, Learning, and Engineering: A Course in the Theory of Games. First edition.

Copyright © 2025 Taka Khoo. All rights reserved.

No part of this book may be reproduced, distributed, or transmitted in any form or by any means without the prior written permission of the author, except for brief quotations in reviews and scholarly work. Permission is granted to reproduce the book, in whole or in part, for non-commercial educational use at the Thayer School of Engineering, Dartmouth College, with attribution. For permission requests, write to the author at takakhoo@gmail.com.

This book is based on the lectures of Professor Peter Chin for ENGS 102, *Game-Theoretic Design, Learning, and Engineering*, taught at the Thayer School of Engineering, Dartmouth College, in the Fall term of 2025. It reorganizes and expands the lecture notes, problem sets, and examinations of that course. The mathematical development follows Thomas S. Ferguson's *Game Theory* and Martin J. Osborne's *An Introduction to Game Theory*, the two texts assigned for the course, supplemented by the classical literature cited throughout. Worked examples, problem statements, and solutions originate in the course materials; where an original solution contained an arithmetic slip, the corrected value is given and the discrepancy noted in the text.

Self-published by the author in Hanover, New Hampshire.

How to cite this book. Khoo, T. (2025). *Game-Theoretic Design, Learning, and Engineering: A Course in the Theory of Games* (1st ed.). Self-published. To cite a single chapter, add its number, for example: Khoo, T. (2025), *Game-Theoretic Design, Learning, and Engineering* (1st ed.), Chapter 4.

Written, typeset, and illustrated by Taka Khoo in Latin Modern with L^AT_EX; the diagrams are drawn with TIKZ and PGFPLOTS. Made as a gift for Professor Peter Chin, and as a record of a term's work.

Version of June 8, 2026.

For Professor Peter Chin,

*who showed me that the whole of a game
lives in what the other player is about to do.*

Preface

Game theory is the mathematics of decisions that depend on other decisions. When a chess player chooses a move, a firm sets a price, a router forwards a packet, or a neural network is trained against an adversary, the quality of any single choice depends on choices made by someone else. This book is about the structures that emerge when several rational agents reason about one another at once, and about the theorems that tell us those structures are not chaos.

The material grew out of *ENGS 102: Game-Theoretic Design, Learning, and Engineering*, taught by Professor Peter Chin at Dartmouth in the autumn of 2025. The course began with a duel of poisoned goblets and ended with the division of a bankrupt estate by a rule from the Talmud. In between it passed through Nim and the Sprague–Grundy theorem, von Neumann’s minimax theorem and its proof by linear programming, Nash’s existence theorem and its proof by a fixed point, and the cooperative theory of Shapley and Schmeidler. The arc is wide. One of the pleasures of the subject is that a single set of ideas, value, equilibrium, and fair division, recurs at every scale, from a child’s game to an auction to a learning algorithm.

What this book tries to be

A complete record. Every topic that appeared on the board, every example that was worked, every problem that was assigned, and every examination question is here, with the algebra carried out in full. Nothing is left to the phrase “it can be shown.” A reader who follows a derivation line by line should never have to supply a missing step from elsewhere. Variables are defined where they first appear and collected in the notation guide. Results are numbered consecutively within each chapter so that any definition, theorem, or example can be found by a single reference, for instance [Theorem 4.8](#) or [Example 1.7](#).

The emphasis is computational in the honest sense: we solve things. The companion to the theory is a set of four programming assignments, written in Python, that turn the theorems into working code, a Nim player that never loses from a winning position, a solver that handles any finite zero-sum game by seven different methods, a recursion that values a stochastic game, and a routine that computes the Shapley value and the nucleolus of a coalition. These appear as a Computational Companion alongside the relevant chapters and in full in [Appendix C](#).

How the material is organized

The ten chapters follow the logical development of the subject rather than the calendar of the term. [Chapter 1](#) sets the stage with the vocabulary of strategic interaction and a preview of the four theorems the book will prove. [Chapter 2](#) treats impartial combinatorial games, where there is no chance and no hidden information, and proves the Sprague–Grundy theorem that reduces every such game to a single pile of Nim. [Chapter 3](#) introduces game trees and backward induction. [Chapters 4 to 6](#) develop the theory of two-person zero-sum games to its conclusion: the minimax theorem, the full toolkit for solving a matrix game by hand, and the duality between game values and linear programs. [Chapter 7](#) extends the zero-sum theory to games that call themselves. [Chapter 8](#) crosses into general-sum games and proves Nash’s theorem. The final two chapters, [Chapter 9](#) and [Chapter 10](#), develop cooperative game theory, from two-person bargaining to the Shapley value, the core, and the nucleolus. A table relating the chapters to the original lectures appears in the section *From the Lectures to the Chapters* of *How to Use This Book*.

Five appendices support the main text: a review of the mathematics the course assumes, the four problem sets with complete solutions, an archive of the examinations, the code, and solutions to selected exercises.

A word on rigor and on errors

The proofs here aim to be complete and elementary. Where a result rests on a deep theorem from outside the course, such as Brouwer’s fixed point theorem, we state that theorem precisely, use it, and give a reference rather than a proof. Where the course’s own worked solutions contained a slip, and a few did, the correct value is given in the text and the discrepancy is flagged plainly, so that the reader learns both the answer and the lesson in checking.

Acknowledgments

The debt is obvious and gladly owned. This book would not exist without Professor Peter Chin, whose lectures supplied its substance, whose examples supplied its color, and whose questions supplied its problems. Thanks are due as well to the course assistants, in particular to Mai, whose review sessions sharpened the solution methods that fill Chapter 5 and Appendix B. Any errors that remain are mine.

Taka Khoo
Hanover, New Hampshire

How to Use This Book

Numbering and cross-references

Within each chapter a single counter numbers every definition, theorem, lemma, proposition, corollary, example, procedure, and remark in the order they appear. Thus Definition 4.1 is followed by Theorem 4.2, then perhaps Example 4.3, and so on. The scheme has one purpose: any statement in the book is located by a single number. When the text refers to a result it names the kind and the number, as in “by Theorem 4.2” or “recall Example 4.3,” and in the electronic version the reference is a live link. Equations are numbered separately, also by chapter, as (4.1), (4.2), and so on. Exercises have their own per-chapter numbering at the end of each chapter.

A reader looking for a specific game, theorem, or term should consult the Index, which lists every named game (Nim, Colonel Blotto, Battle of the Sexes, and the rest), every theorem, and every technical term, with page numbers.

The Computational Companion

Four programming assignments accompanied the course, and their ideas are woven into the chapters wherever a theorem becomes an algorithm. A shaded COMPUTATIONAL COMPANION note marks each such place and points to the relevant notebook. The complete code, with explanation, is collected in [Appendix C](#), and the problem sets that the code answers are stated and solved in [Chapter 11](#).

Paths through the book

The chapters are written to be read in order, but the subject splits into four largely independent strands, and a reader with a particular interest may follow one of them.

- **Combinatorial play.** [Chapter 2](#) alone is self-contained: Nim, the theory of P- and N-positions, and the Sprague–Grundy theorem. It needs nothing but the definitions of [Chapter 1](#).
- **Zero-sum theory.** [Chapters 4](#) to [7](#) form the core of the classical theory: the minimax theorem, hand methods for solving matrix games, linear programming duality, and recursive games. This is the spine of the course.
- **Equilibrium in general games.** [Chapter 3](#) and [Chapter 8](#) treat extensive-form games, backward induction, and the Nash equilibrium with its existence proof.
- **Cooperation.** [Chapter 9](#) and [Chapter 10](#) cover bargaining and coalition games, ending with the Shapley value and the nucleolus.

From the lectures to the chapters

The course was delivered as nineteen lectures over a ten-week term. The book reorders the material into the logical sequence of the subject, so a single lecture may feed more than one chapter, and a single chapter may draw on several lectures. The following table records the correspondence, so that any lecture can be traced to its home in the text.

Lecture	Chapter(s)
1. Introduction; minimax and Nash; Golden Balls	1. Foundations of Strategic Thinking
3, 4. Nim, subtraction games, Grundy values	2. Impartial Combinatorial Games
2. Game trees, backward induction, Centipede	3. Extensive-Form Games
10. Kuhn trees, backward induction (trees part)	3. Extensive-Form Games
4, 5. Saddle points, the minimax theorem, 2×2 and $2 \times n$	4. Two-Person Zero-Sum Games
6, 7. Indifference, invariance, symmetry, Colonel Blotto	5. Solving Zero-Sum Games
9X. Midterm methods review	5. Solving Zero-Sum Games (and B)
8, 9. Linear programming, duality, the simplex method	6. Linear Programming and Duality
10, 11. Repeated games, deduction games, recursive and stochastic games	7. Recursive and Stochastic Games
11, 12. Safety levels, general-sum and bimatrix games	8. General-Sum Games and Nash Equilibrium
13, 14. Nash's existence theorem and its fixed-point proof	8. General-Sum Games and Nash Equilibrium
15, 16, 17. Bargaining, TU and NTU games, the exchange rate	9. Cooperative Games I
17, 18, 19, "Finally". Coalitions, Shapley value, core, nucleolus	10. Cooperative Games II

A few lectures supplied material to two chapters. Lecture 4 introduced both the last of the combinatorial games and the first zero-sum game, so it appears in Chapter 2 and again in Chapter 4. Lecture 10 opened with the tree machinery that belongs to Chapter 3 and continued with the repeated and deduction games of Chapter 7. Lecture 17 bridged bargaining and coalitions and is split accordingly.

Contents

Preface	v
How to Use This Book	vii
List of Figures	xiii
Notation	xv
I Foundations and Games of Perfect Information	1
1 Foundations of Strategic Thinking	3
1.1 What is a game?	3
1.2 A first puzzle: dividing a bankrupt estate	4
1.3 The battle of wits: reasoning about reasoning	4
1.4 The founders and the four theorems	5
1.5 The strategic form	6
1.6 Dominance and the Prisoner's Dilemma	7
1.7 Best response and Nash equilibrium	8
1.8 A taxonomy of games	10
1.9 Where this goes: design, learning, and engineering	11
Notes and References	11
Chapter summary	11
Exercises	11
2 Impartial Combinatorial Games	13
2.1 Combinatorial games and normal play	13
2.2 The take-away game	14
2.3 Nim and Bouton's theorem	14
2.4 The Sprague–Grundy function	16
2.5 Sums of games and the Sprague–Grundy theorem	18
2.6 A catalogue of impartial games	20
2.7 Wythoff's game and the golden ratio	21
2.8 Misère play	23
Notes and References	23
Chapter summary	24
Exercises	24
3 Extensive-Form Games and Backward Induction	27
3.1 Game trees	27
3.2 Strategies in the extensive form	28
3.3 Backward induction and subgame perfection	28
3.4 Backward induction through a chance node	31
3.5 Backward induction with continuous choices: Stackelberg competition	32
3.6 Connection to combinatorial games and Zermelo's theorem	33
3.7 The Centipede game and the limits of backward induction	34
3.8 Imperfect information and behavioral strategies	35
Notes and References	37
Chapter summary	37
Exercises	37

II	Two-Person Zero-Sum Games	39
4	Two-Person Zero-Sum Games and the Minimax Theorem	41
4.1	Zero-sum games and the value	41
4.2	Pure strategies, saddle points, and the two values	41
4.3	Mixed strategies and the minimax theorem	43
4.4	The 2×2 game in closed form	44
4.5	Domination and reduction	46
4.6	The graphical method: $2 \times n$ and $n \times 2$ games	49
4.7	The payoff-value lemma	51
4.8	A recipe for small zero-sum games	51
	Notes and References	52
	Chapter summary	52
	Exercises	52
5	Solving Zero-Sum Games: Indifference, Invariance, and Symmetry	55
5.1	The principle of indifference	55
5.2	The non-degenerate case: solution by matrix inversion	56
5.3	One game, solved five ways	57
5.4	Symmetric games are fair	59
5.5	Invariance	60
5.6	Colonel Blotto	61
5.7	Two lemmas toward the minimax theorem	63
5.8	Learning to play: fictitious play	63
	Notes and References	64
	Chapter summary	64
	Exercises	65
6	Linear Programming and Duality	67
6.1	Linear programs	67
6.2	The geometry of optima	67
6.3	Duality	68
6.4	From a game to a pair of linear programs	69
6.5	The minimax theorem, proved	69
6.6	Complementary slackness is the principle of indifference	70
6.7	The simplex method	71
6.8	Shadow prices	73
	Notes and References	73
	Chapter summary	74
	Exercises	74
7	Recursive and Stochastic Games	75
7.1	The substitution principle	75
7.2	A first recursive game	76
7.3	Symmetry in recursive games	76
7.4	Role-switching recursion	77
7.5	Coupled subgames	77
7.6	Stochastic games	77
7.7	Repeated games and the inspection problem	79
7.8	The repeated Prisoner's Dilemma and the Folk Theorem	80
7.9	A deduction game with subgame states	81
	Notes and References	82
	Chapter summary	82
	Exercises	82
III	General-Sum Games and Nash Equilibrium	85
8	General-Sum Games and Nash Equilibrium	87
8.1	From zero-sum to general-sum	87

8.2	Safety levels	87
8.3	N-person games and Nash equilibrium	88
8.4	Canonical two-person games	88
8.5	Iterated elimination of dominated strategies	91
8.6	The price of anarchy	92
8.7	Cournot competition	93
8.8	Brouwer's fixed-point theorem	94
8.9	Nash's existence theorem	94
8.10	Correlated equilibrium	97
8.11	Evolutionarily stable strategies	98
	Notes and References	98
	Chapter summary	98
	Exercises	99
IV	Cooperative Game Theory	101
9	Cooperative Games I: Bargaining and Transferable Utility	103
9.1	Cooperation and binding agreements	103
9.2	Feasible sets	104
9.3	The transferable-utility solution	104
9.4	The Nash bargaining solution	106
9.5	Shapley's exchange rate	109
	Notes and References	110
	Chapter summary	110
	Exercises	110
10	Cooperative Games II: The Shapley Value, the Core, and the Nucleolus	113
10.1	Coalitional games	113
10.2	Imputations and the core	114
10.3	Balanced games and the Bondareva–Shapley theorem	115
10.4	The Shapley value	116
10.5	The nucleolus	120
10.6	The bankruptcy problem and the Talmud	121
10.7	Shapley value against nucleolus	122
	Notes and References	122
	Chapter summary	123
	Exercises	123
V	The Problem Sets	125
11	Problem Set 1: Combinatorial Games	127
12	Problem Set 2: Two-Person Zero-Sum Games	131
13	Problem Set 3: Recursive Games	135
14	Problem Set 4: N-Person and Cooperative Games	139
VI	Appendices	143
A	Mathematical Review	145
A.1	Sets, functions, and the simplex	145
A.2	Linear algebra	145
A.3	Convexity	146
A.4	Continuity and compactness	146
A.5	Probability and expectation	147
A.6	Fixed-point theorems	147
A.7	Asymptotic notation and binary arithmetic	147

B Examination Archive	149
B.1 Midterm examination	149
B.2 Practice final examination	152
B.3 Additional practice problems	156
B.4 Examination methods at a glance	162
C Computational Companion	163
C.1 The Nim and Grundy toolkit	163
C.2 The universal zero-sum solver	164
C.3 Recursive game solvers	164
C.4 Fixed points and the cooperative solver	165
C.5 Backward induction and fictitious play	165
C.6 Combinatorial and cooperative solvers	166
D Solutions to Selected Exercises	167
E Glossary	181
F A Timeline of Game Theory	185
G Further Reading	187
H A Catalog of Canonical Games	189
I Common Pitfalls	191
Bibliography	193
Index	197

List of Figures

1.1	The poisoned-goblet duel as a tree	5
1.2	Best responses in the Stag Hunt	10
2.1	Grundy values on a game graph	17
2.2	Cold positions of Wythoff's game	23
3.1	A two-stage game	28
3.2	Backward induction as folding	29
3.3	A three-level game	29
3.4	The entry-deterrence game	31
3.5	Entry under an uncertain market	32
3.6	Stackelberg versus Cournot	33
3.7	The Centipede paradox	34
3.8	A poker bluffing game	36
3.9	A behavioral strategy in poker	37
4.1	The payoff surface near a saddle	43
4.2	A 2×4 game and its envelope	50
5.1	Graphical solution of a 2×2 game	59
5.2	A Colonel Blotto contest	62
5.3	Fictitious play converging	64
6.1	An LP feasible polytope	68
7.1	A stochastic coin game	78
8.1	Best responses in the Battle of the Sexes	90
8.2	Pigou's two-road network	92
8.3	Cournot reaction curves	93
9.1	NTU and TU feasible sets	104
9.2	The Nash bargaining solution	107
9.3	Bargaining from safety levels	108
9.4	Nash versus Kalai–Smorodinsky	109
10.1	The imputation triangle and core	115
11.1	The position graph of Problem 3. Each arrow is a legal move; sinks are terminal. Grundy values are computed in reverse topological order by the mex rule. Shaded vertices (v_1, v_4, v_6) are the P-positions, where $g = 0$	129
13.1	The mutually recursive trio: each game's value feeds the next around the cycle. Solving the loop turns the three equations into a single quadratic in v_3	136
14.1	Feasible sets for Problem 3. Shaded: the NTU hull of the cell payoffs. Bold blue: the Pareto frontier. Dashed red: the TU frontier $u_1 + u_2 = V$ at the largest total. In Game (ii) the dominated cells $(0, 2)$ and $(3, 0)$ sit strictly inside the frontier, so bargaining happens only along the short segment $(1, 2)$ to $(3, 1)$	140

A.1	The probability simplex	145
B.1	An extensive-form exam game	152
B.2	Graphical solution of an exam game	153
B.3	A bluffing game (exam)	155

Notation

The symbols below are used consistently throughout the book. Each is defined again where it first appears in context. Vectors are set in bold, $\mathbf{x}$; their components are set in plain italic with a subscript, x_i . All vectors are column vectors unless stated otherwise, and $\mathbf{x}^\top$ denotes the transpose.

Sets and numbers

$\mathbb{R}, \mathbb{Z}, \mathbb{N}, \mathbb{Q}$	the real numbers, integers, natural numbers, rationals
$\mathbb{R}^n$	n -dimensional Euclidean space
Δ_n	the probability simplex $\{\mathbf{p} \in \mathbb{R}^n : p_i \geq 0, \sum_i p_i = 1\}$
$ S $	the number of elements of a finite set S
$\text{conv}(X)$	the convex hull of a set X
2^N	the collection of all subsets of N

Vectors, matrices, and operators

$\mathbf{1}, \mathbf{0}$	the all-ones and all-zeros vectors
$A, \mathbf{x}^\top A \mathbf{y}$	a payoff matrix and the bilinear expected payoff
$A^\top$	transpose of A
$\arg \max, \arg \min$	the maximizing and minimizing arguments
$\mathbb{E}[\cdot], \mathbb{P}(\cdot)$	expectation and probability
$\text{supp}(\mathbf{p})$	the support $\{i : p_i > 0\}$ of a mixed strategy

Strategic games

N	the set of players, usually $\{1, 2, \dots, n\}$
S_i, A_i	the (pure) strategy set and action set of player i
s_i, s_{-i}	a strategy of player i ; the strategies of all players but i
u_i	the payoff (utility) function of player i
$\mathbf{p}, \mathbf{q}$	mixed strategies of the row and column player
$\text{BR}_i(\cdot)$	the best-response correspondence of player i
$V, \underline{V}, \bar{V}$	the value, lower value, and upper value of a zero-sum game

Combinatorial games

$g(x)$	the Grundy (Sprague–Grundy) value, or nimber, of position x
$\text{mex}(T)$	the minimal excludant: the least nonnegative integer not in T
$\oplus$	the nim-sum, that is, bitwise addition modulo 2 (XOR)
P, N	a previous-player-wins and next-player-wins position
$G + H$	the disjunctive sum of two games

Cooperative games

(N, v)	a coalitional game with player set N and characteristic function v
$v(S)$	the worth of coalition $S \subseteq N$
$\varphi_i(v)$	the Shapley value of player i
$\mathcal{C}(v)$	the core of the game
$\mathcal{I}(v)$	the set of imputations
$e(S, \mathbf{x})$	the excess $v(S) - \sum_{i \in S} x_i$ of coalition S at allocation $\mathbf{x}$
$\mathcal{N}(v)$	the nucleolus

A note on a recurring convention. In a two-person zero-sum game the row player (Player I) is the maximizer and the column player (Player II) is the minimizer, and a single number a_{ij} in the payoff matrix is the amount Player II pays Player I when Player I chooses row i and Player II chooses column j . In a general-sum game each cell holds a pair (a_{ij}, b_{ij}) , the payoffs to Players I and II respectively.

PART I

**Foundations and Games of Perfect
Information**

CHAPTER 1

Foundations of Strategic Thinking

*“You must have known I was not a great fool; you would have counted on it,
so I can clearly not choose the wine in front of you.”*

Vizzini, The Princess Bride

A decision is strategic when its quality depends on a decision someone else is making at the same time. A chess player choosing a move, a firm setting a price, a wireless node choosing a transmission power, and a generative model trained against a discriminator are all in the same predicament: the payoff to any action is a function of actions chosen by others, who are reasoning in exactly the same way about us. Game theory is the branch of mathematics that studies these situations and, remarkably, proves that they have structure. This chapter assembles the vocabulary and previews the four theorems that the rest of the book will establish in full.

1.1 What is a game?

Informally, a *game* is any situation involving two or more decision makers, called *players*, each of whom chooses an action, where the outcome and each player’s satisfaction with it depend on the whole profile of chosen actions. Three ingredients are always present.

- **Players.** The set of decision makers, written $N = \{1, 2, \dots, n\}$.
- **Strategies.** For each player i , a set S_i of available choices.
- **Payoffs.** For each player i , a rule u_i that assigns a real number to every combination of choices, measuring how much player i values that outcome.

The defining feature, the one that separates game theory from ordinary optimization, is *strategic interdependence*: player i ’s payoff depends not only on i ’s own choice but on the choices of everyone else. There is no “best action” in isolation. The best action depends on what the others do, which depends on what they expect us to do, which depends on what they expect us to expect, and so on. Taming that regress is the central problem of the subject.

We will make each of these ingredients precise in Section 1.5. First, three examples set the tone: a problem of fair division that opened the course, a duel that illustrates the regress, and a game show that exposes the gap between individual and collective rationality.

1.2 A first puzzle: dividing a bankrupt estate

Before any formal machinery, consider a problem of *fair division* that is two thousand years old and that we will not fully solve until Chapter 10.

Example 1.1 (The contested estate). A venture has failed. Three investors hold claims against what remains:

$$c_1 = 100, \quad c_2 = 200, \quad c_3 = 300,$$

in millions of dollars, so the total claimed is $c_1 + c_2 + c_3 = 600$. After liquidation only an estate of size E remains, with $E < 600$, so the claims cannot all be met. How should the estate be divided?

The proportional rule, giving each claimant the fraction $c_i/600$ of the estate, is one answer, but it is not the answer that a particular ancient source recommends. The Babylonian Talmud records the following allocations for three specific estates, and for nearly two millennia they appeared simply inconsistent:

Estate E	to claimant 1	to claimant 2	to claimant 3
100	$33\frac{1}{3}$	$33\frac{1}{3}$	$33\frac{1}{3}$
200	50	75	75
300	50	100	150

When the estate is small ($E = 100$) the three claimants split it equally, ignoring the sizes of their claims. When the estate is large ($E = 300$) the division is proportional to the claims, $50 : 100 : 150 = 1 : 2 : 3$. The middle case follows neither rule. What single principle produces all three rows? $\diamond$

The resolution, found by Aumann and Maschler [1], is that every row is the *nucleolus* of an associated cooperative game, a solution concept we will define and compute in Chapter 10. For now the example serves two purposes. It shows that “fairness” is a mathematical object with competing definitions, and it plants a question that the machinery of the final chapter will answer completely. The same numbers reappear, fully derived, in Section 10.6.

The idea behind the table is old and intuitive in the two-claimant case, where it is called the *contested-garment* rule. Imagine a garment worth 100, with one person claiming all of it and another claiming half. The claimant of half is, in effect, conceding the other half: they admit that 50 is not theirs. That conceded 50 goes to the first claimant uncontested. Only the remaining 50 is genuinely in dispute, so it is split evenly, 25 each. The first claimant receives $50 + 25 = 75$ and the second receives 25. The rule is: *each claimant first concedes to the other whatever they do not themselves claim, and the contested remainder is divided equally*. The three-claimant rows of the table above are the unique consistent extension of this single rule to more parties, which is the sense in which the Talmud’s rulings, far from being arbitrary, follow one principle.

Historical Note

The bankruptcy passage appears in the Mishnah (Ketubot 93a). For most of its history it was read as a list of disparate case rulings. Robert Aumann, who shared the 2005 Nobel Memorial Prize in Economics, and Michael Maschler showed in 1985 that the three rulings are the consistent output of a single game-theoretic rule, an instance of modern mathematics illuminating an ancient text.

1.3 The battle of wits: reasoning about reasoning

In the 1987 film *The Princess Bride*, the hero Westley challenges the villain Vizzini to a duel of logic. Two goblets of wine sit on a table. Westley has secretly poisoned one of them. Vizzini must choose a goblet to drink; both then drink simultaneously. The loser is the one who swallows the poison.

This is a two-player, simultaneous-move, strictly competitive game with one player holding private information. Vizzini’s celebrated monologue is an attempt to deduce Westley’s choice by reasoning about Westley’s reasoning:

“I think that he thinks that I think that ...”

Each layer of “he thinks that I think” adds a level to a *belief hierarchy*. If the recursion never terminates, pure logic alone cannot single out a choice. The way out, which we will formalize in Chapter 4, is *randomization*: a player who is genuinely unpredictable cannot be outguessed, no matter how deep the opponent’s reasoning runs.

We can already write the strategic skeleton. Let Westley’s hidden choice be the goblet he poisons, A or B , and let Vizzini’s choice be the goblet he drinks. Vizzini does not observe Westley’s choice, so he holds a belief, a probability,

$$\mu(\text{Westley poisoned } A) = p, \quad \mu(\text{Westley poisoned } B) = 1 - p,$$

and he selects the goblet that minimizes his probability of death. Westley, anticipating this, wants to make Vizzini’s belief as useless as possible. The tension between a maximizer and a minimizer who anticipate one another,

$$\max_{\text{Westley}} \min_{\text{Vizzini}} \mathbb{E}[\text{outcome}] \quad \text{against} \quad \min_{\text{Vizzini}} \max_{\text{Westley}} \mathbb{E}[\text{outcome}],$$

is exactly the structure that von Neumann’s minimax theorem resolves.

We can compute the answer now, by hand, and it is worth doing because it shows the whole logic of the book in miniature. Suppose Westley poisons goblet A with probability w , and Vizzini, unable to see, drinks from A with probability d . Vizzini dies exactly when he drinks the poisoned goblet, so his probability of death is

$$\mathbb{P}(\text{Vizzini dies}) = \underbrace{dw}_{\text{drinks } A, A \text{ poisoned}} + \underbrace{(1-d)(1-w)}_{\text{drinks } B, B \text{ poisoned}} = (1-w) + d(2w-1).$$

Read the last expression as a function of Vizzini’s choice d . Its slope in d is $2w-1$. If Westley poisons A more often than half the time ($w > \frac{1}{2}$), the slope is positive and Vizzini drives his death probability down by taking $d = 0$, drinking B ; if $w < \frac{1}{2}$, he takes $d = 1$. Either lopsided choice lets Vizzini push his risk below $\frac{1}{2}$. The one value Westley can choose that denies Vizzini any such handle is

$$w = \frac{1}{2}, \quad \text{whereupon } \mathbb{P}(\text{Vizzini dies}) = (1 - \frac{1}{2}) + d \cdot 0 = \frac{1}{2} \quad \text{for every } d.$$

By poisoning each goblet with equal probability, Westley makes Vizzini *indifferent*: no amount of reasoning about which goblet to drink changes the outcome, which stays a coin flip. The value of the duel is $\frac{1}{2}$, and the optimal strategy is to be maximally unpredictable. This is the principle of indifference of Chapter 5 in its smallest instance, and it is why Vizzini’s tower of “he thinks that I think” leads nowhere: against a player who has randomized correctly, depth of reasoning is worthless.




Figure 1.1. The poisoned-goblet duel as a tree. Westley (W) moves first but in secret, so Vizzini (V) cannot distinguish the two nodes joined by the dashed line. This is an *information set*, the device of Chapter 3.

1.4 The founders and the four theorems

Two mathematicians built the foundations on which everything here rests.

Historical Note

John von Neumann (1903–1957) proved the minimax theorem for two-person zero-sum games in 1928 [63], then with the economist Oskar Morgenstern wrote *Theory of Games and Economic Behavior* [64], the book that founded the field. **John Nash** (1928–2015) generalized equilibrium to any number of players and any payoffs in a 1950 note and a 1951 paper [36, 38], proving existence with a fixed-point theorem. Nash received the 1994 Nobel Memorial Prize in Economics and the 2015 Abel Prize.

The book proves four theorems. We state them informally now, as a map of the territory, and prove each one in its place.

1. **The Sprague–Grundy theorem** (Chapter 2). Every impartial combinatorial game, however complicated, is equivalent to a single heap of Nim. A game’s entire strategic content collapses to one number, its Grundy value.
2. **The minimax theorem** (Chapters 4 and 6). Every finite two-person zero-sum game has a *value* V and optimal mixed strategies for both players:

$$\max_{x \in \Delta_m} \min_{y \in \Delta_n} x^\top A y = \min_{y \in \Delta_n} \max_{x \in \Delta_m} x^\top A y = V. \quad (1.1)$$

The order of optimization does not matter; foresight gives the opponent no advantage.

3. **Nash’s existence theorem** (Chapter 8). Every finite game, with any number of players and any payoffs, has at least one equilibrium in mixed strategies.
4. **The Shapley value** (Chapter 10). There is exactly one way to divide the worth of a cooperative venture among its participants that satisfies four natural fairness axioms, and there is a formula for it.

Each rests on a tool from elsewhere in mathematics: the minimax theorem on linear programming duality, Nash’s theorem on Brouwer’s fixed-point theorem, the Sprague–Grundy theorem on induction over game positions, and the Shapley value on a characterization by axioms. Part of the pleasure of the subject is watching these tools arrive exactly when they are needed.

1.5 The strategic form

We now make the three ingredients precise. The *strategic form*, also called *normal form*, is the most basic way to write a game.

Definition 1.2 (Finite strategic-form game). A *finite strategic-form game* is a triple

$$G = (N, (S_i)_{i \in N}, (u_i)_{i \in N}),$$

consisting of

- a finite set of **players** $N = \{1, \dots, n\}$;
- for each player i , a finite, nonempty set of **pure strategies** S_i ;
- for each player i , a **payoff function** $u_i : S_1 \times S_2 \times \dots \times S_n \rightarrow \mathbb{R}$.

A **strategy profile** is a tuple $s = (s_1, \dots, s_n)$ with $s_i \in S_i$ for every i . We write $S = S_1 \times \dots \times S_n$ for the set of profiles.

It is convenient to separate one player from the rest. Write s_{-i} for the profile of the other players,

$$s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n),$$

so that a full profile is (s_i, s_{-i}) , and $u_i(s_i, s_{-i})$ is player i ’s payoff when i plays s_i and everyone else plays s_{-i} . The notation s_{-i} recurs constantly; it is the formal home of “what everybody else does.”

For two players with finite strategy sets, the game is displayed as a matrix. The row player (Player I) has strategies indexing the rows, the column player (Player II) has strategies indexing the columns, and each cell records the pair of payoffs.

Definition 1.3 (Bimatrix game). A two-player game in which Player I has m strategies and Player II has n strategies is a *bimatrix game*. It is given by two $m \times n$ matrices A and B , where $a_{ij} = u_1(i, j)$ and $b_{ij} = u_2(i, j)$. We display it by writing the pair (a_{ij}, b_{ij}) in row i , column j .

1.5.1 Mixed strategies

A player need not commit to a single action. A player may *randomize*, choosing each action with some probability. This is not a technicality; Section 1.3 already showed that randomization is the only escape from infinite regress, and von Neumann's theorem will show it is always enough.

Definition 1.4 (Mixed strategy). A *mixed strategy* for player i is a probability distribution over the pure strategies S_i . If $S_i = \{1, \dots, m_i\}$, a mixed strategy is a vector

$$\mathbf{p} = (p_1, \dots, p_{m_i}) \in \Delta_{m_i}, \quad \Delta_{m_i} = \left\{ \mathbf{p} \in \mathbb{R}^{m_i} : p_k \geq 0, \sum_k p_k = 1 \right\},$$

where Δ_{m_i} is the **probability simplex**. A pure strategy is the special case in which one component equals 1. The set $\text{supp}(\mathbf{p}) = \{k : p_k > 0\}$ is the **support** of $\mathbf{p}$.

When players randomize independently, payoffs are replaced by their expectations. In a two-player game with payoff matrix A for Player I, if Player I uses $\mathbf{p}$ and Player II uses $\mathbf{q}$, the probability of the cell (i, j) is $p_i q_j$, and Player I's expected payoff is the bilinear form

$$u_1(\mathbf{p}, \mathbf{q}) = \sum_{i=1}^m \sum_{j=1}^n p_i a_{ij} q_j = \mathbf{p}^\top A \mathbf{q}. \quad (1.2)$$

The expression $\mathbf{p}^\top A \mathbf{q}$, linear in $\mathbf{p}$ for fixed $\mathbf{q}$ and linear in $\mathbf{q}$ for fixed $\mathbf{p}$, is the workhorse of the entire zero-sum theory.

Example 1.5 (Computing an expected payoff, two ways). Let Player I's payoff matrix be $A = \begin{bmatrix} 3 & -1 \\ 0 & 2 \end{bmatrix}$, and suppose Player I mixes $\mathbf{p} = (\frac{1}{2}, \frac{1}{2})$ over the two rows while Player II mixes $\mathbf{q} = (\frac{2}{3}, \frac{1}{3})$ over the two columns. What does Player I expect to earn? *Cell by cell.* Because the players randomize independently, the chance of landing in cell (i, j) is $p_i q_j$. Lay the four probabilities and payoffs side by side:

prob $\times$ payoff	$q_1 = \frac{2}{3}$	$q_2 = \frac{1}{3}$
$p_1 = \frac{1}{2}$	$\frac{1}{3} \cdot (3)$	$\frac{1}{6} \cdot (-1)$
$p_2 = \frac{1}{2}$	$\frac{1}{3} \cdot (0)$	$\frac{1}{6} \cdot (2)$

$$\sum = \underbrace{\frac{1}{3}(3)}_1 + \underbrace{\frac{1}{6}(-1)}_{-1/6} + \underbrace{\frac{1}{3}(0)}_0 + \underbrace{\frac{1}{6}(2)}_{1/3} = \frac{7}{6}.$$

Matrix form. The same number drops out of $\mathbf{p}^\top A \mathbf{q}$ if we first apply A to $\mathbf{q}$,

$$A \mathbf{q} = \begin{bmatrix} 3 & -1 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} \frac{2}{3} \\ \frac{1}{3} \end{bmatrix} = \begin{bmatrix} 3 \cdot \frac{2}{3} - 1 \cdot \frac{1}{3} \\ 0 \cdot \frac{2}{3} + 2 \cdot \frac{1}{3} \end{bmatrix} = \begin{bmatrix} \frac{5}{3} \\ \frac{2}{3} \end{bmatrix}, \quad \mathbf{p}^\top (A \mathbf{q}) = \frac{1}{2} \cdot \frac{5}{3} + \frac{1}{2} \cdot \frac{2}{3} = \frac{5}{6} + \frac{2}{6} = \frac{7}{6}.$$

Both routes give $u_1(\mathbf{p}, \mathbf{q}) = \frac{7}{6} \approx 1.17$. The expected payoff is just the payoff in each cell weighted by how likely that cell is, which is why a mixed strategy is best read as a *long-run average*: if these two players faced off many times with these dice, Player I would bank about $\frac{7}{6}$ per round. $\diamond$

Intuition: a mixed strategy is a weather forecast rather than a hedge

Mixing does not mean “playing it safe” or splitting the difference. It means making yourself *unpredictable*, the way a tennis server varies between the body and the wide serve so the returner cannot lean early. The payoff $\mathbf{p}^\top A \mathbf{q}$ is the average outcome over the joint roll of both players' dice, exactly as a forecast of “ $\frac{2}{3}$ chance of sun” is an average over possible days. We will see that against a clever opponent there is an optimal forecast to commit to, and that committing to it openly costs nothing, which is the content of the minimax theorem (Theorem 4.7).

1.6 Dominance and the Prisoner's Dilemma

The first tool for analyzing a game is to ask whether some strategy is simply better than another, whatever the opponents do.

Definition 1.6 (Dominance). Fix player i . A strategy s_i **strictly dominates** a strategy s'_i if

$$u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i}) \quad \text{for every } s_{-i},$$

and **weakly dominates** it if the inequality holds with $\geq$ for every s_{-i} and strictly for at least one s_{-i} . A strategy is **dominant** if it dominates every other strategy of player i .

A rational player never plays a strictly dominated strategy, because another strategy does better regardless of the opponents. This single observation already determines the outcome of one of the most studied games in all of science.

Example 1.7 (Golden Balls: split or steal). The British television program *Golden Balls* ended each episode with the following game. Two contestants face a jackpot of size J . Each secretly chooses a ball marked either SPLIT or STEAL, and the balls are revealed simultaneously. If both choose SPLIT, they share the jackpot equally. If one chooses STEAL while the other chooses SPLIT, the stealer takes the whole jackpot and the splitter gets nothing. If both choose STEAL, both leave with nothing. The bimatrix is

	SPLIT	STEAL
SPLIT	$(\frac{J}{2}, \frac{J}{2})$	$(0, J)$
STEAL	$(J, 0)$	$(0, 0)$

Player I chooses the row, Player II the column.

Analysis. Examine Player I's choice against each choice of Player II.

- If Player II plays SPLIT, Player I gets J by stealing versus $J/2$ by splitting. Stealing is better.
- If Player II plays STEAL, Player I gets 0 either way. Stealing ties.

So STEAL weakly dominates SPLIT for Player I, and by symmetry for Player II. Each player has a weakly dominant strategy, STEAL, and the profile (STEAL, STEAL) results, paying $(0, 0)$. $\diamond$

The conclusion of Example 1.7 is unsettling: two rational players walk away with nothing, even though both prefer the (SPLIT, SPLIT) outcome that pays each $J/2$. Individual rationality and collective benefit point in opposite directions. This is the *Prisoner's Dilemma*, the most famous structure in game theory.

Example 1.8 (The Prisoner's Dilemma). Two suspects are interrogated separately. Each may COOPERATE with the other (stay silent) or DEFECT (testify). The standard payoffs, written so that larger is better, are

	COOPERATE	DEFECT
COOPERATE	$(3, 3)$	$(0, 5)$
DEFECT	$(5, 0)$	$(1, 1)$

For each player, DEFECT strictly dominates COOPERATE: $5 > 3$ when the opponent cooperates, and $1 > 0$ when the opponent defects. Rational play yields (DEFECT, DEFECT) paying $(1, 1)$, although (COOPERATE, COOPERATE) paying $(3, 3)$ is better for both. The gap between the equilibrium outcome and the jointly preferred outcome is the dilemma. $\diamond$

Why, then, do real contestants sometimes split, and real suspects sometimes stay silent? Because the one-shot game with these exact payoffs omits forces that operate in life: reputation, the shadow of future encounters, moral commitment, and the possibility of binding agreements. Each of these is itself a game-theoretic idea. Repetition is the subject of Chapter 7; binding agreements are the subject of Chapters 9 and 10. The lesson to carry forward is that the model must include every force that matters, and that the dilemma is a property of the payoffs rather than a failure of the players.

A Nash equilibrium is *stable*, but it need not be *good*. The unique equilibrium of the Prisoner's Dilemma, (DEFECT, DEFECT) paying $(1, 1)$, is strictly worse for both players than (COOPERATE, COOPERATE) paying $(3, 3)$. Stability against unilateral deviation and collective efficiency are different properties, and they often conflict. Reading "equilibrium" as "optimal" is the most common error in applying the theory.

1.7 Best response and Nash equilibrium

Dominance settles Examples 1.7 and 1.8, but most games have no dominant strategies. The general solution concept replaces "best no matter what" with "best given what the others actually do."

Definition 1.9 (Best response). A strategy $s_i \in S_i$ is a **best response** to the others' profile s_{-i} if

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}) \quad \text{for all } s'_i \in S_i.$$

The set of all best responses to s_{-i} is written $BR_i(s_{-i})$.

Definition 1.10 (Nash equilibrium). A strategy profile $\mathbf{s}^* = (s_1^*, \dots, s_n^*)$ is a **Nash equilibrium** if every player's strategy is a best response to the others':

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*) \quad \text{for every } i \in N \text{ and every } s_i \in S_i. \quad (1.3)$$

Equivalently, $s_i^* \in BR_i(s_{-i}^*)$ for every i . The definition applies verbatim to mixed strategies, with u_i replaced by expected payoff and S_i by Δ_{m_i} .

The interpretation is stability against unilateral deviation. At a Nash equilibrium, no single player can improve by changing strategy alone, holding the others fixed. The concept does not assert that the outcome is good for the players jointly, as (DEFECT, DEFECT) in Example 1.8 makes painfully clear: that profile is the unique Nash equilibrium and it is bad for both. It does not assert that equilibrium is unique, as the next chapters will show. It asserts consistency: an equilibrium is a profile at which everyone's expectation about everyone else is confirmed.

Remark 1.11. A profile in which each player uses a (weakly) dominant strategy is always a Nash equilibrium, since a dominant strategy is a best response to everything. Thus (STEAL, STEAL) and (DEFECT, DEFECT) are Nash equilibria of Examples 1.7 and 1.8. The converse fails: most Nash equilibria do not come from dominance.

Remark 1.12 (Intuition: mutually confirmed beliefs). A Nash equilibrium is best pictured as a set of beliefs that confirm themselves. If each player believes the others will play $\mathbf{s}_{-i}^*$ and best-responds to that belief, the resulting play is exactly $\mathbf{s}^*$, so every belief is borne out and no player regrets their choice. Away from equilibrium, some player is acting on a belief that the others' behavior contradicts, and that player will want to revise. An equilibrium is a fixed point of this process of revising beliefs in the direction of best response, which is precisely why its existence (Chapter 8) is proved with a fixed-point theorem.

Example 1.13 (The Stag Hunt, solved completely). Two hunters each decide, without communicating, whether to hunt a STAG or a HARE. A stag feeds both but can be caught only if *both* commit to it; a hare is caught alone and needs no partner. With payoffs scaled so larger is better, the symmetric bimatrix is

	STAG	HARE
STAG	(4, 4)	(0, 3)
HARE	(3, 0)	(3, 3)

Step 1: pure equilibria by best response. Against a partner who hunts STAG, Player I earns 4 from STAG versus 3 from HARE, so STAG is the best response; against a partner who hunts HARE, Player I earns 0 from STAG versus 3 from HARE, so HARE is the best response. The same holds for Player II by symmetry. The two profiles where both play a mutual best response are (STAG, STAG), paying (4, 4), and (HARE, HARE), paying (3, 3): two pure Nash equilibria.

Step 2: the mixed equilibrium by indifference. Suppose Player I hunts STAG with probability p . Player II is willing to mix only if her two pure strategies pay equally against that p . Her STAG pays $4p + 0(1 - p) = 4p$ (she gets 4 only when Player I also chose STAG); her HARE pays $3p + 3(1 - p) = 3$ (a hare is caught regardless). Setting these equal,

$$4p = 3 \implies p = \frac{3}{4}.$$

By symmetry Player II also hunts STAG with probability $q = \frac{3}{4}$, so the mixed equilibrium is $\mathbf{p}^* = \mathbf{q}^* = (\frac{3}{4}, \frac{1}{4})$. Its payoff to each player is the common value of the two indifferent options, $4 \cdot \frac{3}{4} = 3$, the same (3, 3) as the all-HARE equilibrium.

Step 3: read the picture. Figure 1.2 plots the two best-response correspondences in the square of mixing probabilities. They meet at exactly three points, the three Nash equilibria: the two corners and the interior crossing at $(\frac{3}{4}, \frac{3}{4})$.

Step 4: which equilibrium? The Stag Hunt sharpens the lesson of the Prisoner's Dilemma. Here the efficient outcome (STAG, STAG) is an equilibrium, so cooperation is self-enforcing once expected, this is the **payoff-dominant** equilibrium. Yet (HARE, HARE) is also an equilibrium, and it is **risk-dominant**: hunting hare guarantees 3 no matter what the




Figure 1.2. Best-response correspondences for the Stag Hunt of Example 1.13. Player I's best response (blue) gives p as a function of the partner's q ; Player II's (red) gives q as a function of p . The curves intersect at the three Nash equilibria: the payoff-dominant (STAG, STAG), the risk-dominant (HARE, HARE), and the mixed equilibrium at $(\frac{3}{4}, \frac{3}{4})$.

partner does, while hunting stag risks 0 if the partner defects. A hunter who is unsure of the partner may rationally take the safe hare, and the group is trapped at the worse equilibrium not by a dominant strategy but by mistrust. The gap between payoff dominance and risk dominance is the modern theory of equilibrium selection. $\diamond$

The two questions that organize the rest of the book are now visible. *Does an equilibrium always exist?* Nash's theorem (Chapter 8) answers yes, provided we allow mixed strategies. *How do we find it?* The answer depends on the structure of the game, and supplying it for each structure is the work of the chapters between here and there.

1.8 A taxonomy of games

Games are classified along several axes, and the right tools for a game depend on where it falls. The book is organized largely around this taxonomy.

Zero-sum versus general-sum. A two-player game is **zero-sum** if $u_1(s) + u_2(s) = 0$ for every profile, so one player's gain is exactly the other's loss. Chess, poker, and the goblet duel are zero-sum. When the payoffs do not cancel, the game is **general-sum**, and the players may have partly aligned interests, as in the Prisoner's Dilemma. Zero-sum games are the subject of Chapters 4 to 7; general-sum games of Chapter 8.

Cooperative versus non-cooperative. In a **non-cooperative** game, players choose independently and cannot make binding agreements; the solution concept is equilibrium. In a **cooperative** game, players may form coalitions and sign binding contracts, and the questions become which coalitions form and how they divide their gains. Cooperative theory occupies Chapters 9 and 10.

Perfect versus imperfect information. A game has **perfect information** if every player, when moving, knows the entire history of prior moves. Chess and Nim have perfect information; the goblet duel and poker do not, because a player moves without seeing a relevant earlier choice. The distinction is the theme of Chapter 3.

Simultaneous versus sequential. In a **simultaneous** game players move at once (the strategic form of Definition 1.2); in a **sequential** game they move in turns, and the natural representation is a tree (Chapter 3).

Impartial versus partizan. Within combinatorial games, a game is **impartial** if the available moves depend only on the position and not on whose turn it is. Nim is impartial; chess, where one player moves only the white pieces, is partizan. Impartial games are the entire subject of Chapter 2.

1.9 Where this goes: design, learning, and engineering

The course that produced this book is titled *Game-Theoretic Design, Learning, and Engineering*, and the three words mark three reasons the theory matters beyond the classic economic examples.

Design is the inverse problem. Instead of analyzing a given game, the engineer designs the rules so that self-interested play produces a desired outcome. Auctions that allocate spectrum to those who value it most, and protocols that keep a shared network stable, are designed games. This is mechanism design, previewed in [Chapter 9](#).

Learning concerns how agents reach equilibrium without computing it directly, by adjusting strategies in response to experience. Fictitious play [\[10, 45\]](#) and best-response dynamics are learning rules whose convergence is a game-theoretic question, and the training of a generative adversarial network [\[23\]](#) is literally a two-player minimax game played by gradient descent.

Engineering is the use of these ideas inside systems: worst-case planning by the minimax principle, resource allocation measured against the price of anarchy [\[28\]](#), and adversarial search in board games, the setting of the course's own term project. Each application reduces, when examined, to one of the four theorems of [Section 1.4](#).

Historical Note: game theory and AI, both at home in Hanover

The course's three themes converge on a local coincidence. The term *artificial intelligence* was coined a few hundred yards from this classroom: the 1956 Dartmouth Summer Research Project on Artificial Intelligence, organized by John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon, gave the field its name and its founding agenda. Two of the four theorems of this book sit squarely inside the modern version of that agenda. A generative adversarial network is literally a two-player minimax game trained by gradient play [\[23\]](#), the learning analogue of [Example 5.15's](#) fictitious play; and the superhuman Go and poker programs of [Chapter 3](#) are engineered searches for the optimal strategies that von Neumann, Nash, and Zermelo proved must exist. Game theory supplies the targets; learning and engineering chase them.

Notes and References

The modern subject begins with von Neumann [\[63\]](#) and the treatise of von Neumann and Morgenstern [\[64\]](#); the equilibrium concept and its existence are due to Nash [\[36, 38\]](#). The Prisoner's Dilemma was devised by Merrill Flood and Melvin Dresher at the RAND Corporation in 1950 and given its prison story by Albert Tucker. The bankruptcy table and its nucleolus reading are from Aumann and Maschler [\[1\]](#), who traced the rule to the contested-garment principle of the Mishnah. The two texts assigned for the course, and the ones this book follows most closely, are Ferguson [\[16\]](#) and Osborne [\[39\]](#); broader graduate treatments are Osborne and Rubinstein [\[40\]](#), Myerson [\[35\]](#), and Maschler et al. [\[32\]](#), and Karlin and Peres [\[26\]](#) is a proof-oriented companion. The applications of [Section 1.9](#) trace to Brown [\[10\]](#), Robinson [\[45\]](#) for learning dynamics, Goodfellow et al. [\[23\]](#) for adversarial networks, and Koutsoupas and Papadimitriou [\[28\]](#) for the price of anarchy.

Chapter summary

A game has players, strategies, and payoffs, and its defining feature is strategic interdependence: the value of an action depends on others' actions. The strategic form ([Definition 1.2](#)) records this as payoff functions over strategy profiles, and a two-player game is displayed as a bimatrix. A mixed strategy is a probability distribution over pure strategies ([Definition 1.4](#)), and randomization is essential, both to escape the infinite regress of [Section 1.3](#) and, as later chapters show, to guarantee that equilibria exist. The first analytic tool is dominance ([Definition 1.6](#)), which alone solves the Prisoner's Dilemma and Golden Balls and exposes the conflict between individual and collective rationality. The general solution concept is the Nash equilibrium ([Definition 1.10](#)), a profile stable against unilateral deviation. Games are classified as zero-sum or general-sum, cooperative or non-cooperative, perfect or imperfect information, impartial or partizan, and the book is organized around this taxonomy, building toward four theorems: Sprague–Grundy, minimax, Nash existence, and the Shapley value.

Exercises

Dominance and equilibrium

Exercise 1.1. In Example 1.7, suppose the jackpot is $J = 100,000$ dollars. Write the payoff matrix with these numbers and verify directly from Definition 1.6 that STEAL weakly dominates SPLIT for each player. Is the dominance strict? Explain why the answer matters for predicting behavior.

Exercise 1.2. For each game, find all strictly and weakly dominated strategies, then determine the outcome by iterated removal of strictly dominated strategies.

		L		R			L		C	R
(a)	T	$(2, 1)$		$(0, 0)$	(b)	T	$(4, 3)$		$(2, 7)$	$(0, 4)$
	B	$(1, 1)$		$(1, 2)$		M	$(5, 5)$		$(5, -1)$	$(-4, -2)$
							B		$(3, 5)$	$(1, 5)$
									$(-1, 6)$	

Exercise 1.3. A pure-strategy Nash equilibrium can be found by the best-response method: in each column, underline Player I's best payoff; in each row, underline Player II's best payoff; a cell with both payoffs underlined is an equilibrium. Apply this to game (a) of Exercise 1.2 and to the Prisoner's Dilemma of Example 1.8. How many pure equilibria does each have?

Exercise 1.4. Consider the coordination game

		A	B
A		$(2, 2)$	$(0, 0)$
B		$(0, 0)$	$(1, 1)$

Show that it has two pure-strategy Nash equilibria. Which do the players prefer, and what does this reveal about the limits of the equilibrium concept as a predictor? (The mixed equilibrium of this game is computed in Chapter 8.)

Reading games and applications

Exercise 1.5. Show that a two-player game with payoff matrices A and B is zero-sum if and only if $B = -A$. For such a game, explain why the single matrix A carries all the information, and state what "Player II minimizes" means in terms of A .

Exercise 1.6. In Example 1.1, compute the proportional allocation $E \cdot c_i/600$ for each of the three estates $E \in \{100, 200, 300\}$. Compare with the Talmud table. For which estate do the two rules agree, and why might that estate be special? (You will prove the full result in Chapter 10.)

Exercise 1.7. *Tragedy of the commons.* Each of n herders chooses how many cattle to graze on a shared pasture. If the total grazed is large, each animal is worth less because the pasture degrades. Model this as a game: identify the players, a reasonable strategy set, and a payoff function in which each herder's gain from one more animal is partly borne by the others. Argue informally that the Nash equilibrium overgrazes relative to the outcome that maximizes total value, and relate the structure to the Prisoner's Dilemma.

Exercise 1.8. In the goblet duel of Section 1.3, suppose Westley poisons goblet A with probability w and Vizzini, who cannot observe the choice, drinks from A with probability d . Write Vizzini's probability of drinking poison as a function of w and d . For which value of w is Vizzini's survival probability the same whatever d he chooses? Interpret this value in light of the claim that randomization defeats deeper reasoning.

CHAPTER 2

Impartial Combinatorial Games

“In Nim there are no secrets and no luck. There is only arithmetic, hidden until you learn to see it.”

We begin the real theory with the simplest games imaginable: two players, perfect information, no chance, and a rule that the same moves are available to whoever is to move. These are the *impartial combinatorial games*. They look like puzzles, but they conceal one of the most satisfying theorems in the subject. The Sprague–Grundy theorem says that every such game, no matter how baroque its rules, is secretly a single pile of Nim, and that the disjunctive sum of games is governed by a binary exclusive-or. By the end of the chapter we will be able to take an unfamiliar game, compute one integer for each position, and read off perfect play.

2.1 Combinatorial games and normal play

Definition 2.1 (Impartial combinatorial game). An *impartial combinatorial game* is a finite set X of **positions** together with, for each position x , a set $F(x) \subseteq X$ of **followers**, the positions reachable in one legal move. Two players alternate moves. The available moves $F(x)$ depend only on the position x and are the same whichever player is to move, which is what *impartial* means. We assume the **ending condition**: from every starting position, play reaches a position with no followers in finitely many moves. Under the **normal-play convention**, the player who cannot move (faces a position x with $F(x) = \emptyset$) loses.

The structure (X, F) is a finite directed acyclic graph: the vertices are positions, the edges are moves, and acyclicity is exactly the ending condition. Because the graph is acyclic, we may analyze it by induction, working backward from terminal positions. The *misère* convention, where the last player to move loses, is discussed in [Section 2.8](#); until then “win” and “lose” refer to normal play.

Definition 2.2 (P-positions and N-positions). A position is a **P-position** if the *previous* player can force a win, that is, the player who just moved into it. A position is an **N-position** if the *next* player to move can force a win. The two classes are defined together by backward induction:

1. every terminal position (no followers) is a P-position;
2. a position is an N-position if at least one follower is a P-position;
3. a position is a P-position if every follower is an N-position.

The labels record perfect play. From an N-position the player to move wins by moving to a P-position and repeating; from a P-position every move hands the opponent an N-position, so the player to move loses. The entire problem of solving an impartial game is to determine which positions are P-positions.

2.2 The take-away game

The smallest interesting example has a single pile.

Example 2.3 (Subtract one, two, or three). A pile holds n chips. A move removes 1, 2, or 3 chips. The player who takes the last chip wins. Which n are P-positions?

Compute by induction from $n = 0$. The empty pile is terminal, hence a P-position. A pile is an N-position if some legal move reaches a P-position. The pattern is immediate:

n	0	1	2	3	4	5	6	7	8
type	P	N	N	N	P	N	N	N	P

A pile is a P-position exactly when n is a multiple of 4. The winning strategy from an N-position is to remove $n \bmod 4$ chips, leaving a multiple of 4; whatever the opponent removes (between 1 and 3), you restore a multiple of 4, and eventually you leave 0.

The reasoning generalizes at once. Removing any number from 1 to k partitions the integers into residue classes modulo $k + 1$.

Proposition 2.4 (Subtraction game modulo $k + 1$). In the one-pile game where a move removes between 1 and k chips, under normal play the P-positions are exactly the multiples of $k + 1$:

$$n \text{ is a P-position} \iff n \equiv 0 \pmod{k+1}.$$

Proof. We argue by strong induction on n , verifying the two defining properties of P-positions. The base case $n = 0$ is a multiple of $k + 1$ and is terminal, hence a P-position.

Suppose the claim holds for all piles smaller than n .

If $n \equiv 0 \pmod{k+1}$, then any legal move removes some r with $1 \leq r \leq k$, leaving $n - r$. Since $1 \leq r \leq k$, we have $n - r \not\equiv 0 \pmod{k+1}$, so by the induction hypothesis $n - r$ is an N-position. Every follower is an N-position, so n is a P-position, as required.

If $n \not\equiv 0 \pmod{k+1}$, write $n \equiv r \pmod{k+1}$ with $1 \leq r \leq k$. Removing r chips is legal and leaves $n - r \equiv 0 \pmod{k+1}$, which by the induction hypothesis is a P-position. Some follower is a P-position, so n is an N-position.

Both properties hold, completing the induction. ■

Computational Companion

Problem Set 1 opens with the case $k = 6$: a move removes from 1 to 6 chips. By [Proposition 2.4](#) the P-positions are the multiples of 7. From a pile of 31, since $31 \equiv 3 \pmod{7}$, the winning move removes 3 chips, leaving $28 = 4 \cdot 7$. The assignment also asks for the P-positions under three other subtraction sets, $S_1 = \{1, 3, 5, 7\}$, $S_2 = \{1, 3, 6\}$, and $S_3 = \{1, 2, 4, 8, 16, \dots\}$, whose answers are worked in full in [Chapter 11](#).

2.3 Nim and Bouton's theorem

The central example of the whole theory is Nim.

Definition 2.5 (Nim). *Nim* is played with several piles of chips. A position is a tuple $(x_1, x_2, \dots, x_n)$ of nonnegative pile sizes. A move chooses one pile and removes at least one chip from it (possibly the whole pile). Under normal play, the player who removes the last chip wins, equivalently the player facing $(0, \dots, 0)$ loses.

The solution rests on a binary operation.

Definition 2.6 (Nim-sum). The **nim-sum** of nonnegative integers, written $\oplus$, is bitwise addition modulo 2 (the exclusive-or of the binary representations, with no carrying). To compute it, write the numbers in binary, stack them

with their bits aligned by place value, and in each column write a 1 if that column contains an odd number of ones and a 0 if it contains an even number. For example, to compute $5 \oplus 4 \oplus 13$, write $5 = (0101)_2$, $4 = (0100)_2$, and $13 = (1101)_2$ and stack them:

$$\begin{array}{r}
 8\ 4\ 2\ 1 \\
 \hline
 5 = 0\ 1\ 0\ 1 \\
 4 = 0\ 1\ 0\ 0 \\
 13 = 1\ 1\ 0\ 1 \\
 \hline
 \text{ones per column: } 1\ 3\ 0\ 2 \\
 \text{parity (the XOR): } 1\ 1\ 0\ 0
 \end{array}$$

Reading the parity row, $5 \oplus 4 \oplus 13 = (1100)_2 = 8 + 4 = 12$. Column by column: the 8s column has one 1 (odd, write 1), the 4s column has three 1s (odd, write 1), the 2s column has no 1s (even, write 0), and the 1s column has two 1s (even, write 0). There is no carrying between columns, which is what distinguishes nim-sum from ordinary addition.

The nim-sum is associative and commutative, every element is its own inverse ($a \oplus a = 0$), and 0 is the identity. These are the laws of the additive group of a vector space over the two-element field, and they are the only algebra the next proof needs.

Theorem 2.7 (Bouton, 1901). A Nim position $(x_1, \dots, x_n)$ is a P-position if and only if its nim-sum is zero:

$$x_1 \oplus x_2 \oplus \dots \oplus x_n = 0.$$

Proof. Write $S = x_1 \oplus \dots \oplus x_n$. We verify the two defining properties of P-positions, namely that every position with $S = 0$ has all followers with $S \neq 0$, and that every position with $S \neq 0$ has at least one follower with $S = 0$. Together with the fact that the terminal position $(0, \dots, 0)$ has $S = 0$, induction on the total number of chips then identifies the $S = 0$ positions as exactly the P-positions.

From $S = 0$, every move gives $S \neq 0$. A move replaces some x_i by a strictly smaller $y_i < x_i$, leaving the other piles unchanged. The new nim-sum is

$$S' = S \oplus x_i \oplus y_i = 0 \oplus x_i \oplus y_i = x_i \oplus y_i.$$

Because $x_i \neq y_i$, their nim-sum is nonzero, so $S' \neq 0$.

From $S \neq 0$, some move gives $S = 0$. Let k be the position of the leftmost (most significant) nonzero bit of S . Since that bit of S is 1, an odd number of the piles have a 1 in bit k ; in particular at least one does, say pile i . Set

$$y_i = x_i \oplus S.$$

We claim $y_i < x_i$. Compare x_i and y_i bit by bit, from the most significant down. Above bit k , S has only zeros, so x_i and y_i agree there. At bit k , x_i has a 1 and S has a 1, so y_i has a 0. Thus y_i first differs from x_i at bit k , where it is smaller, giving $y_i < x_i$. Removing $x_i - y_i$ chips from pile i is therefore a legal move, and the resulting nim-sum is

$$S \oplus x_i \oplus y_i = S \oplus x_i \oplus (x_i \oplus S) = 0.$$

So this move reaches a position with nim-sum 0.

Both properties hold, and the induction is complete. ■

Remark 2.8 (Intuition: balancing the binary columns). Write the pile sizes in binary and stack them in a table, one pile per row. The nim-sum is zero exactly when every column holds an even number of ones, a perfectly *balanced* position. Bouton's theorem says the balanced positions are the losing ones: any move you make unbalances some column, and a skilled opponent rebalances it on the next turn. The winning move from an unbalanced position restores balance by adjusting the one pile that controls the leftmost unbalanced column, which is why the recipe changes exactly one pile.

Recipe

To win at Nim from an N-position. Compute the nim-sum $S = x_1 \oplus \dots \oplus x_n$. Locate the leftmost bit in which S equals 1, and pick any pile x_i whose binary form also has a 1 in that bit. Reduce that pile to $y_i = x_i \oplus S$, which is always smaller than x_i . The resulting position has nim-sum 0. Repeat whenever the opponent returns a nonzero nim-sum, and you will take the last chip.

Example 2.9 (Reading off a winning move). Consider Nim with piles (9, 17, 21). Stack the binary forms with bits aligned by place value (16, 8, 4, 2, 1) and take the column parities:

$$\begin{array}{r}
 16 \ 8 \ 4 \ 2 \ 1 \\
 \hline
 9 = 0 \ 1 \ 0 \ 0 \ 1 \\
 17 = 1 \ 0 \ 0 \ 0 \ 1 \\
 21 = 1 \ 0 \ 1 \ 0 \ 1 \\
 \hline
 S = 0 \ 1 \ 1 \ 0 \ 1
 \end{array}$$

The red columns (8, 4, 1) are the *unbalanced* ones, the places where an odd number of piles carry a 1; the balanced columns (16, 2) are already even and need no fixing. The nim-sum collects exactly the unbalanced columns: $S = 9 \oplus 17 \oplus 21 = (01101)_2 = 8 + 4 + 1 = 13 \neq 0$, an N-position. To win, the recipe says: find a pile x_i with $y_i = x_i \oplus S < x_i$ and reduce it to y_i . The leftmost 1 of $S = (01101)_2$ is the 8s bit, so we need a pile with a 1 in its 8s bit. Only $9 = (01001)_2$ qualifies. Test all three anyway:

$$21 \oplus 13 = 24 > 21 \text{ (illegal)}, \quad 17 \oplus 13 = 28 > 17 \text{ (illegal)}, \quad 9 \oplus 13 = 4 < 9 \text{ (legal)}.$$

The one legal reduction is $9 \rightarrow 4$, computed bit by bit as $9 \oplus 13 = (01001)_2 \oplus (01101)_2 = (00100)_2 = 4$, so the winning move removes $9 - 4 = 5$ chips from the first pile. To confirm, the new position (4, 17, 21) has nim-sum

$$4 \oplus 17 \oplus 21 = (00100)_2 \oplus (10001)_2 \oplus (10101)_2 = (00000)_2 = 0,$$

a P-position, so the opponent is now lost. ◇

Historical Note

Because Bouton's winning strategy is a fixed binary computation, Nim was a natural showcase for early computing machinery. The *Nimatron*, an electromechanical relay machine conceived by the physicist Edward Condon and built by Westinghouse, played Nim against visitors at the 1939–1940 New York World's Fair and won the large majority of its games. A decade later the *Nimrod*, built by Ferranti and designed by John Bennett, played Nim at the 1951 Festival of Britain to demonstrate that a digital computer could be programmed for a decision task. Both machines implemented exactly the nim-sum recipe above: read the heaps, compute the exclusive-or, and move to restore a zero nim-sum.

Computational Companion

The first programming task of Problem Set 1 is exactly the construction in the proof of Theorem 2.7: given pile sizes, compute the nim-sum; if it is zero, report a loss; otherwise return every winning move by testing $y_i = x_i \oplus S < x_i$ for each pile. For (1, 2, 3, 4, 5) the nim-sum is 1, and the three winning moves reduce pile 1 to 0, pile 3 to 2, or pile 5 to 4. The full solver appears in Appendix C.

2.4 The Sprague–Grundy function

Bouton's theorem solves Nim. The Sprague–Grundy theory solves everything else, by attaching to each position of any impartial game a number that behaves like a Nim pile size.

Definition 2.10 (Minimum excludant). The **minimum excludant** of a set T of nonnegative integers, written $\text{mex}(T)$, is the smallest nonnegative integer not in T :

$$\text{mex}(T) = \min\{n \geq 0 : n \notin T\}.$$

For instance $\text{mex}(\emptyset) = 0$, $\text{mex}\{0, 1, 3\} = 2$, and $\text{mex}\{0, 1, 2, 3\} = 4$.

Definition 2.11 (Sprague–Grundy function). The **Sprague–Grundy function** of an impartial game (X, F) is the function $g : X \rightarrow \{0, 1, 2, \dots\}$ defined by backward induction,

$$g(x) = \text{mex}\{g(y) : y \in F(x)\}. \tag{2.1}$$

In words, $g(x)$ is the smallest nonnegative integer that is not the Grundy value of any follower of x . A terminal position, having no followers, satisfies $g(x) = \text{mex}(\emptyset) = 0$. The value $g(x)$ is also called the **Grundy value** or **number** of x .

The Grundy function refines the P/N classification, and reduces to it through a single clean statement.

Theorem 2.12 (Grundy value and P-positions). For every position x of an impartial game,

$$x \text{ is a P-position} \iff g(x) = 0.$$

Proof. We check the two defining properties of P-positions for the class $\{x : g(x) = 0\}$, proceeding by induction along the acyclic move graph.

If $g(x) = 0$, then by (2.1) the value 0 is excluded from the followers' Grundy values, so no follower y has $g(y) = 0$. Hence every follower has $g(y) > 0$.

If $g(x) > 0$, then 0 is not excluded, meaning 0 does appear among the followers' Grundy values, so some follower y has $g(y) = 0$.

Thus positions with $g = 0$ have all followers with $g > 0$, and positions with $g > 0$ have some follower with $g = 0$. Since terminal positions have $g = 0$ and are P-positions, the inductive definition of P-positions (Definition 2.2) identifies $\{g = 0\}$ exactly with the P-positions. ■

The Grundy value carries more information than the single bit “P or N.” Its meaning is that the position plays exactly like a Nim pile of $g(x)$ chips, a claim the next section makes precise and exploits.

Example 2.13 (Computing a Grundy function by the mex recurrence). Take the subtraction game with set $\{1, 2, 3\}$ from Example 2.3 and compute its Grundy function from the ground up, one value at a time, using $g(n) = \text{mex}\{g(n-1), g(n-2), g(n-3)\}$ (dropping any term with a negative argument, since those moves are illegal).

$$\begin{aligned} g(0) &= \text{mex}\{\} = 0 && \text{(terminal: no followers)} \\ g(1) &= \text{mex}\{g(0)\} = \text{mex}\{0\} = 1 && \text{(smallest missing from } \{0\}) \\ g(2) &= \text{mex}\{g(1), g(0)\} = \text{mex}\{1, 0\} = 2 \\ g(3) &= \text{mex}\{g(2), g(1), g(0)\} = \text{mex}\{2, 1, 0\} = 3 \\ g(4) &= \text{mex}\{g(3), g(2), g(1)\} = \text{mex}\{3, 2, 1\} = 0 && (0 \text{ is now missing}) \\ g(5) &= \text{mex}\{g(4), g(3), g(2)\} = \text{mex}\{0, 3, 2\} = 1 \\ g(6) &= \text{mex}\{g(5), g(4), g(3)\} = \text{mex}\{1, 0, 3\} = 2 \\ g(7) &= \text{mex}\{g(6), g(5), g(4)\} = \text{mex}\{2, 1, 0\} = 3 \\ g(8) &= \text{mex}\{g(7), g(6), g(5)\} = \text{mex}\{3, 2, 1\} = 0. \end{aligned}$$

The values repeat with period 4: $g(n) = n \bmod 4$. The P-positions, where $g = 0$, are the multiples of 4, matching the P/N table of Example 2.3; but the Grundy function tells us more, namely the exact Nim pile each position imitates, which is what lets us add this game to others in the next section. ◇

The recurrence works on *any* finite directed acyclic graph of positions, not just piles: label the sinks 0 and apply mex to each vertex once all its followers are labelled, working back along the arrows. Figure 2.1 carries this out on a small game graph.




Figure 2.1. Sprague–Grundy values on a directed acyclic game graph (arrows point to followers). The two sinks D, F get $g = 0$. Then $g(C) = \text{mex}\{g(D), g(F)\} = \text{mex}\{0, 0\} = 1$, $g(E) = \text{mex}\{g(F)\} = 1$, $g(B) = \text{mex}\{g(C), g(E)\} = \text{mex}\{1, 1\} = 0$, and $g(A) = \text{mex}\{g(B), g(C), g(D)\} = \text{mex}\{0, 1, 0\} = 2$. The P-positions (shaded, $g = 0$) are B, D, F .

2.5 Sums of games and the Sprague–Grundy theorem

Many games are really several independent games played side by side. On your turn you move in exactly one of them.

Definition 2.14 (Disjunctive sum). The **disjunctive sum** $G = G_1 + G_2 + \cdots + G_n$ of impartial games has positions $(x_1, \dots, x_n)$, one from each component. A legal move selects one component i and replaces x_i by a follower $x'_i \in F_i(x_i)$, leaving the other components unchanged. Play ends when no component admits a move.

Nim itself is the disjunctive sum of one-pile Nim games, where a single pile of size m has Grundy value $g(m) = m$ (its followers are the piles $0, 1, \dots, m-1$, whose Grundy values are $0, 1, \dots, m-1$, and $\text{mex}\{0, \dots, m-1\} = m$). The following theorem says that the Grundy value of a sum is the nim-sum of the Grundy values, which immediately recovers Bouton's theorem and extends it to all impartial games.

Theorem 2.15 (Sprague–Grundy). If g_i is the Sprague–Grundy function of G_i for $i = 1, \dots, n$, then the disjunctive sum $G = G_1 + \cdots + G_n$ has Sprague–Grundy function

$$g(x_1, \dots, x_n) = g_1(x_1) \oplus g_2(x_2) \oplus \cdots \oplus g_n(x_n).$$

Every impartial game is a Nim heap

The Sprague–Grundy theorem reduces all of impartial game theory to Nim. A position x of any impartial game plays exactly like a single Nim pile of $g(x)$ chips, and a disjunctive sum plays like the Nim position with one such pile per component. To analyze a sum, replace each component by its Grundy value and take the nim-sum; the whole is a P-position precisely when that nim-sum is zero. Bouton's theorem (Theorem 2.7) is the special case where each component is already a Nim pile, since $g(m) = m$.

Proof. It suffices to treat the sum of two games, since the general case follows by induction on the number of components. Fix a position $x = (x_1, \dots, x_n)$ and set

$$b = g_1(x_1) \oplus \cdots \oplus g_n(x_n).$$

By Definition 2.11 we must show two things about the value $g(x)$, after which the equality $g(x) = b$ follows because $g(x)$ is by definition the mex of its followers' values:

1. for every nonnegative integer $a < b$, some follower of x has Grundy value a ;
2. no follower of x has Grundy value b .

Proof of (1). Let $d = a \oplus b$. Because $a < b$, the leftmost bit where a and b differ is a position k where b has a 1 and a has a 0; this is the leftmost bit of d , so d has a 1 in bit k and zeros above it. Since $b = g_1(x_1) \oplus \cdots \oplus g_n(x_n)$ has a 1 in bit k , some component, say $g_1(x_1)$, has a 1 in bit k . Then $d \oplus g_1(x_1) < g_1(x_1)$, by the same bit-comparison argument used in Bouton's theorem. By the definition of g_1 as a mex, every value smaller than $g_1(x_1)$ occurs among the Grundy values of the followers of x_1 , so there is a follower x'_1 with $g_1(x'_1) = d \oplus g_1(x_1)$. Moving the first component from x_1 to x'_1 is legal in G , and the resulting Grundy value is

$$g_1(x'_1) \oplus g_2(x_2) \oplus \cdots \oplus g_n(x_n) = d \oplus g_1(x_1) \oplus g_2(x_2) \oplus \cdots \oplus g_n(x_n) = d \oplus b = a.$$

So a follower with Grundy value a exists.

Proof of (2). Suppose, for contradiction, that some follower of x has Grundy value b . A move changes a single component, say the first, from x_1 to x'_1 , giving

$$g_1(x'_1) \oplus g_2(x_2) \oplus \cdots \oplus g_n(x_n) = b = g_1(x_1) \oplus g_2(x_2) \oplus \cdots \oplus g_n(x_n).$$

Cancelling the common terms $g_2(x_2), \dots, g_n(x_n)$ (each is its own inverse under $\oplus$) yields $g_1(x'_1) = g_1(x_1)$. But a position can never have a follower of the same Grundy value, since $g_1(x_1) = \text{mex}\{g_1(y) : y \in F_1(x_1)\}$ excludes its own value. This contradiction proves (2).

By (1) and (2), the followers of x realize every value below b but never b itself, so $g(x) = \text{mex}\{\text{followers' values}\} = b$. ■

Example 2.16 (A sum of three subtraction games). Play three subtraction games at once. In the first, the subtraction set is $\{1, 2, 3\}$ and the pile has 9 chips; in the second, $\{1, 2, 3, 4, 5\}$ with 10 chips; in the third, $\{1, 2, 3, 4, 5, 6, 7\}$ with 14 chips. By Proposition 2.4 the Grundy value of a subtract- $\{1, \dots, k\}$ pile of size m is $m \bmod (k+1)$, so the three component values are

$$g_1 = 9 \bmod 4 = 1, \quad g_2 = 10 \bmod 6 = 4, \quad g_3 = 14 \bmod 8 = 6.$$

By the Sprague-Grundy theorem the whole position behaves like Nim on piles $(1, 4, 6)$. Take the nim-sum in binary:

$$\begin{array}{r} 4 \ 2 \ 1 \\ 1 = 0 \ 0 \ 1 \\ 4 = 1 \ 0 \ 0 \\ 6 = 1 \ 1 \ 0 \\ \hline S = 0 \ 1 \ 1 \end{array} \quad S = g_1 \oplus g_2 \oplus g_3 = (011)_2 = 3 \neq 0,$$

so the position is an N-position. To win, change one component so the three Grundy values nim-sum to 0: component i should move to the target value $g_i \oplus S$, which works whenever that target is actually reachable in component i . Check each:

- Component 1: target $g_1 \oplus S = 1 \oplus 3 = 2$. From the pile of 9 the legal moves reach piles 8, 7, 6 with Grundy values 0, 3, 2; the target 2 is among them (pile 6). So removing 3 chips from the first pile, $9 \rightarrow 6$, is a winning move, after which $2 \oplus 4 \oplus 6 = 0$.
- Component 3: target $g_3 \oplus S = 6 \oplus 3 = 5$. From the pile of 14 the legal moves reach Grundy values 5, 4, 3, 2, 1, 0, 7; the target 5 is among them (pile 13, since $13 \bmod 8 = 5$). So removing 1 chip from the third pile, $14 \rightarrow 13$, is a winning move, after which $1 \oplus 4 \oplus 5 = 0$.
- Component 2: target $g_2 \oplus S = 4 \oplus 3 = 7$, but the second game's Grundy values only range over $0, \dots, 5$, so 7 is unreachable and this component offers no winning move.

There are therefore exactly two winning moves: remove 3 from the first pile, or remove 1 from the third. $\diamond$

Example 2.17 (A sum of three different games). The Sprague-Grundy theorem shines brightest when the parts are unlike one another. Play three games at once: a *Nim heap* of 5 chips, from which any number may be removed; a *subtraction game* with set $\{1, 2, 3\}$ on a pile of 10; and a subtraction game with set $\{2, 3\}$ on a pile of 8. Each collapses to a single Nim pile through its Grundy value.

Component A, the Nim heap of 5. A Nim pile of size m has Grundy value m , so $g_A = 5$.

Component B, subtract $\{1, 2, 3\}$ from 10. By Proposition 2.4, $g_B = 10 \bmod 4 = 2$.

Component C, subtract $\{2, 3\}$ from 8. This set has no one-step modulus, so run the mex recurrence $g(n) = \text{mex}\{g(n-2), g(n-3)\}$ from the bottom up:

$$\begin{array}{c|cccccccc} n & 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 \\ \hline g(n) & 0 & 0 & 1 & 1 & 2 & 0 & 0 & 1 & 1 \end{array}$$

The values settle into the repeating block 0, 0, 1, 1, 2 of length five, and $g_C = g(8) = 1$.

Combine. By Theorem 2.15 the whole position plays exactly like Nim on the piles $(5, 2, 1)$. Add the three Grundy values in binary without carrying:

$$\begin{array}{r} 4 \ 2 \ 1 \\ 5 = 1 \ 0 \ 1 \\ 2 = 0 \ 1 \ 0 \\ 1 = 0 \ 0 \ 1 \\ \hline S = 1 \ 1 \ 0 \end{array} \quad S = g_A \oplus g_B \oplus g_C = (110)_2 = 6 \neq 0,$$

so the starting position is an N-position. To win, drive the nim-sum to zero by moving one component to its target value $g_i \oplus S$:

- Component A: target $5 \oplus 6 = 3$. A Nim heap reaches every smaller Grundy value, and $3 < 5$, so reduce the heap from 5 chips to 3. The new triple $(3, 2, 1)$ has $3 \oplus 2 \oplus 1 = 0$.
- Component B: target $2 \oplus 6 = 4$, but subtract- $\{1, 2, 3\}$ Grundy values never exceed 3, so 4 is out of reach and B offers no winning move.

- Component C: target $1 \oplus 6 = 7$, while subtract- $\{2, 3\}$ values never exceed 2, so C offers none either.

The position has exactly one winning move: remove two chips from the Nim heap. Three games with three different rulebooks collapse to one binary sum, and the whole analysis comes down to three Grundy values and a single exclusive-or. $\diamond$

2.6 A catalogue of impartial games

With Theorem 2.15 in hand, solving a game means tabulating its one-pile Grundy function and then nim-summing. We record the Grundy functions of the games studied in lecture. Each is computed by the recurrence (2.1), where a move that splits a pile into two creates a sum, whose value is the nim-sum of the parts.

2.6.1 Even if not all, all if odd

A move removes any even number of chips, provided it is not the whole pile, or removes the whole pile provided the pile has an odd number of chips. The terminal positions are 0 and 2: from 0 there is nothing to remove, and from 2 the only even number below the pile is none (removing 2 would be the whole pile, which is barred because 2 is even), so neither has a move. The Grundy values are

x	0	1	2	3	4	5	6	7	8	9	10
$g(x)$	0	1	0	2	1	3	2	4	3	5	4

To see where these come from, apply the recurrence. From $x = 3$ (odd) the legal moves are “remove 2” (reaching 1) and “remove all 3” (reaching 0), so $g(3) = \text{mex}\{g(1), g(0)\} = \text{mex}\{1, 0\} = 2$. From $x = 4$ (even) the only legal move is “remove 2” (reaching 2; removing all 4 is barred), so $g(4) = \text{mex}\{g(2)\} = \text{mex}\{0\} = 1$. From $x = 5$ (odd) the moves are “remove 2” $\rightarrow 3$, “remove 4” $\rightarrow 1$, and “remove all 5” $\rightarrow 0$, so $g(5) = \text{mex}\{g(3), g(1), g(0)\} = \text{mex}\{2, 1, 0\} = 3$. Continuing this way fills the table, and the pattern settles into $g(2k) = k - 1$ and $g(2k - 1) = k$ for $k \geq 1$.

2.6.2 Lasker’s Nim

Each move either removes any number of chips from one pile, as in Nim, or splits one pile of size at least 2 into two nonempty piles without removing any chips. The one-pile Grundy values are

x	0	1	2	3	4	5	6	7	8	9	10	11	12
$g(x)$	0	1	2	4	3	5	6	8	7	9	10	12	11

The split move is what makes Lasker’s Nim more than ordinary Nim, and it is handled by the Sprague–Grundy theorem: a pile split into parts a and b becomes the *sum* of two piles, so its value is $g(a) \oplus g(b)$. Take $x = 3$. The Nim-style removals reach piles 2, 1, 0 with values 2, 1, 0, and the one split $3 = 1 + 2$ contributes $g(1) \oplus g(2) = 1 \oplus 2 = 3$. The follower values are $\{2, 1, 0, 3\}$, so $g(3) = \text{mex}\{0, 1, 2, 3\} = 4$. Take $x = 4$. The removals reach 3, 2, 1, 0 with values 4, 2, 1, 0; the split $4 = 1 + 3$ gives $g(1) \oplus g(3) = 1 \oplus 4 = 5$; and the split $4 = 2 + 2$ gives $g(2) \oplus g(2) = 2 \oplus 2 = 0$. The follower values are $\{0, 1, 2, 4, 5\}$, missing 3, so $g(4) = 3$. These satisfy the closed form, valid for all $k \geq 0$,

$$\begin{aligned} g(4k + 1) &= 4k + 1, & g(4k + 2) &= 4k + 2, \\ g(4k + 3) &= 4k + 4, & g(4k + 4) &= 4k + 3. \end{aligned}$$

The values equal x except that each pair $(4k + 3, 4k + 4)$ is swapped, so the sequence reads

$$0, 1, 2, \mathbf{4}, \mathbf{3}, 5, 6, \mathbf{8}, \mathbf{7}, 9, 10, \mathbf{12}, \mathbf{11}, \dots$$

with the swapped pairs in bold. This game appears in Problem Set 2 in the form “remove any number, or split a pile into two,” which is the same rule and gives the same Grundy function.

2.6.3 Kayles

Imagine a row of bowling pins. A move removes one or two adjacent pins, which may split the remaining pins into two separated groups. As a chip game: remove one or two chips from a pile, then optionally split that pile into two nonempty piles. The Grundy values are

x	0	1	2	3	4	5	6	7	8	9	10	11
$g(x)$	0	1	2	3	1	4	3	2	1	4	2	6

The Kayles sequence is eventually periodic with period 12, repeating from $x = 71$ onward (the last value that breaks the pattern sits at $x = 70$). This is a recurring phenomenon: many take-and-break games are eventually periodic, a fact that lets a finite computation settle the game for all sizes.

2.6.4 Dawson's Chess

The rules are: remove one chip if it is the whole pile; or remove two chips from any pile; or remove three chips from any pile and, if desired, split that pile into two nonempty parts. The Grundy values are

x	0	1	2	3	4	5	6	7	8	9	10	11	12	13
$g(x)$	0	1	1	2	0	3	1	1	0	3	3	2	2	4

To see one entry in detail, take $x = 5$. Its followers are: remove two chips to reach 3 with $g(3) = 2$; remove three chips to reach 2 with $g(2) = 1$; or remove three chips and split the remaining two into $(1, 1)$ with value $g(1) \oplus g(1) = 0$. The followers' values are $\{2, 1, 0\}$, so $g(5) = \text{mex}\{0, 1, 2\} = 3$. Dawson's Chess is eventually periodic with period 34.

2.6.5 Grundy's game

A single move splits one pile into two piles of *different* sizes. No chips are removed. Piles of size 0, 1, and 2 are terminal, because they cannot be split into two unequal positive parts. The Grundy values are

x	0	1	2	3	4	5	6	7	8	9	10
$g(x)$	0	0	0	1	0	2	1	0	2	1	0

For example $g(5) = \text{mex}\{g(1) \oplus g(4), g(2) \oplus g(3)\} = \text{mex}\{0 \oplus 0, 0 \oplus 1\} = \text{mex}\{0, 1\} = 2$. Unlike the previous games, Grundy's game is *not* known to be eventually periodic, and whether it ever becomes periodic is a famous open problem, tested without resolution into the billions.

Computational Companion

The final task of Problem Set 1 computes the Grundy function for $x \leq 100$ for three of these games, Kayles, Dawson's Chess, and Grundy's game, and asks whether a period appears. The computation confirms period 12 for Kayles and period 34 for Dawson's Chess, and finds no period for Grundy's game within range. The memoized recursion is in [Appendix C](#).

2.7 Wythoff's game and the golden ratio

One impartial game stands apart because its P-positions, instead of falling into a residue class or a nim-sum condition, are governed by the golden ratio. It is a two-pile game with an unusual extra move.

Definition 2.18 (Wythoff's game). *Wythoff's game* is played with two piles. A move is one of: remove any positive number of chips from *one* pile (an ordinary Nim move), or remove the *same* positive number from *both* piles at once. The player taking the last chip (reaching $(0, 0)$) wins.

The diagonal move couples the two piles, so the disjunctive theory of Section 2.5 does not apply, and we find the P-positions directly. Recall that a position is a P-position iff every move leads to an N-position, and an N-position iff some move leads to a P-position (Definition 2.2).

Example 2.19 (Building the P-positions by hand). List positions (a, b) with $a \leq b$ (the game is symmetric in the two piles) and sweep upward, marking each as P exactly when no single move reaches an earlier P-position. The cold (P) positions emerge in pairs. The first is $(0, 0)$, terminal. The smallest unsolved position is $(1, 2)$. Its legal followers are: remove from the larger pile to reach $(1, 1)$ or $(1, 0)$; remove the whole smaller pile to reach $(0, 2)$; or remove one from each (the diagonal move) to reach $(0, 1)$. So the follower set is $\{(0, 2), (1, 1), (1, 0), (0, 1)\}$. A position is a P-position only if none of its followers is a P-position, and the only P-position found so far is $(0, 0)$; since $(0, 0)$ is not among the followers of $(1, 2)$, the position $(1, 2)$ is itself a P-position. Continuing, the P-positions with $a \leq b$ are

$$(0, 0), (1, 2), (3, 5), (4, 7), (6, 10), (8, 13), (9, 15), (11, 18), \dots$$

There is a clean rule generating them. Writing the k -th pair as (a_k, b_k) ,

$$a_k = \text{mex}\{a_0, b_0, a_1, b_1, \dots, a_{k-1}, b_{k-1}\}, \quad b_k = a_k + k,$$

so a_k is the smallest nonnegative integer not yet used in any earlier pair, and the gap $b_k - a_k = k$ grows by one each time. For instance $a_2 = \text{mex}\{0, 0, 1, 2\} = 3$ and $b_2 = 3 + 2 = 5$; then $a_3 = \text{mex}\{0, 1, 2, 3, 5\} = 4$ and $b_3 = 4 + 3 = 7$. $\diamond$

Theorem 2.20 (Wythoff, 1907). The P-positions of Wythoff's game are exactly the pairs

$$(a_k, b_k) = (\lfloor k\varphi \rfloor, \lfloor k\varphi^2 \rfloor), \quad k = 0, 1, 2, \dots,$$

where $\varphi = (1 + \sqrt{5})/2$ is the golden ratio. Because $\varphi^2 = \varphi + 1$, the second coordinate satisfies $\lfloor k\varphi^2 \rfloor = \lfloor k\varphi \rfloor + k$, matching the gap rule $b_k = a_k + k$.

The two sequences $\lfloor k\varphi \rfloor$ and $\lfloor k\varphi^2 \rfloor$ are *Beatty sequences*: since $\frac{1}{\varphi} + \frac{1}{\varphi^2} = 1$, every positive integer appears in exactly one of them, which is why the mex construction never repeats a value and the P-positions tile the first quadrant along two golden rays.

Proof of Theorem 2.20. Two facts combine: the two sequences partition the positive integers, and the greedy construction of the cold positions reproduces them.

The sequences partition $\{1, 2, 3, \dots\}$ (Beatty's theorem). For any positive integer M the numbers M/φ and M/φ^2 are irrational and sum to M , so their floors sum to $M - 1$. Using $\lfloor k\alpha \rfloor \leq N \iff k\alpha < N + 1$ for irrational $k\alpha$, the count of sequence values not exceeding N is

$$\#\{k \geq 1 : \lfloor k\varphi \rfloor \leq N\} + \#\{k \geq 1 : \lfloor k\varphi^2 \rfloor \leq N\} = \left\lfloor \frac{N+1}{\varphi} \right\rfloor + \left\lfloor \frac{N+1}{\varphi^2} \right\rfloor = N.$$

Subtracting the same count for $N - 1$ leaves $N - (N - 1) = 1$: exactly one sequence value equals N . Every positive integer therefore appears in precisely one of the two sequences.

The greedy cold positions are the Beatty pairs. The cold positions are built greedily by $a_k = \text{mex}\{a_0, b_0, \dots, a_{k-1}, b_{k-1}\}$ and $b_k = a_k + k$. We show by induction that $a_k = \lfloor k\varphi \rfloor$ and $b_k = \lfloor k\varphi^2 \rfloor$. Because $\varphi^2 = \varphi + 1$,

$$\lfloor k\varphi^2 \rfloor = \lfloor k\varphi + k \rfloor = \lfloor k\varphi \rfloor + k,$$

so the gap rule $b_k = a_k + k$ already matches the Beatty pair. Assume the claim for all indices below k . Every integer smaller than $\lfloor k\varphi \rfloor$ lies in one of the two sequences by the partition, so it equals some earlier a_i or b_i (its index is below k because both sequences increase) and is already used; meanwhile $\lfloor k\varphi \rfloor$ itself, belonging to the first sequence alone, has not yet appeared. The smallest unused integer is thus $\lfloor k\varphi \rfloor$, which is exactly what mex returns, so $a_k = \lfloor k\varphi \rfloor$ and $b_k = a_k + k = \lfloor k\varphi^2 \rfloor$. The greedy cold positions and the golden-ratio pairs coincide. $\blacksquare$

Figure 2.2 plots these cold positions; they fall along two rays from the origin, of slopes φ and $1/\varphi$.

Example 2.21 (A winning move in Wythoff's game). From $(5, 8)$, is the mover winning? The nearby P-positions are $(3, 5)$, $(4, 7)$, $(6, 10)$, and $(8, 13)$; since $(5, 8)$ is not among them, it is an N-position, so a winning move exists. Search the reachable P-positions: removing from the larger pile, $(5, 8) \rightarrow (5, b)$ can reach $(5, 3)$ (a P-position, since $(3, 5)$ is) by taking 5 from the second pile, $8 \rightarrow 3$. Indeed $(5, 3)$ is the mirror of $(3, 5)$, a cold position. So the winning move removes 5 chips from the pile of 8, leaving $(5, 3)$, from which the opponent is lost. $\diamond$




Figure 2.2. The P-positions (cold positions) of Wythoff's game, marked in red, lie along the two golden rays $b = \varphi a$ and $b = a/\varphi$. From any other (hot) position the player to move can reach a cold one; from a cold position every move leaves a hot one.

2.8 Misère play

Reversing the winning condition, so that the player who takes the last chip *loses*, seems like a small change, but it can complicate the theory dramatically. For Nim, however, the answer is clean.

Theorem 2.22 (Misère Nim). In misère Nim, where the player who takes the last chip loses, a position is a P-position if and only if either

1. every pile has size at most 1 and the number of nonempty piles is *odd*; or
2. at least one pile has size at least 2 and the nim-sum of all piles is 0.

Proof sketch. While some pile has size at least 2, a player can always keep the play identical to normal Nim, steering by the nim-sum, because the presence of a large pile gives enough freedom to both zero the nim-sum and avoid handing the opponent the last chip. The strategies of normal and misère play therefore coincide until the position consists only of piles of size 1. At that point the game is a pure parity count: with all piles of size 1, the player to move is forced to take exactly one chip, and the player who faces an odd number of single chips takes the last one and loses. A careful induction, splitting on whether exactly one pile or more than one pile exceeds size 1, turns this outline into a full proof; the details are carried out in Chapter 11 with Problem Set 2. ■

The contrast is worth stating. For Nim the misère and normal theories differ only in the endgame of all-ones piles. For general impartial games no such simple description is known, and misère theory remains genuinely harder than normal-play theory.

The clean theory of this chapter is fragile to two changes. Switching to misère play breaks the Sprague–Grundy reduction for general games: there is no misère analogue of the Grundy value that combines by a simple rule, and the disjunctive theory becomes genuinely hard. Adding *partizan* moves, where the two players have different options (as in chess or checkers), breaks it further, and is the subject of a separate and deeper theory of game values. Everything here relies on *impartial, normal* play.

Notes and References

Nim was solved by Bouton [6], in the second series of the *Annals of Mathematics* (volume 3, pages 35–39); it is the first complete mathematical analysis of a nontrivial game. Charles Bouton, a Harvard mathematician, coined the name from the German imperative *nimm* (“take”). The Grundy value and the reduction of every impartial game to a Nim heap were found independently by Sprague [60], in the *Tôhoku Mathematical Journal*, and Grundy [24], in the Cambridge undergraduate journal *Eureka*. Wythoff's game (Section 2.7) is from Wythoff [67]; its golden-ratio P-positions are a classical gem, and the recognition that the two coordinate sequences are complementary Beatty sequences came about

two decades later. The deepest modern development is John Conway’s theory of *surreal numbers* [13], which grew out of analyzing Go endgames and embeds both the real numbers and the ordinals inside a single class built from games. The catalogue of take-and-break games, their eventual periodicity, and the open status of Grundy’s game are treated at length in Berlekamp et al. [4], the definitive reference on combinatorial game theory, which also develops the partizan and misère theories this chapter only mentions. The presentation here follows Ferguson [16], where Nim, the subtraction games, and the take-and-break games appear with the same Grundy-function computations.

Historical Note: John Conway

The deepest theory of combinatorial games is the work of John Horton Conway (1937–2020). From analyzing endgames he built the *surreal numbers*, a single class holding the reals, the ordinals, and the infinitesimals at once, and with Berlekamp and Guy he wrote *Winning Ways*, the field’s encyclopedia. His range was extraordinary: the cellular automaton known as the Game of Life, the sporadic Conway group in the theory of symmetry, the Conway polynomial of a knot, and the Doomsday rule for the day of the week. A fixture first of Cambridge and then of Princeton, he was famous for doing serious mathematics as play. He died of COVID-19 in April 2020.

Chapter summary

An impartial combinatorial game is a finite acyclic graph of positions with the same moves available to both players (Definition 2.1). Under normal play, positions split into P-positions, from which the previous player wins, and N-positions, from which the next player wins (Definition 2.2). For a single subtraction pile the P-positions are periodic (Proposition 2.4). For Nim, Bouton’s theorem identifies the P-positions as those with nim-sum zero, and its proof gives the winning move explicitly (Theorem 2.7). The Sprague–Grundy function $g(x) = \text{mex}\{g(y) : y \in F(x)\}$ assigns each position a number, with $g(x) = 0$ exactly at P-positions (Theorem 2.12), and the Sprague–Grundy theorem computes the Grundy value of a disjunctive sum as the nim-sum of the parts (Theorem 2.15). Every impartial game is therefore equivalent to a single Nim pile. We tabulated the Grundy functions of several take-and-break games, observing eventual periodicity for Kayles and Dawson’s Chess and its apparent absence for Grundy’s game, and we recorded the clean misère theory for Nim (Theorem 2.22).

Exercises

Subtraction games, Grundy values, and disjunctive sums

Exercise 2.1. A pile has n chips and a move removes 1 to 6 chips; the player taking the last chip wins. Identify all P-positions. If $n = 31$, state the winning move. If $n = 28$, do you prefer to move first or second?

Exercise 2.2. For a single pile of 100 chips, find the P-positions and the Grundy value $g(100)$ under each subtraction set, and say whether you want to move first.

$$(a) S_1 = \{1, 3, 5, 7\}, \quad (b) S_2 = \{1, 3, 6\}, \quad (c) S_3 = \{1, 2, 4, 8, 16, 32, 64\}.$$

(Hint for (a): every move changes the parity of the pile. For (c): consider residues modulo 3.) These are parts of Problem Set 1; full solutions are in Chapter 11.

Exercise 2.3. Three piles, each of 100 chips, are played as a disjunctive sum: pile 1 uses subtraction set $S_1 = \{1, 3, 5, 7\}$, pile 2 uses $S_2 = \{1, 3, 6\}$, and pile 3 uses $S_3 = \{1, 2, 4, 8, 16, 32, 64\}$. Using your Grundy values from Exercise 2.2, determine whether the starting position is a P- or N-position, and if you must move first, give a winning move.

Nim, nim-sums, and impartial-game variants

Exercise 2.4. For each Nim position, compute the nim-sum, classify it as P or N, and give all winning first moves: (a) (3, 5, 7); (b) (1, 4, 9, 16); (c) (13, 17, 26).

Exercise 2.5. A position graph is a directed acyclic graph with edges pointing to followers. Define the Grundy value of every vertex by $g(x) = \text{mex}\{g(y) : y \in F(x)\}$, starting from the sinks (which have $g = 0$). Compute the Grundy values for the graph of Problem Set 1, reproduced in [Chapter 11](#), and classify each vertex as P or N. State an algorithm that performs this computation for an arbitrary finite acyclic graph.

Exercise 2.6. A move either removes any number of chips from one pile, as in Nim, or splits one pile of size at least 2 into two nonempty piles. Compute $g(0)$ through $g(8)$ from the recurrence, confirm the values 0, 1, 2, 4, 3, 5, 6, 8, 7, and conjecture the closed form. (This is Lasker's Nim of [Section 2.6.2](#).)

Exercise 2.7. Use [Theorem 2.22](#) to decide the misère outcome of the Nim positions (1, 1, 1), (1, 1, 1, 1), (2, 2), and (1, 2, 3). For each N-position, give a winning move.

Exercise 2.8 (Wythoff's game). Two piles. A move removes any number of chips from one pile, or removes the *same* number from both piles. The player taking the last chip wins. Show by hand that the smallest P-positions are (0, 0), (1, 2), and (3, 5). (The full set of P-positions is $(\lfloor k\phi \rfloor, \lfloor k\phi^2 \rfloor)$ for $k \geq 0$, where $\phi = (1 + \sqrt{5})/2$ is the golden ratio; this is a classical result of Wythoff, 1907, and a good target for a short computational experiment.)

Further practice (worked solutions in the back)

Exercise 2.9. Three Nim heaps hold 7, 9, and 12 chips and the player taking the last chip wins. Compute the nim-sum in binary, classify the position as P or N, and if it is an N-position give every winning move, then verify your move leaves a nim-sum of zero. (Worked in [Appendix D](#).)

Exercise 2.10. A single heap is played with subtraction set $S = \{1, 2, 3\}$: a move removes 1, 2, or 3 chips, last chip wins. (a) Compute the Grundy values $g(0), g(1), \dots, g(8)$ from the mex rule and state the pattern. (b) Use the pattern to find $g(10)$, $g(15)$, and $g(21)$. (c) Three such heaps of sizes 10, 15, and 21 are played as a disjunctive sum; using the Sprague–Grundy theorem, classify the combined position and say whether you want to move first. (Worked in [Appendix D](#).)

CHAPTER 3

Extensive-Form Games and Backward Induction

“In chess, you must anticipate not your move, but the reply, and the reply to the reply, to the very end of the game.”

The strategic form of Chapter 1 flattens a game into a single simultaneous choice. Many games are not like that. Players move in turn, each seeing some of what came before, and the order of play matters. The natural picture is a tree: a root where the first decision is made, branches for the available actions, and leaves carrying payoffs. This is the *extensive form*. For games of perfect information it comes with a complete solution method, backward induction, which both generalizes the P- and N-position analysis of Chapter 2 and exposes a subtlety that the strategic form hides: some Nash equilibria rest on threats no rational player would carry out.

3.1 Game trees

Definition 3.1 (Extensive-form game). A finite *extensive-form game* consists of

- a finite rooted tree whose nodes are the possible states of play and whose edges are **actions**;
- a partition of the non-terminal (decision) nodes among the players, assigning to each decision node the player who moves there;
- for each player i , a payoff u_i attached to every terminal node (leaf);
- a partition of each player’s decision nodes into **information sets**, where the nodes in one information set are indistinguishable to the player when moving, so the same actions must be available at all of them.

The **root** is where play begins, and a **path** from root to leaf is one possible course of play. When chance intervenes, an extra player “Nature” moves at designated nodes according to fixed probabilities.

Definition 3.2 (Perfect information). A game has **perfect information** if every information set is a single node, that is, whenever a player moves, the player knows exactly where in the tree play has arrived, and so knows the entire history of prior moves.

Chess, checkers, and Nim have perfect information. The poisoned-goblet duel of Section 1.3 does not: Vizzini moves without knowing which goblet Westley poisoned, so two nodes lie in one information set, drawn as the dashed link in Figure 1.1. This chapter treats perfect information; the general case, where information sets are nontrivial, is the subject of Chapter 7.

Example 3.3 (A two-stage game). Player I moves first, choosing L or R . Player II observes the choice and then chooses U or D . Figure 3.1 draws the tree; the four paths end in payoff pairs (u_1, u_2) , with Player I’s payoff first:

$$L \rightarrow U : (3, 2), \quad L \rightarrow D : (1, 1), \quad R \rightarrow U : (0, 0), \quad R \rightarrow D : (2, 3).$$




Figure 3.1. The two-stage game of Example 3.3. Player I moves at the root; Player II moves at each of the two interior nodes, having seen Player I's choice.

3.2 Strategies in the extensive form

The word “strategy” means more in a tree than a single action. It is a complete contingency plan.

Definition 3.4 (Pure strategy). A **pure strategy** for player i in an extensive-form game is a function that assigns, to each of player i 's information sets, one of the actions available there. A strategy specifies what the player would do at every decision point, including points that will never be reached if the strategy is followed.

The clause about unreached nodes is not pedantry; it is the whole source of the subtlety to come. In Example 3.3, Player I has two pure strategies, L and R . Player II moves at two nodes and chooses one of two actions at each, so Player II has $2 \times 2 = 4$ pure strategies, which we write as ordered pairs (action after L , action after R):

$$UU, \quad UD, \quad DU, \quad DD.$$

A strategy such as UD means “play U if Player I chose L , play D if Player I chose R .” Once both players fix pure strategies, the path through the tree and hence the payoffs are determined, so the extensive game collapses to a strategic-form game, its **normal form**. For Example 3.3 the normal form is the bimatrix

	UU	UD	DU	DD
L	(3, 2)	(3, 2)	(1, 1)	(1, 1)
R	(0, 0)	(2, 3)	(0, 0)	(2, 3)

where the columns repeat because, for instance, when Player I plays L , the action Player II has planned for the unreached R node makes no difference to the outcome.

3.3 Backward induction and subgame perfection

In a perfect-information tree, the leaves carry payoffs and every interior node begins a self-contained smaller game. This invites a recursive solution.

Definition 3.5 (Subgame). A **subgame** is the tree consisting of a decision node together with everything below it, treated as a game in its own right. In a perfect-information game, every decision node starts a subgame.

Procedure 3.6 (Backward induction). To solve a finite perfect-information game:

1. Find a decision node all of whose children are leaves. The player there chooses the action leading to the leaf that maximizes that player's own payoff.
2. Replace that node by the chosen leaf's payoff vector, turning it into a leaf. The tree shrinks.
3. Repeat until only the root remains. The recorded choices form a strategy for each player, and the path they trace is the predicted play.

Applied to Example 3.3: at the node after L , Player II compares U giving 2 against D giving 1, and chooses U , so that node is worth (3, 2). At the node after R , Player II compares U giving 0 against D giving 3, and chooses D , so

that node is worth (2, 3). Player I now compares L leading to (3, 2), payoff 3, against R leading to (2, 3), payoff 2, and chooses L . The backward-induction solution is

$$s_1^* = L, \quad s_2^* = UD \text{ (play } U \text{ after } L, D \text{ after } R),$$

with outcome (3, 2). Figure 3.2 shows the tree collapsing one level at a time: first each of Player II's nodes is replaced by the payoff vector of the action she prefers there, then Player I chooses at the root among the resulting leaves.




Figure 3.2. Backward induction as folding. *Step 1:* replace each of Player II's decision nodes by the payoff of the branch she prefers (U after L , worth (3, 2); D after R , worth (2, 3)). *Step 2:* Player I, now facing two leaves, compares his own payoffs 3 and 2 and takes the larger, choosing L . The surviving path (highlighted) is the subgame perfect outcome (3, 2).

Example 3.7 (Folding a three-level tree). Backward induction folds any depth, not just two. Player I moves at the root (L or R); Player II responds; and on one branch Player I moves a *third* time. The tree, with payoffs (u_1, u_2) at the leaves, is Figure 3.3.




Figure 3.3. The three-level game of Example 3.7. Player I moves at the root and again at the node reached by L then D ; Player II moves at the two middle nodes.

Fold the deepest node first. The bottom node belongs to Player I (reached by L, D). He compares his own payoffs G : 5 and H : 2 and takes G , so that node is worth (5, 0).

Fold Player II's nodes. At the left node II compares her payoffs from C : 2 (the leaf (4, 2)) and D : 0 (the folded value (5, 0), in which she gets 0); she takes C , worth (4, 2). At the right node she compares E : 1 and F : 3 and takes F , worth (3, 3).

Fold the root. Player I now compares L : 4 (the leaf (4, 2)) against R : 3 (the leaf (3, 3)) and takes L . The subgame perfect equilibrium is

$$s_1^* = (L \text{ at the root, } G \text{ at the third-level node}), \quad s_2^* = (C \text{ after } L, F \text{ after } R),$$

with outcome (4, 2). Notice that Player I's tempting payoff of 5 at leaf G is never reached: Player II, foreseeing it, steers away from D by choosing C , exactly the off-path reasoning that backward induction enforces. $\diamond$

Definition 3.8 (Subgame perfect equilibrium). A strategy profile is a **subgame perfect Nash equilibrium** (SPNE) if it induces a Nash equilibrium in every subgame, the whole game included.

Theorem 3.9 (Existence by backward induction). Every finite extensive-form game of perfect information has a subgame perfect Nash equilibrium in pure strategies, and backward induction (Procedure 3.6) computes one.

Proof. Induct on the height h of the tree, the length of the longest path from root to a leaf.

If $h = 0$ the tree is a single leaf; there are no decisions, and the (empty) profile is trivially subgame perfect.

Suppose the result holds for all perfect-information games of height less than h , and let G have height h with root owned by player j . Each child of the root begins a subgame of height at most $h - 1$. By the induction hypothesis, each such subgame has a subgame perfect equilibrium computed by backward induction; fix one for each, and let w_c be the resulting payoff vector at child c . Now let player j choose at the root an action leading to a child c^* that maximizes j 's own component of w_{c^*} . Combine j 's root choice with the fixed subgame equilibria to form a profile for G .

This profile is subgame perfect. In every proper subgame (one rooted at a child or below), it restricts to the chosen subgame perfect equilibrium, which is a Nash equilibrium there by construction. In the whole game, player j is choosing a best response at the root by the maximization, and no other player moves at the root, so the profile is a Nash equilibrium of G as well. Hence it is subgame perfect, and the construction is exactly backward induction. ■

Backward induction solves perfect information

Every finite game of perfect information has a subgame perfect equilibrium in *pure* strategies, and backward induction finds it: solve the smallest subgames first, replace each by its value, and work up to the root. Two consequences follow at once. Such games never need randomization to be solved, and (in win-lose-draw form) one side can always force its best outcome, which is Zermelo's determinacy theorem below.

3.3.1 Credible and incredible threats

Backward induction does more than find an equilibrium; it selects among the equilibria, discarding those built on empty threats. To see this, find all Nash equilibria of the normal form by the best-response method, underlining each player's best reply to each of the opponent's choices. Player I (the row player) compares the two payoffs in each column:

$$\begin{array}{ll} \text{vs } UU : \max(3, 0) = 3 \rightarrow L, & \text{vs } UD : \max(3, 2) = 3 \rightarrow L, \\ \text{vs } DU : \max(1, 0) = 1 \rightarrow L, & \text{vs } DD : \max(1, 2) = 2 \rightarrow R. \end{array}$$

Player II (the column player) compares her four payoffs in each row:

$$\text{vs } L : \max(2, 2, 1, 1) = 2 \rightarrow \{UU, UD\}, \quad \text{vs } R : \max(0, 3, 0, 3) = 3 \rightarrow \{UD, DD\}.$$

A cell is a Nash equilibrium when both entries are best replies. Player I's best replies sit in columns UU, UD, DU at row L and in column DD at row R ; Player II's best replies sit at $(L, UU), (L, UD), (R, UD), (R, DD)$. The cells appearing in both lists are

$$(L, UU), \quad (L, UD), \quad (R, DD),$$

the three pure Nash equilibria. Only one of them, (L, UD) , is subgame perfect. Consider (R, DD) . Here Player II plans to play D after L , and given that plan Player I prefers R , leading to $(2, 3)$. The plan is a Nash equilibrium because, on the path actually taken, neither player can improve. But it is not credible: the threat to play D after L would, if Player I ever called it by playing L , give Player II a payoff of 1 instead of the 2 that U would bring. A rational Player II would not carry out the threat, so a rational Player I should not believe it. Subgame perfection is exactly the requirement that every planned action be optimal where it is made, which rules out (R, DD) and the similarly incredible (L, UU) , leaving the backward-induction solution (L, UD) alone.

Remark 3.10. This is the central lesson of the chapter. The strategic form remembers only what happens on the equilibrium path and so accepts non-credible threats; the extensive form, read by backward induction, demands rationality off the path as well. Subgame perfection [51] refines Nash equilibrium by insisting that threats be ones the threatener would actually want to execute.

Every subgame perfect equilibrium is a Nash equilibrium, but not conversely. A game tree can have many Nash equilibria, some sustained by threats the threatener would never carry out; backward induction keeps only the credible ones. When a problem asks for the solution of a perfect-information game, it wants the subgame perfect equilibrium, and reporting a plain Nash equilibrium with an empty threat is the usual mistake.

3.3.2 The entry-deterrence game

The canonical economic use of subgame perfection is the *entry game*. A potential **Entrant** (Player I) decides whether to ENTER a market or STAY OUT. If it enters, an **Incumbent** (Player II) decides whether to FIGHT a price war or ACCOMMODATE and share the market. Figure 3.4 draws the tree, where the payoffs $(u_{\text{Entrant}}, u_{\text{Incumbent}})$, in millions of profit, are

$$\text{STAY OUT} : (0, 2), \quad \text{ENTER} \rightarrow \text{ACCOMMODATE} : (1, 1), \quad \text{ENTER} \rightarrow \text{FIGHT} : (-1, -1).$$




Figure 3.4. The entry-deterrence game. The Entrant moves at the root; the Incumbent moves only if entry occurs. Backward induction folds the Incumbent's node first.

Backward induction. Start at the only non-trivial decision node, the Incumbent's. Comparing FIGHT (-1) against ACCOMMODATE (1) , the Incumbent accommodates, so that node is worth $(1, 1)$. Fold it into the root: the Entrant now compares STAY OUT (0) against ENTER leading to $(1, 1)$, payoff 1, and enters. The subgame perfect equilibrium is

$$(\text{ENTER}, \text{ACCOMMODATE}), \quad \text{outcome } (1, 1).$$

The Incumbent would *like* to deter entry by threatening to fight, since (STAY OUT) pays it $2 > 1$. But the threat is not credible: once entry has occurred, fighting pays -1 against accommodating's 1 , so a rational Incumbent never carries it out, and a rational Entrant enters. The profile $(\text{STAY OUT}, \text{FIGHT})$ is a Nash equilibrium of the normal form, yet backward induction discards it as resting on an empty threat. The way an incumbent can deter entry for real is to change the game *before* it is played, for instance by sinking cost into capacity that makes fighting genuinely profitable, which is the theory of commitment. $\diamond$

Historical Note: Reinhard Selten

Subgame perfection is due to Reinhard Selten (1930–2016), who introduced it in 1965 and a decade later added the trembling-hand refinement, asking that an equilibrium survive small mistakes. Selten shared the 1994 Nobel Memorial Prize with Nash and Harsanyi and was a founder of experimental economics, testing in the laboratory whether people actually play the equilibria theory predicts. The son of a Jewish father, he lived through the Nazi years and was uprooted by the war. A lifelong enthusiast of the constructed language Esperanto, he wrote and lectured on game theory in it. He died in 2016.

3.4 Backward induction through a chance node

Uncertainty fits backward induction with no new machinery: at a chance node, fold the *expected* payoff. A startup chooses to LAUNCH or WAIT. If it launches, a rival decides to ENTER or STAY OUT; if the rival enters, the market turns out GOOD for the startup with probability $\frac{2}{3}$ and BAD with probability $\frac{1}{3}$. Figure 3.5 draws the tree, with payoffs (u_s, u_r) to startup and rival at the leaves.

Fold from the leaves. At the chance node the startup's good and bad outcomes average to

$$\left(\frac{2}{3} \cdot 4 + \frac{1}{3} \cdot (-2), \quad \frac{2}{3} \cdot 1 + \frac{1}{3} \cdot 3 \right) = \left(2, \frac{5}{3} \right).$$

The rival, moving just above, compares entering (her share $\frac{5}{3}$) against staying out (her 0) and enters. The startup, at the root, compares launching (its expected 2) against waiting (its 1) and launches. The subgame perfect outcome is LAUNCH then ENTER, worth an expected $(2, \frac{5}{3})$: the startup accepts a contested market because its expected value still beats the safe wait, and the only departure from a deterministic tree is that one comparison is taken in expectation.




Figure 3.5. A startup's entry decision with an uncertain market. Decision nodes are squares (the startup “s” and rival “r”), the chance node “N” (Nature) is a circle, and the bold path is the subgame perfect solution. Backward induction replaces the chance node by its expected payoff $(2, \frac{5}{3})$, then folds the rest as usual.

3.5 Backward induction with continuous choices: Stackelberg competition

Backward induction is not limited to trees with finitely many branches. When a player chooses a *number*, the same logic, solve the last mover's problem first and feed the answer up, runs through calculus instead of a finite comparison. The cleanest economic example is Stackelberg's model of two firms choosing output one after the other.

Two firms sell an identical good. Firm 1 (the *leader*) first picks a quantity q_1 ; firm 2 (the *follower*) sees q_1 and then picks q_2 . The market price falls with total output, $P = a - (q_1 + q_2)$, and to keep the arithmetic clean both firms produce at zero cost, so each firm's profit is its quantity times the price. Take $a = 12$.

Example 3.11 (Solving Stackelberg by backward induction). *Step 1: the follower's problem (the last move).* Given the leader's q_1 , firm 2 chooses q_2 to maximize

$$\pi_2 = q_2 (12 - q_1 - q_2).$$

This is a downward parabola in q_2 ; setting the derivative to zero, $12 - q_1 - 2q_2 = 0$, gives firm 2's **reaction function**

$$q_2 = R_2(q_1) = \frac{12 - q_1}{2}.$$

Step 2: the leader anticipates (the first move). A rational leader knows firm 2 will play $R_2(q_1)$, so it substitutes that response into its own profit before choosing:

$$\pi_1 = q_1 (12 - q_1 - R_2(q_1)) = q_1 \left(12 - q_1 - \frac{12 - q_1}{2} \right) = q_1 \cdot \frac{12 - q_1}{2}.$$

Maximizing, $\frac{d\pi_1}{dq_1} = \frac{12 - 2q_1}{2} = 6 - q_1 = 0$, so $q_1 = 6$. Feeding this back through the follower's reaction function, $q_2 = R_2(6) = \frac{12 - 6}{2} = 3$. Total output is 9, the price is $P = 12 - 9 = 3$, and the profits are

$$\pi_1 = 6 \cdot 3 = 18, \quad \pi_2 = 3 \cdot 3 = 9.$$

Step 3: compare with simultaneous moves. If instead both firms chose at the same time (Cournot competition, Chapter 8), each would solve $q_i = \frac{12 - q_j}{2}$ simultaneously, giving the symmetric solution $q_1 = q_2 = 4$, price $P = 4$, and equal profits $\pi_1 = \pi_2 = 16$. The leader's profit rose from 16 to 18 and the follower's fell from 16 to 9: moving first, and being *known* to have moved first, is worth 2 units of profit here. Figure 3.6 shows the leader sliding along the follower's reaction line to the most profitable point on it, rather than settling at the symmetric Cournot crossing. $\diamond$




Figure 3.6. Stackelberg versus Cournot. Each firm’s reaction line gives its profit-maximizing output as a function of the other’s. Simultaneous play lands at the symmetric crossing, Cournot $(4, 4)$. The Stackelberg leader, choosing first, picks the point on the follower’s red reaction line that maximizes its own profit, $(6, 3)$, producing more and earning more while the follower produces less.

Intuition: the advantage of going first and being seen to commit

The leader does better than under simultaneous play for one reason: by committing to a large output where everyone can see it, it forces the follower to scale back. This is the mirror image of the incumbent’s *incredible* threat above. There, talk before the move changed nothing; here, an *irreversible* first action changes everything. Commitment is power when it cannot be undone, which is why a general burning the bridges behind the army, or a firm building a factory it cannot unbuild, can come out ahead of a more flexible rival. The catch is that going first only helps when the move is observable and binding; a leader who could secretly revise q_1 would gain nothing.

3.6 Connection to combinatorial games and Zermelo’s theorem

Backward induction is the same idea as the P- and N-position labeling of Chapter 2, viewed from the strategic-form side. An impartial game under normal play is a perfect-information extensive game with win/lose payoffs, and the recursive rule “a position wins if some move reaches a losing position” is backward induction with payoffs in $\{\text{win}, \text{lose}\}$. The Grundy function refines this with more than one bit of information, but the engine is identical.

The general principle, the oldest theorem in game theory, predates the field’s modern vocabulary.

Theorem 3.12 (Zermelo, 1913). In a finite two-player game of perfect information with no chance moves and outcomes restricted to win, lose, or draw, exactly one of the following holds: one player has a strategy that forces a win, or each player has a strategy that forces at least a draw.

Proof idea. Apply backward induction with outcomes ordered $\text{win} > \text{draw} > \text{lose}$ from the moving player’s viewpoint. The recursion assigns each node a determined value, and the value of the root is the determined outcome under optimal play. Finiteness and perfect information are what make the recursion well defined. ■

Zermelo’s theorem says chess is determined: either White can force a win, or Black can, or both can force a draw. Which of the three holds is unknown, because the tree, though finite, is far too large to traverse. Determinacy is a statement of existence, and it carries no promise that the strategy can be found in practice. The course’s term project, building a strong player for the game of Go, lives precisely in this gap between a theorem that guarantees optimal play exists and the computation that tries to approximate it.

Historical Note: the gap between existence and computation

Zermelo's theorem guarantees a value for chess and Go; finding it is another matter, and the history of game-playing programs is the history of closing that gap one game at a time. Shannon [53] estimated the chess game tree at about 10^{120} positions, the "Shannon number," to argue brute force is hopeless; Go on a 19×19 board has on the order of 10^{170} legal positions, vastly more. Checkers, with roughly 5×10^{20} positions, was *weakly solved* by the Chinook team after eighteen years of computation: perfect play is a draw [49]. Go fell not to enumeration but to learning: AlphaGo combined deep value and policy networks with Monte Carlo tree search to beat a professional in 2016 [58], and its self-play successors [59], together with the open-source KataGo [66], now exceed the strongest humans. Even imperfect-information games have yielded: heads-up limit poker was essentially solved [7], and superhuman play was reached in no-limit [11] and in six-player poker [12]. Each result computes, for a specific game, the optimal play that Zermelo and von Neumann proved must exist.

3.7 The Centipede game and the limits of backward induction

Backward induction is logically airtight, yet its predictions can feel wrong, and sometimes are wrong about how people actually play. The sharpest illustration is the Centipede game [46].

Example 3.13 (A growing pie). Two players alternate. At each turn the mover may STOP, ending the game, or CONTINUE, passing the move and enlarging the pie. In the four-stage version used in lecture, stopping at successive nodes pays

$$(1, 0), \quad (2, 1), \quad (3, 2), \quad (4, 3),$$

and if the last mover also continues the game ends at $(0, 0)$.

Backward induction runs as follows. At the fourth node Player II compares STOP paying 3 against CONTINUE paying 0, and stops, fixing that node at $(4, 3)$. At the third node Player I compares STOP paying 3 against CONTINUE, which leads to $(4, 3)$ and pays Player I 4, so Player I continues. At the second node Player II compares STOP paying 1 against continuing to $(4, 3)$ paying 3, and continues. At the first node Player I compares STOP paying 1 against continuing to $(4, 3)$ paying 4, and continues. The subgame perfect path is therefore CONTINUE, CONTINUE, CONTINUE, STOP, reaching $(4, 3)$.

Remark 3.14. With the payoffs of Example 3.13 there is no paradox: cooperation to the last node is exactly what backward induction predicts, because each mover's reward for stopping never exceeds what continuing eventually yields. The famous Centipede paradox requires payoffs that tempt each mover to stop early, which the next example supplies.

Example 3.15 (The Centipede paradox). Keep the alternating STOP/CONTINUE structure, but let stopping always pay the mover slightly more than waiting one more round would:

$$\text{stop at node 1: } (2, 0), \quad \text{node 2: } (1, 3), \quad \text{node 3: } (4, 2), \quad \text{node 4: } (3, 5),$$

with a final CONTINUE paying $(6, 4)$. Figure 3.7 draws the chain. Now backward induction unravels cooperation




Figure 3.7. The Centipede paradox of Example 3.15. The pie grows from a total of 2 to a total of 10 along the top, yet backward induction stops at the first node.

completely. At node 4 Player II takes 5 rather than 4 and stops, fixing $(3, 5)$. At node 3 Player I takes 4 rather than the 3 that continuing yields, and stops, fixing $(4, 2)$. At node 2 Player II takes 3 rather than the 2 from continuing, and stops, fixing $(1, 3)$. At node 1 Player I takes 2 rather than the 1 from continuing, and stops. The unique subgame perfect outcome is $(2, 0)$ at the very first move, although both players would be far better off at the cooperative $(6, 4)$.

The paradox is that flawless individual reasoning forecloses an outcome both players prefer. Three observations soften it.

- **Common knowledge of rationality.** Backward induction assumes not only that every player is rational, but that each knows the others are rational, and each knows that each knows, without end. This tower of assumptions is **common knowledge of rationality** (CKR). If a player assigns even a small probability to an opponent who might continue, continuing can become the better choice, and the unravelling stops.
- **Experimental behavior.** In laboratory play of the Centipede, people routinely continue for several rounds, capturing much of the cooperative surplus. In the classic experiment of McKelvey and Palfrey [33], the first mover TOOK (the backward-induction move) on the opening node only about 7% of the time in a four-move game and about 1% in a six-move game, against the 100% that subgame perfection predicts. The deviation is systematic rather than random, and it grows with the number of rounds of cooperation available.
- **The model rather than the mathematics.** The backward-induction outcome is the correct solution of the game as written. When real players cooperate, they are playing a different game, one whose payoffs include reputation, trust, or the expectation of future dealings. This is the same lesson as the Prisoner's Dilemma in Chapter 1: enrich the payoffs to match the situation.

The Centipede motivates the refinements that occupy later research, behavioral game theory and models of bounded rationality, but within the classical framework of this book the subgame perfect equilibrium remains the solution.

3.8 Imperfect information and behavioral strategies

Backward induction needed perfect information: every player, when moving, knew the entire history. Many games break this assumption, the poisoned-goblet duel of Section 1.3 among them, and the extensive form handles them through the information sets of Definition 3.1. This section develops the imperfect-information case, introduces the right notion of a randomized strategy for trees, and states Kuhn's theorem, which makes such strategies manageable. The running example is a small game of poker.

Definition 3.16 (Imperfect information and behavioral strategies). A game has **imperfect information** if some information set contains more than one node, so that a player, when moving there, cannot tell which of those nodes play has reached. A **behavioral strategy** for player i assigns to each of i 's information sets a probability distribution over the actions available there, the player randomizing *independently* at each set. This is different from a mixed strategy, which randomizes once over complete pure plans (Definition 3.4). The game has **perfect recall** if every player remembers their own past moves and observations.

Example 3.17 (A poker bluffing game). Two players ante one dollar each, building a pot of two. From a deck of four cards, three losing and one winning, Nature deals Player I a single card: a winner with probability $\frac{1}{4}$ and a loser with probability $\frac{3}{4}$. Player I sees the card; Player II does not. Player I chooses to BET or CHECK. A check goes straight to a showdown, paying Player I +1 on a winner and -1 on a loser. A bet adds two to the pot and passes the move to Player II, who chooses to CALL or FOLD without seeing the card: a call leads to a showdown paying ± 3 to Player I, while a fold concedes the pot, paying Player I +1 regardless of the card. The game is zero-sum, and Figure 3.8 draws its tree.

A pure strategy for Player I names an action for each card, written (action on a winner, action on a loser), so Player I has four: (b, b) , (b, c) , (c, b) , (c, c) . Player II, who moves only after a bet and in a single information set, has two: call or fold. Averaging each outcome over the deal gives the normal form, with payoffs to Player I:

	call	fold
(b, b)	$\frac{1}{4}(3) + \frac{3}{4}(-3) = -\frac{3}{2}$	$\frac{1}{4}(1) + \frac{3}{4}(1) = 1$
(b, c)	$\frac{1}{4}(3) + \frac{3}{4}(-1) = 0$	$\frac{1}{4}(1) + \frac{3}{4}(-1) = -\frac{1}{2}$
(c, b)	$\frac{1}{4}(1) + \frac{3}{4}(-3) = -2$	$\frac{1}{4}(1) + \frac{3}{4}(1) = 1$
(c, c)	$\frac{1}{4}(1) + \frac{3}{4}(-1) = -\frac{1}{2}$	$\frac{1}{4}(1) + \frac{3}{4}(-1) = -\frac{1}{2}$

Two of Player I's strategies are poor and fall to dominance. Checking a winner while betting a loser, (c, b) , is dominated by (b, b) (its call column $-2 < -\frac{3}{2}$, its fold column tied), and checking always, (c, c) , earns $-\frac{1}{2}$ no matter what, worse




Figure 3.8. The poker bluffing game of Example 3.17. Nature deals; Player I sees the card and bets or checks; Player II, who cannot see the card, calls or folds. The dashed link is Player II's information set: the two nodes after a bet are indistinguishable. Payoffs shown are to Player I.

than strategies that sometimes earn 1. Removing them leaves the 2×2 game on rows $(b, b), (b, c)$ and columns call, fold,

$$A = \begin{bmatrix} -\frac{3}{2} & 1 \\ 0 & -\frac{1}{2} \end{bmatrix}.$$

Its row minima are $\min(-\frac{3}{2}, 1) = -\frac{3}{2}$ and $\min(0, -\frac{1}{2}) = -\frac{1}{2}$, so the maximin is $-\frac{1}{2}$; its column maxima are $\max(-\frac{3}{2}, 0) = 0$ and $\max(1, -\frac{1}{2}) = 1$, so the minimax is 0. Since $-\frac{1}{2} \neq 0$ there is no saddle, and we solve by indifference (Theorem 4.8). Label $a = -\frac{3}{2}, b = 1, c = 0, d = -\frac{1}{2}$, so the denominator is

$$D = a - b - c + d = -\frac{3}{2} - 1 - 0 - \frac{1}{2} = -3.$$

Player I picks the probability p of the top row (b, b) to make Player II indifferent between her columns:

$$p^* = \frac{d - c}{D} = \frac{-\frac{1}{2} - 0}{-3} = \frac{-\frac{1}{2}}{-3} = \frac{1}{6}, \quad 1 - p^* = \frac{5}{6}.$$

Player II picks the probability q of calling to make Player I indifferent between his rows:

$$q^* = \frac{d - b}{D} = \frac{-\frac{1}{2} - 1}{-3} = \frac{-\frac{3}{2}}{-3} = \frac{1}{2}.$$

The value is

$$V = \frac{ad - bc}{D} = \frac{(-\frac{3}{2})(-\frac{1}{2}) - (1)(0)}{-3} = \frac{\frac{3}{4}}{-3} = -\frac{1}{4}.$$

As a check, play $q^* = \frac{1}{2}$ against the top row: $\frac{1}{2}(-\frac{3}{2}) + \frac{1}{2}(1) = -\frac{3}{4} + \frac{1}{2} = -\frac{1}{4}$, and against the bottom row $\frac{1}{2}(0) + \frac{1}{2}(-\frac{1}{2}) = -\frac{1}{4}$, the same number, as indifference requires. The solution is poker wisdom made precise. Player I bets a winning card always (both surviving strategies bet on a winner) and bluffs a losing card with probability $\frac{1}{6}$ (only (b, b) bets on a loser). The value is negative because the deck is unfavorable, but the occasional bluff is what keeps Player II honest and recovers what timid play would surrender. $\diamond$

Player I's solution above was given as a behavioral strategy without comment: bet on a winner with probability 1, bluff on a loser with probability $\frac{1}{6}$, one distribution per information set (Figure 3.9). That this loses nothing compared with mixing over the four pure plans is the content of Kuhn's theorem.

Theorem 3.18 (Kuhn, 1953). In a finite extensive-form game with perfect recall, every mixed strategy is outcome-equivalent to a behavioral strategy, and conversely: the two produce the same probability distribution over terminal nodes against any strategies of the other players.



Figure 3.9. Player I's optimal *behavioral* strategy in the poker game (Example 3.17): one probability distribution over actions at each of his two information sets. Holding a winner he bets for sure; holding a loser he bluffs (bets) only $\frac{1}{6}$ of the time and otherwise checks. By Kuhn's theorem this is outcome-equivalent to mixing over the four pure plans $(b, b), (b, c), (c, b), (c, c)$, but it needs only two numbers instead of three.

Behavioral strategies suffice

Under perfect recall, a player may randomize independently at each information set instead of over whole plans, with no loss. This is what makes large sequential games tractable: a behavioral strategy carries one distribution per information set, while a mixed strategy carries one number per pure plan, and the plans are exponentially many.

Perfect recall is essential to Kuhn's theorem. If a player can forget their own past actions, as in the “absent-minded driver,” behavioral and mixed strategies can give genuinely different sets of achievable outcomes, and the clean equivalence fails. Every game in this book has perfect recall, so we use behavioral strategies freely.

Notes and References

The extensive form, information sets, and the formal apparatus of game trees are due to Kuhn [29]. Determinacy of finite perfect-information games predates the field, in Zermelo [68], often called the first paper in game theory; the careful historical record (and the correction of the common claim that Zermelo used backward induction) is in Schwalbe and Walker's study of that work. Subgame perfection and its refinements are Selten's (Selten 51, 52), work for which he shared the 1994 Nobel Memorial Prize with Nash and Harsanyi. The Centipede game is from Rosenthal [46], and the canonical experimental study, showing how rarely people play the backward-induction move, is McKelvey and Palfrey [33]. The computational game-solving milestones cited above span Bowling et al. [7], Brown and Sandholm [11, 12], Schaeffer et al. [49], Shannon [53], Silver et al. [58, 59]. The treatment here follows Osborne [39], where extensive-form games and backward induction are developed in detail.

Chapter summary

An extensive-form game is a tree of moves with payoffs at the leaves, players assigned to decision nodes, and information sets recording what each mover knows (Definition 3.1). Under perfect information every information set is a single node (Definition 3.2). A pure strategy is a complete contingency plan, specifying an action at every decision point, even unreached ones (Definition 3.4), and fixing strategies reduces the tree to a normal-form game. Backward induction (Procedure 3.6) solves any finite perfect-information game and produces a subgame perfect equilibrium (Theorem 3.9), which refines Nash equilibrium by discarding non-credible threats. The same recursion underlies the P/N analysis of combinatorial games and, in its win-lose-draw form, Zermelo's determinacy theorem (Theorem 3.12). The Centipede game shows that the subgame perfect prediction can diverge from observed human play, a divergence attributable to the failure of common knowledge of rationality or to payoffs the simple model omits.

Exercises

Backward induction and subgame perfection

Exercise 3.1. Solve the following perfect-information game by backward induction. Player I chooses A or B ; after A , Player II chooses C or D ; after B , Player II chooses E or F . Payoffs (u_1, u_2) are $A \rightarrow C : (2, 1)$, $A \rightarrow D : (0, 0)$, $B \rightarrow E : (3, 0)$, $B \rightarrow F : (1, 2)$. State the SPNE strategy profile and the outcome.

Exercise 3.2. Write the normal form of the game in Exercise 3.1. Find all pure Nash equilibria by the best-response method, and identify which are subgame perfect. For each Nash equilibrium that is not subgame perfect, name the non-credible threat that sustains it.

Exercise 3.3. A perfect-information game has Player I moving at the root with three actions; after each, Player II moves with two actions; after each of those, Player I moves again with two actions. How many pure strategies does each player have? (Count information sets and actions; remember that a strategy specifies an action at every decision node of the player.)

Paradoxes, applications, and theory

Exercise 3.4. Verify the backward-induction outcome of the Centipede paradox in Example 3.15 node by node. Then modify the final continuation payoff from $(6, 4)$ to $(6, 6)$ and recompute. Does the outcome change? Explain why or why not in terms of each mover's local comparison.

Exercise 3.5. *The ultimatum game.* Player I proposes a split of one dollar, offering Player II an amount $x \in \{0, 0.25, 0.50, 0.75, 1\}$ and keeping $1 - x$. Player II then accepts, giving payoffs $(1 - x, x)$, or rejects, giving $(0, 0)$. Find the subgame perfect equilibrium. Compare it with how people are observed to play, and relate the difference to the discussion of the Centipede game.

Exercise 3.6. Draw the game tree for Nim with two piles of sizes $(1, 2)$ under normal play, labeling each leaf with the winner. Solve it by backward induction, and confirm that the result agrees with Bouton's theorem (Theorem 2.7): the position $(1, 2)$ has nim-sum $3 \neq 0$, hence is an N-position. What is the winning first move?

Exercise 3.7. Explain why Zermelo's theorem (Theorem 3.12) guarantees that exactly one of the three determinacy outcomes holds for checkers, yet does not tell us which. Contrast this with Nim, where Bouton's theorem gives the answer for every position in closed form. What feature of Nim makes the difference?

PART II

Two-Person Zero-Sum Games

CHAPTER 4

Two-Person Zero-Sum Games and the Minimax Theorem

“There is a value, and you can guarantee it. That is the whole of von Neumann’s theorem, and it took the rest of us twenty years to believe him.”

A two-person zero-sum game is pure conflict: whatever one player wins, the other loses. This is the setting von Neumann solved completely in 1928, and it remains the most satisfying chapter of the theory because the answer is unconditional. Every such game has a *value*, a number the row player can guarantee from below and the column player can guarantee from above, and each has a strategy that achieves it. This chapter states the minimax theorem, proves the part of it that is elementary, and develops the hand methods that solve any small zero-sum game: saddle points, the 2×2 formula, the graphical method for $2 \times n$ and $n \times 2$ games, and reduction by dominance. The remaining half of the theorem, the existence of the value in general, is proved in Chapter 6 through linear programming.

4.1 Zero-sum games and the value

Definition 4.1 (Two-person zero-sum game). A finite *two-person zero-sum game* is a triple (X, Y, A) , where $X = \{1, \dots, m\}$ is the set of Player I’s pure strategies, $Y = \{1, \dots, n\}$ is the set of Player II’s pure strategies, and A is an $m \times n$ **payoff matrix**. The entry a_{ij} is the amount Player II pays Player I when Player I chooses row i and Player II chooses column j . Player I, the **row player**, wants a_{ij} large; Player II, the **column player**, wants it small. Because Player II’s payoff is $-a_{ij}$, the two payoffs sum to zero, and a single matrix A describes the whole game.

The convention is worth fixing in the mind: *rows maximize, columns minimize, and the single number in a cell is a payment from II to I*. A positive value means the game favors Player I, a negative value favors Player II, and a value of zero means the game is *fair*.

Example 4.2 (The even-odd game). Each player secretly chooses 1 or 2. Player I wins the sum from Player II when the sum is odd and pays it when the sum is even. With Player I as the row player, the payoff matrix is

$$A = \begin{bmatrix} -2 & 3 \\ 3 & -4 \end{bmatrix},$$

where, for instance, the top-left entry -2 records that when both choose 1 the sum 2 is even, so Player I pays 2. We will return to this game repeatedly; it is the simplest game with no pure-strategy solution.

4.2 Pure strategies, saddle points, and the two values

Before randomizing, ask what each player can guarantee using a single pure strategy.

Player I reasons defensively: if I commit to row i , the worst that can happen is $\min_j a_{ij}$, since Player II will choose the column that hurts me most. To make this guaranteed floor as high as possible, Player I chooses the row attaining

$$\underline{V} = \max_i \min_j a_{ij}, \tag{4.1}$$

called the **lower value** or **maximin**. Player II reasons symmetrically: if I commit to column j , the worst is $\max_i a_{ij}$, and to make this guaranteed ceiling as low as possible Player II chooses the column attaining

$$\bar{V} = \min_j \max_i a_{ij}, \quad (4.2)$$

the **upper value** or **minimax**.

Proposition 4.3 (Maximin never exceeds minimax). For every payoff matrix, $\underline{V} \leq \bar{V}$. In words, what Player I can guarantee never exceeds what Player II can hold Player I to.

Proof. Let i^* attain the maximin and j^* attain the minimax. For any single entry,

$$\min_j a_{i^*j} \leq a_{i^*j^*} \leq \max_i a_{ij^*},$$

because the left side is a minimum over a set containing $a_{i^*j^*}$ and the right side is a maximum over a set containing $a_{i^*j^*}$. The far left equals $\underline{V}$ by the choice of i^* , and the far right equals $\bar{V}$ by the choice of j^* . Hence $\underline{V} \leq \bar{V}$. ■

When the two values coincide, the common number is achieved by a pair of pure strategies, and neither player can do better.

Definition 4.4 (Saddle point). An entry $a_{i^*j^*}$ is a **saddle point** if it is simultaneously the minimum of its row and the maximum of its column:

$$a_{i^*j^*} = \min_j a_{i^*j} = \max_i a_{ij^*}.$$

Proposition 4.5 (Saddle points are pure-strategy solutions). A game has a saddle point if and only if $\underline{V} = \bar{V}$, and then the common value is the value of the game. At a saddle point $a_{i^*j^*}$, Player I can guarantee at least $a_{i^*j^*}$ by playing row i^* and Player II can guarantee paying at most $a_{i^*j^*}$ by playing column j^* .

Proof. If $a_{i^*j^*}$ is a saddle point then $\bar{V} \leq \max_i a_{ij^*} = a_{i^*j^*} = \min_j a_{i^*j} \leq \underline{V}$, and combined with Proposition 4.3 this forces $\underline{V} = \bar{V} = a_{i^*j^*}$. Conversely, if $\underline{V} = \bar{V}$, the row attaining the maximin and the column attaining the minimax meet at an entry that is its row's minimum and its column's maximum, a saddle point. ■

Example 4.6 (A game with a saddle point). Consider

$$A = \begin{bmatrix} 4 & 1 & -3 \\ 3 & 2 & 5 \\ 0 & 1 & 6 \end{bmatrix}.$$

Row minima (Player I's worst case per row). Taking the smallest entry in each row,

$$\min(4, 1, -3) = -3, \quad \min(3, 2, 5) = 2, \quad \min(0, 1, 6) = 0,$$

so the best of these guaranteed floors is the maximin $\underline{V} = \max(-3, 2, 0) = 2$, attained in row 2. *Column maxima (Player II's worst case per column).* Taking the largest entry in each column,

$$\max(4, 3, 0) = 4, \quad \max(1, 2, 1) = 2, \quad \max(-3, 5, 6) = 6,$$

so the best of these guaranteed ceilings is the minimax $\bar{V} = \min(4, 2, 6) = 2$, attained in column 2. The whole computation fits in one bordered grid: write each row's minimum down the right margin, each column's maximum along the bottom, and look for the entry that is starred in both margins at once.

	c_1	c_2	c_3	row min
r_1	4	1	-3	-3
r_2	3	2	5	2 $\leftarrow \underline{V}$
r_3	0	1	6	0
col max	4	2	6	
		$\uparrow \bar{V}$		

The two red margin numbers are equal, $\underline{V} = \bar{V} = 2$, and they meet at the boxed entry $a_{22} = 2$. That entry is a saddle point: it is the smallest number in its row (so Player II cannot punish row 2 below 2) and the largest in its column (so Player I cannot push column 2 above 2). The value is $V = 2$, Player I plays row 2, and Player II plays column 2. Neither gains by deviating: from a_{22} , every other entry in row 2 is at least 2, so the maximizer has no reason to leave, and every other entry in column 2 is at most 2, so the minimizer has no reason to leave either.

The name is geometric. Imagine the payoff a_{ij} as the height of a surface above the grid of strategy pairs. At a saddle point the surface curves the way a horse's saddle does: walk *along the row* (vary the column) and you are at the *lowest* point of that path, the dip where the rider sits; turn ninety degrees and walk *down the column* (vary the row) and you are at the *highest* point, the rise that runs over the horse's spine. Down one axis it is a valley floor, across the other it is a hilltop, and the saddle point is the single spot that is both at once (Figure 4.1). Geographers call exactly this shape a *mountain pass* or *col*: the lowest gap you can find along the ridge between two peaks, and at the same time the highest point on the trail that crosses from one valley to the next. A traveler going over the range stops climbing right where a hiker walking the ridge stops descending. That meeting point is stable for the same reason the game is: the maximizer (climbing the trail) and the minimizer (walking the ridge down) both find their best stopping point at the very same place, so neither wants to move.




Figure 4.1. The payoff surface near a saddle point. Move along the blue valley line (vary the column, i.e. travel along Player I's row) and the saddle is the lowest point; move along the red ridge line (vary the row, i.e. travel down Player II's column) and the saddle is the highest point. The lone entry that is a row-minimum and a column-maximum at once sits exactly where the two curves cross.

The even-odd game of Example 4.2 has no saddle point: its row minima are -2 and -4 , giving $\underline{V} = -2$, while its column maxima are 3 and 3 , giving $\bar{V} = 3$, and $-2 \neq 3$. When the two values differ, no pure pair is stable, and the players must randomize.

4.3 Mixed strategies and the minimax theorem

Allow each player to randomize. Player I chooses a mixed strategy $\mathbf{p} \in \Delta_m$, Player II chooses $\mathbf{q} \in \Delta_n$, and by (1.2) the expected payoff to Player I is $\mathbf{p}^\top A \mathbf{q}$. The lower and upper values are now taken over mixed strategies,

$$\underline{V} = \max_{\mathbf{p} \in \Delta_m} \min_{\mathbf{q} \in \Delta_n} \mathbf{p}^\top A \mathbf{q}, \quad \bar{V} = \min_{\mathbf{q} \in \Delta_n} \max_{\mathbf{p} \in \Delta_m} \mathbf{p}^\top A \mathbf{q},$$

and the inequality $\underline{V} \leq \bar{V}$ continues to hold by the same argument as Proposition 4.3. Von Neumann's theorem is the remarkable assertion that, once mixing is allowed, the inequality is always an equality.

Theorem 4.7 (Minimax theorem; von Neumann, 1928). For every finite two-person zero-sum game with payoff matrix A , there is a number V , the **value** of the game, and optimal mixed strategies $\mathbf{p}^* \in \Delta_m$ and $\mathbf{q}^* \in \Delta_n$, such that

$$\max_{\mathbf{p} \in \Delta_m} \min_{\mathbf{q} \in \Delta_n} \mathbf{p}^\top A \mathbf{q} = \min_{\mathbf{q} \in \Delta_n} \max_{\mathbf{p} \in \Delta_m} \mathbf{p}^\top A \mathbf{q} = V. \quad (4.3)$$

Equivalently, $\mathbf{p}^*$ guarantees Player I at least V against every column, $(\mathbf{p}^*)^\top A \mathbf{q} \geq V$ for all $\mathbf{q}$, and $\mathbf{q}^*$ guarantees Player II paying at most V against every row, $\mathbf{p}^\top A \mathbf{q}^* \leq V$ for all $\mathbf{p}$. The value V is unique, though the optimal strategies need not be.

Historical Note: John von Neumann

The minimax theorem is the first deep result of game theory, proved by John von Neumann (1903–1957) in 1928. A prodigy born in Budapest, he reshaped field after field: the mathematical foundations of quantum mechanics, ergodic theory, operator algebras, and, with Oskar Morgenstern, the 1944 book that founded game theory itself. During the war he worked on the Manhattan Project, and afterward he set out the design for a stored-program computer, the architecture nearly every machine still follows, along with a theory of self-replicating automata. Few people have left so deep a mark on so many subjects. He died of cancer in 1957.

The minimax theorem

Every finite two-person zero-sum game has a single number V , its value, that the row player can guarantee from below and the column player can guarantee from above:

$$\max_x \min_y x^\top A y = \min_y \max_x x^\top A y = V.$$

There is no advantage to moving second. This is the theorem on which the whole zero-sum theory rests; Chapter 6 proves it from linear programming duality.

The content of (4.3) is that the order of play does not matter. Whether Player I must reveal a mixed strategy first (the left side, where Player II then best-responds) or Player II must reveal first (the right side), the guaranteed payoff is the same number V . Foresight is worthless against an optimal opponent. The proof of existence in the general case is deferred to Chapter 6, where it falls out of linear programming duality. The easy direction, $\underline{V} \leq \bar{V}$, we already have, and for 2×2 games we now prove the theorem by direct construction.

4.4 The 2×2 game in closed form

When a 2×2 game has no saddle point, its value and optimal strategies are given by explicit formulas. Write the matrix in the standard labeling

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix},$$

with Player I mixing $(p, 1 - p)$ over the two rows and Player II mixing $(q, 1 - q)$ over the two columns.

Theorem 4.8 (Solution of a 2×2 game). If the 2×2 game $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$ has no saddle point, then $D := a - b - c + d \neq 0$, and the optimal strategies and value are

$$p^* = \frac{d - c}{D}, \quad q^* = \frac{d - b}{D}, \quad V = \frac{ad - bc}{D} = \frac{\det A}{a - b - c + d}, \quad (4.4)$$

where p^* is the probability Player I plays row 1 and q^* is the probability Player II plays column 1.

Proof. We use the *principle of indifference*. Lacking a saddle point, Player I chooses the mixing probability p so that Player II's two pure columns yield Player I the same expected payoff; then neither column does better than the other, and Player II cannot exploit Player I's choice.

Step 1: Player I's mixture. Player I plays row 1 with probability p and row 2 with probability $1 - p$. Against Player II's column 1 the expected payoff to Player I is

$$U_1(p) = p \cdot a + (1 - p) \cdot c,$$

and against column 2 it is

$$U_2(p) = p \cdot b + (1 - p) \cdot d.$$

Indifference is the equation $U_1(p) = U_2(p)$. Expanding each side by distributing the multiplication,

$$pa + c - pc = pb + d - pd.$$

Move every term containing p to the left and every constant to the right:

$$pa - pc - pb + pd = d - c.$$

Factor p out of the left side:

$$p(a - c - b + d) = d - c, \quad \text{i.e.} \quad p(a - b - c + d) = d - c.$$

Write $D = a - b - c + d$ for the common denominator. Dividing by D , which is nonzero by the last part of the proof,

$$p^* = \frac{d - c}{D}.$$

Step 2: Player II's mixture. By the mirror-image computation, Player II plays column 1 with probability q and column 2 with probability $1 - q$, chosen so that Player I's two *rows* pay equally. Row 1 pays Player I $qa + (1 - q)b$ and row 2 pays $qc + (1 - q)d$; setting these equal,

$$qa + b - qb = qc + d - qd \implies qa - qb - qc + qd = d - b \implies q(a - b - c + d) = d - b,$$

so that

$$q^* = \frac{d - b}{D}.$$

Step 3: the value. Substitute p^* into Player I's column-1 payoff U_1 . We first need $1 - p^*$:

$$1 - p^* = 1 - \frac{d - c}{D} = \frac{D - (d - c)}{D} = \frac{(a - b - c + d) - d + c}{D} = \frac{a - b}{D}.$$

Then

$$V = U_1(p^*) = p^*a + (1 - p^*)c = \frac{(d - c)a}{D} + \frac{(a - b)c}{D} = \frac{ad - ac + ac - bc}{D} = \frac{ad - bc}{D},$$

where the cross terms $-ac$ and $+ac$ cancel. Thus $V = \det A/D$. Substituting into column 2, or evaluating either row at q^* , gives the same number, a useful check in practice.

It remains to show that $D \neq 0$ when there is no saddle point. We prove the contrapositive: if $D = 0$, the game has a saddle point. First note that swapping the two rows or the two columns of A negates D (for instance, the row swap sends $a - b - c + d$ to $c - d - a + b = -D$) and preserves the property of having a saddle point, since it only relabels the players' strategies. So we may relabel until a is the largest of the four entries, which keeps $D = 0$.

Now suppose $D = 0$, that is $a + d = b + c$, with a largest. Solving for the opposite-corner entry,

$$d = b + c - a \leq b \quad (\text{since } c \leq a), \quad d = b + c - a \leq c \quad (\text{since } b \leq a),$$

so d is the smallest entry. Consider the top-right entry b . It is the minimum of its own row $\{a, b\}$, because $b \leq a$; and it is the maximum of its own column $\{b, d\}$, because $b \geq d$. An entry that is its row's minimum and its column's maximum is a saddle point (Definition 4.4). Hence $D = 0$ forces a saddle, and contrapositively, no saddle forces $D \neq 0$, so the formulas above are well defined. ■

Example 4.9 (Solving the even-odd game). We solve the even-odd game $A = \begin{bmatrix} -2 & 3 \\ 3 & -4 \end{bmatrix}$ of Example 4.2 twice: first from scratch by indifference, then as a check against the closed form. Doing it both ways is the habit worth forming.

From scratch. Let Player I play row 1 with probability p and row 2 with probability $1 - p$. Against Player II's column 1, the expected payoff to Player I is

$$U_1(p) = (-2)p + 3(1 - p) = -2p + 3 - 3p = 3 - 5p,$$

and against column 2,

$$U_2(p) = 3p + (-4)(1 - p) = 3p - 4 + 4p = 7p - 4.$$

These are two straight lines in p . Player II, the minimizer, will always pick whichever column is lower, so Player I wants the value of p at which the two lines cross, where neither column can be turned against him. Set $U_1(p) = U_2(p)$:

$$3 - 5p = 7p - 4 \implies 3 + 4 = 7p + 5p \implies 7 = 12p \implies p^* = \frac{7}{12}.$$

So Player I plays $(\frac{7}{12}, \frac{5}{12})$. The value is the common height of the two lines at p^* :

$$V = U_1\left(\frac{7}{12}\right) = 3 - 5 \cdot \frac{7}{12} = \frac{36}{12} - \frac{35}{12} = \frac{1}{12},$$

and the other line agrees, $U_2\left(\frac{7}{12}\right) = 7 \cdot \frac{7}{12} - 4 = \frac{49}{12} - \frac{48}{12} = \frac{1}{12}$, as it must. Repeating the argument for Player II, who mixes $(q, 1 - q)$ to equalize Player I's rows, row 1 pays $-2q + 3(1 - q) = 3 - 5q$ and row 2 pays $3q - 4(1 - q) = 7q - 4$; setting $3 - 5q = 7q - 4$ gives $q^* = \frac{7}{12}$, so Player II also plays $\left(\frac{7}{12}, \frac{5}{12}\right)$.

By the formula. With $a = -2$, $b = 3$, $c = 3$, $d = -4$,

$$D = a - b - c + d = -2 - 3 - 3 - 4 = -12,$$

and

$$\begin{aligned} p^* &= \frac{d - c}{D} = \frac{-4 - 3}{-12} = \frac{-7}{-12} = \frac{7}{12}, & q^* &= \frac{d - b}{D} = \frac{-4 - 3}{-12} = \frac{7}{12}, \\ V &= \frac{ad - bc}{D} = \frac{(-2)(-4) - (3)(3)}{-12} = \frac{8 - 9}{-12} = \frac{-1}{-12} = \frac{1}{12}, \end{aligned}$$

matching the from-scratch values exactly. The game slightly favors Player I, the player who profits from odd sums: he can guarantee an expected gain of $\frac{1}{12}$ per play, and no strategy of Player II can hold him below it. $\diamond$

Remark 4.10. A caution on signs. The value of the even-odd game is $+\frac{1}{12}$, a positive number, so Player I, the player who profits from odd sums, holds the advantage; a value of $-\frac{1}{12}$ would point the wrong way. The closed-form (4.4) is sensitive to the labeling of the matrix, so it is always worth verifying the value by substituting the computed strategy, as we did above. The formula is a convenience; the indifference principle behind it is the real content, and it is developed in full generality in Chapter 5.

Example 4.11 (The penalty kick). A soccer penalty kick is a near-perfect zero-sum game. The kicker chooses to shoot LEFT or RIGHT; the goalkeeper, committing before the ball is struck, dives LEFT or RIGHT. The entry is the kicker's scoring probability (in percent), which the kicker maximizes and the keeper minimizes. A representative matrix for a right-footed kicker, whose natural side is the right, is

$$A = \begin{bmatrix} 40 & 80 \\ 90 & 50 \end{bmatrix} \quad \begin{array}{l} \text{rows: kick L, kick R} \\ \text{cols: dive L, dive R} \end{array}$$

A shot succeeds more often when it goes opposite the dive, and the strong (right) side scores well even when the keeper guesses correctly. The row minima are $\min(40, 80) = 40$ and $\min(90, 50) = 50$, so the maximin is 50; the column maxima are $\max(40, 90) = 90$ and $\max(80, 50) = 80$, so the minimax is 80. Since $50 \neq 80$ there is no saddle, and both players randomize. With $a = 40$, $b = 80$, $c = 90$, $d = 50$ the denominator is $D = a - b - c + d = 40 - 80 - 90 + 50 = -80$, and

$$\begin{aligned} p^* &= \frac{d - c}{D} = \frac{50 - 90}{-80} = \frac{1}{2}, & q^* &= \frac{d - b}{D} = \frac{50 - 80}{-80} = \frac{3}{8}, \\ V &= \frac{ad - bc}{D} = \frac{40 \cdot 50 - 80 \cdot 90}{-80} = \frac{-5200}{-80} = 65. \end{aligned}$$

The kicker should shoot left and right with equal probability, $\mathbf{p}^* = \left(\frac{1}{2}, \frac{1}{2}\right)$; the keeper should dive left only $\frac{3}{8}$ of the time, favoring the kicker's strong side; and the kicker scores 65% of the time at equilibrium. To verify, the keeper's mix $q^* = \frac{3}{8}$ holds both shots to the same value: kick left yields $40 \cdot \frac{3}{8} + 80 \cdot \frac{5}{8} = 15 + 50 = 65$, and kick right yields $90 \cdot \frac{3}{8} + 50 \cdot \frac{5}{8} = 33.75 + 31.25 = 65$. Remarkably, field studies of professional soccer find that real kickers and keepers mix in close agreement with such minimax predictions, one of the cleanest empirical confirmations of the theory [41].

$\diamond$

4.5 Domination and reduction

Large matrices often shrink, because some pure strategies are never worth playing.

Definition 4.12 (Domination of rows and columns). Row i **dominates** row k if $a_{ij} \geq a_{kj}$ for every column j . Column j dominates column ℓ if $a_{ij} \leq a_{i\ell}$ for every row i . More generally, a row is dominated if it is dominated entrywise by some convex combination (a mixed strategy) of the other rows, and similarly for columns.

Player I, the maximizer, never needs a dominated row, and Player II, the minimizer, never needs a dominated column. Deleting them leaves the value unchanged and may reduce the game to a size the closed form handles. The reason the value is preserved is that an optimal strategy can always avoid a dominated pure strategy by shifting its weight to the dominating combination without lowering the payoff.

Example 4.13 (Reducing a 3×3 game by pure dominance). Take

$$A = \begin{bmatrix} 2 & 0 & 4 \\ 1 & 2 & 3 \\ 4 & 1 & 2 \end{bmatrix}.$$

Step 1: a dominated column. Player II minimizes, so she never needs a column whose every entry is at least as large as another column's. Compare columns 2 and 3 entry by entry: $0 \leq 4$, $2 \leq 3$, $1 \leq 2$. Column 2 is at most column 3 in every row (shown in red), so whatever Player I plays, column 2 costs Player II no more than column 3:

$$A = \begin{bmatrix} 2 & 0 & 4 \\ 1 & 2 & 3 \\ 4 & 1 & 2 \end{bmatrix} \xrightarrow{\text{delete col. 3}} \begin{bmatrix} 2 & 0 \\ 1 & 2 \\ 4 & 1 \end{bmatrix}.$$

Step 2: a dominated row. Player I maximizes, so he never needs a row whose every entry is at most another row's. In the matrix just obtained, compare rows 1 and 3: $2 \leq 4$ and $0 \leq 1$, so row 1 (in red) is at most row 3 everywhere, and Player I discards it, leaving the 2×2 game on rows $\{2, 3\}$ and columns $\{1, 2\}$:

$$\begin{bmatrix} 2 & 0 \\ 1 & 2 \\ 4 & 1 \end{bmatrix} \xrightarrow{\text{delete row 1}} \begin{bmatrix} 1 & 2 \\ 4 & 1 \end{bmatrix}.$$

Step 3: solve it. The row minima are $\min(1, 2) = 1$ and $\min(4, 1) = 1$, so the maximin is 1; the column maxima are $\max(1, 4) = 4$ and $\max(2, 1) = 2$, so the minimax is 2; since $1 \neq 2$ there is no saddle. By Theorem 4.8 with $a = 1, b = 2, c = 4, d = 1$ and $D = 1 - 2 - 4 + 1 = -4$,

$$p^* = \frac{1-4}{-4} = \frac{3}{4}, \quad q^* = \frac{1-2}{-4} = \frac{1}{4}, \quad V = \frac{1 \cdot 1 - 2 \cdot 4}{-4} = \frac{-7}{-4} = \frac{7}{4}.$$

Step 4: restore and verify. Restoring the discarded strategies with zero weight, the full solution is

$$p^* = \left(0, \frac{3}{4}, \frac{1}{4}\right), \quad q^* = \left(\frac{1}{4}, \frac{3}{4}, 0\right), \quad V = \frac{7}{4}.$$

The cleanest check is to play p^* against each of the three *original* columns and confirm none pays Player II less than $\frac{7}{4}$:

$$\begin{aligned} \text{column 1: } & 0 \cdot 2 + \frac{3}{4} \cdot 1 + \frac{1}{4} \cdot 4 = \frac{3}{4} + 1 = \frac{7}{4}, \\ \text{column 2: } & 0 \cdot 0 + \frac{3}{4} \cdot 2 + \frac{1}{4} \cdot 1 = \frac{6}{4} + \frac{1}{4} = \frac{7}{4}, \\ \text{column 3: } & 0 \cdot 4 + \frac{3}{4} \cdot 3 + \frac{1}{4} \cdot 2 = \frac{9}{4} + \frac{2}{4} = \frac{11}{4}. \end{aligned}$$

Columns 1 and 2 both pay exactly $\frac{7}{4}$, and the discarded column 3 pays $\frac{11}{4} > \frac{7}{4}$, so Player II gains nothing by ever using it. Player I's mixture p^* guarantees at least $\frac{7}{4}$ no matter what, which matches the value. $\diamond$

Example 4.14 (Reducing by mixed dominance). Sometimes no pure strategy dominates, but a mixture does. Take

$$A = \begin{bmatrix} 0 & 4 & 6 \\ 5 & 7 & 4 \\ 9 & 6 & 3 \end{bmatrix}.$$

Step 0: confirm no pure dominance. Player II compares columns. Column 2 $= (4, 7, 6)^\top$ is not entrywise above column 1 $= (0, 5, 9)^\top$, because $6 < 9$ in the third row; nor above column 3 $= (6, 4, 3)^\top$, because $4 < 6$ in the first row. So no single column dominates column 2, and the pure-dominance trick of Example 4.13 is stuck.

Step 1: a dominating mixture. Try splitting Player II's weight evenly between columns 1 and 3. The resulting mixed column is the average

$$\frac{1}{2}(0, 5, 9)^\top + \frac{1}{2}(6, 4, 3)^\top = \left(\frac{0+6}{2}, \frac{5+4}{2}, \frac{9+3}{2}\right)^\top = (3, 4.5, 6)^\top.$$

Compare it entrywise to column 2 $= (4, 7, 6)^\top$: $3 \leq 4$, $4.5 \leq 7$, and $6 \leq 6$. The mixed column (in red below) is at most column 2 in every row, so for any row Player I picks, the $\frac{1}{2}c_1 + \frac{1}{2}c_3$ mixture costs Player II no more than c_2 does, and she discards c_2 :

$$\begin{bmatrix} 0 & 4 & 6 \\ 5 & 7 & 4 \\ 9 & 6 & 3 \end{bmatrix} \xrightarrow{\text{delete col. 2}} \begin{bmatrix} 0 & 6 \\ 5 & 4 \\ 9 & 3 \end{bmatrix}.$$

Step 2: a dominated row. Now Player I looks for a row beaten by a mixture of the others. Try weighting rows 1 and 3 by $\frac{1}{3}$ and $\frac{2}{3}$:

$$\frac{1}{3}(0, 6) + \frac{2}{3}(9, 3) = \left(\frac{0+18}{3}, \frac{6+6}{3}\right) = (6, 4).$$

Entrywise this beats row 2 = (5, 4), since $6 \geq 5$ and $4 \geq 4$. So row 2 (in red) is dominated by the mixture, Player I discards it, and the game collapses to

$$\begin{bmatrix} 0 & 6 \\ \textcolor{red}{5} & \textcolor{red}{4} \\ 9 & 3 \end{bmatrix} \xrightarrow{\text{delete row 2}} \begin{bmatrix} 0 & 6 \\ 9 & 3 \end{bmatrix} \quad (\text{rows } r_1, r_3; \text{ columns } c_1, c_3).$$

Step 3: solve the 2×2 . With $a = 0, b = 6, c = 9, d = 3$, the denominator is $D = a - b - c + d = 0 - 6 - 9 + 3 = -12$, so

$$p^* = \frac{d - c}{D} = \frac{3 - 9}{-12} = \frac{-6}{-12} = \frac{1}{2}, \quad q^* = \frac{d - b}{D} = \frac{3 - 6}{-12} = \frac{-3}{-12} = \frac{1}{4},$$

$$V = \frac{ad - bc}{D} = \frac{0 \cdot 3 - 6 \cdot 9}{-12} = \frac{-54}{-12} = \frac{9}{2}.$$

Step 4: restore and verify. Restoring the discarded strategies with weight 0,

$$p^* = \left(\frac{1}{2}, 0, \frac{1}{2}\right), \quad q^* = \left(\frac{1}{4}, 0, \frac{3}{4}\right), \quad V = \frac{9}{2}.$$

Play p^* against the original columns: column 1 pays $\frac{1}{2} \cdot 0 + 0 \cdot 5 + \frac{1}{2} \cdot 9 = \frac{9}{2}$; column 3 pays $\frac{1}{2} \cdot 6 + 0 \cdot 4 + \frac{1}{2} \cdot 3 = 3 + \frac{3}{2} = \frac{9}{2}$; and the discarded column 2 pays $\frac{1}{2} \cdot 4 + 0 \cdot 7 + \frac{1}{2} \cdot 6 = 2 + 3 = 5 > \frac{9}{2}$. Every column pays at least $\frac{9}{2}$, confirming both the discard and the value. $\diamond$

Intuition: a dominated strategy is a strictly worse job offer

A dominated row is one you would never play, the way you would never take a job that pays less *and* has worse hours *and* a worse commute than another offer on the table: whatever the world does, the other option is at least as good. Deleting it loses nothing, because an optimal plan would put zero weight on it anyway. The subtle part is that the comparison can be against a *blend* of other options (Example 4.14), not just a single one, just as a salary-plus-benefits package might only be beaten by splitting your week between two other jobs.

Example 4.15 (Two rounds of elimination, shown on the grid). Some games peel away in stages: a column falls, which exposes a row, which leaves a clean 2×2 . Take

$$A = \begin{bmatrix} 4 & 0 & 5 \\ 0 & 3 & 4 \\ 3 & -1 & 4 \end{bmatrix}.$$

Step 0: no saddle. The row minima are 0, 0, -1, so $\underline{V} = 0$; the column maxima are 4, 3, 5, so $\overline{V} = 3$. Since $0 \neq 3$ the game needs mixing, and we shrink it first.

Step 1: a dominated column. Player II pays the column entries and wants them small, so she compares columns entry by entry. Column 3 = (5, 4, 4)^T sits against column 1 = (4, 0, 3)^T row by row: $5 \geq 4$, $4 \geq 0$, and $4 \geq 3$, so column 3 is *everywhere at least* column 1. Whatever row Player I picks, column 3 costs Player II as much or more than column 1, so she discards it. In red, with an arrow to what remains:

$$\begin{bmatrix} 4 & 0 & \textcolor{red}{5} \\ 0 & 3 & \textcolor{red}{4} \\ 3 & -1 & \textcolor{red}{4} \end{bmatrix} \xrightarrow{\text{delete col. 3}} \begin{bmatrix} 4 & 0 \\ 0 & 3 \\ 3 & -1 \end{bmatrix}.$$

Step 2: a dominated row, newly visible. With column 3 gone, look at the rows Player I maximizes. Row 3 = (3, -1) is *everywhere at most* row 1 = (4, 0), since $3 \leq 4$ and $-1 \leq 0$. Player I would never prefer row 3 to row 1, so it goes:

$$\begin{bmatrix} 4 & 0 \\ 0 & 3 \\ \textcolor{red}{3} & \textcolor{red}{-1} \end{bmatrix} \xrightarrow{\text{delete row 3}} \begin{bmatrix} 4 & 0 \\ 0 & 3 \end{bmatrix}.$$

This row only became removable *after* the column left; that is the point of working in rounds.

Step 3: solve the 2×2 . With $a = 4, b = 0, c = 0, d = 3$ the denominator is $D = a - b - c + d = 7$, so

$$p^* = \frac{d - c}{D} = \frac{3}{7}, \quad q^* = \frac{d - b}{D} = \frac{3}{7}, \quad V = \frac{ad - bc}{D} = \frac{12}{7}.$$

Restoring the discarded strategies with weight 0, Player I plays $\mathbf{p}^* = (\frac{3}{7}, \frac{4}{7}, 0)$ and Player II plays $\mathbf{q}^* = (\frac{3}{7}, \frac{4}{7}, 0)$.

Step 4: verify against the original matrix. Play $\mathbf{p}^*$ and read every original column; play $\mathbf{q}^*$ and read every original row:

vs. columns ($\mathbf{p}^*$)	c_1	c_2	c_3		vs. rows ($\mathbf{q}^*$)	r_1	r_2	r_3
pays II	$\frac{12}{7}$	$\frac{12}{7}$	$\frac{31}{7}$		pays I	$\frac{12}{7}$	$\frac{12}{7}$	$\frac{5}{7}$

Both kept strategies pay exactly $\frac{12}{7}$, the discarded column 3 costs Player II *more* ($\frac{31}{7} > \frac{12}{7}$, so she was right to avoid it), and the discarded row 3 earns Player I *less* ($\frac{5}{7} < \frac{12}{7}$, so he was right to avoid it). The value is $V = \frac{12}{7}$. $\diamond$

4.6 The graphical method: $2 \times n$ and $n \times 2$ games

When one player has exactly two pure strategies, the game can be solved by a picture, because the relevant payoffs are straight lines.

4.6.1 The $2 \times n$ case

Suppose Player I has two rows and Player II has n columns. If Player I plays row 1 with probability p , then against Player II's column j the expected payoff is the line

$$U_j(p) = p a_{1j} + (1 - p) a_{2j}, \quad 0 \leq p \leq 1.$$

For each p , Player II will pick the column giving the smallest payoff, so Player I actually receives the *lower envelope* $\min_j U_j(p)$. Player I chooses p to maximize this envelope. The maximum of the lower envelope is the value, and it occurs where two of the lines cross.

Example 4.16 (A 2×4 game). Let

$$A = \begin{bmatrix} 2 & 3 & 1 & 5 \\ 4 & 1 & 6 & 0 \end{bmatrix}.$$

Each column j gives one line $U_j(p) = p a_{1j} + (1 - p) a_{2j}$, read straight off the two entries in that column. Substituting the columns of A one at a time:

$$\begin{aligned} U_1(p) &= p \cdot 2 + (1 - p) \cdot 4 = 4 - 2p, & U_2(p) &= p \cdot 3 + (1 - p) \cdot 1 = 1 + 2p, \\ U_3(p) &= p \cdot 1 + (1 - p) \cdot 6 = 6 - 5p, & U_4(p) &= p \cdot 5 + (1 - p) \cdot 0 = 5p. \end{aligned}$$

Reading the envelope. Player II answers each p with the smallest line, so Player I receives $\min_j U_j(p)$. Check the two ends. At $p = 0$ the lines read $(U_1, U_2, U_3, U_4) = (4, 1, 6, 0)$, so the smallest is $U_4 = 0$. At $p = 1$ they read $(2, 3, 1, 5)$, so the smallest is $U_3 = 1$. As p sweeps from 0 to 1 the binding line therefore changes, and the lower envelope is the thick curve in Figure 4.2: it climbs along U_4 , switches to U_2 , then to U_3 .

Finding the peak. Player I wants the highest point of that envelope. Where it climbs along $U_4 = 5p$ it is increasing, and where it falls along $U_3 = 6 - 5p$ it is decreasing, so the maximum sits at the junction in the middle, where the rising piece U_2 meets the falling piece U_3 . Setting them equal,

$$1 + 2p = 6 - 5p \implies 2p + 5p = 6 - 1 \implies 7p = 5 \implies p^* = \frac{5}{7}.$$

The value is the height there. Using either active line as a check,

$$V = U_2(p^*) = 1 + 2 \cdot \frac{5}{7} = 1 + \frac{10}{7} = \frac{17}{7}, \quad U_3(p^*) = 6 - 5 \cdot \frac{5}{7} = 6 - \frac{25}{7} = \frac{17}{7} \checkmark$$

so Player I plays $\mathbf{p}^* = (\frac{5}{7}, \frac{2}{7})$ for a guaranteed $\frac{17}{7}$.

Player II's mixture. Only the two columns active at the peak, $j = 2$ and $j = 3$, can carry weight; the others (U_1, U_4) lie strictly above the envelope at p^* and are wasted. So Player II mixes columns 2 and 3, say with probabilities $(q, 1 - q)$,




Figure 4.2. The 2×4 game of Example 4.16. The thick curve is the lower envelope $\min_j U_j(p)$; Player I maximizes it at $p^* = 5/7$, where lines U_2 and U_3 cross.

on the 2×2 subgame $\begin{bmatrix} 3 & 1 \\ 1 & 6 \end{bmatrix}$. By the indifference principle of Section 5.1, Player II picks q to make Player I indifferent between his two rows:

$$\underbrace{3q + 1(1-q)}_{\text{row 1}} = \underbrace{1 \cdot q + 6(1-q)}_{\text{row 2}} \implies 1 + 2q = 6 - 5q \implies 7q = 5 \implies q = \frac{5}{7}.$$

Player II plays columns 2, 3 with weights $(\frac{5}{7}, \frac{2}{7})$ and the rest zero, i.e. $\mathbf{q}^* = (0, \frac{5}{7}, \frac{2}{7}, 0)$. As a final check, the common row payoff is $1 + 2 \cdot \frac{5}{7} = \frac{17}{7} = V$, matching Player I's guarantee. $\diamond$

4.6.2 The $n \times 2$ case

Now Player II has two columns and Player I has n rows. If Player II plays column 1 with probability q , the payoff if Player I plays row i is the line $U_i(q) = q a_{i1} + (1-q) a_{i2}$. Player I, given q , picks the largest, so Player I receives the *upper envelope* $\max_i U_i(q)$, and Player II chooses q to minimize it.

Example 4.17 (An $n \times 2$ game with a pure best row). Let

$$A = \begin{bmatrix} 1 & 5 \\ 4 & 4 \\ 6 & 2 \end{bmatrix}.$$

The three lines. Each row i gives one line $U_i(q) = q a_{i1} + (1-q) a_{i2}$, read off that row's two entries:

$$\begin{aligned} U_1(q) &= q \cdot 1 + (1-q) \cdot 5 = 5 - 4q, \\ U_2(q) &= q \cdot 4 + (1-q) \cdot 4 = 4, \\ U_3(q) &= q \cdot 6 + (1-q) \cdot 2 = 2 + 4q. \end{aligned}$$

The upper envelope. Player I answers each q with his best row, so he receives $\max_i U_i(q)$, and Player II picks q to push that maximum down. Here row 2 is flat at $U_2 \equiv 4$, so the maximum can never fall below 4: whatever q is, Player I can always take row 2 and collect exactly 4.

When can Player II hold the maximum at 4? She needs the other two lines to sit at or below 4 simultaneously:

$$5 - 4q \leq 4 \implies 4q \geq 1 \implies q \geq \frac{1}{4}, \quad 2 + 4q \leq 4 \implies 4q \leq 2 \implies q \leq \frac{1}{2}.$$

Both hold precisely for $q \in [\frac{1}{4}, \frac{1}{2}]$. On that interval the envelope equals 4; outside it, one of U_1, U_3 pokes above 4 and the maximum is larger. Hence

$$V = \min_q \max_i U_i(q) = 4.$$

Optimal strategies. Player I guarantees the value with the single pure row 2, so $\mathbf{p}^* = (0, 1, 0)$. Player II achieves it with *any* $\mathbf{q} \in [\frac{1}{4}, \frac{1}{2}]$, for instance $\mathbf{q} = \frac{1}{4}$ (where U_1 just touches 4) or $\mathbf{q} = \frac{1}{2}$ (where U_3 just touches 4). The value is unique, but Player II's optimal strategy is a whole interval, a useful reminder that optimal strategies need not be unique even when the value is. $\diamond$

4.7 The payoff-value lemma

A useful identity confirms a computed solution and underlies the principle of indifference.

Lemma 4.18 (Payoff-value lemma). If $\mathbf{p}^*$ is optimal for Player I and $\mathbf{q}^*$ is optimal for Player II in a game of value V , then the expected payoff under the optimal pair equals the value:

$$(\mathbf{p}^*)^\top A \mathbf{q}^* = V.$$

Proof. Because $\mathbf{p}^*$ is optimal, it guarantees at least V against every column, so in particular $\sum_i p_i^* a_{ij} \geq V$ for each j . Multiplying by $q_j^* \geq 0$ and summing over j , using $\sum_j q_j^* = 1$,

$$(\mathbf{p}^*)^\top A \mathbf{q}^* = \sum_j \left(\sum_i p_i^* a_{ij} \right) q_j^* \geq \sum_j V q_j^* = V.$$

Because $\mathbf{q}^*$ is optimal, it holds Player I to at most V against every row, so $\sum_j a_{ij} q_j^* \leq V$ for each i . Multiplying by $p_i^* \geq 0$ and summing,

$$(\mathbf{p}^*)^\top A \mathbf{q}^* = \sum_i p_i^* \left(\sum_j a_{ij} q_j^* \right) \leq \sum_i p_i^* V = V.$$

The two inequalities force $(\mathbf{p}^*)^\top A \mathbf{q}^* = V$. $\blacksquare$

The lemma also shows that at the optimum a player mixes only over pure strategies that achieve the value: if $p_i^* > 0$ but row i earns strictly less than V against $\mathbf{q}^*$, the weighted sum could not reach V . This is the principle of indifference, and it is the engine of Chapter 5.

4.8 A recipe for small zero-sum games

The methods of this chapter combine into a procedure that solves any small game by hand.

- Procedure 4.19 (Solving a two-person zero-sum game).**
1. **Saddle point.** Compute the row minima and column maxima. If $\max_i \min_j a_{ij} = \min_j \max_i a_{ij}$, that common number is the value and the attaining row and column are optimal pure strategies. Stop.
 2. **Dominance.** Delete any dominated row (for Player I) or column (for Player II), pure or mixed, to shrink the matrix.
 3. **Closed form.** If the game is now 2×2 , apply Theorem 4.8.
 4. **Graphical method.** If one player has exactly two pure strategies, plot the payoff lines and read the value off the appropriate envelope (Section 4.6).
 5. **General case.** Otherwise solve the linear program of Chapter 6.
 6. **Check.** Verify with the payoff-value lemma that $(\mathbf{p}^*)^\top A \mathbf{q}^* = V$.

This recipe handles every example in the chapter and most of the problem-set and examination questions. The methods it leaves open, the principle of indifference for larger games, invariance and symmetry, and the linear program, are the subjects of the next two chapters.

Notes and References

The minimax theorem is due to von Neumann [63], who had first presented it to the Göttingen Mathematical Society on 7 December 1926, two years before the printed paper. Émile Borel introduced the mixed-strategy idea earlier (Borel 5) but doubted that a general minimax theorem could hold; Weinstein [65] shows that the widely repeated claim “Borel believed it false for more than seven strategies” is a mistranslation of the 1953 English edition, in which Borel’s “as soon as n exceeds 3” became “7.” The first *elementary* proof, by convexity and separating hyperplanes rather than a fixed point, is Ville [62], and it is the bridge to the linear-programming view of Chapter 6. The treatise of von Neumann and Morgenstern [64] made the theorem the cornerstone of the new field. The hand methods of this chapter, the saddle test, the 2×2 formula, dominance, and the graphical envelope, follow Ferguson [16]; Osborne [39] and Karlin and Peres [26] give parallel treatments. The proof of the theorem itself is deferred to Chapter 6; a second route through fixed points appears in Chapter 8 as a special case of Nash’s theorem.

Chapter summary

A two-person zero-sum game is described by a single payoff matrix A in which the row player maximizes and the column player minimizes (Definition 4.1). With pure strategies, the row player can guarantee the maximin $\underline{V} = \max_i \min_j a_{ij}$ and the column player the minimax $\bar{V} = \min_j \max_i a_{ij}$, and always $\underline{V} \leq \bar{V}$ (Proposition 4.3). Equality holds exactly when a saddle point exists, giving a pure-strategy solution (Proposition 4.5). When no saddle point exists, mixing is required, and von Neumann’s minimax theorem (Theorem 4.7) guarantees a unique value V with $\max_p \min_q p^\top A q = \min_q \max_p p^\top A q = V$. For 2×2 games the value and optimal strategies have a closed form (Theorem 4.8); larger games reduce by dominance (Definition 4.12) and, when one player has two strategies, by the graphical envelope method (Section 4.6). The payoff-value lemma (Lemma 4.18) verifies solutions and yields the principle of indifference. The full proof of the minimax theorem awaits the linear programming of Chapter 6.

Exercises

The exercises are grouped by method. Those marked (*Worked in Appendix D*) have full solutions in the back; the rest are answered by the same techniques.

Saddle points and the upper and lower values

Exercise 4.1. For each matrix, find the maximin and minimax, decide whether a saddle point exists, and if so give the value and optimal pure strategies.

$$(a) \begin{bmatrix} 3 & 1 \\ 0 & 2 \end{bmatrix}, \quad (b) \begin{bmatrix} -2 & 2 & -1 \\ 5 & 3 & 1 \\ 2 & 0 & 1 \end{bmatrix}, \quad (c) \begin{bmatrix} 1 & 2 & 3 & 4 & 5 \\ 3 & 4 & 5 & 6 & 7 \end{bmatrix}.$$

Exercise 4.2. The game $A = \begin{bmatrix} 6 & 5 \\ 2 & 3 \end{bmatrix}$ looks like it needs mixing. Tabulate the row minima and column maxima, compute the maximin and minimax, and decide whether a saddle point exists. If it does, give the value and the optimal pure strategies, and explain why neither player gains by randomizing. (*Worked in Appendix D.*)

Exercise 4.3. Explain in words why $\underline{V} \leq \bar{V}$ holds for every matrix, but $\underline{V} = \bar{V}$ can fail without mixing, using the even-odd game (Example 4.2) as the example. What does the minimax theorem add once mixed strategies are allowed?

Solving 2×2 games by the closed form

Exercise 4.4. Solve each 2×2 game: check for a saddle point first, and if none exists apply Theorem 4.8. Verify each value by substituting the strategy.

$$(a) \begin{bmatrix} 0 & -10 \\ 1 & 2 \end{bmatrix}, \quad (b) \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}, \quad (c) \begin{bmatrix} 2 & -1 \\ -1 & 1 \end{bmatrix}.$$

Exercise 4.5. Solve $A = \begin{bmatrix} 2 & 5 \\ 6 & 3 \end{bmatrix}$. First rule out a saddle point, then apply the 2×2 closed form to get $\mathbf{p}^*$, $\mathbf{q}^*$, and V , and finally verify the value two ways: by Player I's mix against each column and by Player II's mix against each row. (Worked in Appendix D.)

Dominance and the graphical method

Exercise 4.6. Solve by the graphical method, stating p^* or q^* and the value.

$$(a) \begin{bmatrix} 2 & 3 & 1 & 5 \\ 4 & 1 & 6 & 0 \end{bmatrix}, \quad (b) \begin{bmatrix} 1 & 5 \\ 4 & 4 \\ 6 & 2 \end{bmatrix}.$$

(These are Examples 4.16 and 4.17; reproduce the envelopes and identify which lines are active at the optimum.)

Exercise 4.7. Solve the 2×3 game $A = \begin{bmatrix} 0 & 2 & 5 \\ 4 & 3 & 1 \end{bmatrix}$ by the graphical method. Write Player I's payoff against each of Player II's three columns as a function of the weight p on row 1, identify the lower envelope and its highest point, name the two columns active there, and confirm the third column is slack. Report $\mathbf{p}^*$, $\mathbf{q}^*$, and V . (Worked in Appendix D.)

Exercise 4.8. Reduce each game by dominance to a 2×2 game and solve it.

$$(a) \begin{bmatrix} 2 & 0 & 4 \\ 1 & 2 & 3 \\ 4 & 1 & 2 \end{bmatrix}, \quad (b) \begin{bmatrix} 0 & 4 & 6 \\ 5 & 7 & 4 \\ 9 & 6 & 3 \end{bmatrix}.$$

In (b), exhibit the mixed combination of columns that dominates the deleted column.

Exercise 4.9. Solve $A = \begin{bmatrix} 3 & 2 & 4 \\ 2 & 1 & 3 \\ 5 & 4 & 6 \end{bmatrix}$. Show that one row dominates both of the others, reduce the game accordingly, and explain why the reduced game has a saddle point even though the original matrix has no obvious one. Give V and the optimal pure strategies. (Worked in Appendix D.)

Theory and problem-set games

Exercise 4.10. Prove that the value of a zero-sum game satisfies $\underline{V} \leq V \leq \overline{V}$ and that adding a constant c to every entry of A raises the value by exactly c while leaving the optimal strategies unchanged. (This shift trick is used in Chapter 6 to make a value positive.)

Exercise 4.11. From Problem Set 2. Solve each game by hand, choosing the appropriate method, and give $\mathbf{p}^*$, $\mathbf{q}^*$, and V .

$$(a) \begin{bmatrix} 1 & -2 & 3 & -4 \\ 0 & 1 & -2 & 3 \\ 0 & 0 & 1 & -2 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (b) \begin{bmatrix} 1 & 2 & -1 \\ 2 & -1 & 4 \\ -1 & 4 & -3 \end{bmatrix}, \quad (c) \begin{bmatrix} 10 & 0 & 7 & 0 \\ 0 & 6 & 4 & 0 \\ 0 & 0 & 3 & 3 \end{bmatrix}.$$

(Full solutions appear in Chapter 11. Part (a) yields $V = 1/12$ by the principle of indifference of Chapter 5.)

CHAPTER 5

Solving Zero-Sum Games: Indifference, Invariance, and Symmetry

“If you cannot be predicted, you cannot be beaten; and the way not to be predicted is to be exactly indifferent among your own good moves.”

Chapter 4 solved the smallest zero-sum games by saddle points, the 2×2 formula, and the graphical method. This chapter supplies the tools for the rest: the principle of indifference, which turns the search for optimal mixed strategies into a linear system; the matrix-inversion formula for non-degenerate games; the special structure of symmetric games, which are always fair; and the principle of invariance, which exploits a game’s symmetries to collapse it to a smaller one. The chapter closes with Colonel Blotto, the classic game of resource allocation, and with two lemmas that prepare the full proof of the minimax theorem in Chapter 6.

5.1 The principle of indifference

The payoff-value lemma (Lemma 4.18) showed that an optimal pair earns exactly the value. Its converse, the principle of indifference, is the practical engine for finding optimal strategies: a player puts positive weight only on pure strategies that achieve the value, and so every active strategy can be set equal to the value, yielding equations.

Theorem 5.1 (Principle of indifference, or complementary slackness). Let $(\mathbf{p}^*, \mathbf{q}^*)$ be optimal strategies in a zero-sum game of value V . Then:

1. if $p_i^* > 0$, then row i earns exactly the value against $\mathbf{q}^*$, that is $\sum_j a_{ij}q_j^* = V$;
2. if row i earns strictly less than the value, $\sum_j a_{ij}q_j^* < V$, then $p_i^* = 0$;
3. symmetrically, if $q_j^* > 0$ then column j costs exactly the value, $\sum_i p_i^*a_{ij} = V$, and if column j costs strictly more than the value then $q_j^* = 0$.

Proof. We prove the first two statements; the third is identical with the roles of the players exchanged. Optimality of $\mathbf{q}^*$ gives $\sum_j a_{ij}q_j^* \leq V$ for every row i (Theorem 4.7). Suppose, toward a contradiction, that $p_k^* > 0$ for some row k with $\sum_j a_{kj}q_j^* < V$. Then, weighting these row inequalities by p_i^* and summing,

$$(\mathbf{p}^*)^\top A \mathbf{q}^* = \sum_i p_i^* \left(\sum_j a_{ij}q_j^* \right) < \sum_i p_i^* V = V,$$

where the inequality is strict because the term $i = k$ has $p_k^* > 0$ multiplying a quantity strictly below V , while all other terms are at most $p_i^* V$. This contradicts the payoff-value lemma, which requires $(\mathbf{p}^*)^\top A \mathbf{q}^* = V$. Hence every row with $p_i^* > 0$ earns exactly V , proving (1), and contrapositively a row earning less than V has $p_i^* = 0$, proving (2). ■

The principle of indifference

At an optimal solution a player randomizes only over pure strategies that all earn exactly the value, and never over one that earns less. To solve a game, then, *guess the support*, set each active strategy equal to V , and solve the resulting linear equations together with $\sum_i p_i = 1$. This turns the search for an optimal mixed strategy into ordinary algebra, and it is the workhorse of the chapter.

The name complementary slackness comes from linear programming, where the same condition appears (Chapter 6): a constraint is tight exactly when its dual variable is positive. The practical use is direct. The intuition is worth stating plainly: if one of your active pure strategies earned less than the value, you would shift its probability onto a better one and do strictly better, so at an optimum every move you actually use must be tied for best. The opponent, by playing $\mathbf{q}^*$, has arranged exactly that tie, leaving you *indifferent* and therefore unable to exploit any particular choice.

Example 5.2 (The odd-even game by indifference). Players simultaneously declare “odd” or “even.” With Player I as the row player choosing odd or even, and the payoff matrix

$$A = \begin{bmatrix} 0 & 1 \\ 1 & -2 \end{bmatrix} \quad \begin{array}{ll} \text{rows: odd, even} \\ \text{cols: even, odd} \end{array}$$

its row minima are $\min(0, 1) = 0$ and $\min(1, -2) = -2$, giving maximin 0, while its column maxima are $\max(0, 1) = 1$ and $\max(1, -2) = 1$, giving minimax 1; since $0 \neq 1$ there is no saddle point and both rows are active. Player II mixes $\mathbf{q}^* = (q_1, 1 - q_1)$ over the two columns, and indifference says each of Player I’s rows earns the same value against it. Compute the two row payoffs:

$$\text{row 1 (odd): } 0 \cdot q_1 + 1 \cdot (1 - q_1) = 1 - q_1, \quad \text{row 2 (even): } 1 \cdot q_1 + (-2)(1 - q_1) = 3q_1 - 2.$$

Set them equal and solve for q_1 :

$$1 - q_1 = 3q_1 - 2 \implies 1 + 2 = 3q_1 + q_1 \implies 3 = 4q_1 \implies q_1 = \frac{3}{4}.$$

The value is the common row payoff, $V = 1 - q_1 = 1 - \frac{3}{4} = \frac{1}{4}$ (and the other row agrees: $3 \cdot \frac{3}{4} - 2 = \frac{9}{4} - 2 = \frac{1}{4}$). Repeating for Player I, who mixes $\mathbf{p}^* = (p_1, 1 - p_1)$ to equalize the two columns, column 1 (even) pays $0 \cdot p_1 + 1 \cdot (1 - p_1) = 1 - p_1$ and column 2 (odd) pays $1 \cdot p_1 + (-2)(1 - p_1) = 3p_1 - 2$; setting $1 - p_1 = 3p_1 - 2$ gives $p_1 = \frac{3}{4}$ by the same algebra. So $\mathbf{p}^* = (\frac{3}{4}, \frac{1}{4})$ and $\mathbf{q}^* = (\frac{3}{4}, \frac{1}{4})$, and Player I, the “odd” player, can guarantee an expected gain of $\frac{1}{4}$. $\diamond$

5.2 The non-degenerate case: solution by matrix inversion

When every pure strategy of an $n \times n$ game is active, the indifference equations form a square linear system that the inverse of A solves in closed form.

Proposition 5.3 (Inversion formula for non-degenerate games). Suppose an $n \times n$ game has an optimal pair in which all strategies are active and the value $V \neq 0$. If A is invertible, then

$$V = \frac{1}{\mathbf{1}^\top A^{-1} \mathbf{1}}, \quad \mathbf{q}^* = \frac{A^{-1} \mathbf{1}}{\mathbf{1}^\top A^{-1} \mathbf{1}}, \quad (\mathbf{p}^*)^\top = \frac{\mathbf{1}^\top A^{-1}}{\mathbf{1}^\top A^{-1} \mathbf{1}}, \quad (5.1)$$

where $\mathbf{1}$ is the all-ones vector.

Proof. Because all rows are active, the principle of indifference makes every row earn the value against $\mathbf{q}^*$, which in matrix form is $A\mathbf{q}^* = V\mathbf{1}$. Solving, $\mathbf{q}^* = V A^{-1} \mathbf{1}$. The probabilities sum to one, so

$$\mathbf{1} = \mathbf{1}^\top \mathbf{q}^* = V \mathbf{1}^\top A^{-1} \mathbf{1},$$

which gives $V = 1/(\mathbf{1}^\top A^{-1} \mathbf{1})$, provided the scalar $\mathbf{1}^\top A^{-1} \mathbf{1} \neq 0$. Substituting back yields the formula for $\mathbf{q}^*$. The analogous argument for Player I, using $(\mathbf{p}^*)^\top A = V \mathbf{1}^\top$, gives $(\mathbf{p}^*)^\top = V \mathbf{1}^\top A^{-1}$, which is the stated expression. $\blacksquare$

Remark 5.4. The formula assumes both that A is invertible and that the resulting $\mathbf{p}^*, \mathbf{q}^*$ are genuine probability vectors, with all entries nonnegative. If some entry comes out negative, a strategy that was assumed active is in fact inactive, and one drops it and resolves on the smaller support. The formula also fails when $V = 0$, since then $A\mathbf{q}^* = \mathbf{0}$ and the normalization cannot be recovered this way; for that case, and for symmetric games, the indifference equations are solved directly, as in the next section.

Example 5.5 (Solving a 3×3 game by inversion). Consider

$$A = \begin{bmatrix} 2 & 1 & 1 \\ 1 & 3 & 1 \\ 1 & 1 & 4 \end{bmatrix}.$$

No row dominates another (row 1 beats row 2 in the first column but loses in the second, and similarly for every pair), and the matrix is symmetric, so the same holds for columns; dominance gives no reduction and we expect all three strategies active. Apply [Proposition 5.3](#).

Step 1: the determinant. Expanding along the first row,

$$\det A = 2(3 \cdot 4 - 1 \cdot 1) - 1(1 \cdot 4 - 1 \cdot 1) + 1(1 \cdot 1 - 3 \cdot 1) = 2(11) - 1(3) + 1(-2) = 22 - 3 - 2 = 17.$$

Step 2: the inverse. The cofactors are $C_{11} = 3 \cdot 4 - 1 \cdot 1 = 11$, $C_{12} = -(1 \cdot 4 - 1 \cdot 1) = -3$, $C_{13} = 1 \cdot 1 - 3 \cdot 1 = -2$, and, because A is symmetric, the rest follow by symmetry, giving $C_{22} = 2 \cdot 4 - 1 \cdot 1 = 7$, $C_{23} = -(2 \cdot 1 - 1 \cdot 1) = -1$, $C_{33} = 2 \cdot 3 - 1 \cdot 1 = 5$. The inverse is the transposed cofactor matrix over the determinant,

$$A^{-1} = \frac{1}{17} \begin{bmatrix} 11 & -3 & -2 \\ -3 & 7 & -1 \\ -2 & -1 & 5 \end{bmatrix}.$$

Step 3: apply the formula. Multiply A^{-1} by the all-ones vector, which just sums each row:

$$A^{-1}\mathbf{1} = \frac{1}{17}(11 - 3 - 2, -3 + 7 - 1, -2 - 1 + 5)^\top = \frac{1}{17}(6, 3, 2)^\top.$$

The scalar in the denominator is the sum of these entries,

$$\mathbf{1}^\top A^{-1}\mathbf{1} = \frac{1}{17}(6 + 3 + 2) = \frac{11}{17},$$

so the value is its reciprocal,

$$V = \frac{1}{\mathbf{1}^\top A^{-1}\mathbf{1}} = \frac{17}{11},$$

and the optimal column strategy is

$$\mathbf{q}^* = \frac{A^{-1}\mathbf{1}}{\mathbf{1}^\top A^{-1}\mathbf{1}} = \frac{\frac{1}{17}(6, 3, 2)^\top}{\frac{11}{17}} = \left(\frac{6}{11}, \frac{3}{11}, \frac{2}{11}\right)^\top.$$

Because A is symmetric, $\mathbf{1}^\top A^{-1} = (A^{-1}\mathbf{1})^\top$, so Player I's optimal strategy is the same, $\mathbf{p}^* = \left(\frac{6}{11}, \frac{3}{11}, \frac{2}{11}\right)$.

Step 4: verify. Play $\mathbf{q}^*$ and read off each row's payoff:

$$\text{row 1: } \frac{2 \cdot 6 + 1 \cdot 3 + 1 \cdot 2}{11} = \frac{17}{11}, \quad \text{row 2: } \frac{1 \cdot 6 + 3 \cdot 3 + 1 \cdot 2}{11} = \frac{17}{11}, \quad \text{row 3: } \frac{1 \cdot 6 + 1 \cdot 3 + 4 \cdot 2}{11} = \frac{17}{11}.$$

All three rows earn exactly $\frac{17}{11}$, so every row is a best response and the entries are all positive, confirming the support guess and the value. $\diamond$

5.3 One game, solved five ways

We now have a toolbox: saddle search, the 2×2 closed form, the graphical method, the principle of indifference, and matrix inversion. A natural worry is whether the answer depends on which tool you grab. It does not. A correct method applied to a solvable game returns the same optimal strategies and the same value, every time. To make that concrete, take the single game

$$A = \begin{bmatrix} 1 & 3 \\ 4 & 2 \end{bmatrix}$$

where Player I chooses a row and Player II a column, and the entries are payoffs to the maximizing Player I. We grind it through all five. Watch the matrix reappear at every stage with the active cells marked, so you can see what each method actually looks at.

Method 1: look for a saddle point

Circle each row's minimum and box each column's maximum. A saddle is an entry that is both.

	c_1	c_2	row min
r_1	1	3	<u>1</u>
r_2	4	2	<u>2</u>
col max	4	3	

The largest row minimum (the maximin) is $\max\{1, 2\} = 2$ in row 2. The smallest column maximum (the minimax) is $\min\{4, 3\} = 3$ in column 2. Since $2 \neq 3$, the lower and upper values disagree: *no pure saddle exists*, and the solution must be mixed. Onward.

Method 2: the 2×2 closed form. Label $a = 1$, $b = 3$, $c = 4$, $d = 2$ in $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$. The shared denominator is

$$D = a - b - c + d = 1 - 3 - 4 + 2 = -4,$$

and the closed form from Theorem 4.8 gives, with the very same matrix in view,

$$\begin{bmatrix} 1 & 3 \\ 4 & 2 \end{bmatrix}: \quad p^* = \frac{d - c}{D} = \frac{2 - 4}{-4} = \frac{1}{2}, \quad q^* = \frac{d - b}{D} = \frac{2 - 3}{-4} = \frac{1}{4},$$

$$V = \frac{ad - bc}{D} = \frac{1 \cdot 2 - 3 \cdot 4}{-4} = \frac{-10}{-4} = \frac{5}{2}.$$

So Player I plays $\mathbf{p}^* = (\frac{1}{2}, \frac{1}{2})$, Player II plays $\mathbf{q}^* = (\frac{1}{4}, \frac{3}{4})$, and $v = \frac{5}{2}$.

Method 3: the principle of indifference. Make the opponent indifferent. If Player II splits weight q on c_1 and $1 - q$ on c_2 , Player I's two rows earn

$$\begin{bmatrix} 1 & 3 \\ 4 & 2 \end{bmatrix} \text{ row 1: } 1 \cdot q + 3(1 - q) = 3 - 2q, \quad \begin{bmatrix} 1 & 3 \\ 4 & 2 \end{bmatrix} \text{ row 2: } 4q + 2(1 - q) = 2 + 2q.$$

Setting them equal, $3 - 2q = 2 + 2q \Rightarrow 4q = 1 \Rightarrow q = \frac{1}{4}$, with common value $3 - 2(\frac{1}{4}) = \frac{5}{2}$. Symmetrically, equalizing the columns against Player I's mix $(p, 1 - p)$ gives $p \cdot 1 + (1 - p) \cdot 4 = p \cdot 3 + (1 - p) \cdot 2$, i.e. $4 - 3p = 2 + p \Rightarrow p = \frac{1}{2}$. Same $\mathbf{p}^*, \mathbf{q}^*, v$.

Method 4: the graphical method. Method 3's two row-payoff lines are exactly what the graphical method plots against Player I's mixing weight p (weight on row 1):

$$U_{c_1}(p) = 1 \cdot p + 4(1 - p) = 4 - 3p, \quad U_{c_2}(p) = 3 \cdot p + 2(1 - p) = 2 + p.$$

Player II, the minimizer, holds Player I to the *lower* of the two lines; Player I, the maximizer, picks the p at the highest point of that lower envelope. Figure 5.1 shows the two lines crossing at $p = \frac{1}{2}$, where the envelope peaks at height $\frac{5}{2}$.

$$4 - 3p = 2 + p \implies p = \frac{1}{2}, \quad v = 2 + \frac{1}{2} = \frac{5}{2}.$$

Method 5: matrix inversion. The non-degenerate formula needs A^{-1} . With $\det A = 1 \cdot 2 - 3 \cdot 4 = -10$,

$$\begin{bmatrix} 1 & 3 \\ 4 & 2 \end{bmatrix}^{-1} = \frac{1}{-10} \begin{bmatrix} 2 & -3 \\ -4 & 1 \end{bmatrix} = \begin{bmatrix} -\frac{1}{5} & \frac{3}{10} \\ \frac{2}{5} & -\frac{1}{10} \end{bmatrix}.$$

Sum all four entries: $\mathbf{1}^\top A^{-1} \mathbf{1} = -\frac{1}{5} + \frac{3}{10} + \frac{2}{5} - \frac{1}{10} = \frac{-2+3+4-1}{10} = \frac{4}{10} = \frac{2}{5}$, so

$$v = \frac{1}{\mathbf{1}^\top A^{-1} \mathbf{1}} = \frac{1}{2/5} = \frac{5}{2},$$

$$\mathbf{q}^* = \frac{A^{-1} \mathbf{1}}{\mathbf{1}^\top A^{-1} \mathbf{1}} = \frac{(\frac{1}{10}, \frac{3}{10})}{2/5} = \left(\frac{1}{4}, \frac{3}{4}\right),$$

$$\mathbf{p}^{*\top} = \frac{\mathbf{1}^\top A^{-1}}{\mathbf{1}^\top A^{-1} \mathbf{1}} = \frac{(\frac{1}{5}, \frac{1}{5})}{2/5} = \left(\frac{1}{2}, \frac{1}{2}\right).$$




Figure 5.1. Graphical solution of $A = \begin{pmatrix} 1 & 3 \\ 4 & 2 \end{pmatrix}$. Each gray line is Player I's expected payoff against one of Player II's pure columns, as a function of the weight p on row 1. Player II forces Player I down onto the lower envelope (thick red); Player I maximizes by choosing the peak at $p^* = \frac{1}{2}$, where both lines meet at height $v = \frac{5}{2}$. The crossing is the unique maximin of the envelope.

Five methods, one answer: $\mathbf{p}^* = (\frac{1}{2}, \frac{1}{2})$, $\mathbf{q}^* = (\frac{1}{4}, \frac{3}{4})$, $v = \frac{5}{2}$. The methods differ in cost and scope, rather than in what they conclude. The saddle search is the cheapest screen and rules out a pure solution; the 2×2 form is fastest once you know it is mixed; indifference and the graph build intuition and extend to $2 \times n$ and $n \times 2$ games; inversion handles the general $n \times n$ non-degenerate case; and linear programming (Chapter 6) is the universal hammer when nothing smaller fits. A sixth route, solving the matching pair of linear programs, returns $\frac{5}{2}$ as well.

5.4 Symmetric games are fair

A game can be symmetric in the sense that exchanging the two players leaves it unchanged. For zero-sum games this is captured by skew-symmetry of the payoff matrix.

Definition 5.6 (Symmetric game). A two-person zero-sum game is **symmetric** if its payoff matrix is square and **skew-symmetric**, $A^\top = -A$, that is $a_{ij} = -a_{ji}$ for all i, j . In particular every diagonal entry is zero.

Coin matching, $\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$, is not skew-symmetric, but Rock–Paper–Scissors, $\begin{bmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{bmatrix}$, is. The defining feature of a symmetric game is that the two players are interchangeable, and the following theorem says interchangeable players cannot have an advantage.

Theorem 5.7 (Symmetric games are fair). Every symmetric game has value $V = 0$, and any strategy optimal for one player is optimal for the other, so an optimal pair with $\mathbf{p}^* = \mathbf{q}^*$ exists.

Proof. We use two facts about the bilinear payoff. First, for any vectors the scalar $\mathbf{p}^\top A \mathbf{q}$ equals its own transpose, and with $A^\top = -A$,

$$\mathbf{p}^\top A \mathbf{q} = (\mathbf{p}^\top A \mathbf{q})^\top = \mathbf{q}^\top A^\top \mathbf{p} = -\mathbf{q}^\top A \mathbf{p}. \quad (5.2)$$

Setting $\mathbf{q} = \mathbf{p}$ gives $\mathbf{p}^\top A \mathbf{p} = -\mathbf{p}^\top A \mathbf{p}$, so

$$\mathbf{p}^\top A \mathbf{p} = 0 \quad \text{for every } \mathbf{p}. \quad (5.3)$$

By the minimax theorem the value V exists. Consider the *role-swapped* game in which the two players exchange seats. Its payoff matrix is $-A^\top$, because the original column player, now seated as the maximizer, receives the negatives of the original entries. Swapping the maximizer and the minimizer negates the value, so the role-swapped game has value $-V$. But skew-symmetry gives $-A^\top = -(-A) = A$, so the role-swapped game has the same matrix as the original and therefore the same value V . Hence $V = -V$, which forces $V = 0$.

For the strategies, let $\mathbf{p}^*$ be optimal for Player I, so $(\mathbf{p}^*)^\top A\mathbf{p} \geq V = 0$ for every Player II strategy $\mathbf{p}$. Suppose Player II also adopts the vector $\mathbf{p}^*$. For every Player I strategy $\mathbf{p}$, identity (5.2) gives

$$\mathbf{p}^\top A\mathbf{p}^* = -(\mathbf{p}^*)^\top A\mathbf{p} \leq 0 = V.$$

Thus Player II, by playing $\mathbf{p}^*$, holds Player I to at most the value, so $\mathbf{p}^*$ is optimal for Player II as well, and we may take $\mathbf{q}^* = \mathbf{p}^*$. ■

Example 5.8 (A number-selection game). Each player chooses an integer in $\{1, 2, \dots, 100\}$. The payoffs to Player I are: win 1 if Player I's number is exactly one more than Player II's; lose 2 if it is two or more above; lose 1 if it is exactly one below; win 2 if it is two or more below; and 0 for a tie. The payoff matrix is skew-symmetric. Its first rows are

	1	2	3	4	5
1	0	-1	2	2	2
2	1	0	-1	2	2
3	-2	1	0	-1	2
4	-2	-2	1	0	-1
5	-2	-2	-2	1	0

Strategy k strictly dominates strategy $k+3$, because shifting three or more steps away always lands in the “two or more” regime. Iterated elimination of dominated strategies removes everything but $\{1, 2, 3\}$, leaving the skew-symmetric 3×3 game

$$A = \begin{bmatrix} 0 & -1 & 2 \\ 1 & 0 & -1 \\ -2 & 1 & 0 \end{bmatrix}.$$

By Theorem 5.7, $V = 0$ and $\mathbf{p}^* = \mathbf{q}^*$. Since $V = 0$ we use indifference directly. Writing $\mathbf{p}^* = (p_1, p_2, p_3)$ and setting each row equal to 0 against $\mathbf{p}^*$ (acting as $\mathbf{q}^*$),

$$\text{row 1: } -p_2 + 2p_3 = 0 \Rightarrow p_2 = 2p_3, \quad \text{row 2: } p_1 - p_3 = 0 \Rightarrow p_1 = p_3.$$

With $p_1 + p_2 + p_3 = 1$ this gives $p_3 + 2p_3 + p_3 = 1$, so $p_3 = \frac{1}{4}$ and

$$\mathbf{p}^* = \mathbf{q}^* = \left(\frac{1}{4}, \frac{1}{2}, \frac{1}{4}\right), \quad V = 0.$$

A quick check confirms each row earns 0: row 3 gives $-2 \cdot \frac{1}{4} + 1 \cdot \frac{1}{2} + 0 = 0$.

5.5 Invariance

A game may carry a symmetry that is not full player-exchange but still constrains its solutions. Invariance formalizes “relabeling the strategies leaves the game unchanged,” and its payoff is that some optimal strategy respects the symmetry, which shrinks the search.

Definition 5.9 (Invariance). A zero-sum game (X, Y, A) is **invariant** under a pair of bijections $h : X \rightarrow X$ and $g : Y \rightarrow Y$ if

$$a_{xy} = a_{h(x)g(y)} \quad \text{for all } x \in X, y \in Y.$$

The maps h and g , being bijections of finite sets, are permutations, and they act on mixed strategies by permuting coordinates.

Theorem 5.10 (Invariant optimal strategies exist). If a game is invariant under (h, g) , then there exist optimal strategies $\mathbf{p}^*$ and $\mathbf{q}^*$ that are themselves invariant, $h(\mathbf{p}^*) = \mathbf{p}^*$ and $g(\mathbf{q}^*) = \mathbf{q}^*$.

Proof. We construct an invariant optimal strategy for Player I by averaging over the symmetry; the argument for Player II is identical. First, invariance carries optimal strategies to optimal strategies. If $\mathbf{p}$ is optimal, then for every $\mathbf{q}$,

$$\min_{\mathbf{q}} (h(\mathbf{p}))^\top A\mathbf{q} = \min_{\mathbf{q}} \mathbf{p}^\top A g^{-1}(\mathbf{q}) = \min_{\mathbf{q}'} \mathbf{p}^\top A\mathbf{q}' = V,$$

where the first equality is the invariance identity applied to the bilinear form, and the second uses that g^{-1} permutes the set of all $\mathbf{q}$ onto itself. Hence $h(\mathbf{p})$ is also optimal.

Let $\mathbf{p}$ be any optimal strategy. Since h is a permutation of a finite set, it has finite order k , meaning h^k is the identity. Form the average over the orbit,

$$\bar{\mathbf{p}} = \frac{1}{k} \sum_{j=0}^{k-1} h^j(\mathbf{p}).$$

Each $h^j(\mathbf{p})$ is optimal by the previous paragraph, so each guarantees at least V , $\min_{\mathbf{q}} (h^j(\mathbf{p}))^\top A \mathbf{q} = V$. The guarantee $\mathbf{p} \mapsto \min_{\mathbf{q}} \mathbf{p}^\top A \mathbf{q}$ is a minimum of linear functions, hence concave, so the value of the average is at least the average of the values:

$$\min_{\mathbf{q}} \bar{\mathbf{p}}^\top A \mathbf{q} \geq \frac{1}{k} \sum_{j=0}^{k-1} \min_{\mathbf{q}} (h^j(\mathbf{p}))^\top A \mathbf{q} = V.$$

Therefore $\bar{\mathbf{p}}$ guarantees at least V and is optimal. It is invariant because applying h merely cyclically permutes the terms of the average: $h(\bar{\mathbf{p}}) = \frac{1}{k} \sum_{j=0}^{k-1} h^{j+1}(\mathbf{p}) = \bar{\mathbf{p}}$, using $h^k = h^0$. The same construction with g produces an invariant optimal $\mathbf{q}^*$. ■

The practical content is a reduction. Restrict attention to invariant strategies, which form a lower-dimensional set spanned by the symmetry's orbits, build the smaller payoff matrix among these orbit-strategies, and solve that. The next section carries this out for Colonel Blotto.

5.6 Colonel Blotto

Colonel Blotto is the prototype of games about dividing a resource across several fronts. Two commanders distribute their forces over two posts without seeing each other's deployment, and each post is won by whoever sends more.

Example 5.11 (Blotto with four against three). Colonel Olir (Player I) has four units, Lieutenant Yaxuan (Player II) has three, to divide over two posts. At each post, the side with strictly more units captures it and scores the opponent's units there plus one; the loser scores the negative of its own units there plus one; a tie at a post scores zero to both. The total payoff is the sum over the two posts. Olir's pure strategies are the five splits $X = \{(4, 0), (3, 1), (2, 2), (1, 3), (0, 4)\}$, and Yaxuan's are the four splits $Y = \{(3, 0), (2, 1), (1, 2), (0, 3)\}$. Working out every post contest gives Olir's payoff matrix

$$A = \begin{array}{c|cccc} & (3, 0) & (2, 1) & (1, 2) & (0, 3) \\ \hline (4, 0) & 4 & 2 & 1 & 0 \\ (3, 1) & 1 & 3 & 0 & -1 \\ (2, 2) & -2 & 2 & 2 & -2 \\ (1, 3) & -1 & 0 & 3 & 1 \\ (0, 4) & 0 & 1 & 2 & 4 \end{array}$$

For instance, Olir's (3, 1) against Yaxuan's (2, 1) wins the first post for $2 + 1 = 3$ and ties the second for 0, total 3, the entry in row (3, 1), column (2, 1); Figure 5.2 pictures this contest.

5.6.1 Solving Blotto by invariance

The game has an evident symmetry: relabeling the two posts reverses every split. For Olir this swaps $(4, 0) \leftrightarrow (0, 4)$ and $(3, 1) \leftrightarrow (1, 3)$ while fixing $(2, 2)$; for Yaxuan it swaps $(3, 0) \leftrightarrow (0, 3)$ and $(2, 1) \leftrightarrow (1, 2)$. One checks the payoff matrix is invariant under this pair of permutations. By Theorem 5.10 we may seek an invariant optimal strategy, that is, one giving equal weight to a split and its reverse. The invariant strategies are spanned by the orbit-averages

$$s_1^X = \left(\frac{1}{2}, 0, 0, 0, \frac{1}{2}\right), \quad s_2^X = \left(0, \frac{1}{2}, 0, \frac{1}{2}, 0\right), \quad s_3^X = (0, 0, 1, 0, 0)$$

for Olir, and

$$s_1^Y = \left(\frac{1}{2}, 0, 0, \frac{1}{2}\right), \quad s_2^Y = \left(0, \frac{1}{2}, \frac{1}{2}, 0\right)$$




Figure 5.2. The Colonel Blotto contest of Example 5.11, drawn cleanly: Olir's split $(3, 1)$ meets Yaxuan's $(2, 1)$. Each disc is one unit (Olir in blue, Yaxuan in red). Olir outnumbers Yaxuan 3 to 2 at Post 1 and so captures it for the opponent's units plus one, $2 + 1 = 3$; the posts tie 1 to 1 at Post 2 for 0. The total, 3, is the entry in row $(3, 1)$, column $(2, 1)$ of the payoff matrix.

for Yaxuan. Computing the expected payoff of each orbit-strategy pair reduces the 5×4 game to the 3×2 game

$$B = \begin{array}{c|cc} & s_1^Y & s_2^Y \\ \hline s_1^X & 2 & 3/2 \\ s_2^X & 0 & 3/2 \\ s_3^X & -2 & 2 \end{array}$$

To see one entry, B_{11} is the average of the four corners of A touched by s_1^X and s_1^Y , namely $4, 0, 0, 4$, each with probability $\frac{1}{4}$, giving 2. In B , row s_1^X dominates row s_2^X (since $2 > 0$ and $\frac{3}{2} = \frac{3}{2}$), so s_2^X is discarded, leaving the 2×2 game on rows $\{s_1^X, s_3^X\}$,

$$\begin{bmatrix} 2 & 3/2 \\ -2 & 2 \end{bmatrix}.$$

With $a = 2, b = \frac{3}{2}, c = -2, d = 2$ and $D = a - b - c + d = 2 - \frac{3}{2} + 2 + 2 = \frac{9}{2}$, the formula of Theorem 4.8 gives

$$p^* = \frac{d - c}{D} = \frac{2 + 2}{9/2} = \frac{8}{9}, \quad q^* = \frac{d - b}{D} = \frac{2 - 3/2}{9/2} = \frac{1}{9},$$

so among the orbit-strategies Olir plays $(\frac{8}{9}, 0, \frac{1}{9})$ and Yaxuan plays $(\frac{1}{9}, \frac{8}{9})$. The value is

$$V = p^* \cdot 2 + (1 - p^*)(-2) = \frac{8}{9} \cdot 2 + \frac{1}{9} \cdot (-2) = \frac{14}{9} \approx 1.56.$$

Unpacking the orbit-strategies into the original five and four pure strategies, each orbit splitting its weight evenly between a split and its reverse,

$$\mathbf{p}^* = \left(\frac{4}{9}, 0, \frac{1}{9}, 0, \frac{4}{9}\right), \quad \mathbf{q}^* = \left(\frac{1}{18}, \frac{4}{9}, \frac{4}{9}, \frac{1}{18}\right).$$

Olir, with the extra unit, can guarantee an expected gain of $\frac{14}{9}$.

Historical Note

Colonel Blotto was introduced by Émile Borel in 1921 [5], before von Neumann's theorem, as one of the first formal models of strategic allocation. The general continuous version, dividing a unit of force across many fronts, was solved only much later [44]. The game remains a workhorse wherever a fixed budget must be spread over competing targets: it models a presidential campaign allocating money and visits across battleground states, two firms splitting an advertising budget across markets, and, in the security setting, a defender distributing limited resources across potential targets against a strategic attacker [43]. The lesson is always the one this example shows: the side that is outgunned overall should not spread evenly but should randomize, conceding some fronts to contest others, exactly as Yaxuan abandons one post to fight harder elsewhere.

5.7 Two lemmas toward the minimax theorem

We end with two results that are easy to prove now and that carry the weight of the minimax theorem's proof in Chapter 6. The first says a player facing a known mixed strategy never needs to randomize.

Lemma 5.12 (Best-response purity). For any fixed mixed strategy $\mathbf{q}$ of Player II, Player I has a pure best response:

$$\max_{\mathbf{p} \in \Delta_m} \mathbf{p}^\top A \mathbf{q} = \max_{1 \leq i \leq m} \sum_{j=1}^n a_{ij} q_j.$$

Proof. For any $\mathbf{p}$, the payoff $\mathbf{p}^\top A \mathbf{q} = \sum_i p_i (\sum_j a_{ij} q_j)$ is a convex combination of the row payoffs $r_i = \sum_j a_{ij} q_j$, with weights p_i . A convex combination never exceeds the largest term, so $\mathbf{p}^\top A \mathbf{q} \leq \max_i r_i$, with equality when $\mathbf{p}$ places all weight on a row attaining the maximum. Hence the maximum over mixed $\mathbf{p}$ equals $\max_i r_i$. ■

The second records the easy half of the minimax theorem, already met in Chapter 4, now stated for the record because the hard half is its converse.

Proposition 5.13 (Weak minimax inequality). For every payoff matrix,

$$\max_{\mathbf{p} \in \Delta_m} \min_{\mathbf{q} \in \Delta_n} \mathbf{p}^\top A \mathbf{q} \leq \min_{\mathbf{q} \in \Delta_n} \max_{\mathbf{p} \in \Delta_m} \mathbf{p}^\top A \mathbf{q}.$$

Proof. For any fixed $\mathbf{p}_0$ and $\mathbf{q}_0$, $\min_{\mathbf{q}} \mathbf{p}_0^\top A \mathbf{q} \leq \mathbf{p}_0^\top A \mathbf{q}_0 \leq \max_{\mathbf{p}} \mathbf{p}^\top A \mathbf{q}_0$. The left member is a function of $\mathbf{p}_0$ alone and the right of $\mathbf{q}_0$ alone, so taking the maximum over $\mathbf{p}_0$ on the left and the minimum over $\mathbf{q}_0$ on the right preserves the inequality. ■

The minimax theorem asserts that this inequality is in fact an equality. Two routes prove it: the fixed-point argument that Nash later generalized (Chapter 8), and the linear programming duality of the next chapter, which is the route the course took and which also hands us an algorithm.

5.8 Learning to play: fictitious play

The methods so far compute optimal strategies from the matrix. But players who do *not* know the matrix, and cannot solve it, can still find the value by *learning* from repeated play. In **fictitious play** [10], each player keeps a tally of the opponent's past moves and, each round, best-responds to that empirical frequency as though it were the opponent's mixed strategy. Remarkably, for any zero-sum game the empirical frequencies converge to optimal strategies and the average payoff converges to the value.

Theorem 5.14 (Robinson, 1951). In any finite two-person zero-sum game, the empirical distributions of both players under fictitious play converge to optimal mixed strategies, and the time-average payoff converges to the value V .

Example 5.15 (Fictitious play on Matching Pennies). Take Matching Pennies, $A = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$ (rows and columns H, T ; Player I scores +1 on a match, −1 on a mismatch), whose value is $V = 0$ with optimal mix $(\frac{1}{2}, \frac{1}{2})$ for each. Starting both players at H and breaking ties toward H , each round Player I matches the opponent's more frequent past call and Player II mismatches:

round t	I plays	II plays	payoff	freq. I played H	freq. II played H	avg. payoff
1	H	H	+1	1	1	1
2	H	T	−1	1	$\frac{1}{2}$	0
3	H	T	−1	1	$\frac{1}{3}$	$-\frac{1}{3}$
4	T	T	+1	$\frac{3}{4}$	$\frac{1}{4}$	0
5	T	T	+1	$\frac{3}{5}$	$\frac{1}{5}$	$\frac{1}{5}$
6	T	T	+1	$\frac{1}{2}$	$\frac{1}{6}$	$\frac{1}{3}$
7	T	H	−1	$\frac{3}{7}$	$\frac{2}{7}$	$\frac{1}{7}$
8	T	H	−1	$\frac{3}{8}$	$\frac{3}{8}$	0

The play cycles, but the empirical frequencies are already drifting toward $\frac{1}{2}$ for each player (Figure 5.3), and the average payoff oscillates ever more tightly around the value 0. Continuing the iteration, both frequencies converge to $(\frac{1}{2}, \frac{1}{2})$ and the average payoff to $V = 0$, exactly as Theorem 5.14 promises. Fictitious play is thus a *model-free* route to the minimax solution: no one ever inverts a matrix or solves a linear program; the optimal mixture emerges from adaptation alone. $\diamond$




Figure 5.3. Fictitious play on Matching Pennies (Example 5.15): the empirical frequency with which each player has called H , over the first eight rounds. Both drift toward the optimal $\frac{1}{2}$, Player I (navy) from above and Player II (red) from below, and the running average payoff settles toward the value 0. Convergence is slow, at rate $O(1/\sqrt{t})$, so many more rounds are needed to reach $\frac{1}{2}$ exactly.

This is the simplest instance of the course’s *learning* theme. The same idea, best-responding to an evolving estimate of the opponent, scales up to the gradient play behind training adversarial networks, where a generator and a discriminator improve against each other round after round.

Historical Note: Émile Borel

The first modern study of mixed strategies and the minimax idea belongs to Émile Borel (1871–1956), in a series of notes from 1921 onward. Borel solved small symmetric games and introduced Colonel Blotto, yet he doubted that every game had a value, conjecturing that large ones might not; von Neumann’s 1928 theorem settled the question the other way. Borel was one of the great mathematicians of his age, a founder of measure theory (the Borel sets carry his name), and also a public man: a deputy in the French parliament, Minister of the Navy in 1925, and a member of the Resistance briefly imprisoned under Vichy. He died in 1956.

Notes and References

The principle of indifference, the inversion formula, the fairness of symmetric games, and the invariance reduction are all in Ferguson [16], whose organization this chapter follows; the two symmetry proofs are the content of Problem Set 2. Colonel Blotto originates with Borel [5] and was solved in general only much later by Roberson [44]. The principle of indifference is the combinatorial face of the complementary-slackness conditions of linear programming, made precise in Chapter 6; the inversion formula is the special case in which the optimal solution is fully mixed and the matrix invertible.

Chapter summary

The principle of indifference (Theorem 5.1) states that an optimal strategy puts weight only on pure strategies that achieve the value, turning the search for a solution into a linear system once the support is guessed. For a non-degenerate $n \times n$ game with nonzero value the system is solved in closed form by the inversion formula $V = 1/(\mathbf{1}^\top A^{-1} \mathbf{1})$ (Proposition 5.3). A symmetric game, one with a skew-symmetric matrix, is always fair, $V = 0$, with a common optimal

strategy for both players (Theorem 5.7); the number-selection game illustrates reduction by dominance followed by indifference. The principle of invariance (Theorem 5.10) shows that a game's strategy symmetries are inherited by some optimal strategy, found by averaging over the symmetry's orbit, which reduces Colonel Blotto from a 5×4 game to a 2×2 one with value $\frac{14}{9}$. The best-response purity lemma (Lemma 5.12) and the weak minimax inequality (Proposition 5.13) prepare the full proof of the minimax theorem in Chapter 6.

Exercises

The principle of indifference and matrix inversion

Exercise 5.1. Solve by the principle of indifference, stating the support, the linear system, and the solution.

$$(a) \begin{bmatrix} 0 & 1 \\ 1 & -2 \end{bmatrix}, \quad (b) \begin{bmatrix} 1 & 2 & -1 \\ 2 & -1 & 4 \\ -1 & 4 & -3 \end{bmatrix}.$$

(Part (b) is from Problem Set 2; its full support is all three strategies, and the answer is $\mathbf{p}^* = \mathbf{q}^* = (\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$ with $V = 1$.)

Exercise 5.2. Use the inversion formula (5.1) to solve $A = \begin{bmatrix} \frac{1}{2} & \frac{2}{3} & -1 \\ -1 & \frac{4}{3} & 3 \end{bmatrix}$, assuming all strategies are active. Verify your answer satisfies the indifference equations $A\mathbf{q}^* = V\mathbf{1}$, and confirm the probabilities are nonnegative.

Symmetry and invariance

Exercise 5.3 (Symmetric games are fair: full proof). Let A be skew-symmetric, $A^\top = -A$. (a) Prove the identity $\mathbf{p}^\top A\mathbf{q} = -\mathbf{q}^\top A\mathbf{p}$ for all $\mathbf{p}, \mathbf{q}$, and deduce $\mathbf{p}^\top A\mathbf{p} = 0$. (b) Using the role-swap argument, prove $V = 0$. (c) Prove that any strategy optimal for one player is optimal for the other. (This is Problem 2(a) of Problem Set 2.)

Exercise 5.4. Verify that Rock–Paper–Scissors, $A = \begin{bmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{bmatrix}$, is symmetric, and use Theorem 5.7 together with indifference to show $\mathbf{p}^* = \mathbf{q}^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ and $V = 0$.

Exercise 5.5 (Invariance gives invariant optima: full proof). Let a game be invariant under bijections $h : X \rightarrow X$ and $g : Y \rightarrow Y$, so $a_{xy} = a_{h(x)g(y)}$. (a) Show that if $\mathbf{p}$ is optimal for Player I then so is $h(\mathbf{p})$. (b) Show that the orbit-average $\bar{\mathbf{p}} = \frac{1}{k} \sum_{j=0}^{k-1} h^j(\mathbf{p})$, where k is the order of h , is optimal and satisfies $h(\bar{\mathbf{p}}) = \bar{\mathbf{p}}$. (This is Problem 2(b) of Problem Set 2.)

Exercise 5.6. Reconsider Colonel Blotto with both commanders holding the same number of units, say three each, over two posts. Write the payoff matrix, identify the post-reversal invariance, reduce by invariance, and solve. Is the game fair? Relate your answer to Theorem 5.7.

Edge cases and further practice

Exercise 5.7. Give an example of a 2×2 game with $\det A = 0$ and explain why the inversion formula (5.1) cannot be applied. Solve the game by another method, and state which assumption of Proposition 5.3 fails.

Exercise 5.8. Solve $A = \begin{bmatrix} 3 & 1 \\ 2 & 5 \end{bmatrix}$ by the principle of indifference. Rule out a saddle, set Player II's mix $(q, 1 - q)$ to equalize Player I's two rows, then set Player I's mix $(p, 1 - p)$ to equalize Player II's two columns. Report $\mathbf{p}^*, \mathbf{q}^*, V$ and verify the value both ways. (Worked in Appendix D.)

CHAPTER 6

Linear Programming and Duality

“The minimax theorem and linear programming duality are the same theorem wearing different clothes.”

This chapter pays a debt. In Chapter 4 we stated von Neumann’s minimax theorem, proved its easy half ($\underline{V} \leq \bar{V}$), and solved 2×2 games by hand, but we left the existence of the value in general unproven. The proof comes from linear programming. A zero-sum game turns into a pair of linear programs that are dual to one another, and the strong duality theorem of linear programming, that a primal and its dual share the same optimal value, is exactly the equality $\underline{V} = \bar{V}$ in disguise. Along the way we get an algorithm, the simplex method, that solves any matrix game, and an interpretation of optimal strategies as shadow prices.

6.1 Linear programs

Definition 6.1 (Linear program). A **linear program** (LP) maximizes or minimizes a linear function of several variables subject to linear inequality constraints. In the standard maximization form, with variables $\mathbf{y} \in \mathbb{R}^n$,

$$\text{maximize } \mathbf{c}^\top \mathbf{y} \quad \text{subject to } A\mathbf{y} \leq \mathbf{b}, \mathbf{y} \geq \mathbf{0},$$

where A is a given $m \times n$ matrix, $\mathbf{b} \in \mathbb{R}^m$, and $\mathbf{c} \in \mathbb{R}^n$. The **feasible region** is the set of $\mathbf{y}$ satisfying all constraints, and the **objective** is $\mathbf{c}^\top \mathbf{y}$.

The constraints $A\mathbf{y} \leq \mathbf{b}$, $\mathbf{y} \geq \mathbf{0}$ each cut $\mathbb{R}^n$ with a hyperplane and keep one side. Their intersection is a *convex polytope*, possibly unbounded. Linear programs model any situation in which a linear payoff is optimized against linear limits, from production planning to network flow to, as we will see, the playing of a game.

Example 6.2 (A production program). A firm makes two products in quantities $y_1, y_2 \geq 0$, earning profits c_1, c_2 per unit, and consuming shared resources whose availabilities are the entries of $\mathbf{b}$, with usage matrix A . The firm solves $\max \mathbf{c}^\top \mathbf{y}$ subject to $A\mathbf{y} \leq \mathbf{b}$, $\mathbf{y} \geq \mathbf{0}$: make the most valuable mix that the resources allow. We return to this reading in Section 6.8, where the dual variables turn out to price the resources.

6.2 The geometry of optima

The feasible polytope of an LP has corners, called *vertices* or *extreme points*, and a linear objective is always optimized at one of them.

Proposition 6.3 (Optima at vertices). If a linear program has an optimal solution and its feasible region has at least one vertex, then some vertex is optimal.

Idea of proof. The objective $\mathbf{c}^\top \mathbf{y}$ is linear, hence has no interior maximum on a convex set unless it is constant: from any non-vertex feasible point one can move along a direction inside the polytope that does not decrease the objective until a boundary is hit, and repeating drives the point to a vertex without ever lowering the objective. A linear function restricted to a line segment attains its maximum at an endpoint, which is the engine of each step. ■

Example 6.4 (Solving an LP by checking corners). Maximize $2y_1 + 3y_2$ subject to $y_1 + 2y_2 \leq 8$, $3y_1 + 2y_2 \leq 12$, $y_1, y_2 \geq 0$. The two slanted constraints and the two axes bound a quadrilateral (Figure 6.1). Find its vertices by intersecting the boundary lines in pairs: the origin $(0, 0)$; the point $(4, 0)$ where $3y_1 + 2y_2 = 12$ meets the y_1 -axis; the point $(0, 4)$ where $y_1 + 2y_2 = 8$ meets the y_2 -axis; and the interior corner where the two slanted lines cross, found by subtracting $y_1 + 2y_2 = 8$ from $3y_1 + 2y_2 = 12$ to get $2y_1 = 4$, so $y_1 = 2$ and $y_2 = 3$, the vertex $(2, 3)$. Evaluating $2y_1 + 3y_2$ at the four vertices,

$$(0, 0): 0, \quad (4, 0): 8, \quad (2, 3): 13, \quad (0, 4): 12,$$

so the maximum is 13, attained at the interior corner $(2, 3)$.




Figure 6.1. The feasible polytope of Example 6.4 (shaded). The objective $2y_1 + 3y_2$ increases in the direction of the arrow, so it is maximized at the vertex $(2, 3)$ farthest that way. A linear objective always attains its optimum at a vertex (Proposition 6.3).

When the polytope has many vertices, checking all of them is impractical, and the simplex method of Section 6.7 walks from vertex to neighboring vertex, improving the objective each step, until no improvement is possible.

6.3 Duality

Every linear program has a partner, its *dual*, obtained by transposing the data. The two are bound together by the duality theorems.

Definition 6.5 (Primal and dual). The **primal** maximization program and its **dual** minimization program are

$$(P) \quad \max_{y \geq 0} c^\top y \quad \text{s.t.} \quad Ay \leq b, \quad (D) \quad \min_{x \geq 0} b^\top x \quad \text{s.t.} \quad A^\top x \geq c.$$

The dual has one variable per primal constraint and one constraint per primal variable.

Theorem 6.6 (Weak duality). If y is feasible for (P) and x is feasible for (D), then $c^\top y \leq b^\top x$. Every dual feasible value is an upper bound for every primal feasible value.

Proof. Using $A^\top x \geq c$ and $y \geq 0$ for the first inequality, then $Ay \leq b$ and $x \geq 0$ for the last,

$$c^\top y \leq (A^\top x)^\top y = x^\top Ay \leq x^\top b = b^\top x.$$

■

Weak duality is elementary and already useful: any feasible dual solution certifies that the primal cannot exceed a known bound. The deep fact is that the bound is tight.

Theorem 6.7 (Strong duality). If both (P) and (D) are feasible, then both have optimal solutions and their optimal values are equal:

$$\max_y c^\top y = \min_x b^\top x.$$

Remarks on the proof. Strong duality is the central theorem of linear programming. It can be proved from the *separating hyperplane theorem* or, equivalently, from *Farkas's lemma* on the solvability of linear inequalities: if the primal optimum were strictly below the dual optimum, a separating hyperplane between the achievable objective values and the bound would produce a feasible point violating one of the programs. The result is due in this form to the work of von Neumann, Gale, Kuhn, and Tucker, and the constructive proof is the termination of the simplex method (Section 6.7); see Dantzig [14]. We take strong duality as established and use it to prove the minimax theorem. ■

6.4 From a game to a pair of linear programs

We now translate a zero-sum game into linear programs. Let A be the $m \times n$ payoff matrix. By the shift trick (Exercise 4.10), adding a constant to every entry changes the value by that constant and leaves optimal strategies unchanged, so we may assume every entry of A is positive, which guarantees the value $V > 0$. This positivity is what lets us divide by V below.

Player I's program. Player I seeks a mixed strategy $\mathbf{p}$ and the largest floor v that $\mathbf{p}$ guarantees against every column:

$$\max_{\mathbf{p}, v} \quad v \quad \text{s.t.} \quad \sum_i a_{ij} p_i \geq v \quad (j = 1, \dots, n), \quad \sum_i p_i = 1, \quad \mathbf{p} \geq \mathbf{0}.$$

Introduce the rescaled variables $x_i = p_i/v$, legitimate because $v > 0$. The constraints $\sum_i a_{ij} p_i \geq v$ become $\sum_i a_{ij} x_i \geq 1$, that is $A^\top \mathbf{x} \geq \mathbf{1}$, and the normalization $\sum_i p_i = 1$ becomes $\sum_i x_i = 1/v$. Maximizing v is the same as minimizing $1/v = \sum_i x_i$. So Player I's program is the linear program

$$(I) \quad \min_{\mathbf{x} \geq \mathbf{0}} \mathbf{1}^\top \mathbf{x} \quad \text{s.t.} \quad A^\top \mathbf{x} \geq \mathbf{1}, \quad \text{with } V = \frac{1}{\mathbf{1}^\top \mathbf{x}^*}, \quad \mathbf{p}^* = V \mathbf{x}^*. \quad (6.1)$$

Player II's program. Symmetrically, Player II seeks $\mathbf{q}$ and the smallest ceiling w holding Player I to at most w against every row:

$$\min_{\mathbf{q}, w} \quad w \quad \text{s.t.} \quad \sum_j a_{ij} q_j \leq w \quad (i = 1, \dots, m), \quad \sum_j q_j = 1, \quad \mathbf{q} \geq \mathbf{0}.$$

With $y_j = q_j/w$ this becomes $A\mathbf{y} \leq \mathbf{1}$ and $\sum_j y_j = 1/w$, and minimizing w is maximizing $1/w = \sum_j y_j$. So Player II's program is

$$(II) \quad \max_{\mathbf{y} \geq \mathbf{0}} \mathbf{1}^\top \mathbf{y} \quad \text{s.t.} \quad A\mathbf{y} \leq \mathbf{1}, \quad \text{with } w = \frac{1}{\mathbf{1}^\top \mathbf{y}^*}, \quad \mathbf{q}^* = w \mathbf{y}^*. \quad (6.2)$$

Programs (6.1) and (6.2) are precisely a dual pair in the sense of Definition 6.5, with $\mathbf{b} = \mathbf{c} = \mathbf{1}$: program (II) is the primal $\max \mathbf{1}^\top \mathbf{y}$ subject to $A\mathbf{y} \leq \mathbf{1}$, and program (I) is its dual $\min \mathbf{1}^\top \mathbf{x}$ subject to $A^\top \mathbf{x} \geq \mathbf{1}$.

6.5 The minimax theorem, proved

Everything is now in place to close the argument begun in Chapter 4.

Theorem 6.8 (Minimax theorem, restated). Every finite two-person zero-sum game has a value V and optimal mixed strategies $\mathbf{p}^*, \mathbf{q}^*$ with

$$\max_{\mathbf{p}} \min_{\mathbf{q}} \mathbf{p}^\top A \mathbf{q} = \min_{\mathbf{q}} \max_{\mathbf{p}} \mathbf{p}^\top A \mathbf{q} = V.$$

Proof. Step 1 (reduce to a positive matrix). Add a constant c to every entry of A , large enough that all entries become positive. By Exercise 4.10 this raises the value by c and leaves the optimal strategies unchanged, so it suffices to prove the theorem for $A > 0$ and then subtract c at the end. Positivity guarantees $V > 0$, which is what licenses dividing by V in the rescaling of Section 6.4.

Step 2 (both programs attain optima). Recall the dual pair from Section 6.4: program (I) is $\min \mathbf{1}^\top \mathbf{x}$ subject to $A^\top \mathbf{x} \geq \mathbf{1}, \mathbf{x} \geq \mathbf{0}$, and program (II) is $\max \mathbf{1}^\top \mathbf{y}$ subject to $A\mathbf{y} \leq \mathbf{1}, \mathbf{y} \geq \mathbf{0}$. Program (I) is feasible, since making every x_i large enough forces $A^\top \mathbf{x} \geq \mathbf{1}$; program (II) is feasible, since $\mathbf{y} = \mathbf{0}$ gives $A\mathbf{y} = \mathbf{0} \leq \mathbf{1}$. Program (I) is bounded below by 0, because its objective is a sum of nonnegative numbers, and program (II) is bounded above, because weak

duality (Theorem 6.6) makes every feasible value of (II) at most every feasible value of (I). A feasible and bounded linear program attains its optimum, so optimal $\mathbf{x}^*$ and $\mathbf{y}^*$ exist.

Step 3 (strong duality closes the gap). By Theorem 6.7 the two optimal values are equal; call the common value τ ,

$$\mathbf{1}^\top \mathbf{x}^* = \mathbf{1}^\top \mathbf{y}^* =: \tau.$$

Step 4 (translate back to the game). The construction of Section 6.4 identified the row player's guaranteed value as $\underline{V} = 1/(\mathbf{1}^\top \mathbf{x}^*)$ and the column player's ceiling as $\bar{V} = 1/(\mathbf{1}^\top \mathbf{y}^*)$. Step 3 makes these equal, $\underline{V} = \bar{V} = 1/\tau$, and the common number is the value V . Undoing the rescaling recovers the optimal mixed strategies $\mathbf{p}^* = V \mathbf{x}^*$ and $\mathbf{q}^* = V \mathbf{y}^*$, which are genuine distributions because $\mathbf{1}^\top \mathbf{p}^* = V \mathbf{1}^\top \mathbf{x}^* = V \cdot (1/V) = 1$. Subtracting the constant c from Step 1 gives the value of the original game. Because $\underline{V} = \bar{V}$ is exactly the statement $\max_{\mathbf{p}} \min_{\mathbf{q}} \mathbf{p}^\top A \mathbf{q} = \min_{\mathbf{q}} \max_{\mathbf{p}} \mathbf{p}^\top A \mathbf{q}$, the minimax equality is proved. ■

This is the heart of two-person zero-sum theory. The order of play does not matter, a value exists, and both players have strategies that secure it. The same strong duality theorem that prices resources in a factory guarantees a fair exchange rate in any matrix game.

Duality is minimax

The minimax theorem and the strong duality theorem of linear programming are one statement in two languages. Weak duality gives the easy half, $\underline{V} \leq \bar{V}$; strong duality gives the equality $\underline{V} = \bar{V} = V$. Solving a zero-sum game and solving a linear program are therefore interchangeable: the simplex method that solves the program also solves the game, and the optimal dual variables *are* the opponent's optimal mixed strategy.

Historical Note: “the two problems are equivalent”

The equivalence was discovered in conversation. On 3 October 1947 George Dantzig, who had just invented the simplex method, visited John von Neumann at the Institute for Advanced Study to describe linear programming and ask for a solution method. Instead of answering, von Neumann gave an impromptu lecture connecting linear programs to the theory of games, convex sets, and fixed points, and remarked: “I don't want you to think I am pulling all this out of my sleeve on the spur of the moment like a magician. I have just recently completed a book with Oskar Morgenstern on the theory of games. What I am doing is conjecturing that the two problems are equivalent.” [15]. The conjecture was made precise by Gale et al. [22], who derived the minimax theorem from LP duality. The story is the origin of this chapter's organizing idea.

Example 6.9 (Solving a game through its linear program). Take the game with all-positive matrix

$$A = \begin{bmatrix} 2 & 1 \\ 1 & 3 \end{bmatrix}.$$

Player I's program (I) is $\min x_1 + x_2$ subject to $2x_1 + x_2 \geq 1$, $x_1 + 3x_2 \geq 1$, $\mathbf{x} \geq \mathbf{0}$. At the optimum both constraints bind, giving $2x_1 + x_2 = 1$ and $x_1 + 3x_2 = 1$; solving, $x_1 = \frac{2}{5}$, $x_2 = \frac{1}{5}$, so $\mathbf{1}^\top \mathbf{x}^* = \frac{3}{5}$ and

$$V = \frac{1}{\mathbf{1}^\top \mathbf{x}^*} = \frac{5}{3}, \quad \mathbf{p}^* = V \mathbf{x}^* = \frac{5}{3} \cdot \left(\frac{2}{5}, \frac{1}{5} \right) = \left(\frac{2}{3}, \frac{1}{3} \right).$$

The dual program (II), $\max y_1 + y_2$ subject to $2y_1 + y_2 \leq 1$, $y_1 + 3y_2 \leq 1$, yields $\mathbf{y}^* = (\frac{2}{5}, \frac{1}{5})$ as well, with $\mathbf{1}^\top \mathbf{y}^* = \frac{3}{5}$, so $w = \frac{5}{3} = V$ and $\mathbf{q}^* = (\frac{2}{3}, \frac{1}{3})$. The LP value matches the 2×2 formula of Theorem 4.8, as it must, and strong duality is visible as $\mathbf{1}^\top \mathbf{x}^* = \mathbf{1}^\top \mathbf{y}^* = \frac{3}{5}$.

6.6 Complementary slackness is the principle of indifference

Strong duality has a companion that links the two solutions entry by entry, and in the game setting it is exactly the principle of indifference from Chapter 5.

Theorem 6.10 (Complementary slackness). Let $\mathbf{y}^*$ and $\mathbf{x}^*$ be optimal for the primal (P) and dual (D). Then for each variable and constraint,

$$x_i^* > 0 \Rightarrow (A\mathbf{y}^*)_i = b_i, \quad y_j^* > 0 \Rightarrow (A^\top \mathbf{x}^*)_j = c_j.$$

A positive variable forces its partner constraint to be tight.

Proof. The chain in the proof of weak duality (Theorem 6.6) becomes a chain of equalities at optimality, because strong duality makes its two ends equal. Equality in $\mathbf{c}^\top \mathbf{y}^* \leq (\mathbf{A}^\top \mathbf{x}^*)^\top \mathbf{y}^*$ forces, term by term, that whenever $y_j^* > 0$ the corresponding constraint $(\mathbf{A}^\top \mathbf{x}^*)_j \geq c_j$ holds with equality, and similarly for the other end. ■

Translated through the game-to-LP dictionary, $x_i^* > 0$ means $p_i^* > 0$, and the tight constraint $(\mathbf{A}\mathbf{y}^*)_i = \text{const}$ means row i earns exactly the value. That is precisely Theorem 5.1: a strategy played with positive probability earns exactly the value. The principle of indifference and complementary slackness are one statement seen from two sides.

6.7 The simplex method

The simplex method, due to Dantzig, solves a linear program by walking along the edges of the feasible polytope from vertex to vertex, improving the objective at each step until no neighbor is better, at which point strong duality certifies optimality.

Procedure 6.11 (Simplex method, in outline).

1. **Slack form.** Convert inequalities to equalities by adding nonnegative slack variables, so $\mathbf{A}\mathbf{y} \leq \mathbf{b}$ becomes $\mathbf{A}\mathbf{y} + \mathbf{s} = \mathbf{b}$ with $\mathbf{s} \geq \mathbf{0}$. A **basic feasible solution** sets n of the variables to zero and solves for the rest; geometrically it is a vertex.
2. **Pivot.** Examine the objective's *reduced costs* at the current vertex. If some nonbasic variable would increase the objective, bring it into the basis (let it grow) and remove whichever basic variable hits zero first, by the minimum-ratio test. This moves to an adjacent vertex with a better objective.
3. **Stop.** When no nonbasic variable improves the objective, the current vertex is optimal, and the dual solution read from the final tableau certifies it by strong duality.

For a matrix game, one applies the simplex method to program (I) or (II) of Section 6.4. Program (II) is the convenient one to demonstrate, because the origin is feasible and serves as a starting vertex.

Example 6.12 (A simplex walk on a matrix game). Solve Player II's program for $\mathbf{A} = \begin{bmatrix} 2 & 1 \\ 1 & 3 \end{bmatrix}$ of Example 6.9: maximize $z = y_1 + y_2$ subject to $2y_1 + y_2 \leq 1$, $y_1 + 3y_2 \leq 1$, $\mathbf{y} \geq \mathbf{0}$. Add slacks $s_1, s_2 \geq 0$ to get the equalities

$$2y_1 + y_2 + s_1 = 1, \quad y_1 + 3y_2 + s_2 = 1, \quad z = y_1 + y_2.$$

Vertex 0: the origin. Set the nonbasic $y_1 = y_2 = 0$; then $s_1 = s_2 = 1$ and $z = 0$. Both y_1 and y_2 have reduced cost $+1$ in z , so either entering the basis raises the objective. Bring in y_1 .

Pivot 1 (minimum-ratio test). Increasing y_1 alone, the first constraint allows $y_1 \leq \frac{1}{2}$ (from $2y_1 \leq 1$) and the second allows $y_1 \leq 1$ (from $y_1 \leq 1$). The binding one is the first, $\frac{1}{2} < 1$, so s_1 leaves the basis. Solving the first equation for y_1 gives $y_1 = \frac{1}{2} - \frac{1}{2}y_2 - \frac{1}{2}s_1$, and substituting into z ,

$$z = y_1 + y_2 = \frac{1}{2} + \frac{1}{2}y_2 - \frac{1}{2}s_1.$$

Vertex 1. Now $y_1 = \frac{1}{2}$, $y_2 = 0$, $z = \frac{1}{2}$. The reduced cost of y_2 is still $+\frac{1}{2} > 0$, so we continue; bring in y_2 .

Pivot 2. Substituting $y_1 = \frac{1}{2} - \frac{1}{2}y_2 - \frac{1}{2}s_1$ into the second constraint gives $\frac{5}{2}y_2 - \frac{1}{2}s_1 + s_2 = \frac{1}{2}$, so increasing y_2 alone allows $y_2 \leq \frac{1}{5}$; s_2 leaves. Then $y_2 = \frac{1}{5}$ and back-substitution gives $y_1 = \frac{1}{2} - \frac{1}{2} \cdot \frac{1}{5} = \frac{2}{5}$. The objective is

$$z = \frac{2}{5} + \frac{1}{5} = \frac{3}{5},$$

and rewriting z in the new nonbasic variables yields $z = \frac{3}{5} - \frac{2}{5}s_1 - \frac{1}{5}s_2$.

Vertex 2: optimal. Both remaining reduced costs $-\frac{2}{5}, -\frac{1}{5}$ are negative, so no further entering variable improves z ; the vertex $\mathbf{y}^* = (\frac{2}{5}, \frac{1}{5})$ is optimal with $\mathbf{1}^\top \mathbf{y}^* = \frac{3}{5}$. Rescaling by $w = 1/(\mathbf{1}^\top \mathbf{y}^*) = \frac{5}{3}$ recovers $\mathbf{q}^* = w\mathbf{y}^* = (\frac{2}{3}, \frac{1}{3})$ and $V = w = \frac{5}{3}$, agreeing with Example 6.9. The walk visited three vertices, $(0, 0) \rightarrow (\frac{1}{2}, 0) \rightarrow (\frac{2}{5}, \frac{1}{5})$, improving the objective $0 \rightarrow \frac{1}{2} \rightarrow \frac{3}{5}$ at each step. ◇

For a game that is not already in vertex-walk form, the bookkeeping is cleanest in a *tableau*, which we re-display after every pivot so the reduction is visible at each stage.

Example 6.13 (Solving a 3×3 game by the simplex tableau). Solve

$$A = \begin{bmatrix} 2 & -1 & 6 \\ 0 & 1 & -1 \\ -2 & 2 & 1 \end{bmatrix}.$$

Step 1: no saddle, so make the value positive. The row minima are $-1, -1, -2$ (maximin -1); the column maxima are $2, 2, 6$ (minimax 2); since $-1 \neq 2$ there is no saddle. The tableau method needs a positive value, so add the constant 2 to every entry, which raises the value by 2 and leaves the optimal strategies unchanged (Exercise 4.10):

$$\tilde{A} = A + 2\mathbf{1}\mathbf{1}^\top = \begin{bmatrix} 4 & 1 & 8 \\ 2 & 3 & 1 \\ 0 & 4 & 3 \end{bmatrix}.$$

Step 2: the starting tableau. Border $\tilde{A}$ with a right column of 1s (Player I's normalization), a bottom row of -1 s (the objective), and a 0 corner. Label the rows by Player I's variables x_1, x_2, x_3 and the columns by Player II's y_1, y_2, y_3 :

	y_1	y_2	y_3	1
x_1	4	1	8	1
x_2	2	3	1	1
x_3	0	4	3	1
	-1	-1	-1	0

Step 3: first pivot. A negative entry in the bottom row means the objective can still improve; take column y_1 . The pivot row is the one minimizing the ratio (right entry)/(column entry) among positive entries: $\frac{1}{4}, \frac{1}{2}$, and $\frac{1}{0}$ is excluded, so row x_1 wins with pivot element 4. Pivoting (divide the pivot row by 4, clear the rest of the column, and update each other entry by $t_{ij} \leftarrow t_{ij} - t_{i,y_1}t_{x_1,j}/4$) and swapping the labels $x_1 \leftrightarrow y_1$ gives

	x_1	y_2	y_3	1
y_1	$\frac{1}{4}$	$\frac{1}{4}$	2	$\frac{1}{4}$
x_2	$-\frac{1}{2}$	$\frac{5}{2}$	-3	$\frac{1}{2}$
x_3	0	4	3	1
	$\frac{1}{4}$	$-\frac{3}{4}$	1	$\frac{1}{4}$

Step 4: second pivot. The bottom row still has a negative entry under y_2 . The ratio test $\frac{1/4}{1/4} = 1$, $\frac{1/2}{5/2} = \frac{1}{5}$, $\frac{1}{4} = \frac{1}{4}$ picks row x_2 , pivot element $\frac{5}{2}$. Pivoting and swapping $x_2 \leftrightarrow y_2$ gives

	x_1	x_2	y_3	1
y_1	$\frac{3}{10}$	$-\frac{1}{10}$	$\frac{23}{10}$	$\frac{1}{5}$
y_2	$-\frac{1}{5}$	$\frac{2}{5}$	$-\frac{6}{5}$	$\frac{1}{5}$
x_3	$\frac{4}{5}$	$-\frac{8}{5}$	$\frac{39}{5}$	$\frac{1}{5}$
	$\frac{1}{10}$	$\frac{3}{10}$	$\frac{1}{10}$	$\frac{2}{5}$

Every bottom-row entry is now ≥ 0 , so the tableau is optimal.

Step 5: read off the answer. Let Z be the corner, $Z = \frac{2}{5}$. The value is $1/Z$ minus the shift: $V = \frac{1}{Z} - 2 = \frac{5}{2} - 2 = \frac{1}{2}$. Player I's weights are the bottom-row entries under the swapped x -columns, divided by Z : $x_1 = \frac{1}{10}$, $x_2 = \frac{3}{10}$, and $x_3 = 0$ (never swapped), so

$$\mathbf{p}^* = \frac{1}{Z} \left(\frac{1}{10}, \frac{3}{10}, 0 \right) = \frac{5}{2} \left(\frac{1}{10}, \frac{3}{10}, 0 \right) = \left(\frac{1}{4}, \frac{3}{4}, 0 \right).$$

Player II's weights are the right-column entries on the swapped y -rows, divided by Z : $y_1 = \frac{1}{5}$, $y_2 = \frac{1}{5}$, $y_3 = 0$, so $\mathbf{q}^* = \frac{5}{2} \left(\frac{1}{5}, \frac{1}{5}, 0 \right) = \left(\frac{1}{2}, \frac{1}{2}, 0 \right)$. As a check, play $\mathbf{p}^*$ against the original columns: column 1 pays $\frac{1}{4} \cdot 2 + \frac{3}{4} \cdot 0 + 0 = \frac{1}{2}$, column 2 pays $\frac{1}{4}(-1) + \frac{3}{4}(1) = \frac{1}{2}$, and column 3 pays $\frac{1}{4} \cdot 6 + \frac{3}{4}(-1) = \frac{6}{4} - \frac{3}{4} = \frac{3}{4} \geq \frac{1}{2}$. Both players drop their third strategy, and the value is $\frac{1}{2}$. $\diamond$

The offset that keeps the game value positive, and the reading of both players' strategies from one final tableau, are also exercised in Appendix B, and implemented in code in Appendix C, where Problem Set 2 builds a simplex-based solver that handles any matrix game.

Computational Companion

Problem Set 2 asks for a single program that solves an arbitrary zero-sum game by seven methods: checking a saddle point, the 2×2 formula, the $2 \times n$ and $m \times 2$ graphical method, the principle of indifference, the principle of invariance, the inverse formula for non-degenerate $n \times n$ games, and the linear program. The last method is the simplex routine of this section, and it is the catch-all that solves the games the special methods miss. The solver is cross-checked against an independent reference solver on a hundred random games in [Appendix C](#).

6.8 Shadow prices

The dual variables have a meaning worth stating, because it recurs when games model markets. In the production program of [Example 6.2](#), the dual variable x_i attached to resource i is its **shadow price**: the marginal increase in optimal profit per additional unit of that resource. Complementary slackness says a resource has a positive shadow price only if it is fully used, and a slack resource is free at the margin. In the game setting, the same dual variables are the opponent's optimal strategy, so an opponent's optimal mixed strategy is the "price" the game places on each of one's own pure strategies, charging exactly the value for those one actually uses and more for those one should avoid. The economic and the strategic readings are the same mathematics.

Intuition: a shadow price is the most you would pay to rent one more unit

Picture a workshop that turns scarce resources into profit. The shadow price of, say, an hour of machine time is the extra profit one more hour would unlock, which is exactly the most you should be willing to pay to rent that hour. Two consequences follow and both are complementary slackness in plain words. A resource you do *not* fully use has shadow price zero: another unit sits idle, so you would pay nothing for it. And a resource priced above zero must be used to the last drop, or you could raise profit without buying anything. Reading the same dual through the game lens, the opponent's optimal mix is the schedule of prices the game attaches to your pure strategies: it pays the full value for the moves you genuinely use and overcharges the moves you are right to skip, which is why an optimal player is indifferent among the strategies in their support and strictly prefers to avoid the rest.

Notes and References

Linear programming, the simplex method, and the duality theorem are due to George Dantzig, who developed the simplex method in 1947; the standard reference is Dantzig [14] and his own historical account is Dantzig [15]. The first large LP solved in practice was Stigler's *diet problem* of finding a minimum-cost adequate diet [61]: in 1947 a team at the National Bureau of Standards solved Stigler's 9-constraint, 77-variable model by hand with desk calculators in about 120 person-days, reaching the optimum of \$39.69 per year at 1939 prices. The equivalence between LP duality and the minimax theorem was made explicit by Gale et al. [22] in the same 1950s volumes that contain Shapley [55] and Kuhn [29]. The simplex method is fast in practice but not in the worst case: Klee and Minty [27] built a distorted n -cube with 2^n vertices on which the standard pivot rule visits all of them, an exponential running time that motivated later polynomial methods. Strong duality can be proved from Farkas's lemma or the separating hyperplane theorem (Theorem A.4); we have taken it as given and used it to prove the minimax theorem, repaying the debt left open in Chapter 4. The simplex treatment and the seven-method solver follow Ferguson [16].

Chapter summary

A linear program optimizes a linear objective over a convex polytope (Definition 6.1), and an optimum is always attained at a vertex (Proposition 6.3). Every LP has a dual (Definition 6.5); weak duality bounds the primal by the dual (Theorem 6.6), and strong duality makes the two optimal values equal (Theorem 6.7). A zero-sum game becomes a dual pair of linear programs through the rescaling $\mathbf{x} = \mathbf{p}/V$, $\mathbf{y} = \mathbf{q}/w$ (Section 6.4), and strong duality applied to that pair proves the minimax theorem in full (Theorem 6.8), closing the gap left in Chapter 4. Complementary slackness (Theorem 6.10) is the principle of indifference seen through the duality dictionary. The simplex method (Procedure 6.11) solves any matrix game by walking the vertices of the feasible polytope, and the dual variables carry the economic meaning of shadow prices.

Exercises

Linear programs, geometry, and duality

Exercise 6.1. Solve by checking corners: maximize $3y_1 + 2y_2$ subject to $y_1 + y_2 \leq 4$, $y_1 + 3y_2 \leq 6$, $y_1, y_2 \geq 0$. List the vertices of the feasible region and the objective at each.

Exercise 6.2. State the dual of the linear program in Exercise 6.1 and verify weak duality by exhibiting one feasible dual point whose objective bounds the primal optimum from above.

Exercise 6.3. For the game $A = \begin{bmatrix} 3 & 1 \\ 1 & 2 \end{bmatrix}$, write Player I's linear program (6.1) and Player II's program (6.2), solve both, and check that the value and strategies agree with the 2×2 formula of Theorem 4.8.

Exercise 6.4. The game $A = \begin{bmatrix} -1 & 2 \\ 3 & -2 \end{bmatrix}$ has entries of both signs. Add a constant to make all entries positive, write the resulting LP, solve it, and recover the value of the original game by undoing the shift. Confirm that the optimal strategies are unchanged by the shift.

Exercise 6.5. For the game of Example 6.9, verify complementary slackness directly: check that both rows earn exactly the value $\frac{5}{3}$ against $\mathbf{q}^*$, consistent with both components of $\mathbf{p}^*$ being positive. Relate this to the principle of indifference of Theorem 5.1.

Exercise 6.6. Explain in your own words why strong duality for linear programs is equivalent to the equality $\underline{V} = \overline{V}$ in the minimax theorem. Which direction is weak duality, and which is the substantive content?

The simplex method

Exercise 6.7. Apply the simplex method (Procedure 6.11) to the primal $\max y_1 + y_2$ subject to $2y_1 + y_2 \leq 1$, $y_1 + 3y_2 \leq 1$, $\mathbf{y} \geq \mathbf{0}$ (Player II's program for Example 6.9). Start at the vertex $\mathbf{y} = \mathbf{0}$, perform the pivots, and confirm you reach $\mathbf{y}^* = (\frac{2}{5}, \frac{1}{5})$.

CHAPTER 7

Recursive and Stochastic Games

“The game is worth v because playing it again is worth v . Solve that sentence and you have solved the game.”

The zero-sum theory of Chapters 4 to 6 solves any single matrix game. Real strategic situations, though, often repeat: an inspection is run period after period, a duel of nerves can reset, a coin is flipped again and again. When one outcome of a game is “play this same game again,” the value of that outcome is the value of the whole game, and the problem becomes self-referential. Shapley’s insight [54] is that this self-reference is not a paradox but a fixed-point equation: substitute the unknown value V wherever the rules say “repeat,” solve the resulting one-shot game, and demand that its value equal V . This chapter develops that idea, from a single recursive cell to coupled systems of games to genuinely stochastic games with random transitions between states.

7.1 The substitution principle

A *recursive game* is a zero-sum game in which at least one entry of the payoff matrix, instead of paying a fixed amount, returns play to the same game. To value it, give the whole game an unknown value V and replace the recursive entry by V , because continuing is worth exactly what the game is worth. The result is an ordinary matrix whose entries may contain V , and whose value, computed by the methods of Chapter 4, must come out equal to V . That consistency requirement is a scalar equation for V .

Procedure 7.1 (Solving a recursive game by substitution).

1. Let V denote the value of the entire recursive game. Replace every entry that means “play this game again” by V (by $-V$ if the repeat also switches the players’ roles, and by an expected value if the repeat is reached only with some probability).
2. Compute the value of the resulting matrix $A(V)$ as a function of V , using the saddle-point test and, when there is no saddle, the 2×2 formula $V(A) = (ad - bc)/(a - b - c + d)$.
3. Impose the fixed-point condition $V = V(A(V))$ and solve for V .
4. Select the root that respects the security bounds the player can guarantee by simple pure strategies.

The substitution principle

When a game can “play itself again,” set its unknown value V equal to the value of the matrix that contains V , and solve the resulting fixed-point equation. A role switch on the repeat enters as $-V$; a random transition enters as the expected continuation value. The admissible root is the one the security bounds permit.

The security bounds in the last step matter: the fixed-point equation is often quadratic and returns two roots, only one of which is a value a player could actually achieve.

Substituting V for the recursive cell can produce a value that is not unique. In Example 7.3 the fixed-point condition is satisfied by a whole interval $[1, 5]$, because the self-referential entry, unlike a fixed number, can consistently equal anything between the other column’s payoffs. Always check the candidate roots against the bounds a player can guarantee by simple pure play, and be alert that a recursive game can fail to pin down a single value.

7.2 A first recursive game

Example 7.2 (A recursive 2×2 with a unique value). Consider

$$G = \begin{bmatrix} G & 1 \\ 1 & 0 \end{bmatrix},$$

where the top-left entry G means “play this same game again.” Writing V for the value and substituting $G \mapsto V$ gives the effective matrix $A(V) = \begin{bmatrix} V & 1 \\ 1 & 0 \end{bmatrix}$. There is no saddle for the relevant range, so the 2×2 formula with $a = V$, $b = 1$, $c = 1$, $d = 0$ gives

$$V(A(V)) = \frac{V \cdot 0 - 1 \cdot 1}{V - 1 - 1 + 0} = \frac{-1}{V - 2}.$$

The fixed-point condition $V = -1/(V - 2)$ becomes $V(V - 2) = -1$, that is $V^2 - 2V + 1 = (V - 1)^2 = 0$, so

$$V = 1.$$

The double root signals a unique, stable value. A behavioral check confirms it: if Player I plays $(1 - \varepsilon, \varepsilon)$, Player II’s best reply is column 2, paying Player I $(1 - \varepsilon) \cdot 1 + \varepsilon \cdot 0 = 1 - \varepsilon$, which approaches 1 as $\varepsilon \rightarrow 0$. Player I can guarantee arbitrarily close to 1, and no more.

Recursion can do something a fixed matrix never does: leave the value undetermined within a whole interval.

Example 7.3 (A continuum of self-consistent values). Replace the 1 in the top row by 5:

$$G = \begin{bmatrix} G & 5 \\ 1 & 0 \end{bmatrix}, \quad A(V) = \begin{bmatrix} V & 5 \\ 1 & 0 \end{bmatrix}.$$

Examine the dominance structure as a function of the hypothesized value V . For Player I, the top row reads $(V, 5)$ and the bottom row $(1, 0)$. When $V \geq 1$ the top row weakly dominates, so Player I plays it, and Player II, minimizing over $(V, 5)$, picks the smaller. If also $V \leq 5$, the smaller is V , attained in column 1, so the matrix has a saddle at the top-left with value exactly V . Every $V \in [1, 5]$ is therefore self-consistent. Outside the interval, consistency fails: if $V > 5$ Player II switches to column 2 and caps the value at 5, a contradiction; if $V < 1$ the 2×2 formula gives $V(A(V)) = -5/(V - 6)$, and $V = -5/(V - 6)$ yields $V^2 - 6V + 5 = (V - 1)(V - 5) = 0$, whose roots 1 and 5 are not below 1. The full set of values is

$$V \in [1, 5].$$

The endogenous entry is the source of the non-uniqueness: in an ordinary matrix the value is pinned by exogenous numbers, but here the top-left entry *is* the value, and any number between the other column’s payoffs is consistent with itself.

7.3 Symmetry in recursive games

When a recursive matrix has a symmetry, the invariance principle of [Chapter 5](#) fixes the optimal mixtures and collapses the fixed-point equation to a single average.

Example 7.4 (A cyclic 3×3 recursive game). Let

$$G = \begin{bmatrix} G & 1 & 0 \\ 1 & 0 & G \\ 0 & G & 1 \end{bmatrix}, \quad A(V) = \begin{bmatrix} V & 1 & 0 \\ 1 & 0 & V \\ 0 & V & 1 \end{bmatrix}.$$

Each row and each column contains one V , one 1, and one 0, rotated, so $A(V)$ is invariant under the simultaneous cyclic permutation of rows and columns $1 \rightarrow 2 \rightarrow 3 \rightarrow 1$. By [Theorem 5.10](#) an invariant optimal strategy exists, and the only distribution fixed by a 3-cycle is the uniform one, $\mathbf{p}^* = \mathbf{q}^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Against the uniform column strategy each row earns the average of its entries, $(V + 1 + 0)/3$, so the value satisfies

$$V = \frac{V + 1}{3} \implies 3V = V + 1 \implies V = \frac{1}{2}.$$

A check confirms indifference: at $V = \frac{1}{2}$ each row earns $(\frac{1}{2} + 1)/3 = \frac{1}{2}$ against the uniform mix, so Player I is indifferent among all three rows, and by symmetry so is Player II.

7.4 Role-switching recursion

Sometimes the repeat swaps the players. In a zero-sum game, swapping roles negates the continuation, so the recursive entry is replaced by $-V$ rather than V .

Example 7.5 (Pass or dare). Each round both players choose PASS or DARE. Player I's one-shot payoffs are

	PASS	DARE
PASS	0	1
DARE	3	repeat, roles switched

If both dare, the game repeats with the players exchanged, contributing $-V$ to Player I. Player I can guarantee at least 0 by passing every round (the worst entry in the top row is 0 and the recursive branch is never triggered), so $V \geq 0$. Substituting the role-switch as $-V$,

$$A(V) = \begin{bmatrix} 0 & 1 \\ 3 & -V \end{bmatrix}.$$

The maximin is 0 and the minimax is 1, so there is no saddle and mixing is forced. The 2×2 formula with $a = 0$, $b = 1$, $c = 3$, $d = -V$ gives

$$V(A(V)) = \frac{0 \cdot (-V) - 1 \cdot 3}{0 - 1 - 3 - V} = \frac{-3}{-4 - V} = \frac{3}{V + 4}.$$

The fixed point $V = 3/(V + 4)$ is $V^2 + 4V - 3 = 0$, so $V = -2 \pm \sqrt{7}$. The nonnegativity bound selects

$$V = -2 + \sqrt{7} \approx 0.646.$$

The optimal mixtures, from the indifference formulas with $D = -4 - V$, are $\mathbf{p}^* = (\frac{5-\sqrt{7}}{3}, \frac{\sqrt{7}-2}{3})$ and $\mathbf{q}^* = (3 - \sqrt{7}, \sqrt{7} - 2)$, both valid probability vectors. Moving first confers a genuine but modest edge.

7.5 Coupled subgames

Several games can reference one another, forming a system of value equations.

Example 7.6 (Three linked games). Let

$$G_1 = \begin{bmatrix} G_2 & 0 \\ 0 & G_3 \end{bmatrix}, \quad G_2 = \begin{bmatrix} G_1 & 1 \\ 1 & 0 \end{bmatrix}, \quad G_3 = \begin{bmatrix} G_1 & 2 \\ 2 & 0 \end{bmatrix},$$

with values v_1, v_2, v_3 . Simple pure strategies bound them: in G_2 Player I secures ≥ 0 by the second row and Player II caps the value at ≤ 1 by the second column, so $0 \leq v_2 \leq 1$; likewise $0 \leq v_3 \leq 2$; and in G_1 the minimizer caps the value at $\min\{v_2, v_3\} \leq 1$, so $0 \leq v_1 \leq 1$. Applying the 2×2 formula to each,

$$v_1 = \frac{v_2 v_3}{v_2 + v_3}, \quad v_2 = \frac{1}{2 - v_1}, \quad v_3 = \frac{4}{4 - v_1}.$$

The first relation is the harmonic identity $1/v_1 = 1/v_2 + 1/v_3$. Substituting $1/v_2 = 2 - v_1$ and $1/v_3 = 1 - v_1/4$,

$$\frac{1}{v_1} = (2 - v_1) + \left(1 - \frac{v_1}{4}\right) = 3 - \frac{5}{4}v_1,$$

which clears to $5v_1^2 - 12v_1 + 4 = 0$, with roots $v_1 \in \{\frac{2}{5}, 2\}$. The bound $v_1 \leq 1$ selects $v_1 = \frac{2}{5}$, and back-substitution gives

$$v_1 = \frac{2}{5}, \quad v_2 = \frac{5}{8}, \quad v_3 = \frac{10}{9}.$$

7.6 Stochastic games

Shapley's stochastic games [54] generalize recursion by letting an entry trigger a *random* transition: with stated probabilities, play moves to one of several states or terminates with a payoff. The value of each state then satisfies a Bellman fixed-point equation, and the substitution principle applies once each recursive entry is replaced by the appropriate expected continuation value.

Intuition: a board game whose board changes under you

An ordinary matrix game is one board, played once. A stochastic game is a small archipelago of boards: you play the matrix of the current state, and the cell you land on not only pays out but also rolls dice to decide *which board you play next*, or whether the game ends. Each state's value v_s is what the whole future is worth starting there, so the number written in a cell is its immediate reward *plus* the expected value of wherever you get sent, $\sum_{s'} p_{ss'} v_{s'}$. That circularity, the value of a state depending on the values of the states it can reach (itself included), is why we solve a *system* of fixed-point equations rather than a single matrix. It is the same logic a chess engine uses when it scores a position by the positions it leads to, carried out exactly for games small enough to write down.

Example 7.7 (Coin-flipping between two states). Two states have matrices

$$G_1 = \begin{bmatrix} \frac{1}{2}v_2 + \frac{1}{2} \cdot 0 & 1 \\ 2 & 0 \end{bmatrix}, \quad G_2 = \begin{bmatrix} \frac{2}{3}v_2 + \frac{1}{3}(-2) & 0 \\ 0 & -1 \end{bmatrix},$$

where in G_1 the top-left cell flips a fair coin between continuing in state 2 (worth v_2) and terminating at 0, and in G_2 the top-left cell continues in state 2 with probability $\frac{2}{3}$ and otherwise terminates at -2 . Pure-strategy bounds give $0 \leq v_1 \leq 1$ and $-1 \leq v_2 \leq 0$. The 2×2 formula applied to each effective matrix gives

$$v_1 = \frac{(\frac{1}{2}v_2) \cdot 0 - 1 \cdot 2}{\frac{1}{2}v_2 - 1 - 2 + 0} = \frac{-2}{\frac{1}{2}v_2 - 3} = \frac{4}{6 - v_2}, \quad v_2 = \frac{(\frac{2}{3}v_2 - \frac{2}{3})(-1) - 0}{(\frac{2}{3}v_2 - \frac{2}{3}) - 1} = \frac{2 - 2v_2}{2v_2 - 5}.$$

The second equation involves only v_2 , so solve it first. Cross-multiplying,

$$v_2(2v_2 - 5) = 2 - 2v_2 \implies 2v_2^2 - 5v_2 = 2 - 2v_2 \implies 2v_2^2 - 3v_2 - 2 = 0,$$

which factors as $(2v_2 + 1)(v_2 - 2) = 0$, with roots $v_2 = -\frac{1}{2}$ and $v_2 = 2$. The bound $-1 \leq v_2 \leq 0$ selects $v_2 = -\frac{1}{2}$. Substituting into the first equation,

$$v_1 = \frac{4}{6 - (-\frac{1}{2})} = \frac{4}{13/2} = \frac{8}{13} \approx 0.615.$$

The signs tell the story: starting in state 1 is advantageous ($v_1 > 0$), while state 2 loops back on itself and leaks the occasional loss, so its value is negative. Figure 7.1 diagrams the two states and their transitions.




Figure 7.1. The transition structure of the coin game (Example 7.7). Only the top-left “continue” cell of each state moves play; every other cell ends the game at once. From S_1 that cell flips a fair coin between entering S_2 and terminating at 0; from S_2 it stays in S_2 with probability $\frac{2}{3}$ and terminates at -2 with probability $\frac{1}{3}$. Solving the resulting value equations gives $v_1 = \frac{8}{13}$ and $v_2 = -\frac{1}{2}$.

Example 7.8 (A skewed coin with role-switch). A coin lands heads with probability $\frac{2}{3}$. Each round both players call heads or tails. If exactly one is correct, that player wins and the game ends (Player I collects $+1$ for a correct heads call, $+4$ for the rarer correct tails call, and 0 if Player II is the correct one). If both are correct the game repeats with the same roles; if both are wrong it repeats with roles switched. Averaging each cell over the coin and writing the same-role repeat as V and the switched repeat as $-V$ gives the effective matrix

$$A(V) = \begin{bmatrix} \frac{1}{3}V & \frac{2}{3} \\ \frac{4}{3} & -\frac{1}{3}V \end{bmatrix}.$$

The minimizer can hold the value to at most $\frac{2}{3}$, and safe play gives $V \geq 0$. The 2×2 formula gives

$$V(A(V)) = \frac{(\frac{1}{3}V)(-\frac{1}{3}V) - \frac{2}{3} \cdot \frac{4}{3}}{\frac{1}{3}V - \frac{2}{3} - \frac{4}{3} - \frac{1}{3}V} = \frac{-\frac{V^2}{9} - \frac{8}{9}}{-2} = \frac{V^2 + 8}{18}.$$

The fixed point $V = (V^2 + 8)/18$ is $V^2 - 18V + 8 = 0$, with roots $V = 9 \pm \sqrt{73}$. Only $V = 9 - \sqrt{73} \approx 0.456$ respects the bound $0 \leq V \leq \frac{2}{3}$.

7.7 Repeated games and the inspection problem

A *repeated game* plays a fixed stage game over several periods, and a strategy may condition on the history. The simplest repeated zero-sum games unwind by the same recursive logic, replacing the continuation with the value of the remaining game. The classic example is inspection.

Example 7.9 (The inspector and the thief). An inspector and a thief face off over n remaining time slots. In each slot the inspector chooses to INSPECT or WAIT and the thief to STEAL or WAIT, a zero-sum stage with payoffs to the inspector. Catching a theft, inspecting while the thief steals, is worth 1; an inspection wasted on an idle thief, or a theft that goes unwatched, is worth 0; and if both wait, the horizon shortens to $n - 1$ slots and play continues, worth the continuation value G_{n-1} . Writing G_n for the value with n slots left,

$$M_n = \begin{bmatrix} 1 & 0 \\ 0 & G_{n-1} \end{bmatrix}, \quad G_1 = 1,$$

the boundary $G_1 = 1$ saying that in the final slot the inspector can secure the catch. For $G_{n-1} > 0$ the matrix M_n has no saddle, so the 2×2 formula of Theorem 4.8 with $a = 1$, $b = c = 0$, $d = G_{n-1}$ gives the recursion

$$G_n = \frac{ad - bc}{a - b - c + d} = \frac{1 \cdot G_{n-1} - 0}{1 - 0 - 0 + G_{n-1}} = \frac{G_{n-1}}{1 + G_{n-1}}.$$

Starting from $G_1 = 1$ this telescopes:

$$G_2 = \frac{1}{1+1} = \frac{1}{2}, \quad G_3 = \frac{1/2}{1+1/2} = \frac{1}{3}, \quad \dots, \quad G_n = \frac{1}{n},$$

the last by induction, since $G_{n-1} = 1/(n-1)$ gives $G_n = \frac{1/(n-1)}{1+1/(n-1)} = \frac{1}{n}$. The optimal strategies, from the same formulas with $a = 1$, $b = c = 0$, $d = G_{n-1}$, have the inspector inspect and the thief steal each with probability

$$p_n = q_n = \frac{G_{n-1}}{1 + G_{n-1}} = \frac{1}{n},$$

and wait otherwise. With more time remaining each side acts more rarely, and the value $1/n$ falls: a longer horizon dilutes the single catch the inspector can hope to make. $\diamond$

The inspection game is recursion in its purest form: the bottom-right cell of M_n is the $(n-1)$ -slot game, exactly the “play again” structure of Section 7.1, and the value telescopes to $1/n$.

Application: arms control and treaty verification

Inspection games are not a toy. They are the standard model for verifying arms-control and nuclear-nonproliferation agreements, where an inspectorate with a limited number of on-site visits faces a party that might divert material undetected, and for auditing, customs, and doping control. The recursion is the same as here: each period the inspector decides whether to spend one of its scarce inspections, the inspected party decides whether to cheat, and the continuation is the game with one fewer inspection. The conclusion that both sides should randomize, and that the detection guarantee falls as the horizon lengthens, is exactly what treaty designers must reckon with when fixing how many inspections a regime allows. The game-theoretic literature on this, developed for the Nuclear Non-Proliferation Treaty, is surveyed by Avenhaus et al. [2].

7.8 The repeated Prisoner's Dilemma and the Folk Theorem

The inspection game repeats a zero-sum stage, where the players have nothing to cooperate about. Repetition turns subtle when the stage game is the Prisoner's Dilemma of [Example 1.8](#), whose one-shot equilibrium of mutual defection leaves both worse off than mutual cooperation. Can the prospect of future play hold cooperation together?

Over a *fixed, finite* number of rounds it cannot. The last round is a one-shot Prisoner's Dilemma, so both defect. With the last round settled, the second-to-last round has no future left to protect, so defection dominates there too, and the reasoning unravels back to the first round. The unique subgame perfect equilibrium of the finitely repeated Prisoner's Dilemma is to defect in every round. Backward induction is merciless.

An *infinite* horizon breaks the unraveling, because there is no last round to start the induction from. Let the players discount future payoffs by a factor $\delta \in (0, 1)$ per period, so a stream $u_0, u_1, u_2, \dots$ is worth $\sum_{t \geq 0} \delta^t u_t$. The factor δ measures patience, or equivalently the probability that another round will be played.

Definition 7.10 (Grim-trigger strategy). A player using the **grim trigger** cooperates in the first round and keeps cooperating as long as no one has ever defected; the first time either player defects, they defect in every round afterward.

Theorem 7.11 (Cooperation under repetition). In the infinitely repeated Prisoner's Dilemma with stage payoffs $T > R > P > S$ (temptation, reward, punishment, sucker), both players using the grim trigger is a subgame perfect equilibrium if and only if

$$\delta \geq \frac{T - R}{T - P}.$$

For the payoffs of [Example 1.8](#), with $T = 5$, $R = 3$, $P = 1$, the threshold is $\delta \geq \frac{1}{2}$: cooperation holds once the players value the next round at least half as much as this one.

Proof. Suppose both play the grim trigger. On the cooperative path each earns R every period, worth

$$R + \delta R + \delta^2 R + \dots = \frac{R}{1 - \delta}.$$

Now weigh a one-period deviation while the opponent still cooperates. The best deviation grabs the temptation payoff T that period, but it springs the trigger, so from the next period on both defect and the deviator collects P forever:

$$T + \delta P + \delta^2 P + \dots = T + \frac{\delta P}{1 - \delta}.$$

Cooperation beats deviation exactly when

$$\frac{R}{1 - \delta} \geq T + \frac{\delta P}{1 - \delta}.$$

Multiplying by $1 - \delta$ gives $R \geq T(1 - \delta) + \delta P$, that is $\delta(T - P) \geq T - R$, or $\delta \geq (T - R)/(T - P)$. No single-period deviation pays once this holds, so by the one-shot deviation principle for discounted games the grim trigger is a best response to itself. Its punishment phase, mutual defection, is a Nash equilibrium of the stage game, so the threat is credible and the profile is subgame perfect. For $T = 5$, $R = 3$, $P = 1$ the bound is $\delta \geq 2/4 = \frac{1}{2}$. ■

Why repetition rescues cooperation

A one-shot Prisoner's Dilemma forces defection. Repeated without end, the same game can carry full cooperation, as long as the players are patient enough that the lasting loss from the punishment outweighs the one-time gain from cheating. The future casts a shadow long enough to discipline the present.

The grim trigger is the bluntest enforcer. Softer punishments, such as the forgiving tit-for-tat that defects once in reply and then resets, also sustain cooperation, and tit-for-tat won the round-robin tournaments of Axelrod [3] that made it famous. The result standing behind all of these strategies is the Folk Theorem.

Theorem 7.12 (Folk Theorem, informal). In an infinitely repeated game, fix any payoff profile that is feasible (an average of stage outcomes) and strictly better for every player than the most they can guarantee alone. For every discount factor δ close enough to 1, some subgame perfect equilibrium has that profile as its long-run average payoff.

Remark 7.13. The name records that the result circulated informally long before anyone published it. Its message cuts both ways. Cooperation becomes *possible* when players are patient, which is the encouraging half, but so does nearly everything else: the Folk Theorem yields a vast multiplicity of equilibria, from full cooperation down to perpetual defection, with no word on which the players will actually reach. Existence is generous and prediction is left wanting, the same tension that runs through Chapter 8.

7.9 A deduction game with subgame states

A last example shows recursion in a game of *deduction*, where the state is not a pile size but a configuration of hidden cards and the recursion runs over those configurations. It also unites the bluffing idea of the poker game in Section 3.8 with the substitution principle of this chapter.

7.9.1 Setup, state, and the conjugate relation

Eleven numbered cards are dealt, five to each player and one face down on the table, the *mystery* card. Players alternate. On your turn you may GUESS the table card, winning 1 if right and 0 if wrong, the game ending either way, or ASK the opponent for a number, which they must lay face up if they hold it. Write $G_{m,n}$ for the value, to the player about to move, of the state in which the mover holds m private cards and the opponent holds n , with one mystery card. Because exactly one player eventually guesses correctly, the mover's value and the value once the turn passes, the *conjugate* state $\bar{G}_{m,n}$, sum to one:

$$V(G_{m,n}) + V(\bar{G}_{m,n}) = 1, \quad \text{so} \quad \bar{G}_{m,n} = 1 - G_{m,n}. \quad (7.1)$$

7.9.2 Boundary values

Two states resolve at once. If the opponent holds no cards, $V(G_{m,0}) = 1$: the mover sees their own m cards, so the single number left unseen must be the table card, and they guess it. If the mover holds no cards, $V(G_{0,n}) = 1/(n+1)$: the table card is equally likely to be any of the $n+1$ numbers the mover cannot see, and a raw guess is right with that probability.

7.9.3 The move matrix

On your turn in $G_{m,n}$ you may ask HONESTLY for a number you lack, BLUFF by asking for one you hold, or GUESS; the opponent may IGNORE the ask or CALL it. An honest ask finds the number in the opponent's hand with probability $\frac{n}{n+1}$, whereupon they lay it down and the turn passes to $\bar{G}_{n-1,m}$, and finds it to be the table card with probability $\frac{1}{n+1}$, an immediate win if the ask goes uncontested. Working through the cases gives the payoff matrix to the mover,

	IGNORE	CALL
HONEST	$\frac{n}{n+1} \bar{G}_{n-1,m} + \frac{1}{n+1}$	$\frac{n}{n+1} \bar{G}_{n-1,m}$
BLUFF	$\bar{G}_{n,m-1}$	1
GUESS	$\frac{1}{n+1}$	$\frac{1}{n+1}$

The GUESS row is weakly dominated: a raw guess wins with probability $\frac{1}{m+n+1} \leq \frac{1}{n+1}$, which an honest ask already secures, so an informed player asks rather than guesses while cards remain.

7.9.4 Solving the smallest state

Take $m = n = 1$. The boundary values give $\bar{G}_{0,1} = 1 - \frac{1}{2} = \frac{1}{2}$ and $\bar{G}_{1,0} = 1 - 1 = 0$, and substituting into the matrix, dropping the dominated guess row, leaves the 2×2 game

$$\begin{bmatrix} \frac{3}{4} & \frac{1}{4} \\ 0 & 1 \end{bmatrix} \quad (\text{rows honest, bluff; columns ignore, call}),$$

where, for instance, the honest-ignore entry is $\frac{1}{2}\overline{G}_{0,1} + \frac{1}{2} = \frac{1}{2} \cdot \frac{1}{2} + \frac{1}{2} = \frac{3}{4}$. There is no saddle, and the formula of Theorem 4.8 with $a = \frac{3}{4}$, $b = \frac{1}{4}$, $c = 0$, $d = 1$, so $D = \frac{3}{4} - \frac{1}{4} - 0 + 1 = \frac{3}{2}$, gives

$$V(G_{1,1}) = \frac{\frac{3}{4} \cdot 1 - \frac{1}{4} \cdot 0}{\frac{3}{2}} = \frac{1}{2}, \quad p^* = \frac{1-0}{\frac{3}{2}} = \frac{2}{3}, \quad q^* = \frac{1-\frac{1}{4}}{\frac{3}{2}} = \frac{1}{2}.$$

The mover asks honestly two-thirds of the time and bluffs one-third; the opponent ignores and calls with equal probability; and the value is $\frac{1}{2}$. The bluff is not a flourish but a full third of optimal play, the same lesson as the poker game of Example 3.17, now embedded in a recursion over deduction states and solved by the substitution principle through (7.1). $\diamond$

Notes and References

Stochastic games are due to Shapley [54], published the same year as his value for cooperative games; Shapley proved existence of the value for the *discounted* two-person zero-sum case. That the *undiscounted* (long-run average) value also exists is a much harder theorem, settled only decades later by Mertens and Neyman [34]. The recursive substitution principle and the worked recursive and stochastic examples of this chapter follow Ferguson [16]; the mutually recursive systems and the coin-toss challenge are from Problem Set 3. Inspection games and their application to treaty verification are surveyed by Avenhaus et al. [2]. The repeated Prisoner's Dilemma and the Folk Theorem of Section 7.8 are the surface of a large theory: the grim-trigger result, that patient players can sustain cooperation by Nash reversion, is Friedman [18]; the sharp subgame-perfect Folk Theorem, that essentially any individually rational payoff is an equilibrium average, is Fudenberg and Maskin [19]; and a textbook treatment is Fudenberg and Tirole [20].

Chapter summary

A recursive game contains an outcome meaning “play again,” whose value is the value V of the whole game. The substitution principle (Procedure 7.1) replaces that outcome by V (or $-V$ under a role switch, or an expected value under a random transition), solves the resulting matrix, and imposes $V = V(A(V))$, a scalar fixed-point equation whose admissible root is chosen by security bounds. Worked examples ranged from a unique value (Example 7.2) through a whole continuum of consistent values (Example 7.3), a cyclically symmetric game solved by the uniform mixture (Example 7.4), a role-switching duel (Example 7.5), and a coupled system whose values obey a harmonic identity (Example 7.6). Shapley's stochastic games (Section 7.6) replace recursion by random transitions and satisfy Bellman fixed-point equations, illustrated by two coin games. Repeated games unwind by the same logic, with the inspection game telescoping to the value $1/N$ (Example 7.9). When the stage game is the Prisoner's Dilemma, an unbounded horizon lets patient players hold cooperation in place by the grim trigger, and the Folk Theorem (Theorem 7.12) catalogues the equilibria that patience makes possible.

Exercises

Recursive value equations

Exercise 7.1. Solve the role-switching recursive game whose payoff matrix to Player I is

	C1	C2
R1	repeat, roles switched	2
R2	3	0

where the top-left outcome returns play to the same game with the two players exchanged. Replacing that cell by $-v$ (a role switch negates the continuation), apply the 2×2 formula to $\begin{bmatrix} -v & 2 \\ 3 & 0 \end{bmatrix}$ to get the fixed-point equation $v = 6/(v+5)$, equivalently $v^2 + 5v - 6 = 0$. State both roots and use the security bound (which root can Player I guarantee?) to select the admissible value.

Exercise 7.2. From Problem Set 3. A game has the block structure $G = \begin{bmatrix} 0 & G_1 \\ G_2 & G_3 \end{bmatrix}$ with $G_1 = \begin{bmatrix} 4 & 3 \\ 1 & 2 \end{bmatrix}$, $G_2 = \begin{bmatrix} 0 & 6 \\ 5 & 1 \end{bmatrix}$, and $G_3 = \begin{bmatrix} 0 & -2 \\ -2 & 0 \end{bmatrix}$. Compute the values $v(G_1)$, $v(G_2)$, $v(G_3)$ from the bottom up, then the value of the top-level game. (Answer: $v(G_1) = 3$ by a saddle, $v(G_2) = 3$ by mixing, $v(G_3) = -1$, and the top level is $\begin{bmatrix} 0 & 3 \\ 3 & -1 \end{bmatrix}$ with value $\frac{9}{7}$.)

Exercise 7.3. From Problem Set 3. Three mutually recursive games, $G_1 = \begin{bmatrix} G_2 & 1 \\ 1 & 0 \end{bmatrix}$, $G_2 = \begin{bmatrix} G_3 & 0 \\ 0 & 2 \end{bmatrix}$, $G_3 = \begin{bmatrix} G_1 & 1 \\ 1 & 0 \end{bmatrix}$, lead to a single cubic-by-substitution. Show the elimination reduces to $g_3^2 - 6g_3 + 4 = 0$, so $g_3 = 3 - \sqrt{5}$, and find $(v_1, v_2, v_3) = \left(\frac{5-\sqrt{5}}{4}, \frac{5-\sqrt{5}}{5}, 3 - \sqrt{5}\right) \approx (0.691, 0.553, 0.764)$.

Stochastic, repeated, and continuum games

Exercise 7.4. *Coin-toss challenge (Problem Set 3).* A player flips a coin with heads-probability p ; after any run of $k \geq 1$ heads he may challenge the opponent to match k heads, or continue. (a) For a fair coin $p = \frac{1}{2}$, show the symmetric value is $W = \frac{1}{2}$. (b) For general p , a threshold policy gives value $V(m) = (1 - p^{2m})/(2 - p^m)$; optimizing over the threshold via $x = p^m$ and $f(x) = (1 - x^2)/(2 - x)$, show $f'(x) = 0$ at $x_* = 2 - \sqrt{3}$ and that the limit as $p \rightarrow 1$ is $4 - 2\sqrt{3} \approx 0.536$.

Exercise 7.5. Verify the inspection-game recursion $M_n = \begin{bmatrix} 1 & 0 \\ 0 & G_{n-1} \end{bmatrix}$ with $G_1 = 1$: apply the 2×2 formula to get $G_n = G_{n-1}/(1 + G_{n-1})$, compute $G_2 = \frac{1}{2}$ and $G_3 = \frac{1}{3}$, and prove by induction that $G_n = 1/n$. Confirm that the optimal inspection probability at n slots left is $1/n$.

Exercise 7.6. Explain, with reference to Example 7.3, why a recursive game can have a whole interval of self-consistent values while an ordinary finite matrix game has a unique value. What property of the recursive entry is responsible?

PART III

**General-Sum Games and Nash
Equilibrium**

CHAPTER 8

General-Sum Games and Nash Equilibrium

“I had this idea about equilibrium. It was an obvious idea, but obvious ideas are sometimes the most important.”

John Nash

The zero-sum theory is complete and beautiful, but most real interactions are not zero-sum. Two firms setting prices, two nations negotiating, two drivers at an intersection: their interests are partly opposed and partly aligned. Once the payoffs no longer cancel, the single value of Chapter 4 disappears, and we need a solution concept that survives. That concept is the Nash equilibrium, which we met informally in Chapter 1 and now develop in full for any number of players. The chapter culminates in Nash’s theorem: every finite game has at least one equilibrium in mixed strategies. Its proof, by a fixed point, is the last of the four theorems promised in Section 1.4.

8.1 From zero-sum to general-sum

A two-person general-sum game is given by two payoff matrices.

Definition 8.1 (General-sum bimatrix game). A two-person *general-sum game* is a pair of $m \times n$ matrices (A, B) . When Player I plays row i and Player II plays column j , Player I receives a_{ij} and Player II receives b_{ij} . The game is displayed by writing the pair (a_{ij}, b_{ij}) in each cell. The zero-sum case is $B = -A$.

When $B \neq -A$ the symmetry that gave zero-sum games their single value is gone. The minimax theorem (Theorem 4.7) no longer says anything joint, because the two players are not solving mirror-image problems. What survives for each player individually is the worst-case guarantee.

8.2 Safety levels

Each player can still ask: what can I guarantee myself, whatever the opponent does? The answer is the player’s safety level, computed exactly as a zero-sum value on that player’s own matrix.

Definition 8.2 (Safety level). In a general-sum game (A, B) , the **safety level** of Player I is

$$V_I = \max_{p \in \Delta_m} \min_{1 \leq j \leq n} \sum_{i=1}^m p_i a_{ij} = V(A),$$

the value of Player I’s matrix A played as a zero-sum game. The safety level of Player II is

$$V_{II} = \max_{q \in \Delta_n} \min_{1 \leq i \leq m} \sum_{j=1}^n b_{ij} q_j = V(B^\top),$$

the value of $B^\top$. Each safety level is the payoff a player can secure no matter what the other does.

The safety level treats the opponent as a pure adversary, which is pessimistic in a game where the opponent has their own goals. It is nonetheless the right notion of a *guarantee*, and in Chapter 9 it becomes the disagreement point from which bargaining starts.

Example 8.3 (Computing safety levels). Consider

$$(A, B) = \begin{pmatrix} (2, 0) & (1, 3) \\ (0, 1) & (3, 2) \end{pmatrix}, \quad A = \begin{bmatrix} 2 & 1 \\ 0 & 3 \end{bmatrix}, \quad B = \begin{bmatrix} 0 & 3 \\ 1 & 2 \end{bmatrix}.$$

For Player I, the matrix A has maximin 1 and minimax 2, so there is no saddle, and the 2×2 formula gives $V_I = V(A) = (2 \cdot 3 - 1 \cdot 0) / ((2 - 1) + (3 - 0)) = 6/4 = \frac{3}{2}$, achieved by the equalizing mix $\mathbf{p} = (\frac{3}{4}, \frac{1}{4})$. For Player II, transpose B to $B^T = \begin{bmatrix} 0 & 1 \\ 3 & 2 \end{bmatrix}$; here the second row dominates the first, so Player II's safety strategy is pure (column 2 of the original game) and $V_{II} = 2$. One player needs to randomize to secure a guarantee while the other does not, which is perfectly consistent in a general-sum game.

8.3 N-person games and Nash equilibrium

Nash's framework handles any number of players at once.

Definition 8.4 (Finite n -person game). A finite n -person game has players $1, \dots, n$, each with a finite pure-strategy set $S_i = \{1, \dots, m_i\}$, and for each player a payoff function u_i on the set of pure profiles $S_1 \times \dots \times S_n$. A mixed strategy for player i is a distribution $\mathbf{p}_i \in \Delta_{m_i}$, and the expected payoff under a mixed profile $(\mathbf{p}_1, \dots, \mathbf{p}_n)$ is the multilinear extension

$$u_i(\mathbf{p}_1, \dots, \mathbf{p}_n) = \sum_{a_1 \in S_1} \dots \sum_{a_n \in S_n} \left(\prod_{k=1}^n p_k(a_k) \right) u_i(a_1, \dots, a_n),$$

the expectation of the pure payoff under independent randomization.

Because the expected payoff is a sum of products of probabilities, it is continuous in the whole profile and linear in each player's own mixed strategy separately, two facts the existence proof will use.

Definition 8.5 (Nash equilibrium). A mixed profile $\mathbf{p}^* = (\mathbf{p}_1^*, \dots, \mathbf{p}_n^*)$ is a **Nash equilibrium**, or **strategic equilibrium**, if no player can improve by changing strategy alone: for every player i and every $\mathbf{p}_i \in \Delta_{m_i}$,

$$u_i(\mathbf{p}_i^*, \mathbf{p}_{-i}^*) \geq u_i(\mathbf{p}_i, \mathbf{p}_{-i}^*),$$

where $\mathbf{p}_{-i}^*$ denotes the equilibrium strategies of the other players. Equivalently, each $\mathbf{p}_i^*$ is a best response to $\mathbf{p}_{-i}^*$, so the profile is a fixed point of the combined best-response map.

The definition extends the two-player notion of Definition 1.10 to n players and mixed strategies. Because expected payoff is linear in a player's own strategy, a player is indifferent among all pure strategies in the support of a best response, and each earns the same payoff. This is the principle of indifference again, and it is how mixed equilibria are computed by hand.

The film *A Beautiful Mind* dramatizes Nash inventing equilibrium in a bar: several men face one blonde and several brunettes, and the Nash character proposes that everyone ignore the blonde and pair off with brunettes, calling it the best outcome “for the group.” This is a good illustration of what a Nash equilibrium is *not*. A best-for-the-group plan is irrelevant to the definition, which asks only about unilateral deviations; and the proposed profile is not even an equilibrium, because if all the other men go for brunettes, one man's best response is to deviate to the unattended blonde. Equilibrium is the fixed point of self-interested best response rather than a socially optimal arrangement; conflating the two is the most common misreading of the concept.

8.4 Canonical two-person games

Before the examples, here is the fastest hand method for the *pure* equilibria of any bimatrix.

Finding pure Nash equilibria by marking best responses

1. Go down each *column* and underline Player I's largest first-coordinate payoff: that is Player I's best response to that column.
2. Go across each *row* and underline Player II's largest second-coordinate payoff: that is Player II's best response to that row.
3. Any cell underlined *twice* is a pure Nash equilibrium: each player is best-responding to the other, so neither can profit by deviating.

Example 8.6 (Finding pure equilibria by the marking method). Player I picks a row, Player II a column, and each cell shows (u_1, u_2) :

	<i>L</i>	<i>C</i>	<i>R</i>
<i>T</i>	(3, 2)	(1, 1)	(0, <u>4</u>)
<i>M</i>	(2, 0)	(<u>4</u> , <u>3</u>)	(3, 2)
<i>B</i>	(1, <u>4</u>)	(2, 1)	(<u>5</u> , 3)

Player I's best replies (underline the largest first coordinate in each column). Column *L*: $\max(3, 2, 1) = 3$ at *T*. Column *C*: $\max(1, 4, 2) = 4$ at *M*. Column *R*: $\max(0, 3, 5) = 5$ at *B*. Player II's best replies (underline the largest second coordinate in each row). Row *T*: $\max(2, 1, 4) = 4$ at *R*. Row *M*: $\max(0, 3, 2) = 3$ at *C*. Row *B*: $\max(4, 1, 3) = 4$ at *L*. The only cell carrying *both* underlines is (M, C) , where Player I's 4 meets Player II's 3. So (M, C) with payoff $(4, 3)$ is the unique pure Nash equilibrium: given *C*, Player I's best row is *M*; given *M*, Player II's best column is *C*. The other once-marked cells, such as (B, R) where only Player I is happy, fail because the unmarked player would jump away. $\diamond$

Example 8.7 (A coordination game with a mixed equilibrium). Consider

$$(A, B) = \begin{pmatrix} (3, 3) & (0, 2) \\ (2, 1) & (5, 5) \end{pmatrix}.$$

Two pure equilibria sit on the diagonal. At $(3, 3)$, Player I deviating to row 2 earns $2 < 3$ and Player II deviating to column 2 earns $2 < 3$, so neither moves; at $(5, 5)$, deviations earn $0 < 5$ and $1 < 5$, so it too is an equilibrium. There is also a mixed equilibrium, found by the mutual indifference conditions. To make Player II indifferent, choose Player I's mix $\mathbf{p} = (p, 1 - p)$ so that Player II's two columns pay equally: $3p + 1(1 - p) = 2p + 5(1 - p)$, that is $1 + 2p = 5 - 3p$, giving $p = \frac{4}{5}$. To make Player I indifferent, choose $\mathbf{q} = (q, 1 - q)$ with $3q = 2q + 5(1 - q)$, that is $3q = 5 - 3q$, giving $q = \frac{5}{6}$. The mixed equilibrium is

$$\mathbf{p}^* = \left(\frac{4}{5}, \frac{1}{5}\right), \quad \mathbf{q}^* = \left(\frac{5}{6}, \frac{1}{6}\right),$$

with payoffs $\left(\frac{5}{2}, \frac{13}{5}\right)$, which are exactly the two safety levels. The mixed equilibrium is the worst of the three for both players, while $(5, 5)$ is payoff-dominant. This is the central difficulty general-sum games add: equilibria need not be unique, and the players may disagree about which one to aim for.

Example 8.8 (Battle of the Sexes). Two players prefer to coordinate but on different outcomes:

$$(A, B) = \begin{pmatrix} (2, 1) & (0, 0) \\ (0, 0) & (1, 2) \end{pmatrix},$$

where the rows and columns are two options *F* and *R*. The pure equilibria are (F, F) , which Player I prefers with payoff $(2, 1)$, and (R, R) , which Player II prefers with payoff $(1, 2)$. The mixed equilibrium makes each player indifferent: Player I plays *F* with probability $\frac{2}{3}$ and Player II plays *F* with probability $\frac{1}{3}$,

$$\mathbf{p}^* = \left(\frac{2}{3}, \frac{1}{3}\right), \quad \mathbf{q}^* = \left(\frac{1}{3}, \frac{2}{3}\right),$$

with payoffs $\left(\frac{2}{3}, \frac{2}{3}\right)$, again the safety levels and again worse than either pure equilibrium. The Battle of the Sexes is the standard model of coordination under conflicting preferences, and selecting among its equilibria is a bargaining question deferred to Chapter 9.

The three equilibria are easiest to see by drawing each player's *best-response correspondence* and looking for where they meet (Figure 8.1). Player I prefers *F* exactly when $2q > 1 - q$, that is $q > \frac{1}{3}$; Player II prefers *F* exactly when $p > 2(1 - p)$, that is $p > \frac{2}{3}$. Each best response is therefore a staircase that jumps from one pure choice to the other at the opponent's indifference level, and a Nash equilibrium is any point lying on *both* staircases at once.




Figure 8.1. Best-response correspondences for the Battle of the Sexes. The blue staircase is Player I's best p as a function of Player II's q (jump up at $q = \frac{1}{3}$); the red staircase is Player II's best q as a function of p (jump up at $p = \frac{2}{3}$). They cross at three points, the three Nash equilibria: the pure (R, R) and (F, F) at the corners and the mixed $(\frac{2}{3}, \frac{1}{3})$ in the interior.

Example 8.9 (The Prisoner's Dilemma as a Nash equilibrium). The Prisoner's Dilemma of Example 1.8 has, for each player, a strictly dominant strategy to defect, and a dominant strategy is a best response to everything. The profile (DEFECT, DEFECT) is therefore the unique Nash equilibrium, even though both players strictly prefer mutual cooperation. The equilibrium concept does not promise efficiency; it promises only stability against unilateral deviation, and here the stable outcome is the bad one.

Example 8.10 (The game of Chicken). Two drivers speed toward each other on a one-lane road. Each can SWERVE or drive STRAIGHT. Swerving forfeits the road but is safe; driving straight wins the road if the other yields and ends in a wreck if neither does. Scoring a wreck at -10 and the loss of face from swerving alone at -1 ,

	SWERVE	STRAIGHT
(A, B) = SWERVE	(0, 0)	(<u>-1</u> , <u>1</u>)
STRAIGHT	(<u>1</u> , <u>-1</u>)	(-10, -10)

Pure equilibria by the marking method. Down each column, underline Player I's larger payoff. Column SWERVE: $\max(0, 1) = 1$ at STRAIGHT; column STRAIGHT: $\max(-1, -10) = -1$ at SWERVE. Across each row, underline Player II's larger payoff. Row SWERVE: $\max(0, 1) = 1$ at STRAIGHT; row STRAIGHT: $\max(-1, -10) = -1$ at SWERVE. The two off-diagonal cells come out doubly underlined, so the pure equilibria are (STRAIGHT, SWERVE) and (SWERVE, STRAIGHT): the outcomes in which exactly one driver yields. Each player would rather be the one who drives straight, and the game offers no way to settle who.

The mixed equilibrium. Let each driver go STRAIGHT with probability p , the same p for both by symmetry. Against an opponent who drives straight with probability p , the two pure choices pay

$$u(\text{SWERVE}) = p(-1) + (1-p)0 = -p, \quad u(\text{STRAIGHT}) = p(-10) + (1-p)1 = 1 - 11p.$$

Setting them equal, $-p = 1 - 11p$, gives $10p = 1$, so $p = \frac{1}{10}$: each driver goes straight one time in ten. The symmetric mixed equilibrium pays $-p = -\frac{1}{10}$ to each, and a wreck happens with probability $p^2 = \frac{1}{100}$. $\diamond$

Remark 8.11. Chicken is the anti-coordination twin of the Battle of the Sexes. Both games have two pure equilibria and one mixed, but the Battle of the Sexes rewards players for matching while Chicken rewards them for mismatching. The mixed equilibrium is the cost of symmetry: with no way to agree on who yields, both gamble, and the gamble carries a one-percent wreck. The game is famous for the value of visible commitment. A driver who conspicuously rips out the steering wheel can no longer swerve, and the opponent, seeing it, yields; the reckless-looking move secures the pure equilibrium its maker prefers.

The four social games above, together with the strictly competitive even-odd game (Matching Pennies) of Chapter 4 and the pure coordination game of Exercise 1.4, exhaust the qualitatively different shapes a two-by-two game can

take. Table 8.1 lines them up by the feature that governs prediction: how many equilibria there are, and whether the players are trying to match, to mismatch, or to outguess.

Table 8.1. The canonical two-by-two games, sorted by the shape of their equilibria. Every finite game has at least one equilibrium (Theorem 8.15); what changes from game to game is how many there are and what the players are trying to do to one another.

Game	Pure equilibria	Mixed	The players try to
Prisoner's Dilemma	one (dominant)	none	defect no matter what
Stag Hunt	two (the diagonal)	one	coordinate, but safely
Battle of the Sexes	two (the matches)	one	coordinate, yet differ on which
Chicken	two (the mismatches)	one	mismatch
Matching Pennies	none	one	outguess each other
Pure coordination	two (the diagonal)	one	simply agree

8.5 Iterated elimination of dominated strategies

Before hunting for equilibria, it pays to throw away strategies no rational player would ever use, and then to throw away the ones that become useless *once the first batch is gone*. This is **iterated elimination of strictly dominated strategies** (IEDS). A pure strategy is *strictly dominated* if some other strategy of the same player earns strictly more against every choice of the opponents.

Intuition: cross out the no-brainers, then look again

IEDS is how you reason in a haggle or a chess opening: first rule out the moves that are bad no matter what the other side does, then notice that some of *their* moves are now pointless, rule those out, and keep going. Each deletion can expose the next. What survives is what is left after everyone has applied “I would never do that, and I know you would never do that, and you know that I know...” as far as it goes.

Example 8.12 (IEDS solves a 3×3 game). Player I chooses a row $\{T, M, B\}$, Player II a column $\{L, C, R\}$, and each cell lists (u_1, u_2) :

	L	C	R
T	(4, 3)	(5, 1)	(6, 2)
M	(2, 1)	(8, 4)	(3, 6)
B	(3, 0)	(9, 6)	(2, 8)

Round 1 (Player II drops a column). Player II reads only the *second* coordinates. Her payoffs are $L = (3, 1, 0)$, $C = (1, 4, 6)$, $R = (2, 6, 8)$ down the three rows. Compare R against C row by row: $2 > 1$, $6 > 4$, $8 > 6$. Column R pays her strictly more than column C in *every* row, so C is strictly dominated and she will never play it:

	L	C	R			L	R
T	(4, 3)	(5, 1)	(6, 2)	$\xrightarrow{\text{delete } C}$	T	(4, 3)	(6, 2)
M	(2, 1)	(8, 4)	(3, 6)		M	(2, 1)	(3, 6)
B	(3, 0)	(9, 6)	(2, 8)		B	(3, 0)	(2, 8)

Round 2 (Player I drops two rows). With C gone, Player I reads the *first* coordinates: $T = (4, 6)$, $M = (2, 3)$, $B = (3, 2)$ across the surviving columns L, R . Now T beats M ($4 > 2$, $6 > 3$) and T beats B ($4 > 3$, $6 > 2$), so both M and B are strictly dominated by T , which only became clear after C left:

	L	R			L	R
T	(4, 3)	(6, 2)	$\xrightarrow{\text{delete } M, B}$	T	(4, 3)	(6, 2)
M	(2, 1)	(3, 6)				
B	(3, 0)	(2, 8)				

Round 3 (Player II chooses). Only row T remains, where Player II compares L (pays her 3) against R (pays her 2) and takes L . IEDS has collapsed the game to the single cell $(T, L) = (4, 3)$.

It is the Nash equilibrium. Check directly: if Player II plays L , Player I's column is $(4, 2, 3)$ and T is the best reply; if Player I plays T , Player II's row pays $(3, 1, 2)$ and L is the best reply. Neither can profitably deviate, so (T, L) is

a Nash equilibrium, and because IEDS left nothing else, it is the *only* one. A strategy that survives iterated strict dominance can fail to be unique in general, but when IEDS ends at a single cell, that cell is always the unique Nash equilibrium. $\diamond$

8.6 The price of anarchy

The Prisoner's Dilemma shows that equilibrium can be inefficient for two players. The same effect appears at scale, and it can be measured. The **price of anarchy** is the ratio of the social cost at the worst Nash equilibrium to the social cost of the best coordinated outcome; it quantifies how much selfishness costs. The term was coined by Papadimitriou [42], sharpening the worst-case-equilibrium ratio of Koutsoupias and Papadimitriou [28].

Example 8.13 (Selfish routing: Pigou's network). One unit of traffic must travel from s to t along one of two parallel roads (Figure 8.2). The top road always takes 1 hour, however many use it. The bottom road takes x hours when a fraction x of the traffic uses it: empty it is instant, full it takes an hour.




Figure 8.2. Pigou's two-road network of Example 8.13. The bottom road's travel time grows with its own congestion x ; the top road is constant. Selfish routing overloads the bottom road.

The social optimum. Send a fraction x of the traffic along the bottom road and $1 - x$ along the top. The total travel time, summed over all drivers, is

$$C(x) = \underbrace{(1-x) \cdot 1}_{\text{top}} + \underbrace{x \cdot x}_{\text{bottom}} = 1 - x + x^2.$$

Minimizing, $C'(x) = -1 + 2x = 0$ gives $x = \frac{1}{2}$, and the least possible total time is $C(\frac{1}{2}) = 1 - \frac{1}{2} + \frac{1}{4} = \frac{3}{4}$.

The selfish equilibrium. Now let each driver pick the faster road. While the bottom road carries a fraction $x < 1$, its time x is below the top road's 1, so every driver prefers it and traffic keeps shifting down until $x = 1$, where the two roads are tied at time 1 and no one wants to switch. That is the unique Nash equilibrium, and its total time is

$$C(1) = 0 \cdot 1 + 1 \cdot 1 = 1.$$

The price of anarchy. The ratio is

$$\text{PoA} = \frac{\text{selfish cost}}{\text{optimal cost}} = \frac{1}{3/4} = \frac{4}{3}.$$

Selfishness makes the commute one-third longer than necessary. This $\frac{4}{3}$ is not special to this network: Roughgarden and Tardos [48] proved that for *any* road network with linear congestion costs the price of anarchy is at most $\frac{4}{3}$, and Pigou's example attains it. $\diamond$

Adding a road can make everyone slower. In *Braess's paradox* [8], opening a fast shortcut in a congested network shifts every driver's best response in a way that raises all travel times at equilibrium; removing the shortcut restores the faster outcome. The effect is not hypothetical, it has been reported when streets were closed in Manhattan (42nd Street, 1990) and Seoul (the Cheonggyecheon expressway removal, 2003–2005). More roads expand the strategy set, and a larger strategy set can move the equilibrium to a worse place: equilibrium tracks individual incentives rather than collective good.

8.7 Cournot competition

Nash equilibrium is the workhorse of economic modelling, and its first and most famous application predates Nash by over a century: Cournot's 1838 account of two firms choosing how much to produce. It is a general-sum game with a continuum of strategies, solved not by a payoff grid but by each player's *reaction curve*, the continuous version of a best response.

Two firms produce quantities $q_1, q_2 \geq 0$ of an identical good. The market price falls with total output, $P = a - (q_1 + q_2)$, and each unit costs c to make, with $a > c$. Writing $A = a - c$ for the gap between the choke price and cost, firm i earns the profit

$$\pi_i(q_1, q_2) = q_i (P - c) = q_i (A - q_i - q_j).$$

Holding the rival's output q_j fixed, firm i faces a downward parabola in q_i ; the first-order condition $A - 2q_i - q_j = 0$ gives its **reaction curve**

$$q_i = \frac{A - q_j}{2}.$$

A Nash equilibrium is a pair sitting on *both* reaction curves at once. By symmetry put $q_1 = q_2 = q^*$, so $q^* = (A - q^*)/2$, giving $q^* = A/3$. Each firm makes $A/3$, total output is $2A/3$, the price settles at $a - 2A/3 = (a + 2c)/3$, and each earns $\pi^* = (A/3)(A/3) = A^2/9$.




Figure 8.3. Cournot reaction curves for $A = 12$ (here $a = 12$, $c = 0$). Firm 1's reaction $q_1 = (12 - q_2)/2$ is the steep navy line R_1 ; firm 2's $q_2 = (12 - q_1)/2$ is the shallow red line R_2 . They cross at the Cournot–Nash equilibrium $(4, 4)$, where each firm produces 4 and earns 16. The symmetric collusion point $(3, 3)$, which would pay each firm 18, lies off both reaction curves: better for the pair, but unstable, since from there each firm is tempted to expand toward its own curve.

A numerical case makes the forces visible. Take $a = 12$ and $c = 0$, so $A = 12$ and the reaction curves are $q_1 = (12 - q_2)/2$ and $q_2 = (12 - q_1)/2$, drawn in Figure 8.3. They cross at $q^* = 4$: each firm makes 4, the price is $12 - 8 = 4$, and each profits 16. Two benchmarks frame the result. Were the firms to *collude* for joint profit, they would together make the monopoly quantity $A/2 = 6$, that is 3 apiece, sell at price 6, and earn 18 each, strictly more than the 16 of competition. Were they instead to compete until price met cost, total output would reach $A = 12$ and profit would vanish. Cournot lands between the two, and the comparison exposes the familiar trap: the collusive split $(3, 3)$ pays more, but each firm, taking the other's 3 as given, reads off its reaction curve $q_i = (12 - 3)/2 = 4.5$ and wants to produce more. Only at $(4, 4)$ does neither want to move. The quantity-setting duopoly is a Prisoner's Dilemma with a continuum of strategies, where individually rational expansion pushes both firms below the profit they could jointly hold.

Letting the firms move in sequence rather than at once tilts the outcome toward whoever commits first. That sequential variant is Stackelberg competition, solved by backward induction in Section 3.5: the leader, knowing the follower will accommodate, produces 6 to the follower's 3 and earns 18 against 9. The simultaneous Cournot crossing at $(4, 4)$ and the committed Stackelberg point are the two faces of the same duopoly, one symmetric and one sequential.

8.8 Brouwer's fixed-point theorem

The existence of equilibrium rests on a topological fact about continuous maps.

Theorem 8.14 (Brouwer, 1911). Let $D \subseteq \mathbb{R}^k$ be nonempty, compact, and convex, and let $f : D \rightarrow D$ be continuous. Then f has a fixed point: there is a point $x^* \in D$ with $f(x^*) = x^*$.

Discussion. A full proof belongs to topology; see Nash [36] for the role it plays here, and any topology text for a proof. The one-dimensional case is the intuition: a continuous function $f : [0, 1] \rightarrow [0, 1]$ must cross the diagonal, because $g(x) = f(x) - x$ satisfies $g(0) \geq 0$ and $g(1) \leq 0$, so by the intermediate value theorem g vanishes somewhere, giving $f(x^*) = x^*$. In two dimensions the picture is a disk that is continuously stirred back into itself: some point returns exactly to where it started. A continuous map on a compact convex set cannot move every point, much as a “You Are Here” map laid flat on the wall it depicts has exactly one dot lying directly over the real location it marks. ■

The conditions matter. On the open interval $(0, 1)$ the continuous map $f(x) = x/2 + 1/2$ has its only fixed point at 1, which is not in the set, so compactness is essential. Convexity rules out holes, around which a map could rotate every point.

8.9 Nash's existence theorem

Nash's theorem is the assertion that mixed strategies are always enough for an equilibrium.

Theorem 8.15 (Nash, 1951). Every finite n -person game has at least one Nash equilibrium in mixed strategies.

Nash's existence theorem

Every finite game, with any number of players and any payoffs, has at least one Nash equilibrium once mixed strategies are allowed. The proof recasts the equilibrium condition as a fixed point of a continuous self-map of the strategy space and invokes Brouwer's theorem. Existence is guaranteed; uniqueness, efficiency, and ease of computation are not.

Proof. The strategy space is the product of simplices $X = \Delta_{m_1} \times \cdots \times \Delta_{m_n}$. Each simplex is nonempty, compact, and convex, and so is their product, which sits in $\mathbb{R}^{m_1 + \cdots + m_n}$. We construct a continuous map $F : X \rightarrow X$ whose fixed points are exactly the Nash equilibria, then apply Brouwer.

The deviation gains. Fix a profile $\mathbf{x} = (x_1, \dots, x_n) \in X$. For each player i and each pure strategy $a \in S_i$, define the **gain** from shifting toward a ,

$$\varphi_{i,a}(\mathbf{x}) = \max \left\{ 0, u_i(a, \mathbf{x}_{-i}) - u_i(\mathbf{x}) \right\},$$

the amount by which playing the pure strategy a instead of the current mix $\mathbf{x}_i$ would raise player i 's payoff, counted as zero if the deviation does not help. Because u_i is multilinear, hence continuous, and the maximum with 0 is continuous, each $\varphi_{i,a}$ is a continuous function of $\mathbf{x}$.

The map. Define $F(\mathbf{x}) = \mathbf{y}$ component by component: for each player i and each pure strategy $a \in S_i$,

$$y_i(a) = \frac{x_i(a) + \varphi_{i,a}(\mathbf{x})}{1 + \sum_{b \in S_i} \varphi_{i,b}(\mathbf{x})}. \quad (8.1)$$

The numerator is nonnegative and the denominator is positive, and summing the numerator over $a \in S_i$ gives $1 + \sum_b \varphi_{i,b}(\mathbf{x})$, the denominator, so $\sum_a y_i(a) = 1$. Thus $\mathbf{y}_i \in \Delta_{m_i}$ and F maps X into X . Each $y_i(a)$ is a ratio of continuous functions with nonvanishing denominator, so F is continuous.

Brouwer. By Theorem 8.14, F has a fixed point $\mathbf{x}^* = F(\mathbf{x}^*)$.

Fixed points are equilibria. At the fixed point, (8.1) reads, for every i and a ,

$$x_i^*(a) \left(1 + \sum_b \varphi_{i,b}(\mathbf{x}^*) \right) = x_i^*(a) + \varphi_{i,a}(\mathbf{x}^*),$$

which simplifies to

$$x_i^*(a) \sum_{b \in S_i} \varphi_{i,b}(\mathbf{x}^*) = \varphi_{i,a}(\mathbf{x}^*). \quad (8.2)$$

Now fix a player i . Player i 's equilibrium payoff $u_i(\mathbf{x}^*)$ is the average, weighted by x_i^* , of the pure-strategy payoffs $u_i(a, \mathbf{x}_{-i}^*)$ over a in the support of $\mathbf{x}_i^*$. An average cannot be strictly below all of its terms, so at least one pure strategy a_0 with $x_i^*(a_0) > 0$ earns at most the average, $u_i(a_0, \mathbf{x}_{-i}^*) \leq u_i(\mathbf{x}^*)$. For that a_0 the gain is $\varphi_{i,a_0}(\mathbf{x}^*) = \max\{0, \leq 0\} = 0$. Apply (8.2) to a_0 : since $x_i^*(a_0) > 0$ and the right side is 0,

$$\sum_{b \in S_i} \varphi_{i,b}(\mathbf{x}^*) = 0.$$

Every term in this sum is nonnegative, so every gain vanishes: $\varphi_{i,b}(\mathbf{x}^*) = 0$ for all $b \in S_i$. By the definition of the gain, $u_i(b, \mathbf{x}_{-i}^*) \leq u_i(\mathbf{x}^*)$ for every pure strategy b . Since the best response over mixed strategies is attained at a pure strategy (the payoff is linear in $\mathbf{x}_i$), no mixed deviation beats $\mathbf{x}_i^*$ either. This holds for every player i , so $\mathbf{x}^*$ is a Nash equilibrium. ■

Example 8.16 (Nash's map in action). To make the construction concrete, evaluate the map (8.1) on the Battle of the Sexes (Example 8.8), $A = \begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix}$, $B = \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix}$, at the non-equilibrium profile where both players mix uniformly, $\mathbf{x} = ((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$.

Current payoffs. Averaging over the four cells with weight $\frac{1}{4}$ each,

$$u_I(\mathbf{x}) = \frac{1}{4}(2 + 0 + 0 + 1) = \frac{3}{4}, \quad u_{II}(\mathbf{x}) = \frac{1}{4}(1 + 0 + 0 + 2) = \frac{3}{4}.$$

Player I's gains. Against Player II's $(\frac{1}{2}, \frac{1}{2})$, pure row 1 earns $\frac{1}{2} \cdot 2 + \frac{1}{2} \cdot 0 = 1$ and pure row 2 earns $\frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 1 = \frac{1}{2}$. The gains over the current $\frac{3}{4}$ are

$$\varphi_{I,1} = \max\{0, 1 - \frac{3}{4}\} = \frac{1}{4}, \quad \varphi_{I,2} = \max\{0, \frac{1}{2} - \frac{3}{4}\} = 0.$$

Player I's update. With $\sum_b \varphi_{I,b} = \frac{1}{4}$, the map gives

$$y_I(1) = \frac{\frac{1}{2} + \frac{1}{4}}{1 + \frac{1}{4}} = \frac{3/4}{5/4} = \frac{3}{5}, \quad y_I(2) = \frac{\frac{1}{2} + 0}{1 + \frac{1}{4}} = \frac{1/2}{5/4} = \frac{2}{5}.$$

Player II's gains and update. Symmetrically, pure column 1 earns $\frac{1}{2}$ and column 2 earns 1, so $\varphi_{II,1} = 0$, $\varphi_{II,2} = \frac{1}{4}$, and

$$y_{II}(1) = \frac{\frac{1}{2} + 0}{5/4} = \frac{2}{5}, \quad y_{II}(2) = \frac{\frac{1}{2} + \frac{1}{4}}{5/4} = \frac{3}{5}.$$

So $F(\mathbf{x}) = ((\frac{3}{5}, \frac{2}{5}), (\frac{2}{5}, \frac{3}{5})) \neq \mathbf{x}$: the map shifts each player toward their profitable pure strategy (row 1 for Player I, column 2 for Player II), so $\mathbf{x}$ is not a fixed point and hence not an equilibrium. Iterating F drives the profile toward a fixed point, where all gains vanish; the mixed equilibrium $((\frac{2}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{2}{3}))$ of Example 8.8 is one such point, at which every φ is zero and F holds the profile still. ◇

Remark 8.17. The construction is due to Nash. The map (8.1) nudges each player's mix toward the pure strategies that currently look profitable, and a fixed point is a profile where nothing looks profitable to shift toward, which is exactly equilibrium. Nash's original 1950 note used Kakutani's fixed-point theorem for the set-valued best-response map; the single-valued construction above, often credited to Gale, needs only Brouwer.

8.9.1 A second proof via the best-response correspondence

The map (8.1) is one of two classical roads to Nash's theorem. The other, the route Nash took in his 1950 announcement, works with the best-response *correspondence* directly. The set-valued road repays the detour: it exposes the exact obstacle that the gain-map was built to dodge.

For player i facing a fixed profile $\mathbf{x}_{-i}$ of the others, the set of best replies is

$$\text{BR}_i(\mathbf{x}_{-i}) = \arg \max_{\mathbf{s}_i \in \Delta_{m_i}} u_i(\mathbf{s}_i, \mathbf{x}_{-i}) \subseteq \Delta_{m_i}.$$

Collecting these across players defines a correspondence on the whole strategy space $X = \Delta_{m_1} \times \cdots \times \Delta_{m_n}$,

$$\text{BR}(\mathbf{x}) = \text{BR}_1(\mathbf{x}_{-1}) \times \cdots \times \text{BR}_n(\mathbf{x}_{-n}) \subseteq X.$$

By the definition of equilibrium, a profile $\mathbf{x}^*$ is a Nash equilibrium exactly when every player already best-responds, that is, $\mathbf{x}_i^* \in \text{BR}_i(\mathbf{x}_{-i}^*)$ for each i , which is the single condition $\mathbf{x}^* \in \text{BR}(\mathbf{x}^*)$. Equilibria are the fixed points of BR, the profiles the correspondence carries back onto themselves. The staircase plot of Figure 8.1 shows the picture for a 2×2 game: two step-shaped correspondences whose crossings are the equilibria.

Second proof of Theorem 8.15, via Kakutani's theorem. The obstacle is that BR is not a function. Where a player is indifferent among several pure replies, BR_i returns a whole face of the simplex rather than a single point, so the step curves of Figure 8.1 have flat treads. Brouwer's theorem speaks only of single-valued maps and does not apply to BR. Kakutani's theorem (Theorem A.7) is the set-valued replacement: it grants a fixed point to any correspondence on a nonempty compact convex set that is nonempty-valued, convex-valued, and closed-graphed. We verify the three conditions for BR.

Nonempty values. For fixed $\mathbf{x}_{-i}$ the payoff $u_i(\cdot, \mathbf{x}_{-i})$ is linear, hence continuous, on the compact simplex Δ_{m_i} . A continuous function on a compact set attains its maximum, so $\text{BR}_i(\mathbf{x}_{-i})$ is never empty, and neither is the product $\text{BR}(\mathbf{x})$.

Convex values. Since u_i is linear in $\mathbf{s}_i$, its maximizers form a face of the simplex. If $\mathbf{s}$ and $\mathbf{s}'$ both attain the maximum $M = u_i(\mathbf{s}, \mathbf{x}_{-i}) = u_i(\mathbf{s}', \mathbf{x}_{-i})$, then for any $\lambda \in [0, 1]$,

$$u_i(\lambda \mathbf{s} + (1 - \lambda) \mathbf{s}', \mathbf{x}_{-i}) = \lambda u_i(\mathbf{s}, \mathbf{x}_{-i}) + (1 - \lambda) u_i(\mathbf{s}', \mathbf{x}_{-i}) = M,$$

so the entire segment maximizes as well. Each $\text{BR}_i(\mathbf{x}_{-i})$ is convex, and a product of convex sets is convex.

Closed graph. Suppose $\mathbf{x}^k \rightarrow \mathbf{x}$ and $\mathbf{y}^k \rightarrow \mathbf{y}$ with $\mathbf{y}^k \in \text{BR}(\mathbf{x}^k)$. For each player i and each $\mathbf{s}_i \in \Delta_{m_i}$, the best-response property gives $u_i(\mathbf{y}_i^k, \mathbf{x}_{-i}^k) \geq u_i(\mathbf{s}_i, \mathbf{x}_{-i}^k)$. The payoff is a polynomial in the mixing probabilities, so it is continuous, and letting $k \rightarrow \infty$ preserves the inequality: $u_i(\mathbf{y}_i, \mathbf{x}_{-i}) \geq u_i(\mathbf{s}_i, \mathbf{x}_{-i})$. This holds for every $\mathbf{s}_i$, so $\mathbf{y}_i \in \text{BR}_i(\mathbf{x}_{-i})$ for each i , hence $\mathbf{y} \in \text{BR}(\mathbf{x})$. The graph is closed.

The strategy space X is a product of simplices, so it is nonempty, compact, and convex inside $\mathbb{R}^{m_1 + \cdots + m_n}$. Kakutani's theorem applies to the correspondence BR on X and yields a fixed point $\mathbf{x}^* \in \text{BR}(\mathbf{x}^*)$. By the equivalence above, $\mathbf{x}^*$ is a Nash equilibrium. ■

Remark 8.18 (Two roads, one theorem). The proofs trade the same difficulty in opposite directions. The best-response correspondence is the natural object, but it is set-valued, so it calls for Kakutani. Nash's gain-map (8.1) is an honest function, so it needs only Brouwer, at the price of replacing “jump straight to a best reply” with “edge a little toward whatever currently looks profitable.” Each route ends at a fixed point, and a fixed point of either is a profile in which no player gains by moving alone.

Computational Companion

Problem Set 4 turns this proof into computation. One task proves the theorem in the form above; another approximates a fixed point of a continuous self-map of the square by discretizing the domain, evaluating the displacement $f(z) - z$ on a grid, and refining where it changes sign. The method finds the fixed point of a contraction toward the center of the square at the center, as expected. The full routine appears in Appendix C.

Remark 8.19 (What existence does and does not give). Nash's theorem guarantees an equilibrium exists; it does not say there is only one, nor that the players can agree on which to play, nor that the equilibrium is efficient. The coordination game and the Battle of the Sexes have several equilibria, and the Prisoner's Dilemma has a unique but inefficient one. Existence is the foundation; selection, efficiency, and computation are separate questions, some of which the cooperative theory of the next chapters addresses.

Historical Note: John Nash

Nash proved the theorem of this section in his Princeton doctoral thesis of 1950, a document of fewer than thirty pages, when he was twenty-one. The equilibrium was one of several deep contributions: he also solved standing problems in real algebraic geometry and in the embedding of Riemannian manifolds, and his work on nonlinear partial differential equations brought him the Abel Prize in 2015. From the late 1950s a severe schizophrenia interrupted his career for decades, an illness and recovery later told in the biography and film *A Beautiful Mind*. He shared the 1994 Nobel Memorial Prize with Harsanyi and Selten. Days after receiving the Abel Prize in Norway, he and his wife Alicia died in a car accident on the New Jersey Turnpike in May 2015.

8.10 Correlated equilibrium

Nash equilibrium lets the players randomize independently. Aumann's 1974 *correlated equilibrium* allows their randomizations to be *coupled* by a shared signal, and the extra coordination can leave everyone better off. Picture a trusted device that privately whispers an action to each player, drawing the pair of recommendations from a fixed joint distribution both players know. That distribution is a **correlated equilibrium** if no player, on hearing their own recommendation, expects to do better by disobeying it, given what the recommendation reveals about the other's.

The gain is sharpest in Chicken (Example 8.10), whose mixed equilibrium wastes payoff on the occasional wreck. Install a device that picks one of the three safe profiles (SWERVE, STRAIGHT), (STRAIGHT, SWERVE), (SWERVE, SWERVE), each with probability $\frac{1}{3}$, and never recommends the crash. Obedience is optimal. Told to go STRAIGHT, a player knows the device must have drawn (STRAIGHT, SWERVE), the only safe profile in which they hold the road, so the opponent is surely swerving and driving straight wins, $1 > 0$. Told to SWERVE, the player is equally likely to be in (SWERVE, STRAIGHT) or (SWERVE, SWERVE), so the opponent goes straight with probability $\frac{1}{2}$; swerving then pays $\frac{1}{2}(-1) + \frac{1}{2}(0) = -\frac{1}{2}$, while deviating to STRAIGHT pays $\frac{1}{2}(-10) + \frac{1}{2}(1) = -\frac{9}{2}$, far worse for the crash risk. Neither recommendation is worth disobeying, so the distribution is a correlated equilibrium.

Its payoff splits $\frac{1}{3}$ each over $(-1, 1)$, $(1, -1)$, $(0, 0)$, giving each driver an expected 0, strictly above the $-\frac{1}{10}$ of the mixed Nash equilibrium and with no chance of a wreck. A public coordinating signal, a traffic light at the one-lane bridge, lifts both players past anything they could reach on their own.

Correlated equilibrium

A correlated equilibrium is a joint distribution over action profiles such that, when a referee privately recommends each player their part of a draw, no one wishes to deviate. Every Nash equilibrium is one (let the device randomize independently), but the converse fails: correlation reaches payoffs no Nash equilibrium attains. The correlated equilibria of a game form a convex polytope cut out by linear inequalities, so they are computable by linear programming, easier to find than the Nash equilibria that sit inside.

Historical Note: Robert Aumann

Correlated equilibrium is the work of Robert Aumann (born 1930 in Frankfurt, his family fleeing to the United States in 1938), among the deepest figures in the field. Beyond the idea of this section he shaped the theory of repeated games and, with Michael Maschler, read the Talmud's ancient bankruptcy rule as the modern nucleolus (Example 10.21). He shared the 2005 Nobel Memorial Prize with Thomas Schelling for showing how conflict and cooperation are illuminated by game theory. An Orthodox Jew and a serious student of the Talmud, he brought the two halves of his life together in that work.

8.11 Evolutionarily stable strategies

Biology gave game theory a second notion of equilibrium. Picture not two reasoning players but a large population, each member programmed to play a fixed strategy and repeatedly matched at random to play a symmetric game. A strategy is **evolutionarily stable** if a population playing it cannot be invaded by a small band of mutants playing anything else: against the incumbent the mutant scores strictly less, or ties and then scores strictly less against itself. Maynard Smith and Price introduced the idea in 1973. Every evolutionarily stable strategy is a Nash equilibrium, carrying an extra margin of stability on top.

The standard illustration is *Hawk–Dove*, the biologist’s Chicken. Two animals contest a resource worth V . A HAWK escalates; a DOVE displays and retreats. Two hawks fight and split the value net of an injury cost C ; a hawk meeting a dove takes everything; two doves share:

	HAWK	DOVE
HAWK	$\frac{V-C}{2}$	V
DOVE	0	$\frac{V}{2}$

When injury is costly, $C > V$, neither pure strategy is stable: an all-hawk population earns the negative $\frac{V-C}{2}$ and is invaded by doves, while an all-dove population is invaded by a lone hawk who takes V at every meeting. The stable population plays HAWK with probability $p^* = V/C$, from the indifference condition $p \frac{V-C}{2} + (1-p)V = (1-p)\frac{V}{2}$. With $V = 2$ and $C = 4$ it settles at half hawks and half doves, the share of aggression climbing with the prize and falling with the cost of a fight. The arithmetic that mixes a Nash equilibrium also predicts the make-up of a population, which is why these ideas reach well past economics into evolutionary biology.

Historical Note: John Maynard Smith

The evolutionarily stable strategy is due to John Maynard Smith (1920–2004) and George Price, in a 1973 paper in *Nature*. Maynard Smith reached biology by an unusual road: trained as an aeronautical engineer, he designed military aircraft during the war before turning to genetics under J. B. S. Haldane. His book *Evolution and the Theory of Games* (1982) carried game theory permanently into biology, and the idea that a population, rather than a reasoning player, can rest at an equilibrium reshaped how biologists explain animal conflict, sex ratios, and signalling. He died in 2004.

Notes and References

Nash announced the existence of equilibrium in Nash [38] and gave the full development in Nash [36], using Kakutani’s fixed-point theorem (Theorem A.7); he shared the 1994 Nobel Memorial Prize with Harsanyi and Selten. The single-valued construction used here, which needs only Brouwer’s theorem (Brouwer 9), is due to David Gale, who also suggested the Kakutani route for the original note. Safety levels and the general-sum examples follow Ferguson [16] and Osborne [39]. The Prisoner’s Dilemma was first studied experimentally by Flood and Dresher at RAND in 1950 [17], and the success of the simple tit-for-tat rule in Axelrod’s iterated-game tournaments [3] is the classic case for cooperation emerging from repetition. The price of anarchy originates with Koutsoupias and Papadimitriou [28] and was named by Papadimitriou [42]; the $\frac{4}{3}$ bound for linear routing is Roughgarden and Tardos [48], and the counterintuitive congestion effect is Braess [8]. The equilibrium-selection problem posed by the Battle of the Sexes and the coordination game is addressed by the theory of focal points and by the cooperative theory of the next two chapters.

Chapter summary

A general-sum game is a pair of payoff matrices (Definition 8.1), and once $B \neq -A$ the single value of zero-sum theory is gone. Each player still has a safety level, the zero-sum value of their own matrix (Definition 8.2), which is what they can guarantee. The solution concept is the Nash equilibrium (Definition 8.5): a profile in which no player gains by deviating alone, equivalently a fixed point of the best-response map. Worked examples showed multiple equilibria in a coordination game and in the Battle of the Sexes, and a unique inefficient equilibrium in the Prisoner’s Dilemma. The chapter’s main theorem is Nash’s (Theorem 8.15): every finite game has a mixed equilibrium. The proof builds a continuous self-map of the product of strategy simplices whose fixed points, guaranteed by Brouwer’s theorem (Theorem 8.14), are exactly the equilibria. Existence is assured; uniqueness, selection, and efficiency are not.

Exercises

Safety levels and Nash equilibria

Exercise 8.1. For the game $(A, B) = \begin{bmatrix} (1,2) & (3,1) \\ (2,0) & (0,3) \end{bmatrix}$, compute both safety levels $V_I = V(A)$ and $V_{II} = V(B^\top)$, and the safety strategies that achieve them.

Exercise 8.2. Find all Nash equilibria, pure and mixed, of the coordination game of Example 8.7. Verify that the mixed-equilibrium payoffs equal the safety levels, and explain why that is not a coincidence for a 2×2 game with a fully mixed equilibrium.

Exercise 8.3. For the Battle of the Sexes (Example 8.8), confirm the two pure equilibria and derive the mixed equilibrium from the indifference conditions. Compute the expected payoffs at the mixed equilibrium and compare them with the pure-equilibrium payoffs.

Exercise 8.4. Find the mixed Nash equilibrium of the coordination game of Exercise 1.4, $\begin{bmatrix} (2,2) & (0,0) \\ (0,0) & (1,1) \end{bmatrix}$, and the expected payoff. Which of the three equilibria would you predict the players reach, and why is prediction hard?

Fixed points and existence

Exercise 8.5. Prove the one-dimensional Brouwer theorem in detail: if $f : [0, 1] \rightarrow [0, 1]$ is continuous, then f has a fixed point. (Apply the intermediate value theorem to $g(x) = f(x) - x$.) Give a continuous f on the half-open interval $(0, 1]$ with no fixed point, and identify which hypothesis of Theorem 8.14 fails.

Exercise 8.6. For a 2×2 game, write out Nash's map (8.1) explicitly for one player as a function of the mixing probability, and verify by hand that a profile is a fixed point if and only if neither player has a profitable pure deviation.

Exercise 8.7. *From Problem Set 4.* State and prove that every finite n -person game has a Nash equilibrium in mixed strategies, following the gain-function construction of Theorem 8.15. Identify precisely where continuity of the payoff and compactness and convexity of the strategy space are used.

PART IV

Cooperative Game Theory

CHAPTER 9

Cooperative Games I: Bargaining and Transferable Utility

“Two people can do better by agreeing than by threatening, but the threat is what decides the terms of the agreement.”

Everything so far has been non-cooperative: players choose independently and cannot make binding promises. Cooperative game theory changes the rules. Players may now communicate, form coalitions, and sign binding agreements before they act. Two questions replace the search for equilibrium. What total can the players achieve by cooperating, and how should they divide it? This chapter answers both for two players, first when utility is freely transferable through side payments, then when it is not and the split must be bargained. The tools are the disagreement point inherited from the safety levels of Chapter 8, the Nash bargaining solution, and Shapley’s exchange rate that converts a non-transferable problem into a transferable one.

9.1 Cooperation and binding agreements

A cooperative two-person game still starts from a bimatrix (A, B) , but the players may now agree on a joint choice of strategies, and possibly on a transfer of payoff, before playing. Two regimes arise.

- In a **transferable utility** (TU) game, payoffs are in a common currency and a player may pay part of their payoff to the other. Side payments are allowed.
- In a **non-transferable utility** (NTU) game, no such transfer is possible, because the payoffs are not in comparable units (one player values time, the other money) or simply cannot be moved between players. The players may still agree on a joint randomization, but each keeps the payoff the agreement produces.

The TU case is simpler because the only question is how to split a single number, the largest total the players can produce. The NTU case is genuine bargaining over a region of outcomes.




Figure 9.1. The NTU feasible set (dark triangle, the convex hull of the four payoff pairs) and the TU feasible set (light, the whole region under $u_1 + u_2 = 7$) for $\begin{bmatrix} (4,3) & (0,0) \\ (2,2) & (1,4) \end{bmatrix}$. Side payments slide outcomes along the dashed slope-1 lines, so transferable utility widens the NTU set into the larger one; the TU Pareto frontier is the single line of greatest total.

9.2 Feasible sets

The set of achievable payoff pairs differs in the two regimes.

Definition 9.1 (Feasible sets). For a bimatrix game (A, B) , the **NTU feasible set** is the convex hull of the mn payoff pairs (a_{ij}, b_{ij}) , the outcomes reachable by joint randomization over pure profiles. The **TU feasible set** is the convex hull of all pairs $(a_{ij} + s, b_{ij} - s)$ for $s \in \mathbb{R}$, that is, the NTU set widened along every line of slope -1 , since a side payment of s moves payoff from one player to the other without changing the total.

Figure 9.1 shows the two sets for the bimatrix $\begin{bmatrix} (4,3) & (0,0) \\ (2,2) & (1,4) \end{bmatrix}$. The NTU set is the triangle spanned by the payoff pairs; the TU set is that triangle smeared along every slope-1 line until it fills out the whole region below the line $u_1 + u_2 = \sigma$, where $\sigma = 7$ is the largest cell total (here at $(4,3)$).

Definition 9.2 (Pareto optimality). A feasible pair (v_1, v_2) is **Pareto optimal** if no other feasible pair gives both players at least as much and one of them strictly more. The Pareto optimal pairs form the upper-right boundary of the feasible set.

In the TU set, because every slope -1 line is allowed, the Pareto boundary is a single line of slope -1 through the outcome of greatest total. All efficiency questions in the TU case reduce to maximizing the sum, then dividing it.

9.3 The transferable-utility solution

Suppose the players will cooperate. Efficiency dictates they choose the joint pure outcome of greatest total payoff.

Definition 9.3 (Cooperative strategy and surplus). The **cooperative value** is the largest achievable total,

$$\sigma = \max_{i,j} (a_{ij} + b_{ij}),$$

attained at a **cooperative strategy** $\langle i_0, j_0 \rangle$. The players will realize σ jointly; the only remaining question is how to split it.

How the surplus is divided depends on what each player could do if talks broke down. That fallback is set by *threats*.

Definition 9.4 (Threats and the disagreement point). If the players fail to agree, each carries out a threat strategy, Player I a mix $\mathbf{p}$ and Player II a mix $\mathbf{q}$, producing the **disagreement point**

$$D = (D_1, D_2) = (\mathbf{p}^\top A \mathbf{q}, \mathbf{p}^\top B \mathbf{q}).$$

Each player chooses a threat to improve their own position in the split. Since the division will turn on the *difference* $D_1 - D_2 = \mathbf{p}^\top (A - B)\mathbf{q}$, the optimal threats are the optimal strategies of the zero-sum game with matrix $A - B$, and the relevant number is its value

$$\delta = V(A - B).$$

The TU solution then splits the surplus so that the difference between the players' payoffs equals the difference δ their threats can enforce.

Theorem 9.5 (The TU value). In a two-person TU game with cooperative value σ and threat value $\delta = V(A - B)$, the solution awards

$$\varphi = (\varphi_1, \varphi_2) = \left(\frac{\sigma + \delta}{2}, \frac{\sigma - \delta}{2} \right).$$

This is the point on the efficient line $x + y = \sigma$ midway between the two extreme demands the threats justify. To reach it from the cooperative strategy's payoff $(a_{i_0 j_0}, b_{i_0 j_0})$, Player II pays Player I a side payment of $\varphi_1 - a_{i_0 j_0}$ (a transfer in the other direction if this number is negative).

The transferable-utility solution

When utility transfers freely, the cooperating players first maximize the total $\sigma = \max_{i,j} (a_{ij} + b_{ij})$, then split it so their payoffs differ by exactly what their threats can enforce, $\delta = V(A - B)$:

$$\varphi = \left(\frac{\sigma + \delta}{2}, \frac{\sigma - \delta}{2} \right).$$

Efficiency fixes the sum of the payoffs; the threat game fixes their difference.

Derivation. Efficiency fixes $\varphi_1 + \varphi_2 = \sigma$. The threat game $A - B$ guarantees Player I a lead of at least δ over Player II in the difference of payoffs, and guarantees Player II that the lead is no more than δ , so fairness sets $\varphi_1 - \varphi_2 = \delta$. Solving the two linear equations gives $\varphi_1 = (\sigma + \delta)/2$ and $\varphi_2 = (\sigma - \delta)/2$. The side payment is whatever moves the realized payoff $a_{i_0 j_0}$ to φ_1 . ■

The solution is efficient (the payoffs sum to σ), individually rational (each gets at least the disagreement payoff), and uniquely determined once the disagreement point is fixed.

Example 9.6 (A TU division). Take

$$(A, B) = \begin{pmatrix} (0, 0) & (6, 2) & (-1, 2) \\ (4, -1) & (3, 6) & (5, 5) \end{pmatrix}.$$

The greatest total is at the cell $(5, 5)$, so $\sigma = 10$ with cooperative strategy $(2, 3)$. The threat game is

$$A - B = \begin{pmatrix} 0 & 4 & -3 \\ 5 & -3 & 0 \end{pmatrix}.$$

Column 1 is dominated by column 3 (for the minimizing Player II, $-3 \leq 0$ and $0 \leq 5$), so delete it and solve the 2×2 game $\begin{bmatrix} 4 & -3 \\ -3 & 0 \end{bmatrix}$. It has no saddle, and with $D = 4 + 3 + 3 + 0 = 10$ the value is $\delta = (4 \cdot 0 - (-3)(-3))/10 = -9/10$, with threats $\mathbf{p} = (0.3, 0.7)$ and $\mathbf{q} = (0, 0.3, 0.7)$. The TU solution is

$$\varphi = \left(\frac{10 - 0.9}{2}, \frac{10 + 0.9}{2} \right) = (4.55, 5.45).$$

Starting from the cooperative payoff $(5, 5)$, Player I pays Player II 0.45 to reach φ .

Example 9.7 (Transferable utility in the Prisoner's Dilemma). Give the Prisoner's Dilemma the cooperative reading

$$(A, B) = \begin{pmatrix} (1, 1) & (4, 0) \\ (2, 4) & (3, 3) \end{pmatrix}.$$

The greatest total is $\sigma = 6$ at $(3, 3)$. The threat game is

$$A - B = \begin{bmatrix} 1-1 & 4-0 \\ 2-4 & 3-3 \end{bmatrix} = \begin{bmatrix} 0 & 4 \\ -2 & 0 \end{bmatrix}.$$

Its row minima are $\min(0, 4) = 0$ and $\min(-2, 0) = -2$, giving maximin 0, and its column maxima are $\max(0, -2) = 0$ and $\max(4, 0) = 4$, giving minimax 0. The two agree, so the top-left entry is a saddle point and the threat value is $\delta = 0$: each player's threat is the pure strategy through that saddle (Player I row 1, Player II column 1), and neither can tilt the difference of payoffs. The TU solution is therefore $\varphi = \left(\frac{6+0}{2}, \frac{6-0}{2}\right) = (3, 3)$ with no side payment. Cooperation lets the players reach the efficient $(3, 3)$ that non-cooperative play destroyed in Example 8.9, and the balanced threat game splits the surplus evenly.

9.4 The Nash bargaining solution

When utility is not transferable, the players cannot reduce the problem to splitting one number. They must agree on a point in the feasible region itself, starting from a disagreement point (u^*, v^*) that each can guarantee. Nash proposed a unique solution [37].

Definition 9.8 (Nash bargaining solution). Given a convex feasible set S and a disagreement point $(u^*, v^*) \in S$, the **Nash bargaining solution** is the feasible point that maximizes the product of gains over the disagreement point,

$$(\bar{u}, \bar{v}) = \arg \max_{(u,v) \in S, u \geq u^*, v \geq v^*} (u - u^*)(v - v^*).$$

The maximizer is the Pareto-optimal point where a hyperbola $(u - u^*)(v - v^*) = c$ is tangent to the Pareto frontier.

Nash justified this product by four axioms: the solution should be Pareto optimal, unchanged by rescaling either player's utility, symmetric when the problem is symmetric, and independent of irrelevant alternatives (removing unchosen options does not change it). The remarkable fact is that the product rule is the *only* solution satisfying all four, and the proof is short enough to give in full.

Theorem 9.9 (Nash's bargaining theorem). Let f assign to each bargaining problem $(S, (u^*, v^*))$, with S convex, compact, and containing a point strictly above the disagreement point, a feasible outcome $f(S, (u^*, v^*))$. If f satisfies

1. **(PAR)** Pareto optimality: no feasible point is strictly larger in both coordinates;
2. **(INV)** invariance to positive affine rescaling of either player's utility;
3. **(SYM)** symmetry: if the problem is unchanged by swapping the two players, the two players receive equal payoffs;
4. **(IIA)** independence of irrelevant alternatives: if $T \subseteq S$, $f(S, \cdot) \in T$, then $f(T, \cdot) = f(S, \cdot)$;

then f is exactly the Nash bargaining solution of Definition 9.8: the maximizer of $(u - u^*)(v - v^*)$.

Proof. The Nash point obeys the axioms. Write $z = (\bar{u}, \bar{v})$ for the product maximizer. It is Pareto optimal because the product strictly increases when either coordinate rises, so a point above z in both coordinates would give a larger product (PAR). Rescaling $u \mapsto \alpha u + \beta$ with $\alpha > 0$ multiplies the gain $u - u^*$ by α , hence multiplies the whole product by a positive constant and leaves its maximizer where it was (INV). If the problem is symmetric, the product is a symmetric function of (u, v) , so its unique maximizer lies on the diagonal and gives both players the same amount (SYM). And shrinking S to any T that still contains z cannot raise the maximum, so z stays the maximizer (IIA). The Nash point therefore satisfies all four axioms; the work is showing it is the *only* such rule.

Uniqueness. Suppose f satisfies the four axioms; we prove $f(S, (u^*, v^*)) = z$. By nondegeneracy $\bar{u} > u^*$ and $\bar{v} > v^*$, so the affine map

$$u' = \frac{u - u^*}{\bar{u} - u^*}, \quad v' = \frac{v - v^*}{\bar{v} - v^*}$$

is a positive rescaling of each axis; it sends the disagreement point to $(0, 0)$ and the Nash point z to $(1, 1)$. By (INV) it is enough to prove $f(S', (0, 0)) = (1, 1)$ for the transformed set S' , whose Nash point is now $(1, 1)$.

Step 1: S' lies under the line $u + v = 2$. Suppose some $w = (w_1, w_2) \in S'$ had $w_1 + w_2 > 2$. Since $(1, 1) \in S'$ and S' is convex, the segment to w stays in S' ; a point a fraction t of the way along it is $(1 + t(w_1 - 1), 1 + t(w_2 - 1))$, with gain product

$$(1 + t(w_1 - 1))(1 + t(w_2 - 1)) = 1 + t(w_1 + w_2 - 2) + O(t^2) > 1$$

for small $t > 0$, because $w_1 + w_2 - 2 > 0$. That beats the product 1 at $(1, 1)$, contradicting the fact that $(1, 1)$ maximizes the product on S' . Hence every point of S' satisfies $u + v \leq 2$.

Step 2: enclose S' in a symmetric set. S' is bounded, so choose M large enough that $S' \subseteq T := \{(u, v) : u + v \leq 2, u \leq M, v \leq M\}$. The set T is symmetric across the diagonal and has disagreement point $(0, 0)$. By (PAR) the solution $f(T, (0, 0))$ lies on the Pareto edge $u + v = 2$, and by (SYM) it lies on the diagonal $u = v$; the only such point is $(1, 1)$, so $f(T, (0, 0)) = (1, 1)$.

Step 3: shrink back by (IIA). We have $S' \subseteq T$, and the solution $f(T, (0, 0)) = (1, 1)$ belongs to S' . By (IIA), removing the alternatives in $T \setminus S'$ does not move the solution, so $f(S', (0, 0)) = (1, 1)$. Undoing the affine map gives $f(S, (u^*, v^*)) = z$, the Nash point. ■

The axioms, not the formula, are the content, and changing one axiom changes the solution. The most debated is *independence of irrelevant alternatives*. Kalai and Smorodinsky [25] replaced it with a *monotonicity* axiom (enlarging the feasible set in one player's favor should not hurt that player) and derived a different rule, the Kalai–Smorodinsky solution, which equalizes each player's fraction of their best attainable gain rather than maximizing the product. On a given problem the two solutions can disagree, so “the fair split” is only as well defined as the axioms one is willing to assume.

Example 9.10 (Bargaining over a triangle). Let the feasible set be the triangle with vertices $(0, 0)$, $(3, 0)$, and $(0, 1)$, with disagreement point $(u^*, v^*) = (0, 0)$; Figure 9.2 draws it. The Pareto frontier is the segment from $(3, 0)$ to $(0, 1)$, the line $v = -\frac{1}{3}u + 1$. Maximizing uv subject to this line, the tangency condition equates the slope of the level curve $uv = c$ to the slope $-\frac{1}{3}$ of the frontier, which gives $\bar{u} = 3\bar{v}$. Substituting into the frontier, $\bar{v} = -\frac{1}{3}(3\bar{v}) + 1 = -\bar{v} + 1$, so $2\bar{v} = 1$ and

$$(\bar{u}, \bar{v}) = \left(\frac{3}{2}, \frac{1}{2}\right).$$

The player with the wider range of outcomes, here Player I who can reach 3, receives the larger share, but not all of it: the product rule balances the two gains.




Figure 9.2. The Nash bargaining solution of Example 9.10: the hyperbola $uv = c$ is tangent to the Pareto frontier at $(\frac{3}{2}, \frac{1}{2})$.

The previous example handed us the disagreement point. In a game it must be *computed*, and the natural choice is the pair of safety levels from Chapter 8: what each player can guarantee if talks collapse. The next example carries out the whole pipeline.

Example 9.11 (Bargaining from safety levels). Two partners can split a project's value in two ways but disagree on which, a bargaining game with bimatrix

$$(A, B) = \begin{pmatrix} (6, 2) & (0, 0) \\ (0, 0) & (2, 6) \end{pmatrix}.$$

Step 1: the disagreement point. Each player's safety level is the zero-sum value of their own matrix. For Player I, $A = \begin{bmatrix} 6 & 0 \\ 0 & 2 \end{bmatrix}$ has no saddle (maximin 0, minimax 2), so $V_I = \frac{6 \cdot 2 - 0}{6 - 0 - 0 + 2} = \frac{12}{8} = \frac{3}{2}$. For Player II, $B^\top = \begin{bmatrix} 2 & 0 \\ 0 & 6 \end{bmatrix}$ gives $V_{II} = \frac{12}{8} = \frac{3}{2}$ likewise. The disagreement point is $(u^*, v^*) = (\frac{3}{2}, \frac{3}{2})$.

Step 2: the feasible set and frontier. The NTU feasible set is the convex hull of the payoff pairs $(6, 2)$, $(0, 0)$, $(2, 6)$, a triangle. Its Pareto frontier is the segment from $(6, 2)$ to $(2, 6)$, which has slope $\frac{6-2}{2-6} = -1$ and so lies on the line $u + v = 8$ for $u \in [2, 6]$.

Step 3: maximize the Nash product. Substitute $v = 8 - u$ into the product of gains:

$$N(u) = (u - \frac{3}{2})(8 - u - \frac{3}{2}) = (u - \frac{3}{2})(\frac{13}{2} - u), \quad N'(u) = \frac{13}{2} - u - (u - \frac{3}{2}) = 8 - 2u.$$

Setting $N'(u) = 0$ gives $u = 4$, hence $v = 4$, which lies in $[2, 6]$. The Nash bargaining solution is $(\bar{u}, \bar{v}) = (4, 4)$: each partner secures their guaranteed $\frac{3}{2}$ plus an equal share $\frac{5}{2}$ of the cooperative surplus. $\diamond$




Figure 9.3. Bargaining from safety levels (Example 9.11). The disagreement point d is the pair of safety levels; the Nash solution is where the gain-product hyperbola is tangent to the frontier, at $(4, 4)$.

Example 9.12 (A merger, and the rule that falls out). Two firms consider merging. Their combined operation will earn \$10M of profit to divide, so the feasible frontier is the line $u_1 + u_2 = 10$ (a dollar handed to one is a dollar not handed to the other). If they walk away, firm 1 earns \$4M on its own and firm 2 earns \$1M, so the disagreement point is $d = (4, 1)$. The Nash bargaining solution maximizes the product of gains $N(u_1, u_2) = (u_1 - 4)(u_2 - 1)$ on the frontier. Substituting $u_2 = 10 - u_1$,

$$N(u_1) = (u_1 - 4)(9 - u_1), \quad N'(u_1) = (9 - u_1) - (u_1 - 4) = 13 - 2u_1 = 0 \implies u_1 = \frac{13}{2} = 6.5,$$

so $u_2 = 10 - 6.5 = 3.5$ and the split is $(\bar{u}_1, \bar{u}_2) = (6.5, 3.5)$, with gain product $N = (2.5)(2.5) = 6.25$. Look at where the money went: the surplus created by merging is $s = 10 - 4 - 1 = 5$, and each firm receives its standalone profit plus exactly *half* the surplus,

$$\bar{u}_1 = 4 + \frac{s}{2} = 4 + 2.5 = 6.5, \quad \bar{u}_2 = 1 + \frac{s}{2} = 1 + 2.5 = 3.5.$$

Nash bargaining on a straight frontier: take your floor, split the rest

On a transferable frontier $u_1 + u_2 = M$ with disagreement point $d = (d_1, d_2)$, the Nash bargaining solution is always

$$\bar{u}_i = d_i + \frac{1}{2}(M - d_1 - d_2), \quad i = 1, 2:$$

each party first keeps what it could get alone, then the two split the leftover surplus *equally*. This is the same arithmetic as roommates claiming a bulk discount: each first pays what the item would have cost them individually, and they share the savings fifty-fifty. The disagreement point matters as much as the size of the pie. A firm that can earn more on its own walks away with more, not because it bargains harder, but because its threat to leave is worth more (Figure 9.3 shows the same geometry for a different floor).

Example 9.13 (Nash against Kalai–Smorodinsky on one problem). The subtletybox above warned that the axioms, not the formula, are the content. Here is a single problem on which Nash’s rule and the Kalai–Smorodinsky rule disagree, so the warning has teeth. Let the feasible set be the triangle with vertices $(0, 0)$, $(4, 0)$, $(1, 3)$, with disagreement $d = (0, 0)$. The Pareto frontier is the upper edge from $(4, 0)$ to $(1, 3)$, which we parametrize as

$$(u_1, u_2) = (4 - 3t, 3t), \quad t \in [0, 1].$$

Nash. Maximize the product $u_1 u_2 = (4 - 3t)(3t) = 12t - 9t^2$. Setting the derivative $12 - 18t = 0$ gives $t = \frac{2}{3}$, so the Nash solution is $(\bar{u}_1, \bar{u}_2) = (2, 2)$.

Kalai–Smorodinsky. First find the *ideal point*, each player's best attainable payoff: u_1 peaks at 4 (the corner $(4, 0)$) and u_2 peaks at 3 (the corner $(1, 3)$), so the ideal is $(4, 3)$, generally unreachable at once. The KS rule picks the frontier point where both players get the *same fraction* of their ideal gain, $u_1/4 = u_2/3$. On the frontier this reads $(4 - 3t)/4 = 3t/3 = t$, i.e. $4 - 3t = 4t$, so $t = \frac{4}{7}$ and the KS solution is $(\frac{16}{7}, \frac{12}{7}) \approx (2.29, 1.71)$. Equivalently, it is the point where the straight line from d to the ideal pierces the frontier (Figure 9.4).

The two solutions differ: Nash awards $(2, 2)$, KS awards $(\frac{16}{7}, \frac{12}{7})$, nudging the split toward Player 1, whose reachable peak (4) is the larger. Which is “fair” depends on whether you accept independence of irrelevant alternatives (Nash) or monotonicity in the size of the pie (Kalai–Smorodinsky); the mathematics cannot choose for you.




Figure 9.4. Two “fair” splits of the same problem. The Nash solution (blue) maximizes the product of gains; the Kalai–Smorodinsky solution (red) sits where the line from the disagreement point d to the ideal corner $(4, 3)$ crosses the Pareto frontier, equalizing each player’s fraction of their best attainable gain. They are different points.

9.5 Shapley’s exchange rate

The TU and NTU theories meet through an exchange rate. If the two players’ payoffs are in different currencies, an exchange rate $\lambda > 0$ converts Player I’s units into Player II’s, turning an NTU problem into a TU one. Player I’s payoff a_{ij} becomes λa_{ij} in the common currency, and the combined-payoff problem reads

$$\sigma(\lambda) = \max_{i,j} (\lambda a_{ij} + b_{ij}).$$

One solves the TU problem in the common currency, then converts back. The TU split awards a half-difference of $\frac{\delta}{2}$ in common units, which is $\frac{\delta}{2\lambda}$ in Player I’s own units, so the TU and NTU solutions read

$$\text{TU: } \left(\sigma + \frac{\delta}{2}, \sigma - \frac{\delta}{2} \right), \quad \text{NTU: } \left(\sigma + \frac{\delta}{2\lambda}, \sigma - \frac{\delta}{2\lambda} \right),$$

with the same threat value δ measured in the common currency. The correct exchange rate is the one that makes the resulting split consistent with the geometry of the NTU feasible set, the λ at which the rescaled TU boundary is tangent to the NTU frontier.

Example 9.14 (Solving an NTU game by the exchange rate). Take the non-transferable game

$$(A, B) = \left(\begin{pmatrix} (-1, 1) & (1, 3) \\ (0, 0) & (3, -1) \end{pmatrix}, \quad A = \begin{bmatrix} -1 & 1 \\ 0 & 3 \end{bmatrix}, \quad B = \begin{bmatrix} 1 & 3 \\ 0 & -1 \end{bmatrix} \right).$$

Its four payoff pairs are $(-1, 1), (1, 3), (0, 0), (3, -1)$; the Pareto frontier of their hull is the segment from $(1, 3)$ to $(3, -1)$, the line $u_2 = -2u_1 + 5$.

Step 1: the threat value, for every rate. Convert Player I’s payoffs to Player II’s currency at rate λ . The threat game is the zero-sum spine $M(\lambda) = \lambda A - B$,

$$M(\lambda) = \begin{bmatrix} -\lambda - 1 & \lambda - 3 \\ 0 & 3\lambda + 1 \end{bmatrix}.$$

For every $\lambda > 0$ the entry $M_{21} = 0$ is the minimum of its row (the other entry $3\lambda + 1 > 0$) and the maximum of its column (the other entry $-\lambda - 1 < 0$), hence a saddle. So the threat value is $\delta(\lambda) = 0$ throughout, and the split is the even one $(\frac{\sigma}{2\lambda}, \frac{\sigma}{2})$ in the players' own currencies.

Step 2: the surplus. The combined surplus is $\sigma(\lambda) = \max_{i,j}(\lambda a_{ij} + b_{ij}) = \max\{-\lambda + 1, \lambda + 3, 0, 3\lambda - 1\}$. For $\lambda \leq 2$ the winner is $\lambda + 3$ (the cell $(1, 3)$); for $\lambda \geq 2$ it is $3\lambda - 1$ (the cell $(3, -1)$); the two tie at $\lambda = 2$.

Step 3: pin the rate by the frontier. On the branch $\sigma = \lambda + 3$ the candidate split is $(\frac{\lambda+3}{2\lambda}, \frac{\lambda+3}{2})$. It must land on the NTU frontier $u_2 = -2u_1 + 5$:

$$\frac{\lambda+3}{2} = -2 \cdot \frac{\lambda+3}{2\lambda} + 5 \implies \lambda(\lambda+3) = -2(\lambda+3) + 10\lambda \implies \lambda^2 - 5\lambda + 6 = 0,$$

so $(\lambda-2)(\lambda-3) = 0$. Only $\lambda = 2$ respects this branch's range $\lambda \leq 2$ (and the other branch yields no admissible rate), so $\lambda^* = 2$.

Step 4: read the solution. At $\lambda^* = 2$, $\sigma = 5$, and the split is

$$(u_1, u_2) = \left(\frac{\sigma}{2\lambda^*}, \frac{\sigma}{2}\right) = \left(\frac{5}{4}, \frac{5}{2}\right),$$

which indeed lies on the frontier $(-2 \cdot \frac{5}{4} + 5 = \frac{5}{2})$. The right exchange rate, two of Player I's units per one of Player II's, is exactly the one that makes the transferable split land back on the non-transferable frontier. $\diamond$

The full exchange-rate construction, and its extension to coalitions of more than two players, opens [Chapter 10](#).

Remark 9.15. The exchange rate is what makes “fair division” meaningful when the goods are incomparable. A cow is not a number of sheep until a rate is set, and once it is, the bargaining problem becomes arithmetic. Shapley's insight is that the right rate is not arbitrary; it is pinned down by requiring the transferable solution, converted back, to lie on the non-transferable frontier.

Notes and References

The bargaining solution and its four axioms are from Nash [37], one of the founding papers of cooperative game theory; the best-known alternative, obtained by replacing independence of irrelevant alternatives with monotonicity, is Kalai and Smorodinsky [25]. The transferable-utility value, the use of the threat game $A - B$, and the exchange-rate reduction of NTU to TU problems follow Ferguson [16]; the exchange-rate idea is Shapley's. Two-person cooperative theory and the axiomatic approach to bargaining are developed at length in Osborne and Rubinstein [40] and Myerson [35].

Chapter summary

Cooperative game theory lets players make binding agreements and asks what they can jointly achieve and how to divide it. With transferable utility, side payments widen the feasible set along slope -1 lines ([Definition 9.1](#)), the efficient outcomes lie on one line of slope -1 , and the solution maximizes the total σ then splits it by the threat value $\delta = V(A - B)$, giving $\varphi = ((\sigma + \delta)/2, (\sigma - \delta)/2)$ ([Theorem 9.5](#)). Without transferable utility, the players bargain over the feasible region from a disagreement point, and the Nash bargaining solution ([Definition 9.8](#)) selects the Pareto point maximizing the product of gains, characterized by four axioms and computed by a tangency condition. Shapley's exchange rate ([Section 9.5](#)) reconciles the two by converting NTU payoffs into a common currency, the bridge to the coalitional theory of [Chapter 10](#).

Exercises

Transferable utility and feasible sets

Exercise 9.1. For the TU game $(A, B) = \left[\begin{array}{ccc} (1,5) & (2,2) & (0,1) \\ (4,2) & (1,0) & (2,1) \\ (5,0) & (2,3) & (0,0) \end{array} \right]$, find the cooperative value σ , the threat value $\delta = V(A - B)$, the TU solution φ , and the side payment from each cooperative cell that achieves it.

Exercise 9.2. Verify the TU solution of the cooperative Prisoner's Dilemma of Example 9.7: confirm $\sigma = 6$, show $A - B = \begin{bmatrix} 0 & 4 \\ -2 & 0 \end{bmatrix}$ has a saddle point so $\delta = 0$, and conclude $\varphi = (3, 3)$ with no side payment. Contrast this with the non-cooperative outcome of Example 8.9.

Exercise 9.3 (NTU and TU feasible sets, from Problem Set 4). For each bimatrix, draw the NTU feasible set (the convex hull of the payoff pairs), the TU feasible set (widened along slope -1 lines), and identify the Pareto frontier.

$$(a) \begin{pmatrix} (0, 4) & (3, 2) \\ (4, 0) & (2, 3) \end{pmatrix}, \quad (b) \begin{pmatrix} (3, 1) & (0, 2) \\ (1, 2) & (3, 0) \end{pmatrix}.$$

For (a), find the range of achievable totals; for (b), find the Pareto frontier as a parametrized segment.

The Nash bargaining solution

Exercise 9.4. Solve the Nash bargaining problem on the triangle of Example 9.10 with the disagreement point moved to $(u^*, v^*) = (1, 0)$. Set up the maximization of $(u - 1)(v)$ on the frontier $v = -\frac{1}{3}u + 1$, find the solution, and explain how raising Player I's disagreement payoff shifts the split.

Exercise 9.5. State the four axioms that characterize the Nash bargaining solution, and verify that the solution of Example 9.10 satisfies symmetry would the problem be symmetric, and is unchanged if Player I's payoffs are doubled (invariance to rescaling).

Exercise 9.6. In a two-currency NTU game, explain why the exchange rate λ cannot be chosen arbitrarily. Using the formulas of Section 9.5, describe how the NTU split $(\sigma + \frac{\delta}{2\lambda}, \sigma - \frac{\delta}{2\lambda})$ depends on λ , and what condition pins down the correct rate.

Exercise 9.7. Two divisions can pool their budgets for a joint project worth \$12M, so the transferable frontier is $u_1 + u_2 = 12$. Apart, division 1 can deliver \$5M of value and division 2 can deliver \$1M, giving disagreement point $d = (5, 1)$. Find the Nash bargaining split by maximizing the gain product, then re-derive it with the "floor plus half the surplus" rule of Example 9.12, and confirm the two agree. (Worked in Appendix D.)

CHAPTER 10

Cooperative Games II: The Shapley Value, the Core, and the Nucleolus

“The question is no longer what one player should do, but what a coalition is worth, and how its worth should be shared.”

The bargaining of Chapter 9 was between two players. With more players the central object changes: what matters is the worth of each *coalition*, every subset of players that might band together, and the problem becomes how to divide the worth of the grand coalition among all the players. This chapter develops the coalitional, or characteristic-function, form. It introduces the core, the set of divisions no coalition can object to, then the two great single-valued solutions: the Shapley value, fixed by four fairness axioms, and the nucleolus, which minimizes the worst complaint. The chapter closes the book by returning to the bankruptcy problem of Chapter 1 and showing that the Talmud’s ancient divisions are exactly the nucleolus.

10.1 Coalitional games

Definition 10.1 (Coalitional game). A **coalitional game** is a pair (N, v) where $N = \{1, \dots, n\}$ is the set of players, the **grand coalition**, and the **characteristic function** $v : 2^N \rightarrow \mathbb{R}$ assigns to each coalition $S \subseteq N$ a worth $v(S)$, the total payoff the members of S can secure by cooperating among themselves. It satisfies

$$v(\emptyset) = 0, \quad v(S \cup T) \geq v(S) + v(T) \text{ for disjoint } S, T,$$

the second condition being **super-additivity**: merged coalitions are worth at least the sum of the parts.

A characteristic function can be built from an ordinary strategic game: the worth $v(S)$ is the value of the two-person zero-sum game in which coalition S plays jointly against its complement $N \setminus S$, each side maximizing or minimizing the total payoff to S . This is the reduction that connects Chapter 8 to the present setting. Super-additivity then expresses that cooperation never hurts, which is why the grand coalition tends to form, and $v(N)$ is the total to be divided.

Example 10.2 (Building v from a three-player strategic game). Three players each choose between two pure strategies, written 1 and 2. The payoff triples (u_1, u_2, u_3) are

		$z=1$	$z=2$			$z=1$	$z=2$
Player 1 plays 1 :	$y=1$	(0, 3, 1)	(2, 1, 1)	Player 1 plays 2 :	$y=1$	(1, 0, 0)	(1, 1, 1)
	$y=2$	(4, 2, 3)	(1, 0, 0)		$y=2$	(0, 0, 1)	(0, 0, 1)

where y is Player 2’s choice and z is Player 3’s. To get $v(S)$ we read off the total payoff to S and solve the zero-sum game of S (maximizing) against the complement (minimizing).

$v(\{1\})$. Player 1 alone has 2 rows; the coalition $\{2, 3\}$ has 4 joint columns $(y, z) \in \{11, 12, 21, 22\}$. Reading off u_1 ,

$$\begin{bmatrix} 0 & 2 & 4 & 1 \\ 1 & 1 & 0 & 0 \end{bmatrix} \xrightarrow{\{2,3\} \text{ drop dominated cols } 12,21} \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix},$$

because for the minimizer column $(1, 0)^\top$ is entrywise below columns $(2, 1)^\top$ and $(4, 0)^\top$. The surviving $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ has no saddle, and by Theorem 4.8 its value is $\frac{0 \cdot 0 - 1 \cdot 1}{0 - 1 - 1 + 0} = \frac{1}{2}$. So $v(\{1\}) = \frac{1}{2}$.

$v(\{1, 3\})$, a *cautious case*. Now $\{1, 3\}$ has 4 rows (x, z) and $\{2\}$ has 2 columns. Reading off $u_1 + u_3$,

$$A_{\{1,3\}} = \begin{bmatrix} 1 & 7 \\ 3 & 1 \\ 1 & 1 \\ 2 & 1 \end{bmatrix}.$$

The row minima are all 1, so the coalition can *guarantee* 1 with a pure strategy (maximin = 1); but the column maxima are 3 and 7, so the minimax is $3 \neq 1$ and there is no saddle. The value is therefore not 1: by mixing, $\{1, 3\}$ does better. Only rows (1, 7) and (3, 1) are active at the optimum; letting Player 2 play column 1 with probability q , the coalition is indifferent between them when $7 - 6q = 1 + 2q$, i.e. $q = \frac{3}{4}$, giving value $1 + 2 \cdot \frac{3}{4} = \frac{5}{2}$. So $v(\{1, 3\}) = \frac{5}{2}$ rather than the naive maximin 1, a reminder that the characteristic value is the full game value and exceeds what pure play alone guarantees.

The rest. The same procedure gives $v(\{2\}) = 0$ and $v(\{3\}) = \frac{3}{4}$ (each a 2×4 solve), $v(\{1, 2\}) = 3$ and $v(\{2, 3\}) = 2$ (each a 4×2 with a saddle), and $v(N) = 9$ (the grand coalition just maximizes the total $u_1 + u_2 + u_3$, attained at the cell (4, 2, 3)). Collecting,

S	$\{1\}$	$\{2\}$	$\{3\}$	$\{1, 2\}$	$\{1, 3\}$	$\{2, 3\}$	N
$v(S)$	$\frac{1}{2}$	0	$\frac{3}{4}$	3	$\frac{5}{2}$	2	9

One checks super-additivity (e.g. $v(\{1, 2\}) = 3 \geq v(\{1\}) + v(\{2\}) = \frac{1}{2}$), so the grand coalition is worth forming. This v is now the input to the imputation, core, Shapley, and nucleolus analyses of the rest of the chapter. $\diamond$

The four solution concepts of cooperative theory are the imputation, the core, the Shapley value, and the nucleolus. The first two are sets of acceptable divisions; the last two single out one division each.

10.2 Imputations and the core

A division of $v(N)$ is a payment vector $\mathbf{x} = (x_1, \dots, x_n)$. Two minimal demands make it acceptable.

Definition 10.3 (Imputation). A payment vector $\mathbf{x}$ is an **imputation** for (N, v) if it is

- **individually rational:** $x_i \geq v(\{i\})$ for every player, so no one accepts less than they could secure alone, and
- **group rational (efficient):** $\sum_{i \in N} x_i = v(N)$, so the entire worth of the grand coalition is distributed.

Individual rationality is the weakest stability requirement: it asks only that no single player prefers to walk away. The core strengthens it to every coalition.

Definition 10.4 (Core). The **core** of (N, v) is the set of imputations no coalition can improve upon,

$$\mathcal{C}(v) = \left\{ \mathbf{x} : \sum_{i \in N} x_i = v(N), \sum_{i \in S} x_i \geq v(S) \text{ for every } S \subseteq N \right\}.$$

A division is in the core when every coalition already receives at least its worth, so no group of players can break away and do better. The core may be empty.

Example 10.5 (Computing a core). Let $N = \{1, 2, 3\}$ with worths

$$v(\{1\}) = 1, v(\{2\}) = 0, v(\{3\}) = 1, v(\{1, 2\}) = 4, v(\{1, 3\}) = 3, v(\{2, 3\}) = 5, v(N) = 8.$$

The imputations satisfy $x_1 + x_2 + x_3 = 8$ with $x_1 \geq 1, x_2 \geq 0, x_3 \geq 1$, a triangle. The core adds the three pair constraints, each of which caps one player by way of the efficiency equation:

$$x_1 + x_2 \geq 4 \Rightarrow x_3 \leq 4, \quad x_1 + x_3 \geq 3 \Rightarrow x_2 \leq 5, \quad x_2 + x_3 \geq 5 \Rightarrow x_1 \leq 3.$$

These inequalities cut a smaller region out of the imputation triangle, shown in Figure 10.1; that region is the core.




Figure 10.1. The imputation triangle for Example 10.5, redrawn as a vector figure. Each vertex is the imputation that gives one player their maximum; the three dashed lines are the coalition constraints $x_1 \leq 3$, $x_2 \leq 5$, $x_3 \leq 4$, which slice off the corners. The shaded pentagon that survives is the core, the divisions no coalition can improve upon, and the Shapley value φ lies inside it.

10.3 Balanced games and the Bondareva–Shapley theorem

The core can be empty. When is it not? The answer, found independently by Olga Bondareva and Lloyd Shapley in the 1960s, is a clean duality: the core is nonempty exactly when the game is *balanced*. The condition reads best as a statement about fractional coalitions.

Definition 10.6 (Balanced collection). A collection $\mathcal{B}$ of nonempty coalitions, each carrying a weight $\lambda_S \geq 0$, is **balanced** if every player's weights sum to one:

$$\sum_{S \in \mathcal{B}: i \in S} \lambda_S = 1 \quad \text{for every } i \in N.$$

Equivalently the weighted indicator vectors add to the all-ones vector, $\sum_S \lambda_S \mathbf{1}_S = \mathbf{1}_N$. Any partition of N is balanced with all weights equal to 1. The three pairs of a three-player set are balanced with all weights $\frac{1}{2}$, since each player sits in two of the three pairs.

A balanced collection is a way to cover each player exactly once using fractional coalitions. Picture each player dividing a unit of time among the coalitions that contain them according to the weights; balancedness says every player is fully scheduled. The worth such a fractional schedule generates is $\sum_S \lambda_S v(S)$. The grand coalition can always reproduce a genuine partition; the substance of the next definition is that it must keep pace with every fractional cover as well.

Definition 10.7 (Balanced game). A game (N, v) is **balanced** if every balanced collection $\mathcal{B}$ with weights (λ_S) satisfies

$$\sum_{S \in \mathcal{B}} \lambda_S v(S) \leq v(N).$$

Theorem 10.8 (Bondareva, 1963; Shapley, 1967). The core of (N, v) is nonempty if and only if the game is balanced.

When the core is nonempty

The core is the solution set of a system of linear inequalities, so asking whether it is empty is a linear-programming feasibility question. The dual of that program has one variable per coalition, and its feasible points are exactly

the balanced collections. The core is nonempty precisely when no fractional cover of the players can manufacture more worth than the grand coalition produces on its own.

Proof. Whether the core is nonempty is the question of whether the linear program

$$(P) \quad \min_{\mathbf{x} \in \mathbb{R}^n} \sum_{i \in N} x_i \quad \text{subject to} \quad \sum_{i \in S} x_i \geq v(S) \quad \text{for every } \emptyset \neq S \subseteq N$$

attains the value $v(N)$. The constraint for $S = N$ forces $\sum_i x_i \geq v(N)$, so the optimum of (P) is at least $v(N)$. Should the optimum equal $v(N)$, an optimal $\mathbf{x}^*$ has $\sum_i x_i^* = v(N)$ together with $\sum_{i \in S} x_i^* \geq v(S)$ for all S , which is membership in the core, the singleton constraints supplying individual rationality. Conversely every core vector is feasible for (P) with objective $v(N)$. So the core is nonempty if and only if the optimum of (P) equals $v(N)$.

Now form the dual. Attach a multiplier $\lambda_S \geq 0$ to each constraint $\sum_{i \in S} x_i \geq v(S)$. Since $\mathbf{x}$ ranges freely, each coordinate i contributes a dual equality, and the dual program is

$$(D) \quad \max_{\lambda \geq 0} \sum_S \lambda_S v(S) \quad \text{subject to} \quad \sum_{S: i \in S} \lambda_S = 1 \quad \text{for every } i \in N.$$

Its feasible vectors λ are precisely the balanced collections of [Definition 10.6](#). Both programs are feasible: (P) by taking every x_i large, (D) by the single weight $\lambda_N = 1$ with all others zero, the grand coalition standing alone. Strong duality for linear programs ([Theorem 6.7](#)) equates their optima,

$$\min(P) = \max(D) = \max_{\text{balanced } \lambda} \sum_S \lambda_S v(S).$$

The choice $\lambda_N = 1$ shows this maximum is always at least $v(N)$. Therefore $\min(P) = v(N)$ holds if and only if $\sum_S \lambda_S v(S) \leq v(N)$ for every balanced collection, which is the definition of a balanced game. With the first paragraph, the core is nonempty if and only if the game is balanced. ■

The theorem turns an existence question into an arithmetic one. For three players the arithmetic is short, because the balanced collections are few.

Example 10.9 (Checking balancedness for three players). On $N = \{1, 2, 3\}$ the minimal balanced collections are the partitions, each with unit weights, and one more: the three pairs $\{1, 2\}, \{1, 3\}, \{2, 3\}$ at weight $\frac{1}{2}$, balanced because every player lies in two of the three. Testing those collections settles the core. Take first the game of [Example 10.5](#), with singletons 1, 0, 1, pairs 4, 3, 5, and $v(N) = 8$:

$$\begin{aligned} \{1\}\{2\}\{3\} : 1 + 0 + 1 &= 2 \leq 8, & \{3\}\{1, 2\} : 1 + 4 &= 5 \leq 8, \\ \{1\}\{2, 3\} : 1 + 5 &= 6 \leq 8, & \{2\}\{1, 3\} : 0 + 3 &= 3 \leq 8, \\ \{1, 2\}\{1, 3\}\{2, 3\} &\text{ at } \frac{1}{2} : \frac{1}{2}(4 + 3 + 5) &= 6 \leq 8. \end{aligned}$$

Every balanced collection yields at most $8 = v(N)$, so the game is balanced and the core is nonempty, which confirms the shaded pentagon of [Figure 10.1](#). Now flatten the game to all singletons 0, all pairs 1, and $v(N) = 1$. The pairs at weight $\frac{1}{2}$ give

$$\frac{1}{2}(v(\{1, 2\}) + v(\{1, 3\}) + v(\{2, 3\})) = \frac{1}{2}(1 + 1 + 1) = \frac{3}{2} > 1 = v(N),$$

so balancedness fails and the core is empty. Three players who pair up profitably but cannot all agree at once are the smallest empty core, and the weight $\frac{1}{2}$ is exactly the “two of three pairs” double count that leaves the grand coalition short. ◇

10.4 The Shapley value

The core is often a large set, sometimes empty, and gives no single recommendation. The Shapley value [\[55\]](#) singles out one division, justified by four axioms of fairness.

Definition 10.10 (Shapley axioms). A value is a rule $v \mapsto \varphi(v) = (\varphi_1, \dots, \varphi_n)$ assigning a payment to each player. The **Shapley axioms** are:

1. **Efficiency:** $\sum_i \varphi_i(v) = v(N)$.
2. **Symmetry:** if players i and j satisfy $v(S \cup \{i\}) = v(S \cup \{j\})$ for every S containing neither, then $\varphi_i = \varphi_j$.
3. **Dummy:** if $v(S \cup \{i\}) = v(S)$ for every S not containing i , so i adds nothing anywhere, then $\varphi_i = 0$.
4. **Additivity:** $\varphi(v + v') = \varphi(v) + \varphi(v')$ for any two games on N .

Theorem 10.11 (Shapley, 1953). There is exactly one value satisfying the four axioms, the **Shapley value**, given by the marginal-contribution formula

$$\varphi_i(v) = \sum_{S \ni i} \frac{(|S| - 1)! (n - |S|)!}{n!} [v(S) - v(S \setminus \{i\})], \quad (10.1)$$

where the sum runs over coalitions S containing i .

Historical Note: Lloyd Shapley

Lloyd Shapley (1923–2016) introduced the value that bears his name in his Princeton thesis of 1953, and went on to shape nearly every branch of cooperative theory: stochastic games, the core, the Bondareva–Shapley theorem, and, with David Gale, the deferred-acceptance algorithm that finds stable matchings. The son of the astronomer Harlow Shapley, he served in the war and spent much of his career at the RAND Corporation and later at UCLA. In 2012 he shared the Nobel Memorial Prize with Alvin Roth for the theory of stable allocations and market design, the cooperative ideas of this chapter turned to the practical business of matching students to schools and donors to patients.

The Shapley value

There is exactly one division of a coalition's worth that satisfies efficiency, symmetry, the dummy axiom, and additivity. It pays each player their average marginal contribution over all orders in which the grand coalition might assemble:

$$\varphi_i(v) = \sum_{S \ni i} \frac{(|S| - 1)! (n - |S|)!}{n!} [v(S) - v(S \setminus \{i\})].$$

Construction and uniqueness. The proof builds the value on a basis of simple games and extends by additivity. For each nonempty $T \subseteq N$, the **elementary game** W_T has

$$W_T(S) = \begin{cases} 1 & T \subseteq S, \\ 0 & \text{otherwise,} \end{cases}$$

a unit of worth that appears exactly when all of T is present. The axioms force its value: every player outside T is a dummy, so receives 0; every pair inside T is symmetric, so they share equally; and efficiency distributes the single unit. Hence

$$\varphi_i(W_T) = \begin{cases} 1/|T| & i \in T, \\ 0 & i \notin T. \end{cases}$$

Every characteristic function decomposes uniquely as a linear combination $v = \sum_T c_T W_T$, where the coefficients are computed recursively by $c_T = v(T) - \sum_{S \subsetneq T} c_S$. By additivity,

$$\varphi_i(v) = \sum_{T \ni i} \frac{c_T}{|T|}, \quad (10.2)$$

which determines the value uniquely and proves existence. A counting argument over the orders in which players join the grand coalition rewrites (10.2) as (10.1): the weight $(|S| - 1)! (n - |S|)!/n!$ is the probability that, in a uniformly random ordering of the n players, the members of $S \setminus \{i\}$ all precede i and the rest follow, and $v(S) - v(S \setminus \{i\})$ is i 's marginal contribution upon joining. The Shapley value is thus the average marginal contribution of player i over all arrival orders. ■

Example 10.12 (Computing the Shapley value). Take the game of Example 10.5. The decomposition coefficients, from $c_T = v(T) - \sum_{S \subsetneq T} c_S$, are

$$c_{\{1\}} = 1, \quad c_{\{2\}} = 0, \quad c_{\{3\}} = 1, \quad c_{\{1,2\}} = 3, \quad c_{\{1,3\}} = 1, \quad c_{\{2,3\}} = 4, \quad c_{\{1,2,3\}} = -2,$$

where for instance $c_{\{1,2\}} = 4 - 1 - 0 = 3$ and $c_N = 8 - 3 - 1 - 4 - 1 - 0 - 1 = -2$ (coefficients may be negative). By (10.2),

$$\varphi_1 = \frac{c_{\{1\}}}{1} + \frac{c_{\{1,2\}}}{2} + \frac{c_{\{1,3\}}}{2} + \frac{c_N}{3} = 1 + \frac{3}{2} + \frac{1}{2} - \frac{2}{3} = \frac{7}{3},$$

and similarly $\varphi_2 = \varphi_3 = \frac{17}{6}$. The same numbers come from the marginal-contribution formula (10.1); for player 1,

$$\varphi_1 = \frac{2}{6}[1 - 0] + \frac{1}{6}[4 - 0] + \frac{1}{6}[3 - 1] + \frac{2}{6}[8 - 5] = \frac{1}{3} + \frac{2}{3} + \frac{1}{3} + 1 = \frac{7}{3}.$$

The value is

$$\varphi(v) = \left(\frac{7}{3}, \frac{17}{6}, \frac{17}{6} \right), \quad \frac{14}{6} + \frac{17}{6} + \frac{17}{6} = 8 = v(N).$$

Players 2 and 3 receive more than player 1 despite player 1's larger standalone worth, because they contribute more to the coalitions on average. This Shapley value lies inside the core of Example 10.5, as a check against its inequalities confirms.

Example 10.13 (The Shapley value as an average over arrival orders). The marginal-contribution formula (10.1) is, literally, an average over the $n! = 6$ orders in which the three players might arrive to form the grand coalition. It is worth seeing every term. For each ordering, a player's marginal contribution is the worth they add when they join, $v(\text{those present, including them}) - v(\text{those already present})$.

arrival order	player 1 adds	player 2 adds	player 3 adds	checks sum = 8
1, 2, 3	1	$4 - 1 = 3$	$8 - 4 = 4$	$1 + 3 + 4 = 8$
1, 3, 2	1	$8 - 3 = 5$	$3 - 1 = 2$	$1 + 5 + 2 = 8$
2, 1, 3	$4 - 0 = 4$	0	$8 - 4 = 4$	$4 + 0 + 4 = 8$
2, 3, 1	$8 - 5 = 3$	0	$5 - 0 = 5$	$3 + 0 + 5 = 8$
3, 1, 2	$3 - 1 = 2$	$8 - 3 = 5$	1	$2 + 5 + 1 = 8$
3, 2, 1	$8 - 5 = 3$	$5 - 1 = 4$	1	$3 + 4 + 1 = 8$
column total	14	17	17	

Averaging each column over the six orders,

$$\varphi_1 = \frac{14}{6} = \frac{7}{3}, \quad \varphi_2 = \frac{17}{6}, \quad \varphi_3 = \frac{17}{6},$$

exactly the values found before. Every row sums to $v(N) = 8$, which is why the Shapley value is automatically efficient: no matter the order, the three marginal contributions partition the grand coalition's worth. $\diamond$

Example 10.14 (The glove market: scarcity is power). Three traders hold gloves. Players 1 and 2 each hold a *left* glove; player 3 holds a single *right* glove. Only a complete pair (one left, one right) has value, worth 1; an unmatched glove is worthless. A coalition's worth is the number of pairs it can assemble, $v(S) = \min(\#\text{left in } S, \#\text{right in } S)$. The worths are

$$v(\{1\}) = v(\{2\}) = v(\{3\}) = 0, \quad v(\{1, 2\}) = 0, \quad v(\{1, 3\}) = v(\{2, 3\}) = 1, \quad v(N) = 1,$$

since $\{1, 2\}$ holds two lefts and no right ($\min(2, 0) = 0$), while $\{1, 3\}$ and $\{2, 3\}$ each hold one of each ($\min(1, 1) = 1$), and the grand coalition can form one pair ($\min(2, 1) = 1$). Average each player's marginal contribution over the $3! = 6$ arrival orders:

order	player 1	player 2	player 3
1, 2, 3	0	0	1
1, 3, 2	0	0	1
2, 1, 3	0	0	1
2, 3, 1	0	0	1
3, 1, 2	1	0	0
3, 2, 1	0	1	0
total	1	1	4

For instance, in the order 3, 1, 2 player 3 arrives first and completes no pair (worth 0), then player 1 arrives and completes the pair with 3's right glove (marginal 1), then player 2 adds nothing. The Shapley value is

$$\varphi = \left(\frac{1}{6}, \frac{1}{6}, \frac{4}{6}\right) = \left(\frac{1}{6}, \frac{1}{6}, \frac{2}{3}\right).$$

The lone holder of the scarce right glove captures two-thirds of the surplus, four times what each left-glove holder gets, even though all three are needed for the single pair. Scarcity, not quantity, is what the Shapley value rewards: the two left gloves compete with each other and bid their own value down. $\diamond$

Application: the Shapley value in the wild

The Shapley value is one of the most widely applied ideas in this book. In a weighted voting body it becomes the **Shapley–Shubik power index** [57]: a member's power is the fraction of orderings in which that member is *pivotal*, the one whose vote turns a losing coalition into a winning one, which can differ sharply from their share of the votes. In cost sharing it splits a joint expense by marginal contribution; the long-used rule for airport runway landing fees, charging each aircraft for the runway length only it requires, turns out to be exactly the Shapley value of the associated cost game [30]. Most recently it has entered machine learning as **SHAP** [31]: to explain a prediction, treat the input features as players in a cooperative game whose worth is the model's output, and attribute the prediction to each feature by its Shapley value. The same average-marginal-contribution formula (10.1) underlies all three.

Example 10.15 (Voting power is not vote share). Three partners hold 3, 2, and 1 votes, and a motion passes with a quota of 4 votes. The characteristic function of this *weighted voting game* is $v(S) = 1$ if the votes in S reach 4 and 0 otherwise; the winning coalitions are $\{1, 2\}$ (five votes), $\{1, 3\}$ (four), and the grand coalition. A partner's **Shapley–Shubik power** is the fraction of the $3! = 6$ arrival orders in which that partner is *pivotal*, the one whose entry first carries the running tally to the quota. Tabulating all six orders (the pivot in bold):

order	running tally	reaches 4 at	pivot
1, 2, 3	3, 5, 6	second	2
1, 3, 2	3, 4, 6	second	3
2, 1, 3	2, 5, 6	second	1
2, 3, 1	2, 3, 6	third	1
3, 1, 2	1, 4, 6	second	1
3, 2, 1	1, 3, 6	third	1

Partner 1 is pivotal in four of the six orders, partners 2 and 3 in one each, so the power index is

$$\left(\frac{4}{6}, \frac{1}{6}, \frac{1}{6}\right) = \left(\frac{2}{3}, \frac{1}{6}, \frac{1}{6}\right).$$

Compare this to the raw vote shares $\left(\frac{3}{6}, \frac{2}{6}, \frac{1}{6}\right) = \left(\frac{1}{2}, \frac{1}{3}, \frac{1}{6}\right)$. Two distortions jump out. Partner 1 commands two-thirds of the *power* from only half of the *votes*, because almost every coalition needs him. And partners 2 and 3 wield *equal* power despite 2 holding twice the votes of 3: with partner 1 present either of them completes a quota, and without partner 1 neither can, so their extra vote never matters. Power flows from being pivotal rather than from the size of one's holding, which is why the index, and not the vote count, is the honest measure of influence in a committee. $\diamond$

Example 10.16 (The Banzhaf index: a second measure of power). The Shapley–Shubik index counts pivots over *orderings*; the **Banzhaf index** instead counts *swings*, the coalitions a player turns from losing to winning by joining. For the same weighted vote $[4; 3, 2, 1]$ the winning coalitions are $\{1, 2\}$, $\{1, 3\}$, and $\{1, 2, 3\}$. Partner 1 swings all three, since removing him breaks each; partner 2 swings only $\{1, 2\}$, because $\{1, 2, 3\}$ still wins without him through $\{1, 3\}$; and partner 3 swings only $\{1, 3\}$. With swing counts 3, 1, 1 totalling 5, the normalized Banzhaf index is

$$\left(\frac{3}{5}, \frac{1}{5}, \frac{1}{5}\right) = (0.6, 0.2, 0.2),$$

against the Shapley–Shubik $\left(\frac{2}{3}, \frac{1}{6}, \frac{1}{6}\right) \approx (0.667, 0.167, 0.167)$. The two indices agree on the ranking and on the equal power of partners 2 and 3, yet they divide the total differently, weighting a player’s pivotal coalitions in different ways. Both honestly answer “how much power?”; which to use is a modeling choice about whether the order of joining should count. $\diamond$

Example 10.17 (Splitting a shared taxi: cost-sharing by Shapley value). Three friends share one taxi home along a single straight road. Alone, Ada lives nearest and her fare would be \$6, Bo is farther at \$12, and Cy is farthest at \$18. Riding together the meter only charges for the farthest rider, so the whole trip costs \$18. How should they split it fairly? Model it as a *cost game* whose worth of a coalition S is the fare that group would pay alone, namely the cost of its farthest member:

$$c(\{A\}) = 6, \quad c(\{B\}) = 12, \quad c(\{C\}) = 18,$$

$$c(\{A, B\}) = 12, \quad c(\{A, C\}) = c(\{B, C\}) = c(\{A, B, C\}) = 18.$$

The Shapley value charges each rider their *average marginal cost* over all $3! = 6$ boarding orders: the marginal cost of a rider is how much the fare rises when they get in, given who is already aboard. Tabulating every order (the new rider’s marginal cost in their column):

boarding order	A	B	C
A, B, C	6	6	6
A, C, B	6	0	12
B, A, C	0	12	6
B, C, A	0	12	6
C, A, B	0	0	18
C, B, A	0	0	18
sum	12	30	66
$\div 6$	2	5	11

Read one row to see the logic: in order C, A, B , Cy boards first into an empty cab and bears the full \$18; Ada and Bo then add nothing, because the cab is already going all the way to Cy’s house. In order A, B, C the cost builds up \$6, \$6, \$6 as each extra leg is added. Averaging the columns gives the Shapley split

$$\varphi(A) = \frac{12}{6} = 2, \quad \varphi(B) = \frac{30}{6} = 5, \quad \varphi(C) = \frac{66}{6} = 11,$$

which sums to $2 + 5 + 11 = 18$, the whole fare (efficiency). Ada pays \$2 rather than her solo \$6, Bo pays \$5 rather than \$12, and Cy pays \$11 rather than \$18: each rider saves by sharing the legs of the trip they have in common, and nobody is charged for road they do not travel. This is exactly the airport-runway landing-fee rule of the preceding note, with “length of runway each plane needs” replaced by “length of road each rider travels.” $\diamond$

10.5 The nucleolus

The Shapley value averages contributions. The nucleolus [50] takes a different view of fairness: make the most aggrieved coalition as content as possible.

Definition 10.18 (Excess). For an efficient allocation $\mathbf{x}$ with $\sum_i x_i = v(N)$ and a coalition S , the **excess** is

$$e(S, \mathbf{x}) = v(S) - \sum_{i \in S} x_i,$$

the gap between what S is worth and what it receives. A large excess is a strong complaint: the coalition believes it deserves much more than the allocation grants it.

Definition 10.19 (Nucleolus). For each efficient allocation $\mathbf{x}$, form the vector of all 2^n excesses sorted into decreasing order. The **nucleolus** is the efficient allocation whose sorted excess vector is **lexicographically smallest**: it minimizes the largest excess; among allocations that achieve that minimum, it minimizes the second-largest; and so on. The nucleolus exists, is unique, and lies in the core whenever the core is nonempty.

The nucleolus has no closed-form formula. It is computed by a sequence of linear programs, each driving down the current largest excess until the allocation is pinned. For small games one can balance the binding excesses by hand.

Example 10.20 (A bankruptcy nucleolus). Three creditors A, B, C share an estate. The characteristic function is

$$v(\{B, C\}) = 100, \quad v(N) = 200,$$

with every other proper coalition worth 0. Start from the equal split $\mathbf{x} = (\frac{200}{3}, \frac{200}{3}, \frac{200}{3})$ and read off the excesses; the largest is $e(\{B, C\}, \mathbf{x}) = 100 - \frac{400}{3} = -\frac{100}{3}$. To reduce the largest complaint, balance the two excesses that move oppositely as x_1 changes. Using $x_1 + x_2 + x_3 = 200$,

$$e(\{A\}, \mathbf{x}) = -x_1, \quad e(\{B, C\}, \mathbf{x}) = 100 - (200 - x_1) = x_1 - 100,$$

and setting $-x_1 = x_1 - 100$ gives $x_1 = 50$. With x_1 fixed, $x_2 + x_3 = 150$, and the symmetry of B and C forces $x_2 = x_3 = 75$. The nucleolus is

$$\mathcal{N}(v) = (50, 75, 75).$$

The Shapley value of the same game is $(\frac{100}{3}, \frac{250}{3}, \frac{250}{3}) \approx (33.3, 83.3, 83.3)$, which differs: the nucleolus gives more to the weak player A . The two solutions embody different ideas of fairness, and they need not agree.

10.6 The bankruptcy problem and the Talmud

We can now keep the promise made in [Example 1.1](#). A **bankruptcy problem** has an estate E and claims $c_1, \dots, c_n$ with $\sum_i c_i > E$. Cast it as a coalitional game by giving each coalition what is left of the estate after the outside claims are paid in full,

$$v(S) = \max\left\{0, E - \sum_{i \notin S} c_i\right\}. \quad (10.3)$$

A coalition can count only on what no one outside it can claim. The nucleolus of this game is the allocation the Talmud prescribes [1].

Example 10.21 (The contested estate, resolved). Return to the three claims $c_1 = 100$, $c_2 = 200$, $c_3 = 300$ of [Example 1.1](#), and take the middle estate $E = 200$. Formula (10.3) gives

$$v(\{2, 3\}) = \max\{0, 200 - 100\} = 100, \quad v(N) = 200,$$

and every other proper coalition is worth 0, because each excludes a claim large enough to exhaust the estate. This is exactly the game of [Example 10.20](#), whose nucleolus is $(50, 75, 75)$. That is the Talmud's ruling for $E = 200$. For the small estate $E = 100$, every proper coalition is worth 0, the game is symmetric, and the nucleolus is the equal split $(33\frac{1}{3}, 33\frac{1}{3}, 33\frac{1}{3})$; for the large estate $E = 300$, the same computation yields $(50, 100, 150)$. The three rows of the Talmud table,

E	claimant 1	claimant 2	claimant 3
100	$33\frac{1}{3}$	$33\frac{1}{3}$	$33\frac{1}{3}$
200	50	75	75
300	50	100	150

which looked like three unrelated rulings in [Chapter 1](#), are the nucleolus of the bankruptcy game at three estates. The two-thousand-year-old divisions follow from a single twentieth-century principle: minimize the largest grievance.

Historical Note

That the Talmudic numbers are the nucleolus was discovered by Robert Aumann and Michael Maschler in 1985 [1]. The same rule had been stated for two claimants in the Mishnah as the contested-garment principle: each claimant concedes the part of the estate the other could fully claim, and the contested remainder is split evenly. Aumann and Maschler showed that the three-claimant rulings are exactly the consistent extension of that principle, and that it coincides with the nucleolus of (10.3).

10.7 Shapley value against nucleolus

The two single-valued solutions answer different questions.

- The **Shapley value** rewards each player their average marginal contribution over all orders of coalition formation. It has a closed-form formula (Equation (10.1)), is easy to compute, and is characterized by the four fairness axioms, but it can fall outside the core.
- The **nucleolus** minimizes the largest coalitional complaint, then the next, and so on. It always lies in the core when the core is nonempty, but it has no formula and is computed by linear programming.

Computational Companion

Problem Set 4 closes the course with these computations. One task computes both the Shapley value and the nucleolus of a bankruptcy game, for instance three investors with claims 100, 300, 500 and estate 500, where both solutions come out to (50, 150, 300) and so agree. They need not: at a more contested estate the two diverge, as in Example 10.20, where the Shapley value $(\frac{100}{3}, \frac{250}{3}, \frac{250}{3})$ and the nucleolus (50, 75, 75) differ. Comparing the two principles of fairness, and seeing when they coincide, is the point. The routines for coalition values, the Shapley sum, and the lexicographic nucleolus search appear in Appendix C.

Both belong to the toolkit, and which to use is a modeling choice rather than a mathematical one. The Shapley value asks what each player contributes; the nucleolus asks how to silence the loudest objection. The bankruptcy problem, where they differ, is a fitting place to end: a question of fair division older than the field itself, answered now in two precise and defensible ways.

Notes and References

The Shapley value is from Shapley [55] and the nucleolus from Schmeidler [50]; the reading of the Talmud's bankruptcy rulings as the nucleolus is the result of Aumann and Maschler [1]. The Shapley value's reach is wide: the power index is Shapley and Shubik [57], the airport cost-sharing identity is Littlechild and Owen [30], and its use for explaining machine-learning predictions is Lundberg and Lee [31]. The core is older, with roots in Edgeworth's contract curve and Gillies's 1953 thesis; the precise condition for it to be nonempty is the Bondareva–Shapley theorem on balanced games [56]. The cooperative idea of a stable, unblocked outcome also drives *market design*: the deferred-acceptance algorithm of Gale and Shapley [21] produces stable matchings, and Alvin Roth's redesign of real markets, including the medical-residency match and kidney-exchange clearinghouses [47], applies it at scale, work for which Roth and Shapley shared the 2012 Nobel Memorial Prize. Comprehensive treatments of coalitional games appear in Myerson [35] and Maschler et al. [32]; this chapter follows Ferguson [16] for the Shapley computation, the nucleolus, and the bankruptcy examples.

Chapter summary

A coalitional game (N, v) assigns a worth to every coalition through a super-additive characteristic function (Definition 10.1). A division is an imputation if it is individually and group rational (Definition 10.3), and it lies in the core if no coalition can improve on it (Definition 10.4); the core may be empty. The Shapley value is the unique division satisfying efficiency, symmetry, the dummy axiom, and additivity (Theorem 10.11); it equals each player's average marginal contribution and is computed by (10.1) or through the decomposition into elementary games. The nucleolus (Definition 10.19) instead minimizes the sorted vector of coalition excesses lexicographically; it is unique and sits in the core when that is nonempty. Applied to the bankruptcy game with characteristic function (10.3), the nucleolus reproduces the Talmud's divisions (Example 10.21), resolving the puzzle that opened the book and showing two distinct, principled answers to the oldest question in the subject: how to divide fairly.

Exercises

The core, the nucleolus, and bankruptcy

Exercise 10.1. For the game of Example 10.5, verify that the Shapley value $(\frac{7}{3}, \frac{17}{6}, \frac{17}{6})$ satisfies all the core inequalities, so that it lies in the core. Then find a point of the core that is not the Shapley value.

Exercise 10.2. Give an example of a three-player coalitional game whose core is empty. (Hint: try $v(\{i\}) = 0$, $v(S) = 1$ for every two-player S , and $v(N) = 1$, and show the core inequalities are inconsistent.) Compute the nucleolus of your game and observe that, with an empty core, it need not satisfy every coalition.

Exercise 10.3. For the claims $c = (100, 200, 300)$ and estate $E = 300$, write out the characteristic function (10.3) for all coalitions, and verify that the nucleolus is $(50, 100, 150)$, the third row of the Talmud table. Show also that for $E = 100$ every proper coalition is worth 0, so the nucleolus is the equal split.

Exercise 10.4 (ClosedAI bankruptcy, from Problem Set 4). Three investors hold claims 100, 300, and 500 against an estate of $E = 500$. Build the characteristic function (10.3) (so that, for instance, $v(\{2, 3\}) = \max\{0, 500 - 100\} = 400$ and $v(\{3\}) = \max\{0, 500 - 400\} = 100$). Compute the Shapley value and the nucleolus, confirming that both equal $(50, 150, 300)$, so the two solutions coincide for this estate. Then lower the estate to $E = 200$ and recompute: show that now the Shapley value $(\frac{100}{3}, \frac{250}{3}, \frac{250}{3})$ and the nucleolus $(50, 75, 75)$ differ, and explain why a scarcer estate makes the two principles diverge.

The Shapley value

Exercise 10.5. Compute the Shapley value of the three-player game with $v(\{i\}) = 0$ for each i , $v(\{1, 2\}) = v(\{1, 3\}) = 0$, $v(\{2, 3\}) = 100$, and $v(N) = 200$, using the marginal-contribution formula (10.1). Confirm you obtain $(\frac{100}{3}, \frac{250}{3}, \frac{250}{3})$, and compare with the nucleolus $(50, 75, 75)$ of Example 10.20.

Exercise 10.6. Prove, directly from the axioms, that for the elementary game W_T the Shapley value is $\varphi_i(W_T) = 1/|T|$ for $i \in T$ and 0 otherwise. Identify exactly which axiom forces each part of the answer.

Exercise 10.7. A three-player cooperative game has $v(\emptyset) = v(\{i\}) = 0$ for each i , and

$$v(\{1, 2\}) = 4, \quad v(\{1, 3\}) = 6, \quad v(\{2, 3\}) = 8, \quad v(\{1, 2, 3\}) = 12.$$

(a) Compute the Shapley value by averaging each player's marginal contribution over all six arrival orders, laying out the order table. (b) Check that the result is efficient and lies in the core (every coalition gets at least its worth). (Worked in Appendix D.)

PART V

The Problem Sets

CHAPTER 11

Problem Set 1: Combinatorial Games

The first problem set is about *impartial combinatorial games*: two players alternate, the same moves are available to whoever is to move, and (in normal play) the player who cannot move loses. The whole toolkit is built in Chapter 2, and here we put every piece to work and carry the arithmetic all the way to the end. Three ideas do all the lifting. A position is a **P-position** if the *previous* player (the one who just moved) can force a win, and an **N-position** if the *next* player can; terminal positions are P-positions. The **Sprague–Grundy value** $g(x) = \text{mex}\{g(y) : y \in F(x)\}$ refines this, where mex is the smallest nonnegative integer not among the followers’ values, and x is a P-position exactly when $g(x) = 0$. Finally the **Sprague–Grundy theorem** says a disjunctive sum behaves like a single Nim heap whose size is the **nim-sum** (binary XOR, written $\oplus$) of the components’ Grundy values. Each solution below names its variables, shows the table it is built from, and states the winning move explicitly.

Problem 1: Generalized take-away

A single pile holds n chips. A move removes from 1 to 6 chips, and the player who takes the last chip wins. (a) Find the winning strategy and all P-positions. (b) From 31 chips, give a winning move. (c) From 28 chips, do you want to move first or second?

Solution. (a) Let us classify positions from the bottom up. The empty pile $n = 0$ is terminal: the player facing it cannot move and loses, so $n = 0$ is a **P-position** (the previous player won). Now the key observation: from any pile of size n , the reachable piles are $n - 1, n - 2, \dots, n - 6$, a window of the six positions just below n .

Claim: the P-positions are exactly the multiples of 7. We check the two defining properties.

- From a multiple of 7, every move lands on a non-multiple. If $n = 7q$, then removing $r \in \{1, \dots, 6\}$ leaves $7q - r$, and since $0 < r < 7$, this is strictly between $7(q - 1)$ and $7q$, so it is not a multiple of 7. A P-position can only move to N-positions: good.
- From a non-multiple of 7, some move lands on a multiple. Write $n = 7q + s$ with remainder $s \in \{1, \dots, 6\}$. Removing exactly s chips leaves $7q$, a multiple of 7. An N-position can move to a P-position: good.

Together these say: the multiples of 7 are the losing positions, and from anywhere else the winning move is to remove $n \bmod 7$ chips, restoring a multiple of 7. The first several positions confirm the pattern:

n	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
type	P	N	N	N	N	N	N	P	N	N	N	N	N	N	P

This is the same arithmetic as Nim played modulo 7: think of the chips as a car odometer where the loser is whoever is parked on a multiple of 7, and each turn you may advance the count by 1 to 6. The winner keeps handing the opponent a multiple of 7.

(b) From $n = 31$: since $31 = 7 \cdot 4 + 3$, the remainder is $s = 3$. Remove 3 chips, leaving $28 = 7 \cdot 4$, a P-position. The opponent is now stuck on a multiple of 7.

(c) From $n = 28 = 7 \cdot 4$, the position is already a P-position. The player *to move* loses under optimal play, so you want to move **second**. Whatever your opponent removes, you reply by restoring the next lower multiple of 7 (21, then 14, then 7, then 0), and you take the last chip.

Problem 2: Three subtraction sets, their sum, and two modified rules

For a single pile of 100 chips, with subtraction sets $S_1 = \{1, 3, 5, 7\}$, $S_2 = \{1, 3, 6\}$, and $S_3 = \{1, 2, 4, 8, 16, 32, 64\}$: (a) give the P-positions and Grundy values, and decide first or second. (b) Play the three-pile disjunctive sum, one set per pile, and decide first or second, with a winning move if forced to move first. (c) Replace the piles by three variant games and repeat.

Solution. Everything rests on the recurrence $g(n) = \text{mex}\{g(n-s) : s \in S, s \leq n\}$, built from the ground up. We do each set in turn.

Set $S_1 = \{1, 3, 5, 7\}$ (all odd). Every legal removal is odd, so each move flips the parity of the pile. Building the table,

n	0	1	2	3	4	5	6	7	8
$g_1(n)$	0	1	0	1	0	1	0	1	0

For instance $g_1(2) = \text{mex}\{g_1(1)\} = \text{mex}\{1\} = 0$ (only $s = 1$ fits) and $g_1(3) = \text{mex}\{g_1(2), g_1(0)\} = \text{mex}\{0, 0\} = 1$. The pattern is $g_1(n) = n \bmod 2$, so the P-positions are the even numbers. Since 100 is even, $g_1(100) = 0$: a **P-position**, and alone you would move *second*.

Set $S_2 = \{1, 3, 6\}$. Now $g_2(n) = \text{mex}\{g_2(n-1), g_2(n-3), g_2(n-6)\}$, dropping out-of-range terms:

n	0	1	2	3	4	5	6	7	8	9
$g_2(n)$	0	1	0	1	0	1	2	3	2	0

Three entries are worth seeing, and the cycle then closes:

$$g_2(6) = \text{mex}\{1, 1, 0\} = 2, \quad g_2(7) = \text{mex}\{2, 0, 1\} = 3, \quad g_2(9) = \text{mex}\{2, 2, 1\} = 0.$$

From $g_2(9) = 0$ onward the block 0, 1, 0, 1, 0, 1, 2, 3, 2 repeats with **period** 9. Since $100 = 11 \cdot 9 + 1$, we have $g_2(100) = g_2(1) = 1$: an **N-position**, prefer to move *first*.

Set $S_3 = \{1, 2, 4, 8, \dots\}$ (powers of two). Here $g_3(0) = 0$, $g_3(1) = \text{mex}\{g_3(0)\} = 1$, $g_3(2) = \text{mex}\{g_3(1), g_3(0)\} = 2$, and $g_3(3) = \text{mex}\{g_3(2), g_3(1)\} = \text{mex}\{2, 1\} = 0$ (from 3 only the removals 1 and 2 are powers of two that fit):

n	0	1	2	3	4	5	6	7	8
$g_3(n)$	0	1	2	0	1	2	0	1	2

The values cycle 0, 1, 2 with **period** 3, so $g_3(n) = n \bmod 3$ and the P-positions are the multiples of 3. As $100 = 33 \cdot 3 + 1$, $g_3(100) = 1$: an **N-position**.

(b) The three-pile sum. Playing all three piles of 100 at once is a disjunctive sum, so by the Sprague–Grundy theorem (Theorem 2.15) the combined Grundy value is the nim-sum of the parts:

$$g_1(100) \oplus g_2(100) \oplus g_3(100) = 0 \oplus 1 \oplus 1 = 0.$$

The start is a **P-position**: move *second*, and if forced to move first there is no winning move (every move you make leaves a nonzero nim-sum, which the opponent zeroes again).

(c) Two modified rules. Replace the three components by:

- *Pile 1, “remove at least half.”* From n you may leave any k with $0 \leq k \leq \lfloor n/2 \rfloor$. One checks the Grundy value is $g(n) = \lfloor \log_2 n \rfloor + 1$ (each doubling of the pile adds one reachable Grundy level), so $g(100) = \lfloor \log_2 100 \rfloor + 1 = 6 + 1 = 7$.
- *Pile 2, “even if not all, all if odd.”* A partial removal must be an even number that leaves a nonempty pile, and the whole pile may be taken only when it is odd. Working the recurrence gives $g(2k) = k-1$ and $g(2k+1) = k+1$, so $g(100) = g(2 \cdot 50) = 50 - 1 = 49$.
- *Pile 3, set S_1 .* As above, $g(100) = 0$.

The triple of Grundy values is (7, 49, 0). Align them in binary and add the columns modulo 2:

	32	16	8	4	2	1
7 =	0	0	0	1	1	1
49 =	1	1	0	0	0	1
S =	1	1	0	1	1	0

so $S = 7 \oplus 49 = (110110)_2 = 32 + 16 + 4 + 2 = 54 \neq 0$, an **N-position**. The leftmost unbalanced column is the 32s bit, carried only by pile 2, so pile 2 is the one to play. Its target Grundy value is $49 \oplus 54 = (110001)_2 \oplus (110110)_2 = (000111)_2 = 7$, and a heap of Grundy value 7 in the even-rule game is a heap of 16 (since $g(16) = g(2 \cdot 8) = 8 - 1 = 7$). So the winning move takes pile 2 from 100 down to 16, removing 84 chips, after which the three Grundy values 7, 7, 0 nim-sum to 0 and the opponent is lost.

Problem 3: Grundy values of a directed acyclic graph

A game is given as a directed acyclic graph: each vertex is a position and each edge a legal move. Compute the Sprague–Grundy value of every vertex and classify each as a P- or N-position.

Solution. A position with no outgoing edges is terminal, so it gets $g = 0$. Every other vertex is evaluated only after all of its followers, by $g(x) = \text{mex}\{g(y) : x \rightarrow y\}$. We work the graph of Figure 11.1 in reverse topological order.

$$\begin{aligned}
 g(v_6) &= \text{mex}\{\} = 0 && \text{(sink)} \\
 g(v_4) &= \text{mex}\{\} = 0 && \text{(sink)} \\
 g(v_5) &= \text{mex}\{g(v_6)\} = \text{mex}\{0\} = 1 \\
 g(v_2) &= \text{mex}\{g(v_4), g(v_5)\} = \text{mex}\{0, 1\} = 2 \\
 g(v_3) &= \text{mex}\{g(v_5), g(v_6)\} = \text{mex}\{1, 0\} = 2 \\
 g(v_1) &= \text{mex}\{g(v_2), g(v_3)\} = \text{mex}\{2, 2\} = 0.
 \end{aligned}$$

The last line is the instructive one: both followers of v_1 have value 2, so 0 is missing and $\text{mex} = 0$. The P-positions (value 0) are v_1, v_4, v_6 ; the N-positions are v_2, v_3 (value 2) and v_5 (value 1). From the root v_1 , a P-position, the player to move loses; from any N-position the winning move goes to a follower of value 0 (for example from v_5 move to v_6).




Figure 11.1. The position graph of Problem 3. Each arrow is a legal move; sinks are terminal. Grundy values are computed in reverse topological order by the mex rule. Shaded vertices (v_1, v_4, v_6) are the P-positions, where $g = 0$.

Problem 4: Programming a Nim player

Given pile sizes $(c_1, \dots, c_n)$ for ordinary Nim, compute a winning first move, or report that none exists.

Solution. By Bouton’s theorem (Theorem 2.7), the position loses exactly when its nim-sum is zero, and otherwise a winning move exists. The algorithm is three lines:

1. Compute $S = c_1 \oplus c_2 \oplus \dots \oplus c_n$.
2. If $S = 0$, report “no winning move” (a P-position).
3. Otherwise, for each pile i form the target $y_i = c_i \oplus S$; wherever $y_i < c_i$, the move reducing pile i to y_i wins.

The reduction is legal precisely when $y_i < c_i$, which happens for at least one pile whenever $S \neq 0$ (the pile carrying the leftmost set bit of S). *Worked instance* $(1, 2, 3, 4, 5)$: align the binary digits and add columns modulo 2,

$$\begin{array}{r}
 421 \\
 \hline
 1 = 001 \\
 2 = 010 \\
 3 = 011 \\
 4 = 100 \\
 5 = 101 \\
 \hline
 S = 00\mathbf{1}
 \end{array}$$

so $S = 1$. The only unbalanced column is the 1s bit, carried by the three odd piles 1, 3, 5, so there are three winning moves: $1 \rightarrow 0$, $3 \rightarrow 2$, or $5 \rightarrow 4$, each leaving nim-sum 0. The implementation is in Appendix C.

Problem 5: Programming a Grundy-value calculator

Represent a finite acyclic game as a map from each position to its list of followers, and compute the Grundy value of every position, classifying each as P or N.

Solution. Memoize the recurrence $g(x) = \text{mex}\{g(y) : y \in F(x)\}$, evaluating a vertex only after its followers (a depth-first post-order, or any reverse topological order, as in Problem 3). A position is a P-position iff $g(x) = 0$. The mex of a finite set is found by scanning $0, 1, 2, \dots$ until a gap appears. Memoization makes the whole graph linear in its number of edges. The routine is in Appendix C, and run on the graph of Figure 11.1 it reproduces the labels $g(v_1), \dots, g(v_6) = 0, 2, 2, 0, 1, 0$ computed there by hand.

Problem 6: Programming take-and-break games

Compute the Grundy values for $x \leq 100$ and look for periodicity in (a) Kayles, (b) Dawson’s Chess, and (c) Grundy’s game.

Solution. Each is the memoized mex recurrence of Problem 5, with one new ingredient: a move that *splits* a heap into two independent heaps contributes the *nim-sum* of the two parts’ Grundy values, by the Sprague–Grundy theorem.

- *Kayles* (a row of pins; remove one pin or two adjacent pins, splitting the row): the Grundy values open as

n	0	1	2	3	4	5	6	7	8	9	10	11	12
$g(n)$	0	1	2	3	1	4	3	2	1	4	2	6	4

and become **periodic with period 12** once $n \geq 71$ (the last exceptional value is at $n = 70$). This is the celebrated eventual periodicity of Kayles.

- *Dawson’s Chess* (a take-and-break variant where removing three may split the remainder) is **eventually periodic with period 34**.
- *Grundy’s game* (split any heap into two *unequal* nonempty heaps): the values open $g(0), \dots, g(7) = 0, 0, 0, 1, 0, 2, 1, 0$ and show **no periodicity** within the computed range, and none is known to this day, a small open problem hiding in a children’s game. The code and the full tables are in Appendix C.

CHAPTER 12

Problem Set 2: Two-Person Zero-Sum Games

The second problem set is the computational heart of the course: given a payoff matrix A , find the value V and a pair of optimal mixed strategies $\mathbf{p}^*$ (rows, Player I, the maximizer) and $\mathbf{q}^*$ (columns, Player II, the minimizer), so that $V = (\mathbf{p}^*)^\top A \mathbf{q}^*$. The machinery is built in Chapters 4 to 6; here we run each game all the way to a checked answer, re-displaying the matrix at every step and marking the entries that decide it.

Problem 1: A universal zero-sum solver

Write a program whose input is a payoff matrix and whose output is $(\mathbf{p}, \mathbf{q}, V)$, using in turn the seven methods of the course: the saddle point, the 2×2 formula, the $2 \times n$ and $m \times 2$ graphical method, the principle of indifference, the principle of invariance, the $n \times n$ inverse formula, and the linear program. Test it on 100 random matrices against an independent solver.

Solution. The seven methods form a ladder of increasing cost, and the solver climbs it only as far as it must, exactly as Procedure 4.19 prescribes:

1. **Saddle point.** Compare $\max_i \min_j a_{ij}$ with $\min_j \max_i a_{ij}$; if they agree the game is pure and we stop (cost $O(mn)$).
2. **Dominance.** Delete any row dominated by another row, or column by column, shrinking the matrix before any real work.
3. **Closed forms.** If the surviving game is 2×2 , apply the formula of Theorem 4.8; if it is $2 \times n$ or $m \times 2$, read the value off the graphical envelope (Section 4.6); if it is square and non-degenerate, use the inverse formula $V = 1/(\mathbf{1}^\top A^{-1} \mathbf{1})$.
4. **Linear program.** Otherwise solve the pair of dual linear programs of Chapter 6, which always works.

The implementation is in Appendix C; it agrees with the reference linear-programming solver on all 100 random tests, which is the real point of the exercise: every method, when it applies, gives the same $(\mathbf{p}^*, \mathbf{q}^*, V)$, so the cheap screens never cost accuracy.

Problem 2: Five games by hand

Solve each game, showing all work:

$$\begin{aligned} \text{(i)} \quad & \begin{bmatrix} 1 & 2 & 3 & 4 & 5 \\ 3 & 4 & 5 & 6 & 7 \end{bmatrix}, \quad \text{(ii)} \quad \begin{bmatrix} 1 & -2 & 3 & -4 \\ 0 & 1 & -2 & 3 \\ 0 & 0 & 1 & -2 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \\ \text{(iii)} \quad & \begin{bmatrix} 1 & 2 & -1 \\ 2 & -1 & 4 \\ -1 & 4 & -3 \end{bmatrix}, \quad \text{(iv)} \quad \begin{bmatrix} 0 & 2 & 1 \\ -2 & 0 & -4 \\ -1 & 4 & 0 \end{bmatrix}, \quad \text{(v)} \quad \begin{bmatrix} 10 & 0 & 7 & 0 \\ 0 & 6 & 4 & 0 \\ 0 & 0 & 3 & 3 \end{bmatrix}. \end{aligned}$$

(i) A saddle hiding behind dominance

First the cheap screen. Write each row's minimum down the right margin and each column's maximum along the bottom:

	c_1	c_2	c_3	c_4	c_5	row min
r_1	1	2	3	4	5	1
r_2	3	4	5	6	7	3
col max	3	4	5	6	7	

The maximin is $\max(1, 3) = 3$ (row 2) and the minimax is $\min(3, 4, 5, 6, 7) = 3$ (column 1); they agree, so the boxed entry $a_{21} = 3$ is a saddle point. The value is $V = 3$, with the pure pair $\mathbf{p}^* = (0, 1)$ and $\mathbf{q}^* = (1, 0, 0, 0, 0)$. The reason is visible in the matrix: row 2 dominates row 1 entry by entry ($3 > 1, 4 > 2, \dots$), so the maximizer never leaves row 2, and along row 2 the minimizer takes the smallest entry, in column 1.

(ii) An upper-triangular game, solved bottom-up

The row minima are $-4, -2, -2, 0$ (maximin 0) and the column maxima $1, 1, 3, 3$ (minimax 1); they are attained at different cells, so there is no saddle and all four strategies are live. By the principle of indifference (Theorem 5.1), Player II's mix $\mathbf{q}^* = (q_1, q_2, q_3, q_4)$ must make every one of Player I's rows earn the same value V . Because the matrix is upper triangular, the equations unzip from the bottom row up, each introducing one new unknown:

$$\begin{aligned} \text{row 4: } q_4 &= V, \\ \text{row 3: } q_3 - 2q_4 &= V \Rightarrow q_3 = V + 2V = 3V, \\ \text{row 2: } q_2 - 2q_3 + 3q_4 &= V \Rightarrow q_2 = V + 6V - 3V = 4V, \\ \text{row 1: } q_1 - 2q_2 + 3q_3 - 4q_4 &= V \Rightarrow q_1 = V + 8V - 9V + 4V = 4V. \end{aligned}$$

Normalizing, $q_1 + q_2 + q_3 + q_4 = (4 + 4 + 3 + 1)V = 12V = 1$, so $V = \frac{1}{12}$ and $\mathbf{q}^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{4}, \frac{1}{12})$. The mirror computation for Player I (now each *column* must cost V , and the matrix is lower triangular from Player I's side) gives $\mathbf{p}^* = (\frac{1}{12}, \frac{1}{4}, \frac{1}{3}, \frac{1}{3})$. Check: against column 1, Player I earns $\frac{1}{12} \cdot 1 = \frac{1}{12} = V$; the slim positive value $V = \frac{1}{12}$ shows the game barely favors Player I.

(iii) A symmetric game

The matrix equals its own transpose, so by symmetry $\mathbf{p}^* = \mathbf{q}^*$ and we look for a $\mathbf{p}^*$ solving $A\mathbf{p}^* = V\mathbf{1}$ with $p_1 + p_2 + p_3 = 1$. Guess full support and the value $V = 1$; the candidate $\mathbf{p}^* = (\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$ gives the three row payoffs

$$\text{row 1: } \frac{1}{4} + 2 \cdot \frac{1}{2} - \frac{1}{4} = 1, \quad \text{row 2: } 2 \cdot \frac{1}{4} - \frac{1}{2} + 4 \cdot \frac{1}{4} = 1, \quad \text{row 3: } -\frac{1}{4} + 4 \cdot \frac{1}{2} - 3 \cdot \frac{1}{4} = 1,$$

all equal to $V = 1$. The weights are positive and sum to 1, so the guess is confirmed: $\mathbf{p}^* = \mathbf{q}^* = (\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$ and $V = 1$.

(iv) A saddle, and a second reason for it

Tabulate the margins again:

	c_1	c_2	c_3	row min
r_1	0	2	1	0
r_2	-2	0	-4	-4
r_3	-1	4	0	-1
col max	0	4	0	

The entry $a_{11} = 0$ is the minimum of its row $\{0, 2, 1\}$ and the maximum of its column $\{0, -2, -1\}$, hence a saddle: $V = 0$ with $\mathbf{p}^* = \mathbf{q}^* = (1, 0, 0)$. There is a second, independent reason the value is zero: A is skew-symmetric ($A^\top = -A$), and every such game is fair by Theorem 5.7.

(v) A dominated compromise column

In this 3×4 game each of columns 1, 2, 4 threatens a single row, while column 3 $= (7, 4, 3)$ is a costly compromise. Guess that Player II discards column 3 and that all three rows stay active. Player I's $\mathbf{p}^*$ must then make the three

live columns cost the same amount V :

$$\underbrace{10p_1}_{c_1} = \underbrace{6p_2}_{c_2} = \underbrace{3p_3}_{c_4} = V \implies p_1 = \frac{V}{10}, p_2 = \frac{V}{6}, p_3 = \frac{V}{3}.$$

Normalizing,

$$\frac{V}{10} + \frac{V}{6} + \frac{V}{3} = V \cdot \frac{3+5+10}{30} = \frac{18V}{30} = \frac{3V}{5} = 1 \implies V = \frac{5}{3},$$

so $\mathbf{p}^* = (\frac{1}{6}, \frac{5}{18}, \frac{5}{9})$. *Confirm the discard:* against this $\mathbf{p}^*$, column 3 costs Player II

$$7 \cdot \frac{1}{6} + 4 \cdot \frac{5}{18} + 3 \cdot \frac{5}{9} = \frac{21}{18} + \frac{20}{18} + \frac{30}{18} = \frac{71}{18} \approx 3.94 > \frac{5}{3},$$

strictly worse than the live columns, so Player II was right to drop it. The mirror computation ($10q_1 = 6q_2 = 3q_4 = V$) gives $\mathbf{q}^* = (\frac{1}{6}, \frac{5}{18}, 0, \frac{5}{9})$, with the 0 confirming column 3 unused.

Problem 3: Two proofs

(a) *Prove every symmetric game (one with a skew-symmetric payoff matrix $A^\top = -A$) has value 0. (b) Prove that if a game is invariant under bijections f of the rows and g of the columns, with $A(x, y) = A(f(x), g(y))$, then there exist optimal strategies that are themselves invariant.*

Solution. (a) Let V be the value, which exists by the minimax theorem. Consider the *role-swapped* game: let Player I now use the columns and Player II the rows. Its payoff matrix, from the new row player's viewpoint, is $-A^\top$ (the minimizer's old payoffs negated). Swapping the two roles negates the value, so the role-swapped game has value $-V$. But $-A^\top = -(-A) = A$, so the role-swapped game has the *same* matrix as the original and hence the same value V . Therefore $V = -V$, which forces $V = 0$. (A skew-symmetric matrix has zero diagonal, so playing any pure strategy against its own copy already scores 0, which is the intuition.)

(b) Let $\mathbf{p}^*$ be any optimal strategy for Player I, and let m be the order of the permutation f (so f^m is the identity). Average $\mathbf{p}^*$ over the orbit of f :

$$\bar{\mathbf{p}} = \frac{1}{m} \sum_{k=0}^{m-1} f^k(\mathbf{p}^*).$$

This $\bar{\mathbf{p}}$ is invariant by construction ($f(\bar{\mathbf{p}}) = \bar{\mathbf{p}}$, since f just cycles the terms). It remains optimal because invariance of the payoff gives $f^k(\mathbf{p}^*)^\top A \mathbf{q} = (\mathbf{p}^*)^\top A g^{-k}(\mathbf{q})$ for every $\mathbf{q}$, so each $f^k(\mathbf{p}^*)$ guarantees the value V , and an average of strategies that each guarantee V also guarantees V (the guarantee $\min_{\mathbf{q}} \mathbf{p}^\top A \mathbf{q}$ is concave in $\mathbf{p}$). The same orbit-average over g produces an invariant optimal $\bar{\mathbf{q}}$ for Player II. These are [Theorem 5.7](#) and [Theorem 5.10](#).

Problem 4: Take-and-break and misère Nim

(a) *In the take-and-break game where a move removes any number of chips from a pile or splits a pile into two nonempty piles, find the Grundy function up to $g(100)$. (b) Explain how to play ordinary Nim under the misère convention, where taking the last chip loses.*

Solution. (a) This is Lasker's Nim ([Section 2.6.2](#)). The take moves alone would give $g(n) = n$ (ordinary Nim); the splitting move perturbs this only slightly. Working the mex recurrence, the Grundy values equal n except that within each block of four the values for $4k+3$ and $4k+4$ are swapped:

n	0	1	2	3	4	5	6	7	8
$g(n)$	0	1	2	4	3	5	6	8	7

so $g(100)$ sits in the block $\{97, 98, 99, 100\}$ with $100 = 4 \cdot 25$, giving $g(100) = 99$. (b) By [Theorem 2.22](#), play exactly as in ordinary Nim, keeping the nim-sum at zero, until your move would leave only heaps of size one. At that moment switch rules: leave your opponent an *odd* number of single chips. They are then forced to take chips one at a time until they take the last, and lose. The whole difference between normal and misère play is this single end-game correction.

CHAPTER 13

Problem Set 3: Recursive Games

The third problem set studies games whose cells may say “play this game again,” so the value of a state is defined in terms of other states, possibly including itself. The method throughout is the **substitution principle** of Chapter 7: write an unknown for each state’s value, replace every recursive cell by the value it stands for, solve the resulting ordinary matrix games, and pin the unknowns down from the consistency equations. We carry each game to a closed form and check it.

Problem 1: A recursive card game (programming)

A zero-sum game on a state (m, n) reduces in each state to a 2×2 subgame built from the values of the smaller states $(m, n-1)$ and $(m-1, n)$. Solve it for all (m, n) with $1 \leq m, n \leq 6$ by memoized recursion.

Solution. This is the deduction game of Section 7.9. The base values are $G(m, 0) = 1$ and $G(0, n) = \frac{1}{n+1}$, and each interior state (m, n) assembles a 2×2 matrix from the conjugate values of its two smaller neighbours and solves it by the saddle test and the 2×2 formula (Theorem 4.8). Memoizing $G(m, n)$ over a table computes all 36 states in one sweep. The smallest entry on the diagonal is $G(1, 1) = \frac{1}{2}$, and the diagonal hovers just above one half:

$G(k, k)$	$k = 1$	2	3	4	5	6
value	$\frac{1}{2}$	$\frac{5}{9}$	0.548	0.543	0.539	0.535

The pull toward $\frac{1}{2}$ is the signature of a near-fair game of pure deduction. The code and the full 6×6 table are in Appendix C.

Problem 2: Recursive value equations

Solve the recursive games in which a cell may mean “play this game again”: (i) the nested game $G = \begin{bmatrix} 0 & G_1 \\ G_2 & G_3 \end{bmatrix}$ with $G_1 = \begin{bmatrix} 4 & 3 \\ 1 & 2 \end{bmatrix}$, $G_2 = \begin{bmatrix} 0 & 6 \\ 5 & 1 \end{bmatrix}$, $G_3 = \begin{bmatrix} 0 & -2 \\ -2 & 0 \end{bmatrix}$; and (ii) the mutually recursive trio $G_1 = \begin{bmatrix} G_2 & 1 \\ 1 & 0 \end{bmatrix}$, $G_2 = \begin{bmatrix} G_3 & 0 \\ 0 & 2 \end{bmatrix}$, $G_3 = \begin{bmatrix} G_1 & 1 \\ 1 & 0 \end{bmatrix}$.

(i) Nested blocks, solved bottom up

Here the recursion bottoms out: G_1, G_2, G_3 are ordinary 2×2 games, so we value them first and then substitute into the top level.

- $G_1 = \begin{bmatrix} 4 & 3 \\ 1 & 2 \end{bmatrix}$. Row minima 3, 1 (maximin 3); column maxima 4, 3 (minimax 3). They agree, so the entry $a_{12} = 3$ is a saddle and $v(G_1) = 3$.
- $G_2 = \begin{bmatrix} 0 & 6 \\ 5 & 1 \end{bmatrix}$. Maximin $\max(0, 1) = 1 \neq \min(5, 6) = 5$, no saddle. With $a, b, c, d = 0, 6, 5, 1$ the denominator is $D = a - b - c + d = -10$, so

$$v(G_2) = \frac{ad - bc}{D} = \frac{0 \cdot 1 - 6 \cdot 5}{-10} = \frac{-30}{-10} = 3.$$

- $G_3 = \begin{bmatrix} 0 & -2 \\ -2 & 0 \end{bmatrix}$. No saddle; $D = 0 + 2 + 2 + 0 = 4$, so $v(G_3) = \frac{0 \cdot 0 - (-2)(-2)}{4} = \frac{-4}{4} = -1$.

Substituting these three values into the top-level layout $G = \begin{bmatrix} 0 & v(G_1) \\ v(G_2) & v(G_3) \end{bmatrix}$ turns it into an ordinary game:

$$G = \begin{bmatrix} 0 & G_1 \\ G_2 & G_3 \end{bmatrix} \xrightarrow{\text{substitute the values}} \begin{bmatrix} 0 & 3 \\ 3 & -1 \end{bmatrix}.$$

This has no saddle ($\max(0, -1) = 0 \neq \min(3, 3) = 3$), so with $a, b, c, d = 0, 3, 3, -1$ and $D = 0 - 3 - 3 - 1 = -7$,

$$v(G) = \frac{ad - bc}{D} = \frac{0 \cdot (-1) - 3 \cdot 3}{-7} = \frac{-9}{-7} = \frac{9}{7}, \quad \mathbf{p}^* = \left(\frac{d-c}{D}, \frac{a-b}{D} \right) = \left(\frac{4}{7}, \frac{3}{7} \right) = \mathbf{q}^*.$$

So $v(G) = \frac{9}{7} \approx 1.29$, with both players mixing $(\frac{4}{7}, \frac{3}{7})$.

(ii) A mutually recursive trio

Now the three games refer to one another in a cycle (Figure 13.1): G_1 contains G_2 , which contains G_3 , which contains G_1 again. Write v_1, v_2, v_3 for the three values and substitute each recursive cell. Each game is then a 2×2 solved by the formula:

$$v_1 = \frac{v_2 \cdot 0 - 1 \cdot 1}{v_2 - 1 - 1 + 0} = \frac{-1}{v_2 - 2} = \frac{1}{2 - v_2}, \quad v_2 = \frac{2v_3}{v_3 + 2}, \quad v_3 = \frac{1}{2 - v_1}.$$

Close the loop by eliminating v_1 and v_2 . First




Figure 13.1. The mutually recursive trio: each game's value feeds the next around the cycle. Solving the loop turns the three equations into a single quadratic in v_3 .

$$2 - v_2 = 2 - \frac{2v_3}{v_3 + 2} = \frac{2(v_3 + 2) - 2v_3}{v_3 + 2} = \frac{4}{v_3 + 2}, \quad \text{so} \quad v_1 = \frac{1}{2 - v_2} = \frac{v_3 + 2}{4}.$$

Substituting into $v_3 = \frac{1}{2 - v_1}$,

$$v_3 = \frac{1}{2 - \frac{v_3 + 2}{4}} = \frac{4}{8 - (v_3 + 2)} = \frac{4}{6 - v_3} \implies v_3(6 - v_3) = 4 \implies v_3^2 - 6v_3 + 4 = 0.$$

The quadratic formula gives $v_3 = 3 \pm \sqrt{5}$. The bottom row of G_3 caps its value at 1 (its larger entry is 1), so the admissible root is

$$v_3 = 3 - \sqrt{5} \approx 0.764.$$

Back-substituting, $v_1 = \frac{v_3 + 2}{4} = \frac{5 - \sqrt{5}}{4} \approx 0.691$ and $v_2 = \frac{2v_3}{v_3 + 2} = \frac{5 - \sqrt{5}}{5} \approx 0.553$, matching Exercise 7.3. A 2×2 game whose value involves $\sqrt{5}$ is the smallest reminder that recursion can manufacture irrational values from integer payoffs.

Problem 3: Coin-toss challenge

A player flips a coin with heads-probability p and, after a run of $k \geq 1$ heads, may challenge the opponent to match k heads or continue. (a) For $p = \frac{1}{2}$, find the value. (b) For general p , find the optimal threshold policy and the limit of the value as $p \rightarrow 1$.

Solution. (a) For a fair coin, challenging immediately at $k = 1$ is optimal. Let W be the probability the player to act eventually wins. On the next flip, with probability $\frac{1}{2}$ the coin shows the head needed to win the challenge outright (worth $\frac{1}{2}$ once the matching chance is folded in), and with probability $\frac{1}{2}$ a tail passes the move to the opponent, whose winning chance is then W , so the original player wins with probability $1 - W$:

$$W = \frac{1}{2} \cdot \frac{1}{2} + \frac{1}{2} (1 - W) \implies W + \frac{1}{2}W = \frac{3}{4} \implies \frac{3}{2}W = \frac{3}{4} \implies W = \frac{1}{2}.$$

The game is fair, as symmetry demands.

(b) For general p , a *threshold policy* “continue while the run is below m , then challenge” has value

$$V(m) = \frac{1 - p^{2m}}{2 - p^m}.$$

Put $x = p^m \in (0, 1]$ and maximize $f(x) = \frac{1-x^2}{2-x}$. The derivative,

$$f'(x) = \frac{-2x(2-x) - (1-x^2)(-1)}{(2-x)^2} = \frac{x^2 - 4x + 1}{(2-x)^2},$$

vanishes when $x^2 - 4x + 1 = 0$, i.e. $x_\star = 2 - \sqrt{3} \approx 0.268$ (the other root exceeds 1). The optimal threshold is therefore the integer m closest to $\frac{\ln(2-\sqrt{3})}{\ln p}$, and as $p \rightarrow 1$ the player can hit $x_\star$ almost exactly, so the value tends to

$$f(2 - \sqrt{3}) = 4 - 2\sqrt{3} \approx 0.536,$$

matching [Exercise 7.4](#). Even with a near-certain coin, the challenger wins only about 54% of the time: the right to choose *when* to strike is worth surprisingly little.

Problem 4: Reading

Read Nash's doctoral thesis and its existence proof via Kakutani's fixed-point theorem. The construction and a self-contained proof appear in [Chapter 8](#).

CHAPTER 14

Problem Set 4: N-Person and Cooperative Games

The final problem set leaves the two-person zero-sum world. It proves that equilibria always exist, builds a constructive fixed-point finder, and then turns to cooperation: feasible sets, the Nash equilibria of a bimatrix game, transferable utility, and the bankruptcy problem solved by the nucleolus. The theory is developed in Chapters 8 to 10; here we draw the pictures and carry the divisions to the last cent.

Problem 1: Existence of Nash equilibrium

Prove that every finite n -person game has a Nash equilibrium in mixed strategies.

Solution. This is Theorem 8.15; the constructive heart is a continuous self-map whose fixed points are exactly the equilibria. Let each player i mix over m_i pure strategies, so a profile is a point $\mathbf{x} = (x_1, \dots, x_n)$ in the compact convex product of simplices $X = \prod_i \Delta(S_i)$. For player i and pure strategy s , define the **gain** from switching to s ,

$$\varphi_{i,s}(\mathbf{x}) = \max\{0, u_i(s, \mathbf{x}_{-i}) - u_i(\mathbf{x})\},$$

the amount (if any) by which the pure move s beats player i 's current mixed payoff. Now tilt each player's mix toward its profitable deviations and renormalize:

$$T_{i,s}(\mathbf{x}) = \frac{x_{i,s} + \varphi_{i,s}(\mathbf{x})}{1 + \sum_r \varphi_{i,r}(\mathbf{x})}.$$

The map $T = (T_{i,s})$ is continuous and sends X into X , so by Brouwer's fixed-point theorem (Theorem 8.14) it has a fixed point $\mathbf{x}^* = T(\mathbf{x}^*)$. At a fixed point all the gains $\varphi_{i,s}(\mathbf{x}^*)$ must vanish: if some were positive, the renormalization would strictly shift weight onto the high-gain strategies and away from a strategy already in the support that earns only the average, contradicting fixedness. Zero gains everywhere is precisely the statement that no player can improve by any pure deviation, which is the Nash condition. Hence $\mathbf{x}^*$ is an equilibrium. ■

Problem 2: Approximate fixed point (programming)

Design and implement an algorithm that approximately locates a fixed point of a continuous map $f : B^2 \rightarrow B^2$ of the square, that is, a point with $\|f(z) - z\| \leq \varepsilon$.

Solution. Brouwer guarantees a fixed point but not a formula, so we hunt for one. Form the displacement field $g(z) = f(z) - z$; a zero of g is a fixed point of f . Lay an $n \times n$ grid on the square with spacing $h = 1/n$ and evaluate g at every grid point, returning immediately if some $\|g(z)\| \leq \varepsilon$. Otherwise look for a grid edge across which a component of g changes sign (a one-dimensional intermediate-value test, the discrete shadow of the topological reason a zero must exist), pick the edge minimizing $\|g_{\text{left}}\| + \|g_{\text{right}}\|$, and refine by ternary search on $\phi(t) = \|f(z(t)) - z(t)\|$ along it. If no sign change is found, halve h and repeat. For a contraction toward the centre the method walks straight in and returns the centre. The code, with $n = 100$ and $\varepsilon = 10^{-3}$, is in Appendix C.

Problem 3: NTU and TU feasible sets

For each bimatrix, describe the NTU feasible set, the TU feasible set, and the Pareto frontier: (i) $\begin{bmatrix} (0,4) & (3,2) \\ (4,0) & (2,3) \end{bmatrix}$; (ii) $\begin{bmatrix} (3,1) & (0,2) \\ (1,2) & (3,0) \end{bmatrix}$.

Solution. A cell of a bimatrix is a payoff pair (u_1, u_2) . Without transfers the players can reach any *lottery* over cells, so the **NTU feasible set** is the convex hull of the cell points. With freely transferable utility they can also slide payoff along a 45° line, so the **TU feasible set** is bounded by the line $u_1 + u_2 = V$ at the largest attainable total V . The **Pareto frontier** is the upper-right boundary of the NTU hull, the outcomes that cannot be improved for one player without hurting the other.

(i) The four cells are $(0, 4)$, $(3, 2)$, $(4, 0)$, $(2, 3)$, with totals 4, 5, 4, 5. Their hull is the quadrilateral of Figure 14.1 (left); since $(2, 3)$ lies *above* the segment from $(0, 4)$ to $(3, 2)$, all four points are corners. The Pareto frontier is the upper chain $(0, 4) \rightarrow (2, 3) \rightarrow (3, 2) \rightarrow (4, 0)$. The largest total is $V = 5$, attained at $(3, 2)$ and $(2, 3)$, so the TU frontier is $u_1 + u_2 = 5$.

(ii) The cells $(3, 1)$, $(0, 2)$, $(1, 2)$, $(3, 0)$ have totals 4, 2, 3, 3. Now two cells are dominated: $(1, 2)$ beats $(0, 2)$ (same u_2 , more u_1) and $(3, 1)$ beats $(3, 0)$ (same u_1 , more u_2). The Pareto frontier collapses to the single segment from $(1, 2)$ to $(3, 1)$, parametrized $(2\alpha + 1, 2 - \alpha)$ for $\alpha \in [0, 1]$ (Figure 14.1, right). The largest total is $V = 4$ at $(3, 1)$, so the TU frontier is $u_1 + u_2 = 4$.




Figure 14.1. Feasible sets for Problem 3. Shaded: the NTU hull of the cell payoffs. Bold blue: the Pareto frontier. Dashed red: the TU frontier $u_1 + u_2 = V$ at the largest total. In Game (ii) the dominated cells $(0, 2)$ and $(3, 0)$ sit strictly inside the frontier, so bargaining happens only along the short segment $(1, 2)$ to $(3, 1)$.

Problem 4: A bimatrix by best response

Find the pure Nash equilibria of $\begin{bmatrix} (1,5) & (2,2) & (0,1) \\ (4,2) & (1,0) & (2,1) \\ (5,0) & (2,3) & (0,0) \end{bmatrix}$.

Solution. Underline each player's best reply: Player I underlines the largest *first* coordinate in each column, Player II the largest *second* coordinate in each row. A cell underlined twice is a pure equilibrium.

	C_1	C_2	C_3
R_1	<u>(1, 5)</u>	(2, 2)	(0, 1)
R_2	(4, <u>2</u>)	(1, 0)	(2, 1)
R_3	(5, 0)	(2, <u>3</u>)	(0, 0)

Player I's marks: 5 at (R_3, C_1) , 2 at (R_1, C_2) and (R_3, C_2) , 2 at (R_2, C_3) . Player II's marks: 5 at (R_1, C_1) , 2 at (R_2, C_1) , 3 at (R_3, C_2) . The only cell carrying both underlines is (R_3, C_2) , with payoff $(2, 3)$: the unique pure Nash equilibrium. (It also survives iterated dominance: R_3 weakly dominates R_1 , and once R_1 is gone C_3 falls to C_1 , leaving R_3 against C_2 .)

Problem 5: Transferable-utility solver (programming)

Solve any two-person TU game: given (A, B) , find the cooperative value and an optimal division.

Solution. The cooperative value is the largest joint total, $\sigma = V_{\max}(A + B)$, found by maximizing over the combined matrix $C = A + B$. To split it, score each player's bargaining strength by the *threat game* $A - B$: its value $\delta = V(A - B)$ measures how much Player I can swing the difference. The TU division is then $\varphi = \left(\frac{\sigma + \delta}{2}, \frac{\sigma - \delta}{2}\right)$ as in [Theorem 9.5](#): split the pie evenly, then shade it by the threat advantage. The implementation reuses the zero-sum solver of Problem Set 2 (once for C , once for $A - B$) and appears in [Appendix C](#).

Problem 6: Bankruptcy and the nucleolus

Express the bankruptcy problem, with estate E and claims $c_1, \dots, c_n$, as a coalitional game, and state its solution as a nucleolus. Work the claims $(100, 200, 300)$.

Solution. A coalition can count only on what the estate leaves after the outside claims are paid, so set

$$v(S) = \max\left\{0, E - \sum_{i \notin S} c_i\right\}$$

as in (10.3). The nucleolus of this game is the bankruptcy division. With claims $(100, 200, 300)$ we solve all three classical estates by balancing the two coalitions whose excesses move oppositely with x_1 .

Estate $E = 200$. Every singleton excludes two claims already summing past 200, so $v(\{i\}) = 0$; the only positive worth is $v(\{2, 3\}) = \max\{0, 200 - 100\} = 100$, and $v(N) = 200$. Using efficiency $x_1 + x_2 + x_3 = 200$, the complaint of $\{1\}$ is $v(\{1\}) - x_1 = -x_1$ and the complaint of $\{2, 3\}$ is $100 - (x_2 + x_3) = 100 - (200 - x_1) = x_1 - 100$. Balancing, $-x_1 = x_1 - 100$ gives $x_1 = 50$, and the symmetry of claimants 2, 3 forces $x_2 = x_3 = 75$. The division is $(50, 75, 75)$.

Estate $E = 100$. Now even the pair $\{2, 3\}$ is worth $\max\{0, 100 - 100\} = 0$, and every proper coalition has worth 0. The game is symmetric, so the nucleolus splits equally: $(33\frac{1}{3}, 33\frac{1}{3}, 33\frac{1}{3})$.

Estate $E = 300$. Here $v(\{1, 3\}) = \max\{0, 300 - 200\} = 100$, $v(\{2, 3\}) = \max\{0, 300 - 100\} = 200$, $v(\{3\}) = \max\{0, 300 - 300\} = 0$, and other singletons and $\{1, 2\}$ are 0. Balancing $\{1\}$ (excess $-x_1$) against $\{2, 3\}$ (excess $200 - (300 - x_1) = x_1 - 100$) gives $x_1 = 50$; then balancing $\{2\}$ (excess $-x_2$) against $\{1, 3\}$ (excess $100 - (50 + x_3) = x_2 - 200$, using $x_2 + x_3 = 250$) gives $x_2 = 100$, hence $x_3 = 150$. The division is $(50, 100, 150)$.

estate E	claimant 1	claimant 2	claimant 3
100	$33\frac{1}{3}$	$33\frac{1}{3}$	$33\frac{1}{3}$
200	50	75	75
300	50	100	150

These three rows are exactly the Talmud's rulings for an estate owed $(100, 200, 300)$, the two-thousand-year-old table that Aumann and Maschler [\[1\]](#) identified as the nucleolus and that closes [Chapter 10](#).

PART VI

Appendices

APPENDIX A

Mathematical Review

This appendix collects the mathematics the main text assumes. It is a reference rather than a course: each notion is stated precisely, with just enough detail to make the chapters self-contained. The headings follow the order in which the ideas are first needed.

A.1 Sets, functions, and the simplex

For a finite set S , the number of elements is $|S|$, and 2^S denotes the collection of all subsets of S , of which there are $2^{|S|}$. A function $f : X \rightarrow Y$ is **injective** if distinct inputs give distinct outputs, **surjective** if every $y \in Y$ is some $f(x)$, and **bijective** if both; a bijection of a finite set onto itself is a **permutation**.

The single most-used set in the book is the probability simplex.

Definition A.1 (Probability simplex). The **simplex** in $\mathbb{R}^n$ is

$$\Delta_n = \left\{ \mathbf{p} = (p_1, \dots, p_n) \in \mathbb{R}^n : p_i \geq 0, \sum_{i=1}^n p_i = 1 \right\},$$

the set of probability distributions over n outcomes. Its **vertices** are the unit vectors $\mathbf{e}_i$ (the pure strategies), and it is the convex hull of those vertices.

For $n = 2$ the simplex Δ_2 is the segment from $(1, 0)$ to $(0, 1)$, parametrized by the single number $p_1 = p \in [0, 1]$; this is why a two-strategy mixed strategy is just “play row 1 with probability p .” For $n = 3$ it is a triangle (Figure A.1), with the three pure strategies at the corners and every mixed strategy inside. For instance $\mathbf{p} = (\frac{1}{2}, \frac{1}{3}, \frac{1}{6})$ is a valid point of Δ_3 because its entries are nonnegative and $\frac{1}{2} + \frac{1}{3} + \frac{1}{6} = \frac{3+2+1}{6} = 1$; it sits inside the triangle, closest to the corner $\mathbf{e}_1$ it weights most.




Figure A.1. The simplex Δ_3 of distributions over three outcomes: a triangle whose corners are the pure strategies $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$. A mixed strategy $\mathbf{p}$ is the convex combination $p_1\mathbf{e}_1 + p_2\mathbf{e}_2 + p_3\mathbf{e}_3$, here pulled toward $\mathbf{e}_1$ because $p_1 = \frac{1}{2}$ is the largest weight.

A.2 Linear algebra

Vectors are columns in $\mathbb{R}^n$; $\mathbf{x}^\top$ is the row transpose, and $\mathbf{x}^\top \mathbf{y} = \sum_i x_i y_i$ is the inner product. A matrix $A \in \mathbb{R}^{m \times n}$ acts by $(A\mathbf{y})_i = \sum_j a_{ij} y_j$, and $(A^\top)_{ij} = a_{ji}$. The **bilinear form** $\mathbf{x}^\top A \mathbf{y} = \sum_{i,j} x_i a_{ij} y_j$ is linear in $\mathbf{x}$ for fixed $\mathbf{y}$ and linear in $\mathbf{y}$ for fixed $\mathbf{x}$; it is the expected payoff of Chapter 4.

For a square matrix A , the **determinant** $\det A$ is nonzero exactly when A is **invertible**, in which case A^{-1} satisfies $AA^{-1} = A^{-1}A = I$. For the 2×2 case,

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}, \quad \det A = ad - bc, \quad A^{-1} = \frac{1}{ad - bc} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix},$$

the formulas behind the 2×2 game solution of [Theorem 4.8](#) and the inversion method of [Proposition 5.3](#). For a worked instance, take $A = \begin{bmatrix} 2 & 1 \\ 1 & 3 \end{bmatrix}$: then $\det A = 2 \cdot 3 - 1 \cdot 1 = 5$, so

$$A^{-1} = \frac{1}{5} \begin{bmatrix} 3 & -1 \\ -1 & 2 \end{bmatrix}, \quad \text{and} \quad AA^{-1} = \frac{1}{5} \begin{bmatrix} 2 & 1 \\ 1 & 3 \end{bmatrix} \begin{bmatrix} 3 & -1 \\ -1 & 2 \end{bmatrix} = \frac{1}{5} \begin{bmatrix} 5 & 0 \\ 0 & 5 \end{bmatrix} = I,$$

confirming the inverse. The **all-ones vector** $\mathbf{1}$ has every entry 1, so $\mathbf{1}^\top \mathbf{x} = \sum_i x_i$.

A.3 Convexity

Convexity underlies the existence theorems.

Definition A.2 (Convex set, convex hull). A set $C \subseteq \mathbb{R}^n$ is **convex** if it contains the segment between any two of its points: $\mathbf{x}, \mathbf{y} \in C$ and $\lambda \in [0, 1]$ imply $\lambda \mathbf{x} + (1 - \lambda) \mathbf{y} \in C$. The **convex hull** $\text{conv}(X)$ of a set X is the smallest convex set containing X , equivalently the set of all finite convex combinations $\sum_k \lambda_k \mathbf{x}_k$ with $\mathbf{x}_k \in X$, $\lambda_k \geq 0$, $\sum_k \lambda_k = 1$.

For a concrete convex combination, the three points $(0, 0), (2, 0), (0, 3)$ in $\mathbb{R}^2$ have convex hull the filled triangle they span; the combination $\frac{1}{2}(2, 0) + \frac{1}{2}(0, 3) = (1, \frac{3}{2})$ is the midpoint of one edge, and $\frac{1}{3}(0, 0) + \frac{1}{3}(2, 0) + \frac{1}{3}(0, 3) = (\frac{2}{3}, 1)$ is the centroid, both inside the hull because their weights are nonnegative and sum to 1.

The simplex ([Definition A.1](#)) is convex and compact, and so is any product of simplices, a fact the proof of Nash's theorem ([Theorem 8.15](#)) needs. The feasible sets of [Chapter 9](#) are convex hulls of finitely many payoff points, hence convex polytopes.

Definition A.3 (Convex and concave functions). A function f on a convex set is **convex** if $f(\lambda \mathbf{x} + (1 - \lambda) \mathbf{y}) \leq \lambda f(\mathbf{x}) + (1 - \lambda) f(\mathbf{y})$ for all $\lambda \in [0, 1]$, and **concave** if $-f$ is convex. A maximum (minimum) of linear functions is convex (concave); this is why the upper envelope in the graphical method of [Section 4.6](#) is convex and the lower envelope concave.

Theorem A.4 (Separating hyperplane). If C and D are disjoint nonempty convex sets in $\mathbb{R}^n$ and C is compact and D closed, then there is a vector $\mathbf{a}$ and a scalar b with $\mathbf{a}^\top \mathbf{x} \leq b$ for all $\mathbf{x} \in C$ and $\mathbf{a}^\top \mathbf{y} \geq b$ for all $\mathbf{y} \in D$. A hyperplane separates the two sets.

The separating hyperplane theorem, equivalently *Farkas's lemma* on systems of linear inequalities, is the engine behind linear programming duality ([Theorem 6.7](#)) and hence behind the minimax theorem.

A.4 Continuity and compactness

A set in $\mathbb{R}^n$ is **compact** if it is closed and bounded. The relevant consequence is the guarantee that optima are attained.

Theorem A.5 (Weierstrass extreme value theorem). A continuous real-valued function on a nonempty compact set attains a maximum and a minimum.

Every existence statement in the book reduces to this: the simplex is compact, expected payoff is continuous, so a best response exists ([Lemma 5.12](#)), the value of a game is attained ([Theorem 4.7](#)), and the maps in the fixed-point proofs are continuous on compact domains. The hypothesis cannot be dropped: $f(x) = x$ on the open interval $(0, 1)$ has supremum 1 but no maximum, the same failure that makes compactness essential in Brouwer's theorem.

A.5 Probability and expectation

A **probability distribution** over a finite set $\{1, \dots, n\}$ is a point of the simplex Δ_n . The **expectation** of a quantity g that takes value $g(k)$ with probability p_k is $\mathbb{E}[g] = \sum_k p_k g(k)$. For example, a payoff that is 3 with probability $\frac{1}{4}$, -1 with probability $\frac{1}{2}$, and 2 with probability $\frac{1}{4}$ has expectation $\frac{1}{4} \cdot 3 + \frac{1}{2} \cdot (-1) + \frac{1}{4} \cdot 2 = \frac{3}{4} - \frac{1}{2} + \frac{1}{2} = \frac{3}{4}$. When two players randomize independently with $\mathbf{p}$ and $\mathbf{q}$, the probability of the joint pure outcome (i, j) is $p_i q_j$, and the expected payoff is the bilinear form $\sum_{i,j} p_i a_{ij} q_j = \mathbf{p}^\top \mathbf{A} \mathbf{q}$. Expectation is **linear**: a convex combination of strategies yields the same convex combination of expected payoffs, which is why a best response can always be taken pure (a linear function on the simplex is maximized at a vertex).

A.6 Fixed-point theorems

Two fixed-point theorems are used or cited. Both are stated in [Chapter 8](#); we collect them here for reference.

Theorem A.6 (Brouwer). A continuous map $f : D \rightarrow D$ on a nonempty compact convex set $D \subseteq \mathbb{R}^n$ has a fixed point x^* with $f(x^*) = x^*$.

Theorem A.7 (Kakutani). Let $D \subseteq \mathbb{R}^n$ be nonempty, compact, and convex, and let F assign to each $x \in D$ a nonempty convex subset $F(x) \subseteq D$, with closed graph (if $x_k \rightarrow x$, $y_k \rightarrow y$, and $y_k \in F(x_k)$, then $y \in F(x)$). Then F has a fixed point $x^* \in F(x^*)$.

Brouwer's theorem proves Nash's existence theorem through the single-valued construction of [Theorem 8.15](#). Kakutani's theorem, a set-valued generalization, gives the same conclusion applied directly to the best-response correspondence, which is the route Nash took in his original announcement.

A.7 Asymptotic notation and binary arithmetic

Two small tools appear in the combinatorial and computational parts. The **nim-sum** $a \oplus b$ is the bitwise exclusive-or of the binary representations of a and b , that is, addition in each binary place modulo 2 with no carrying ([Definition 2.6](#)); it is associative and commutative with $a \oplus a = 0$. For algorithms, $O(g(n))$ denotes a quantity bounded by a constant multiple of $g(n)$ for large n , the standard measure of how the cost of a computation grows with its input.

APPENDIX B

Examination Archive

This appendix reproduces the in-class midterm and a full practice final, with worked solutions. The problems test the same methods as the chapters, applied without hints, so the solutions are brief and point to the relevant results. Several original solution keys contained arithmetic slips; the corrected value is given and the discrepancy is noted, in keeping with the book's policy of teaching the answer and the check together.

B.1 Midterm examination

The midterm had six problems: a set of true/false statements, four small zero-sum games, the Grundy analysis of three impartial games, a Nim and take-or-break problem, four matrix games, and an extensive-form game.

Problem 1: True or false, with justification

Problem. Decide each statement, with a one-line reason.

Solution.

1. *The value of a finite two-person zero-sum game is unique.* **True:** the minimax theorem (Theorem 4.7) gives a single number $V = \max_{\mathbf{p}} \min_{\mathbf{q}} \mathbf{p}^\top A \mathbf{q}$.
2. *The optimal strategies are unique.* **False:** in the all-ones matrix every entry is 1, so $V = 1$ and *every* mixed strategy is optimal.
3. *A skew-symmetric game ($A^\top = -A$) lets either player avoid an expected loss.* **True:** its value is 0 (Theorem 5.7), so each player can guarantee at least 0.
4. *For an invertible $n \times n$ game with $V \neq 0$, the row optimum is $\mathbf{p}^{*\top} = V \mathbf{1}^\top A^{-1}$.* **True:** this is the inversion formula (Proposition 5.3), valid exactly when A^{-1} exists and the solution is fully mixed.
5. *A best response to a known mixed strategy can always be taken pure.* **True:** the payoff is linear in one's own mix, so a vertex (a pure strategy) attains the maximum (Lemma 5.12).
6. *A game can have both a pure and a genuinely mixed equilibrium.* **True:** a 2×2 coordination game has two pure equilibria and one mixed one.
7. *From an N-position some move reaches a P-position; from a P-position every move reaches an N-position.* **True** on both counts: this is the definition of P- and N-positions (Definition 2.2).
8. *In an invariant game, some optimal strategies are invariant but others need not be.* **True:** orbit-averaging produces an invariant optimum (Theorem 5.10), yet a non-symmetric optimum can also exist.

Problem 2: Small zero-sum games

(a) *Guess the number.* Player I hides $i \in \{1, 2, 3\}$, Player II guesses j ; a miss costs Player II $|i - j|$, so Player I (who wants the miss large) is the row maximizer with matrix $M = \begin{bmatrix} 0 & 1 & 2 \\ 1 & 0 & 1 \\ 2 & 1 & 0 \end{bmatrix}$ (entry $|i - j|$). The row minima are 0, 0, 0, so the maximin is 0; the column maxima are 2, 1, 2, so the minimax is 1; since $0 \neq 1$ there is no saddle. Try Player I's mix $\mathbf{p}^* = (\frac{1}{2}, 0, \frac{1}{2})$: against column 1 it pays $\frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 2 = 1$, against column 2 it pays $\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 1 = 1$, and against column 3 it pays $\frac{1}{2} \cdot 2 + \frac{1}{2} \cdot 0 = 1$, so Player I guarantees $V = 1$. Player II's optimal strategy is not unique: any $\mathbf{q} = (q, 1 - 2q, q)$ with $0 \leq q \leq \frac{1}{2}$ holds Player I to 1, because then rows 1 and 3 each cost 1 and row 2 costs $2q \leq 1$.

(b) *Weighted Rock–Paper–Scissors.* Rock beats scissors by 3, paper beats rock by 2, scissors beats paper by 1. Ordering the strategies (rock, paper, scissors), the skew-symmetric matrix is $M = \begin{bmatrix} 0 & -2 & 3 \\ 2 & 0 & -1 \\ -3 & 1 & 0 \end{bmatrix}$, so by Theorem 5.7 the value is $V = 0$ and $\mathbf{p}^* = \mathbf{q}^*$. Equalize each row against $\mathbf{q}^* = (q_1, q_2, q_3)$ to 0: row 1 gives $-2q_2 + 3q_3 = 0$, row 2 gives $2q_1 - q_3 = 0$, and row 3 gives $-3q_1 + q_2 = 0$. From the second, $q_3 = 2q_1$; from the third, $q_2 = 3q_1$; the first is then satisfied ($-6q_1 + 6q_1 = 0$). Normalizing $q_1 + 3q_1 + 2q_1 = 6q_1 = 1$ gives $\mathbf{p}^* = \mathbf{q}^* = (\frac{1}{6}, \frac{1}{2}, \frac{1}{3})$.

(c) *The five-number game.* Players pick $i, j \in \{1, \dots, 5\}$; choosing exactly one more than the opponent wins 2, choosing two or more more loses 1, and a tie scores 0. Reading off the payoff to the row player gives the skew-symmetric matrix

$$M = \begin{bmatrix} 0 & -2 & 1 & 1 & 1 \\ 2 & 0 & -2 & 1 & 1 \\ -1 & 2 & 0 & -2 & 1 \\ -1 & -1 & 2 & 0 & -2 \\ -1 & -1 & -1 & 2 & 0 \end{bmatrix}$$

(each row i : a +2 just left of the diagonal, a -2 just right, -1 further left, +1 further right). Since $M^T = -M$, the value is $V = 0$ and $\mathbf{p}^* = \mathbf{q}^*$ (Theorem 5.7). The symmetric optimum, weighting the middle numbers heavily and the extremes lightly, is

$$\mathbf{p}^* = \mathbf{q}^* = \left(\frac{1}{16}, \frac{5}{16}, \frac{1}{4}, \frac{5}{16}, \frac{1}{16} \right).$$

Check: playing $\mathbf{p}^*$ against each column pays 0 in every column, so it guarantees the value 0; by symmetry the same holds for Player II. *Correction.* The original key printed a vector summing to $\frac{19}{16}$; the probabilities above are the corrected ones, and they sum to 1.

(d) *Colonel Blotto, three against two.* Colonel Blotto has 3 units and Lieutenant Kije has 2, and each spreads them over two posts. With the exam's scoring the payoffs to Blotto (rows his splits (3, 0), (2, 1), (1, 2), (0, 3); columns Kije's (2, 0), (1, 1), (0, 2)) are

$$A = \begin{bmatrix} 3 & 1 & 0 \\ 1 & 2 & -1 \\ -1 & 2 & 1 \\ 0 & 1 & 3 \end{bmatrix}.$$

Reflecting the two posts is a symmetry of the game: it swaps Blotto's (3, 0) $\leftrightarrow$ (0, 3) and (2, 1) $\leftrightarrow$ (1, 2) and Kije's (2, 0) $\leftrightarrow$ (0, 2), while fixing (1, 1). Playing each swapped pair with equal weight collapses the contest to the invariant 2×2

$$\tilde{A} = \begin{bmatrix} \frac{3+0}{2} & \frac{1+1}{2} \\ \frac{1-1}{2} & \frac{2+2}{2} \end{bmatrix} = \begin{bmatrix} \frac{3}{2} & 1 \\ 0 & 2 \end{bmatrix}, \quad D = \frac{3}{2} + 2 - 1 - 0 = \frac{5}{2}, \quad V = \frac{3-0}{5/2} = \frac{6}{5}.$$

The reduced game puts weight $p^* = \frac{d-c}{D} = \frac{2}{5/2} = \frac{4}{5}$ on Blotto's symmetric pair and $q^* = \frac{d-b}{D} = \frac{1}{5/2} = \frac{2}{5}$ on Kije's. Lifting back, Blotto plays (3, 0) and (0, 3) each with probability $\frac{2}{5}$ and (2, 1), (1, 2) each with $\frac{1}{10}$; Kije plays (2, 0), (0, 2) each with $\frac{1}{5}$ and (1, 1) with $\frac{3}{5}$. The value is $V = \frac{6}{5}$.

Problem 3: Sprague–Grundy values

Problem. Compute the stated Grundy values for three impartial games, then analyze the composite position with one pile of each, of sizes (5, 7, 9).

Solution. *Lasker's Nim* (remove any number, or split a heap): its Grundy values equal n except that $4k+3$ and $4k+4$ swap, $g = 0, 1, 2, 4, 3, 5, 6, 8, 7, \dots$, so $g(5) = 5$. *Subtraction* $S = \{1, 4\}$: build $g(n) = \text{mex}\{g(n-1), g(n-4)\}$ from the bottom,

n	0	1	2	3	4	5	6	7	8
$g(n)$	0	1	0	1	2	0	1	0	1

so $g(7) = 0$. *Divisor-removal* (remove a divisor of the pile): $g(n) = \nu_2(n) + 1$, where $\nu_2(n)$ is the power of two dividing n ; since 9 is odd, $\nu_2(9) = 0$ and $g(9) = 1$.

The composite (5, 7, 9). By the Sprague–Grundy theorem the combined value is the nim-sum

$$g = 5 \oplus 0 \oplus 1 = (101)_2 \oplus (000)_2 \oplus (001)_2 = (100)_2 = 4 \neq 0,$$

an N-position. The leftmost set bit of 4 (the 4s bit) is carried only by the Lasker pile, whose value 5 must drop to $5 \oplus 4 = 1$. A Lasker heap of Grundy value 1 is a heap of size 1 (since $g(1) = 1$), so the winning move takes the Lasker pile from 5 down to 1. After it the three Grundy values 1, 0, 1 nim-sum to 0.

Problem 4: Nim and take-or-break

Problem. Solve Nim with heaps (9, 17, 21), and describe the Grundy structure of the take-or-break game (remove one or two chips, or split a heap into two nonempty heaps).

Solution. Stack the heaps in binary and add the columns modulo 2:

$$\begin{array}{r} 16 \ 8 \ 4 \ 2 \ 1 \\ \hline 9 = 0 \ 1 \ 0 \ 0 \ 1 \\ 17 = 1 \ 0 \ 0 \ 0 \ 1 \\ 21 = 1 \ 0 \ 1 \ 0 \ 1 \\ \hline S = 0 \ 1 \ 1 \ 0 \ 1 \end{array}$$

The nim-sum is $S = 9 \oplus 17 \oplus 21 = (01101)_2 = 13 \neq 0$, an N-position. The leftmost unbalanced column is the 8s bit, carried only by the heap 9, whose target is $9 \oplus 13 = (01001)_2 \oplus (01101)_2 = (00100)_2 = 4$. So the unique winning move reduces the heap of 9 to 4 (removing 5); then $4 \oplus 17 \oplus 21 = 0$. *Take-or-break.* Adding the splitting move to a subtraction game makes a heap of size n able to move either to $n - 1$, $n - 2$ or to a disjunctive sum $i + (n - i)$, so $g(n) = \text{mex}(\{g(n - 1), g(n - 2)\} \cup \{g(i) \oplus g(n - i)\})$; the resulting Grundy sequence is eventually periodic with period 4.

Problem 5: Matrix games

Problem. Solve each game (Player I maximizes over rows, Player II minimizes over columns). (a) $\begin{bmatrix} 4 & 4 & 1 & 0 \\ 4 & 3 & 2 & -1 \\ 1 & -1 & 4 & 3 \\ 0 & -3 & 1 & 2 \end{bmatrix}$, (b)

$$\begin{bmatrix} 1 & 3 & 5 \\ 5 & 1 & 3 \\ 3 & 5 & 1 \end{bmatrix}, \quad \text{(c)} \begin{bmatrix} -2 & 2 & -1 \\ 5 & 3 & 1 \\ 2 & 0 & 1 \end{bmatrix}, \quad \text{(d)} \begin{bmatrix} 0 & -1 & 2 \\ -1 & 0 & -1 \\ 2 & -1 & 0 \end{bmatrix}.$$

Solution. (a) **Support reduction to a 2×2 .** Solving the 4×4 game, the optimal play falls on rows $\{1, 3\}$ and columns $\{2, 4\}$; on that support the induced game is

$$\tilde{A} = \begin{bmatrix} 4 & 0 \\ -1 & 3 \end{bmatrix}, \quad D = 4 - 0 - (-1) + 3 = 8, \quad V = \frac{4 \cdot 3 - 0 \cdot (-1)}{8} = \frac{12}{8} = \frac{3}{2},$$

with $p^* = \frac{d-c}{D} = \frac{4}{8} = \frac{1}{2}$ and $q^* = \frac{d-b}{D} = \frac{3}{8}$. Lifting back, $p^* = (\frac{1}{2}, 0, \frac{1}{2}, 0)$ and $q^* = (0, \frac{3}{8}, 0, \frac{5}{8})$. *Check:* against p^* the four columns pay $\frac{5}{2}, \frac{3}{2}, \frac{5}{2}, \frac{3}{2}$ (minimum $\frac{3}{2}$); against q^* the four rows pay $\frac{3}{2}, \frac{1}{2}, \frac{3}{2}, \frac{1}{8}$ (maximum $\frac{3}{2}$); the two bounds meet at $V = \frac{3}{2}$.

(b) **A circulant.** Each row is the one above cycled one place, so the game is unchanged by a cyclic relabelling of the strategies and the uniform mix is optimal. Every column sums to $1 + 5 + 3 = 9$, so against $p^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ each column costs $\frac{9}{3} = 3$. Hence $V = 3$ with $p^* = q^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.

(c) **A saddle.** The row minima are $-2, 1, 0$ (maximin 1, the second row) and the column maxima are $5, 3, 1$ (minimax 1, the third column); they agree at the entry $a_{23} = \boxed{1}$, a saddle. So $V = 1$ with the pure pair $p^* = (0, 1, 0)$, $q^* = (0, 0, 1)$.

(d) **An even-odd symmetric game.** Here $M^\top = M$, so $p^* = q^*$; equalizing the columns (solving $p^{*\top} M = V \mathbf{1}^\top$ with $\sum_i p_i = 1$) gives $p^* = q^* = (\frac{1}{6}, \frac{2}{3}, \frac{1}{6})$. *Check:* playing p^* holds every column to $-\frac{1}{3}$ (column 1: $\frac{1}{6}(0) + \frac{2}{3}(-1) + \frac{1}{6}(2) = -\frac{1}{3}$; column 2: $\frac{1}{6}(-1) + \frac{2}{3}(0) + \frac{1}{6}(-1) = -\frac{1}{3}$; column 3: $\frac{1}{6}(2) + \frac{2}{3}(-1) + \frac{1}{6}(0) = -\frac{1}{3}$), so $V = -\frac{1}{3}$ to the row maximizer (and $+\frac{1}{3}$ with the roles reversed).

Problem 6: An extensive-form game

Problem. The midterm's Kuhn tree (Figure B.1) is a two-player zero-sum game in which Nature moves first, taking each of its two branches with probability $\frac{1}{2}$, after which Players I and II move and Player II cannot tell which branch Nature chose. Player I has four pure strategies AC, AD, BC, BD and Player II has two, a and b ; all payoffs are to Player I. Find the value and optimal strategies.

Solution. Average each strategy's two leaf payoffs over Nature's fair coin to read off the strategic form. For instance $u_1(BC, b) = \frac{1}{2}(8) + \frac{1}{2}(2) = 5$; doing this for all eight cells gives the 4×2 game

	a	b
AC	5	2
AD	4	3
BC	3	5
BD	3	5

Rows BC and BD coincide, and row AD is never a best reply, so the contest is between AC and BC . If Player II plays column a with probability q , those two rows pay $2 + 3q$ (row AC) and $5 - 2q$ (row BC); they cross at $q = \frac{3}{5}$, where both equal $\frac{19}{5}$ while AD trails at $\frac{18}{5}$. So $\mathbf{q}^* = (\frac{3}{5}, \frac{2}{5})$. Matching on the other side, Player I plays AC with probability p and BC with $1 - p$; the columns then pay $3 + 2p$ and $5 - 3p$, equal at $p = \frac{2}{5}$. Hence

$$V = \frac{19}{5}, \quad \mathbf{p}^* = (\frac{2}{5}, 0, \frac{3}{5}, 0), \quad \mathbf{q}^* = (\frac{3}{5}, \frac{2}{5}),$$

and both securities check: $\mathbf{p}^*$ makes each column pay exactly $\frac{19}{5}$, and $\mathbf{q}^*$ holds every row to at most $\frac{19}{5}$.




Figure B.1. The midterm's Kuhn tree. Nature opens with a fair coin, sending play into one of two subgames in which Players I and II move; Player II acts without learning Nature's choice, so the game has imperfect information. Averaging each leaf pair over Nature's coin collapses the tree to the 4×2 strategic form solved above, whose value is $V = \frac{19}{5}$.

B.2 Practice final examination

The practice final had eight problems spanning the whole course. Each is restated and worked in full below, with the method named and the arithmetic carried to a checked answer.

Problem 1: Impartial games

Problem. Three Nim heaps hold 13, 17, and 26 chips; the player taking the last chip wins. Determine who wins under optimal play and give a winning first move. State briefly how you would handle (b) a game given as a directed acyclic graph and (c) a “remove or split a heap” game.

Solution. (a) Stack the heaps in binary and add each column modulo 2:

$$\begin{array}{r}
 16 \ 8 \ 4 \ 2 \ 1 \\
 13 = 0 \ 1 \ 1 \ 0 \ 1 \\
 17 = 1 \ 0 \ 0 \ 0 \ 1 \\
 26 = 1 \ 1 \ 0 \ 1 \ 0 \\
 \hline
 S = 0 \ 0 \ 1 \ 1 \ 0
 \end{array}$$

The nim-sum is $S = 13 \oplus 17 \oplus 26 = (00110)_2 = 6 \neq 0$, so the position is an N-position and the player to move wins. The leftmost unbalanced column is the 4s bit, carried only by the heap 13 (the heaps 17 and 26 have a 0 there). Its

target is $13 \oplus S = (01101)_2 \oplus (00110)_2 = (01011)_2 = 11$, so the unique winning move reduces the heap of 13 to 11, removing 2 chips. Check: $11 \oplus 17 \oplus 26 = 0$, a P-position handed to the opponent.

(b) A game presented as a directed acyclic graph is solved by labelling every vertex with its Grundy value $g(x) = \text{mex}\{g(y) : x \rightarrow y\}$, working in reverse topological order from the sinks (which get $g = 0$); a vertex is a P-position iff $g(x) = 0$, as in Chapter 11. (c) A “remove any number, or split a heap into two” game is Lasker’s Nim: evaluate it by the take-and-break recurrence of Section 2.6, in which a split contributes the *nim-sum* of the two parts’ Grundy values.

Problem 2: Zero-sum matrix games

Problem. Find the value and optimal strategies of each game (Player I chooses a row and maximizes): (a) $\begin{bmatrix} 6 & 5 & 7 \\ 3 & 4 & 2 \\ 1 & 5 & 0 \end{bmatrix}$, (b) $\begin{bmatrix} 2 & 5 \\ 6 & 3 \end{bmatrix}$, (c) $\begin{bmatrix} 0 & 2 & 5 \\ 4 & 3 & 1 \end{bmatrix}$.

Solution. (a) **Saddle point.** Border the matrix with row minima and column maxima:

	c_1	c_2	c_3	row min
r_1	6	5	7	5
r_2	3	4	2	2
r_3	1	5	0	1
col max	6	5	7	

The maximin $\max(5, 2, 1) = 5$ and the minimax $\min(6, 5, 7) = 5$ agree, so $a_{12} = 5$ is a saddle: $V = 5$, with $\mathbf{p}^* = (1, 0, 0)$ and $\mathbf{q}^* = (0, 1, 0)$.

(b) **The 2×2 formula.** Row minima 2, 3 give maximin 3; column maxima 6, 5 give minimax 5; since $3 \neq 5$ the game is mixed. With $a, b, c, d = 2, 5, 6, 3$ and $D = a - b - c + d = -6$,

$$p^* = \frac{d - c}{D} = \frac{3 - 6}{-6} = \frac{1}{2}, \quad q^* = \frac{d - b}{D} = \frac{3 - 5}{-6} = \frac{1}{3}, \quad V = \frac{ad - bc}{D} = \frac{6 - 30}{-6} = 4.$$

So $\mathbf{p}^* = (\frac{1}{2}, \frac{1}{2})$, $\mathbf{q}^* = (\frac{1}{3}, \frac{2}{3})$, $V = 4$. Check: column 1 costs $\frac{1}{2}(2) + \frac{1}{2}(6) = 4$ and column 2 costs $\frac{1}{2}(5) + \frac{1}{2}(3) = 4$, both equal to V .

(c) **The graphical method.** Player I mixes $(p, 1 - p)$ over the two rows; against each column the payoff is a line in p :

$$U_{c_1}(p) = 4 - 4p, \quad U_{c_2}(p) = 3 - p, \quad U_{c_3}(p) = 1 + 4p.$$

Player II forces Player I down to the lower envelope, and Player I picks its highest point (Figure B.2). The falling line U_{c_1} meets the rising line U_{c_3} at $4 - 4p = 1 + 4p$, i.e. $p^* = \frac{3}{8}$, where $V = 1 + 4 \cdot \frac{3}{8} = \frac{5}{2}$. The middle line there is $U_{c_2}(\frac{3}{8}) = \frac{21}{8} > \frac{5}{2}$, so column 2 is slack and $q_2^* = 0$; solving the remaining 2×2 on columns 1, 3 gives $\mathbf{q}^* = (\frac{1}{2}, 0, \frac{1}{2})$. Thus $\mathbf{p}^* = (\frac{3}{8}, \frac{5}{8})$, $\mathbf{q}^* = (\frac{1}{2}, 0, \frac{1}{2})$, $V = \frac{5}{2}$.




Figure B.2. Graphical solution of game (c). Player II holds Player I to the lower envelope (thick red); its peak at $p^* = \frac{3}{8}$, where columns 1 and 3 cross, gives the value $V = \frac{5}{2}$. Column 2 never touches the envelope, so Player II does not use it.

Problem 3: Special zero-sum games

Problem. (a) Solve the symmetric game $A = \begin{bmatrix} 0 & 1 & -2 \\ 1 & 0 & 1 \\ -2 & 1 & 0 \end{bmatrix}$. (b) Two players each name a positive integer and the larger name wins a fixed prize; what is the best a player can guarantee?

Solution. (a) The matrix is symmetric ($A = A^\top$), so the two players face the same problem and $\mathbf{p}^* = \mathbf{q}^*$. We look for a fully mixed $\mathbf{p}^* = (p_1, p_2, p_3)$ making every row earn the common value V , that is $A\mathbf{p}^* = V\mathbf{1}$:

$$\begin{aligned} \text{row 1: } & 0p_1 + 1p_2 - 2p_3 = V, \\ \text{row 2: } & 1p_1 + 0p_2 + 1p_3 = V, \\ \text{row 3: } & -2p_1 + 1p_2 + 0p_3 = V. \end{aligned}$$

Subtracting row 3 from row 1 gives $2p_1 - 2p_3 = 0$, so $p_1 = p_3$; the symmetry of the problem is already visible. Row 2 then reads $2p_1 + 0 = V$, so $p_2 = 1 - 2p_1$ from normalization. Substituting $p_1 = p_3$ into row 1, $p_2 - 2p_1 = V = 2p_1$, so $p_2 = 4p_1$. With $p_1 + p_2 + p_3 = 2p_1 + 4p_1 = 6p_1 = 1$, we get $p_1 = \frac{1}{6}$, hence

$$\mathbf{p}^* = \mathbf{q}^* = \left(\frac{1}{6}, \frac{2}{3}, \frac{1}{6}\right), \quad V = 2p_1 = \frac{1}{3}.$$

Check each row against $\mathbf{p}^*$: row 1 earns $\frac{2}{3} - 2 \cdot \frac{1}{6} = \frac{1}{3}$, row 2 earns $\frac{1}{6} + \frac{1}{6} = \frac{1}{3}$, row 3 earns $-2 \cdot \frac{1}{6} + \frac{2}{3} = \frac{1}{3}$, all equal to $V = \frac{1}{3}$, and the weights are positive, confirming the solution.

(b) There is no optimal strategy: whatever integer n you commit to, the opponent can name $n + 1$ and beat you. The supremum of what you can guarantee is therefore the losing payoff, and the “least bad” deterministic choice is to name the smallest legal integer, 1, which guarantees a payoff of -1 (you lose unless the opponent also names 1). The game has no value among pure strategies because the strategy space is unbounded, the standard cautionary example against dropping compactness.

Problem 4: A card game in extensive form

Problem. Nature deals Player I a card that is HIGH or LOW, each with probability $\frac{1}{2}$. Player I sees it and either BETS or CHECKS. On a check there is a showdown worth $+1$ to Player I if the card is high and -1 if low. On a bet, Player II (who has *not* seen the card) either CALLS or FOLDS: a fold pays Player I $+1$; a call pays ± 2 according to the card. Find the value of the game and the optimal strategies (Figure B.3).

Solution. With a high card, betting weakly dominates checking ($+2$ or $+1$ versus $+1$), so Player I always bets the high card. The only real choices are how often Player I *bluffs* the low card (bet with probability β) and how often Player II *calls* a bet (probability c). Player I's expected payoff is

$$f(\beta, c) = \underbrace{\frac{1}{2}(2c + (1 - c))}_{\text{high: bet}} + \underbrace{\frac{1}{2}(\beta(-2c + (1 - c)) + (1 - \beta)(-1))}_{\text{low: bluff}} = \frac{c}{2} + \frac{\beta(2 - 3c)}{2}.$$

Player I maximizes over β , Player II minimizes over c , and at an interior equilibrium each makes the other indifferent. Player I is indifferent when the coefficient of β vanishes, $2 - 3c = 0$, so $c^* = \frac{2}{3}$; Player II is indifferent when the coefficient of c in $f = \beta + c \cdot \frac{1 - 3\beta}{2}$ vanishes, $1 - 3\beta = 0$, so $\beta^* = \frac{1}{3}$. The value is $f(\frac{1}{3}, \frac{2}{3}) = \frac{1}{3}$. Player I bets every high card and bluffs one low card in three; Player II calls two bets in three; and the game is worth $V = \frac{1}{3}$ to Player I. The bluff is what makes the honest bet pay: without it Player II would always fold to a bet.

Problem 5: Recursive games

Problem. Solve the recursive game $G = \begin{bmatrix} G & 1 & 0 \\ 1 & 0 & G \\ 0 & G & 1 \end{bmatrix}$, in which the diagonal cells mean “play G again.”

Solution. Let V be the value. By the substitution principle, replace each recursive cell G by V , turning the game into the ordinary matrix

$$M(V) = \begin{bmatrix} V & 1 & 0 \\ 1 & 0 & V \\ 0 & V & 1 \end{bmatrix},$$

whose value must again be V for consistency. The matrix is a *circulant*: each row is the previous one shifted one place, so the game is unchanged by cyclically relabelling the strategies of both players together. By that symmetry




Figure B.3. The bluffing game of Problem 4. Player I (circles) sees the card; Player II (squares) does not, so the two bet nodes lie in one information set (dashed). Payoffs are to Player I. The equilibrium bets every high card, bluffs a third of low cards, and is called two-thirds of the time, for value $\frac{1}{3}$.

(orbit-averaging, Theorem 5.10) the uniform mix $\mathbf{p}^* = \mathbf{q}^* = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$ is optimal. Every column then costs the same amount, namely one-third of its column sum, and each column of $M(V)$ sums to $V + 1$, so the uniform strategy guarantees

$$\frac{V + 1}{3}.$$

Consistency demands this equal V :

$$\frac{V + 1}{3} = V \implies V + 1 = 3V \implies V = \frac{1}{2}.$$

So $V = \frac{1}{2}$, with both players mixing uniformly. (Check: at $V = \frac{1}{2}$ each column of $M(\frac{1}{2})$ sums to $\frac{3}{2}$; and $\frac{3}{2}/3 = \frac{1}{2} = V$.)

Problem 6: Bimatrix games

Problem. For the bimatrix $(u_1, u_2) = \begin{bmatrix} (3,3) & (0,5) \\ (5,0) & (1,1) \end{bmatrix}$, find the pure Nash equilibria, then the cooperative transferable-utility value and an optimal division.

Solution. Non-cooperative. Underline best replies (Player I the larger first coordinate down each column, Player II the larger second coordinate across each row):

	C_1	C_2
R_1	<u>(3, 3)</u>	(0, <u>5</u>)
R_2	(<u>5</u> , 0)	(<u>1</u> , <u>1</u>)

Only (R_2, C_2) carries both underlines, so the unique pure Nash equilibrium is mutual defection with payoff $(1, 1)$. This is a Prisoner's Dilemma: each player has a dominant move toward the bad corner.

Cooperative (TU). With binding agreements and side payments the partners maximize the joint total, which lives in $A + B = \begin{bmatrix} 6 & 5 \\ 5 & 2 \end{bmatrix}$; its largest entry is $\sigma = 6$, at the cooperative cell (R_1, C_1) . To split the surplus fairly we score the threat game $A - B = \begin{bmatrix} 0 & -5 \\ 5 & 0 \end{bmatrix}$, which is skew-symmetric, so its value is $\delta = 0$. By Theorem 9.5 the division is

$$\varphi = \left(\frac{\sigma + \delta}{2}, \frac{\sigma - \delta}{2} \right) = (3, 3).$$

The cooperative cell already pays $(3, 3)$, so no side payment is needed: cooperation lifts both players from the equilibrium $(1, 1)$ to the efficient $(3, 3)$, the whole point of allowing binding agreements.

Problem 7: Bankruptcy, fully worked

Problem. Three venture investors hold claims 100, 300, 500 against a firm whose estate is only $E = 500$. Model the wind-up as a coalitional game and find both the Shapley value and the nucleolus. Contrast them with the scarcer estate $E = 200$.

Solution. A coalition can count only on what the estate leaves after the outside claims are paid, so $v(S) = \max\{0, E - \sum_{i \notin S} c_i\}$. At $E = 500$ this gives

$$v(\{3\}) = 100, \quad v(\{1, 3\}) = 200, \quad v(\{2, 3\}) = 400, \quad v(N) = 500,$$

with every other coalition worth 0. *Shapley value.* Average each investor's marginal contribution over all $3! = 6$ arrival orders:

order	1	2	3
1, 2, 3	0	0	500
1, 3, 2	0	300	200
2, 1, 3	0	0	500
2, 3, 1	100	0	400
3, 1, 2	100	300	100
3, 2, 1	100	300	100
sum	300	900	1800
$\div 6$	50	150	300

so $\varphi = (50, 150, 300)$. *Nucleolus.* Efficiency fixes $x_1 + x_2 + x_3 = 500$. The two strongest opposing complaints are coalition $\{1\}$, with excess $v(\{1\}) - x_1 = -x_1$, and coalition $\{2, 3\}$, with excess $400 - (x_2 + x_3) = 400 - (500 - x_1) = x_1 - 100$. Balancing, $-x_1 = x_1 - 100$ gives $x_1 = 50$ and a top excess of -50 . With $x_1 = 50$ fixed, balance $\{2\}$ (excess $-x_2$) against $\{1, 3\}$ (excess $200 - (50 + x_3) = x_2 - 300$, using $x_2 + x_3 = 450$): $-x_2 = x_2 - 300$ gives $x_2 = 150$, hence $x_3 = 300$. So $\mathcal{N} = (50, 150, 300)$, the same as the Shapley value here.

Contrast at $E = 200$. Now the estate is scarcer: $v(\{2, 3\}) = \max\{0, 200 - 100\} = 100$, $v(\{1, 3\}) = \max\{0, 200 - 300\} = 0$, and only $\{3\}$ among smaller sets has worth $v(\{3\}) = \max\{0, 200 - 400\} = 0$. The Shapley value, still averaging marginals, comes out $(\frac{100}{3}, \frac{250}{3}, \frac{250}{3}) \approx (33.3, 83.3, 83.3)$, but the nucleolus is $(50, 75, 75)$. The two fair-division rules now *disagree*: the nucleolus protects the small claimant more, the contrast drawn in [Section 10.7](#).

Problem 8: Nash, minimax, and fixed points

Problem. State the minimax theorem, the Nash existence theorem, and Brouwer's fixed-point theorem, and prove that the fixed points of Nash's gain map are exactly the Nash equilibria.

Solution. *The three theorems.* **Minimax** ([Theorem 4.7](#)): every finite two-person zero-sum game has a value V with $\max_{\mathbf{p}} \min_{\mathbf{q}} \mathbf{p}^\top \mathbf{A} \mathbf{q} = \min_{\mathbf{q}} \max_{\mathbf{p}} \mathbf{p}^\top \mathbf{A} \mathbf{q} = V$. **Nash** ([Theorem 8.15](#)): every finite n -person game has at least one Nash equilibrium in mixed strategies. **Brouwer** ([Theorem 8.14](#)): every continuous map of a nonempty compact convex set into itself has a fixed point.

The gain map and its fixed points. For a profile $\mathbf{x}$ define, for each player i and pure strategy s , the gain $\varphi_{i,s}(\mathbf{x}) = \max\{0, u_i(s, \mathbf{x}_{-i}) - u_i(\mathbf{x})\}$ and the map

$$T_{i,s}(\mathbf{x}) = \frac{x_{i,s} + \varphi_{i,s}(\mathbf{x})}{1 + \sum_r \varphi_{i,r}(\mathbf{x})}.$$

We show $\mathbf{x}^* = T(\mathbf{x}^*)$ if and only if $\mathbf{x}^*$ is a Nash equilibrium.

$(\Leftarrow)$ *Equilibrium $\Rightarrow$ fixed point.* If $\mathbf{x}^*$ is an equilibrium, no pure deviation beats the current payoff, so every gain $\varphi_{i,s}(\mathbf{x}^*) = 0$. Then each $T_{i,s}(\mathbf{x}^*) = x_{i,s}^*/1 = x_{i,s}^*$, so $\mathbf{x}^*$ is fixed.

$(\Rightarrow)$ *Fixed point $\Rightarrow$ equilibrium.* Suppose $\mathbf{x}^*$ is fixed but, for contradiction, some player i has a profitable deviation, so $\sum_r \varphi_{i,r}(\mathbf{x}^*) > 0$. Player i 's mixed payoff is the average of the pure payoffs over the support of $\mathbf{x}_i^*$, so at least one pure strategy t in that support earns *no more* than the average, giving $\varphi_{i,t}(\mathbf{x}^*) = 0$ while $x_{i,t}^* > 0$. For that t ,

$$T_{i,t}(\mathbf{x}^*) = \frac{x_{i,t}^* + 0}{1 + \sum_r \varphi_{i,r}(\mathbf{x}^*)} < x_{i,t}^*,$$

because the denominator exceeds 1. So $T(\mathbf{x}^*) \neq \mathbf{x}^*$, contradicting fixedness. Hence no profitable deviation exists and $\mathbf{x}^*$ is an equilibrium. Existence then follows: T is continuous on the compact convex product of simplices, so Brouwer supplies a fixed point, which by the equivalence just proved is a Nash equilibrium. ■

B.3 Additional practice problems

A short cross-section of exam-style problems with full solutions, one for each principal method.

Practice 1: A Nim position

Problem. Three heaps hold 3, 5, and 9 chips; the player taking the last chip wins. Is the position a win for the mover, and if so what is a winning move? *Solution.* Take the nim-sum in binary:

$$\begin{array}{r} 8\ 4\ 2\ 1 \\ 3 = 0\ 0\ 1\ 1 \\ 5 = 0\ 1\ 0\ 1 \\ 9 = 1\ 0\ 0\ 1 \\ \hline S = 1\ 1\ 1\ 1 \end{array}$$

The nim-sum is $S = (1111)_2 = 15 \neq 0$, so the position is an N-position and the mover wins. A winning move drives some heap h to $h \oplus S$: the heap of 9 goes to $9 \oplus 15 = 6$, legal because $6 < 9$, whereas the heaps of 3 and 5 would have to grow (to 12 and 10) and cannot. So remove 3 chips from the heap of 9, leaving (3, 5, 6) with nim-sum $3 \oplus 5 \oplus 6 = 0$.

Practice 2: A 2×2 zero-sum game

Problem. Solve $A = \begin{bmatrix} 1 & 4 \\ 5 & 2 \end{bmatrix}$. *Solution.* Row minima 1, 2 give maximin 2; column maxima 5, 4 give minimax 4. They differ, so there is no saddle and the game is mixed. With $a = 1, b = 4, c = 5, d = 2$ and $D = a + d - b - c = 1 + 2 - 4 - 5 = -6$, the 2×2 formula gives

$$V = \frac{ad - bc}{D} = \frac{2 - 20}{-6} = 3, \quad p^* = \left(\frac{d-c}{D}, \frac{a-b}{D} \right) = \left(\frac{1}{2}, \frac{1}{2} \right), \quad q^* = \left(\frac{d-b}{D}, \frac{a-c}{D} \right) = \left(\frac{1}{3}, \frac{2}{3} \right).$$

Check: against q^* both rows pay $1 \cdot \frac{1}{3} + 4 \cdot \frac{2}{3} = 3$ and $5 \cdot \frac{1}{3} + 2 \cdot \frac{2}{3} = 3$, confirming $V = 3$.

Practice 3: Backward induction

Problem. Player I chooses L or R . After L , Player II chooses a or b ; after R , she chooses c or d . The payoffs (u_1, u_2) are $La = (2, 2)$, $Lb = (4, 1)$, $Rc = (3, 3)$, and $Rd = (1, 4)$. Find the subgame perfect equilibrium. *Solution.* Fold from the leaves. At the L node Player II weighs her payoffs 2 (from a) and 1 (from b) and picks a , so L leads to $(2, 2)$. At the R node she weighs 3 (from c) and 4 (from d) and picks d , so R leads to $(1, 4)$. Player I then compares 2 from L against 1 from R and picks L . The subgame perfect equilibrium is $(L, (a, d))$ with outcome $(2, 2)$. The jointly better $(3, 3)$ at Rc is out of reach: Player II would never honor c , preferring d , so the credible continuation after R is the worse $(1, 4)$, and Player I steers away from it.

Practice 4: All equilibria of a bimatrix

Problem. Find every Nash equilibrium of $(u_1, u_2) = \begin{bmatrix} (2,1) & (0,0) \\ (0,0) & (1,2) \end{bmatrix}$. *Solution.* Marking best replies, the diagonal cells are both pure equilibria, $(R_1, C_1) = (2, 1)$ and $(R_2, C_2) = (1, 2)$. For the mixed equilibrium let Player I play row 1 with probability p and Player II column 1 with probability q . Player II is indifferent when $p = 2(1 - p)$, so $p = \frac{2}{3}$; Player I is indifferent when $2q = 1 - q$, so $q = \frac{1}{3}$. The mixed equilibrium $((\frac{2}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{2}{3}))$ pays each player $\frac{2}{3}$. Three equilibria, two pure and one mixed, the signature of a coordination game with conflicting preferences.

Practice 5: A Shapley value

Problem. Three firms can form a joint venture. Singletons earn 0; the pairs earn $v(\{1, 2\}) = v(\{1, 3\}) = 60$ and $v(\{2, 3\}) = 0$; all three together earn $v(N) = 72$. Find the Shapley value. *Solution.* Average each firm's marginal contribution over the six arrival orders. Firm 1 belongs to both valuable pairs, firms 2 and 3 to only one each, and the order table yields column sums 264, 84, 84, so

$$\varphi = \left(\frac{264}{6}, \frac{84}{6}, \frac{84}{6} \right) = (44, 14, 14),$$

summing to $72 = v(N)$. Firm 1, indispensable to every productive coalition, collects 44; firms 2 and 3 are symmetric (swapping them leaves every worth unchanged) and split the remainder equally at 14 each.

Practice 6: The graphical method

Problem. Solve the 2×3 game $A = \begin{bmatrix} 3 & 0 & 4 \\ 1 & 5 & 2 \end{bmatrix}$ graphically. *Solution.* Let p be Player I's weight on row 1. Against p the three columns pay $c_1 = 1 + 2p$, $c_2 = 5 - 5p$, $c_3 = 2 + 2p$. Since $c_1 = 1 + 2p$ is below $c_3 = 2 + 2p$ for every p , column 3 never sets the lower envelope, which is $\min(1 + 2p, 5 - 5p)$. Player I maximizes it; the two active lines meet where $1 + 2p = 5 - 5p$, that is $p = \frac{4}{7}$, giving $V = 1 + 2 \cdot \frac{4}{7} = \frac{15}{7}$. So $\mathbf{p}^* = (\frac{4}{7}, \frac{3}{7})$, and Player II mixes the two active columns c_1, c_2 to make the rows indifferent: $3q = 5 - 4q$ gives $q = \frac{5}{7}$, so $\mathbf{q}^* = (\frac{5}{7}, \frac{2}{7}, 0)$, with column 3 slack. Both rows then pay $\frac{15}{7}$.

Practice 7: Nash bargaining

Problem. Two partners split a surplus along the frontier $u + v = 10$, with disagreement point $d = (2, 1)$. Find the Nash bargaining solution. *Solution.* Maximize the gain product $(u - 2)(v - 1)$ on $v = 10 - u$, that is $(u - 2)(9 - u)$; its derivative $11 - 2u$ vanishes at $u = \frac{11}{2}$, so $v = \frac{9}{2}$. The split is $(\frac{11}{2}, \frac{9}{2})$. Equivalently, the total surplus above disagreement is $10 - (2 + 1) = 7$, and the symmetric Nash solution gives each partner their disagreement payoff plus half the surplus: $2 + \frac{7}{2} = \frac{11}{2}$ and $1 + \frac{7}{2} = \frac{9}{2}$.

Practice 8: The core of a three-player game

Problem. For $v(\{i\}) = 0$, $v(\{i, j\}) = 2$ for every pair, and $v(N) = 3$, find the core. *Solution.* An imputation has $x_1 + x_2 + x_3 = 3$ with each $x_i \geq 0$. The three pair constraints $x_i + x_j \geq 2$ each cap the excluded player: $x_1 + x_2 \geq 2$ forces $x_3 \leq 1$, and likewise $x_1 \leq 1$, $x_2 \leq 1$. With the total fixed at 3 and each coordinate at most 1, every coordinate must equal exactly 1. The core is the single point $(1, 1, 1)$, which also equals the Shapley value because the three players are symmetric.

Practice 9: A subtraction game

Problem. A pile of 20 chips is played with subtraction set $\{1, 4\}$, last chip wins. Compute the Grundy values and classify the position. *Solution.* The recurrence $g(n) = \text{mex}\{g(n - 1), g(n - 4)\}$ gives

$$g(0), \dots, g(10) = 0, 1, 0, 1, 2, 0, 1, 0, 1, 2, 0,$$

which settles into the period-5 block $0, 1, 0, 1, 2$. Since $20 \equiv 0 \pmod{5}$, $g(20) = 0$, so the pile of 20 is a P-position: the player to move loses under best play.

Practice 10: Spotting a saddle point

Problem. Decide whether $A = \begin{bmatrix} 7 & 4 & 6 \\ 3 & 4 & 2 \\ 8 & 4 & 5 \end{bmatrix}$ has a saddle, and if so solve it. *Solution.* Border the grid with row minima and column maxima:

	C_1	C_2	C_3	row min
R_1	7	4	6	4
R_2	3	4	2	2
R_3	8	4	5	4
col max	8	4	6	

The maximin is $\max(4, 2, 4) = 4$ and the minimax is $\min(8, 4, 6) = 4$. They agree, so the entry $a_{12} = 4$ (boxed) is a saddle: it is the smallest in its row and the largest in its column. The value is $V = 4$, with Player I playing row 1 and Player II column 2; neither gains by randomizing, since column 2 holds Player I to 4 whatever he does and row 1 guarantees Player I at least 4.

Practice 11: Misère Nim

Problem. Under misère play, where the player taking the last chip *loses*, who wins from $(1, 1, 1)$ and from $(2, 3, 5)$? *Solution.* By Theorem 2.22, a position with some heap of size ≥ 2 keeps its normal-play classification, while a position of all ones reverses it. From $(2, 3, 5)$ a heap exceeds 1 and the nim-sum is $2 \oplus 3 \oplus 5 = 4 \neq 0$, so the mover wins exactly as in normal play; the move $5 \rightarrow 1$ leaves $(2, 3, 1)$ with nim-sum 0, and a heap of size ≥ 2 remains, so it is legal under misère. From $(1, 1, 1)$ every heap is a single chip; normal play would call it an N-position (nim-sum 1), but misère

reverses all-ones positions, so the mover *loses*: any move leaves $(1, 1)$, from which the opponent takes one and forces the mover to take the last chip.

Practice 12: Wythoff's game

Problem. Is $(4, 7)$ a cold position of Wythoff's game? From $(10, 15)$, find a winning move. *Solution.* The cold (P-) positions are $(\lfloor k\varphi \rfloor, \lfloor k\varphi^2 \rfloor)$. For $k = 3$, $\lfloor 3\varphi \rfloor = 4$ and $\lfloor 3\varphi^2 \rfloor = 7$, so $(4, 7)$ is cold and its mover loses. For $k = 6$, $\lfloor 6\varphi \rfloor = 9$ and $\lfloor 6\varphi^2 \rfloor = 15$, so $(9, 15)$ is cold. From $(10, 15)$ remove one chip from the larger-indexed pile to reach $(9, 15)$, a cold position, after which the opponent is lost.

Practice 13: Cournot equilibrium

Problem. Two firms face inverse demand $P = 20 - Q$ where $Q = q_1 + q_2$, each producing at unit cost 2. Find the Cournot equilibrium quantities and profits. *Solution.* With $A = a - c = 20 - 2 = 18$, each firm's reaction is $q_i = (18 - q_j)/2$, and the symmetric solution is $q^* = A/3 = 6$. Total output is 12, the price is $P = 20 - 12 = 8$, and each firm earns $q^*(P - c) = 6(8 - 2) = 36$. (Collusion would cut each to $A/4 = 4.5$ and lift profit to $A^2/8 = 40.5$ apiece, the usual gap.)

Practice 14: Safety levels

Problem. For the bimatrix $(u_1, u_2) = \begin{bmatrix} (2,3) & (5,1) \\ (1,4) & (0,2) \end{bmatrix}$, compute both safety levels. *Solution.* Player I's safety level is $V(A)$ for $A = \begin{bmatrix} 2 & 5 \\ 1 & 0 \end{bmatrix}$: the row minima are 2, 0 and the column maxima 2, 5, so the entry $a_{11} = 2$ is a saddle and $V_I = 2$. Player II's safety level is $V(B^\top)$ for $B^\top = \begin{bmatrix} 3 & 4 \\ 1 & 2 \end{bmatrix}$: row minima 3, 1 and column maxima 3, 4 give a saddle at the 3, so $V_{II} = 3$. The safety levels are $(2, 3)$, the payoffs each player can guarantee alone whatever the other does.

Practice 15: Bankruptcy by the contested garment

Problem. Two creditors are owed 50 and 100 by an estate worth 90. Divide it by the contested-garment rule. *Solution.* Each creditor first concedes to the other whatever the estate covers beyond their own claim. Creditor 1 concedes $\max(0, 90 - 50) = 40$ to creditor 2; creditor 2 concedes $\max(0, 90 - 100) = 0$ to creditor 1. The remaining contested amount is $90 - 40 - 0 = 50$, split evenly at 25 each. Final awards are $0 + 25 = 25$ to creditor 1 and $40 + 25 = 65$ to creditor 2, totalling 90. This is the two-creditor case of the Talmud's bankruptcy rule, the nucleolus of the associated coalitional game.

Practice 16: Grim trigger in the repeated Prisoner's Dilemma

Problem. The stage Prisoner's Dilemma has $T = 5$, $R = 3$, $P = 1$, $S = 0$. For which discount factors δ does the grim trigger sustain cooperation? Check $\delta = 0.4$ and $\delta = 0.6$. *Solution.* Cooperation holds when $\delta \geq (T - R)/(T - P) = 2/4 = \frac{1}{2}$. At $\delta = 0.4$ cooperating is worth $R/(1 - \delta) = 3/0.6 = 5$, while a one-time defection is worth $T + \delta P/(1 - \delta) = 5 + 0.4/0.6 \approx 5.67 > 5$, so cooperation collapses. At $\delta = 0.6$ cooperating is worth $3/0.4 = 7.5$ against a deviation worth $5 + 0.6/0.4 = 6.5 < 7.5$, so cooperation holds. Patience past the threshold $\frac{1}{2}$ is what lets the future discipline the present.

Practice 17: A symmetric 3×3 by indifference

Problem. Solve $A = \begin{bmatrix} 3 & 1 & 4 \\ 1 & 5 & 2 \\ 4 & 2 & 1 \end{bmatrix}$ assuming all strategies are active. *Solution.* The matrix is symmetric, so $\mathbf{p}^* = \mathbf{q}^*$, found by making every row pay the same against $\mathbf{q}$. Solving $A\mathbf{q} = V\mathbf{1}$ with $q_1 + q_2 + q_3 = 1$ gives $\mathbf{q}^* = (\frac{5}{12}, \frac{1}{3}, \frac{1}{4})$, and each row then pays $V = \frac{31}{12}$: row 1 gives $3 \cdot \frac{5}{12} + \frac{1}{3} + 4 \cdot \frac{1}{4} = \frac{31}{12}$, and rows 2 and 3 likewise. All three weights are positive, so the full-support assumption holds.

Practice 18: Grundy values on a game graph

Problem. A position graph has moves $a \rightarrow \{b, c\}$, $b \rightarrow \{d\}$, $c \rightarrow \{d, e\}$, with d and e terminal. Compute the Grundy values and classify a . *Solution.* Work up from the sinks: $g(d) = g(e) = 0$; then $g(b) = \text{mex}\{g(d)\} = \text{mex}\{0\} = 1$ and $g(c) = \text{mex}\{g(d), g(e)\} = \text{mex}\{0, 0\} = 1$; finally $g(a) = \text{mex}\{g(b), g(c)\} = \text{mex}\{1, 1\} = 0$. Since $g(a) = 0$, the starting position a is a P-position and the player to move loses.

Practice 19: A 2×2 with mixed signs

Problem. Solve $A = \begin{bmatrix} -1 & 2 \\ 3 & -4 \end{bmatrix}$. *Solution.* Row minima $-1, -4$ give maximin -1 ; column maxima $3, 2$ give minimax 2 . No saddle, so the game is mixed. With $D = a + d - b - c = -1 - 4 - 2 - 3 = -10$,

$$V = \frac{(-1)(-4) - 2 \cdot 3}{-10} = \frac{-2}{-10} = \frac{1}{5}, \quad \mathbf{p}^* = \left(\frac{7}{10}, \frac{3}{10}\right), \quad \mathbf{q}^* = \left(\frac{3}{5}, \frac{2}{5}\right).$$

Check: against $\mathbf{q}^*$ row 1 pays $-1 \cdot \frac{3}{5} + 2 \cdot \frac{2}{5} = \frac{1}{5}$ and row 2 pays $3 \cdot \frac{3}{5} - 4 \cdot \frac{2}{5} = \frac{1}{5}$, confirming $V = \frac{1}{5}$.

Practice 20: Airport cost-sharing

Problem. Three planes need a runway of length 1, 2, and 4; building length ℓ costs ℓ . Share the cost of the length-4 runway by the Shapley value. *Solution.* Split the runway into segments and charge each segment equally to the planes that use it. Segment $[0, 1]$ is needed by all three planes (cost 1, so $\frac{1}{3}$ each); segment $[1, 2]$ by the two longer planes (cost 1, so $\frac{1}{2}$ each); segment $[2, 4]$ by the longest plane alone (cost 2). The Shapley cost shares are

$$\left(\frac{1}{3}, \frac{1}{3} + \frac{1}{2}, \frac{1}{3} + \frac{1}{2} + 2\right) = \left(\frac{1}{3}, \frac{5}{6}, \frac{17}{6}\right),$$

which total 4, the whole cost. Each plane pays for the stretch it needs, sharing the common stretches evenly, which is exactly what the Shapley value prescribes for this cost game.

Practice 21: A five-heap Nim

Problem. Heaps of 1, 2, 3, 4, 5 chips are played as Nim, last to move winning. Who wins, and how? *Solution.* The nim-sum is $1 \oplus 2 \oplus 3 \oplus 4 \oplus 5$. Combining left to right, $1 \oplus 2 = 3$, then $3 \oplus 3 = 0$, $0 \oplus 4 = 4$, $4 \oplus 5 = 1$. The total is $1 \neq 0$, so the mover wins. Removing the heap of 1 leaves $(2, 3, 4, 5)$ with nim-sum $2 \oplus 3 \oplus 4 \oplus 5 = 0$, a P-position, so taking the whole first heap wins.

Practice 22: Another 2×2

Problem. Solve $A = \begin{bmatrix} 7 & 2 \\ 3 & 6 \end{bmatrix}$. *Solution.* Row minima 2, 3 give maximin 3; column maxima 7, 6 give minimax 6. No saddle. With $D = 7 + 6 - 2 - 3 = 8$, the formula gives $V = (42 - 6)/8 = \frac{9}{2}$, $\mathbf{p}^* = (\frac{3}{8}, \frac{5}{8})$, and $\mathbf{q}^* = (\frac{1}{2}, \frac{1}{2})$. Check: against $\mathbf{q}^*$ both rows pay $\frac{9}{2}$.

Practice 23: A Shapley value

Problem. A three-firm venture has singleton worths 0, pair worths $v(\{1, 2\}) = 40$, $v(\{1, 3\}) = 50$, $v(\{2, 3\}) = 30$, and $v(N) = 90$. Find the Shapley value. *Solution.* Averaging marginal contributions over the six arrival orders gives column sums 210, 150, 180, so

$$\varphi = \left(\frac{210}{6}, \frac{150}{6}, \frac{180}{6}\right) = (35, 25, 30),$$

which sums to $90 = v(N)$. Firm 1, in the two most valuable pairs, earns the most.

Practice 24: A core

Problem. For singleton worths 0, pair worths $v(\{i, j\}) = 4$, and $v(N) = 6$, find the core. *Solution.* Each pair constraint $x_i + x_j \geq 4$ caps the excluded player at $x_k \leq 2$. With the total fixed at 6 and each coordinate at most 2, all three equal 2. The core is the single point $(2, 2, 2)$, again the Shapley value by symmetry.

Practice 25: A Cournot duopoly

Problem. Two firms face $P = 30 - Q$ with $Q = q_1 + q_2$ and unit cost 6. Find the Cournot quantities and profits.

Solution. Here $A = 30 - 6 = 24$, so $q^* = A/3 = 8$ each. Total output is 16, the price is $30 - 16 = 14$, and each firm earns $8(14 - 6) = 64$.

Practice 26: Wythoff again

Problem. Is $(6, 10)$ a cold position of Wythoff's game? From $(7, 10)$, give a winning move. *Solution.* For $k = 4$, $\lfloor 4\varphi \rfloor = 6$ and $\lfloor 4\varphi^2 \rfloor = 10$, so $(6, 10)$ is cold and loses for the mover. From $(7, 10)$ remove one chip from the heap of 7 to reach the cold $(6, 10)$, leaving the opponent lost.

Practice 27: A subtraction game

Problem. A pile of 12 is played with subtraction set $\{2, 3\}$, last chip wins. Classify it and give a winning move if one exists. *Solution.* The recurrence $g(n) = \text{mex}\{g(n-2), g(n-3)\}$ gives the period-5 block $0, 0, 1, 1, 2$, so $g(12) = g(2) = 1 \neq 0$: an N-position, the mover wins. Removing 2 chips reaches the pile of 10, where $g(10) = 0$, a P-position, so taking 2 wins.

Practice 28: A small backward induction

Problem. Player I chooses A , ending the game at $(3, 1)$, or B , after which Player II chooses $c \rightarrow (2, 2)$ or $d \rightarrow (0, 4)$. Find the subgame perfect equilibrium. *Solution.* At the B node Player II compares 2 (from c) with 4 (from d) and picks d , so B leads to $(0, 4)$. Player I then compares A 's payoff 3 with B 's 0 and picks A . The subgame perfect equilibrium is (A, d) with outcome $(3, 1)$.

Practice 29: All equilibria of a bimatrix

Problem. Find every Nash equilibrium of $(u_1, u_2) = \begin{bmatrix} (4,4) & (0,3) \\ (3,0) & (2,2) \end{bmatrix}$. *Solution.* Both diagonal cells are pure equilibria: at $(4, 4)$ a deviation pays $3 < 4$ for each, and at $(2, 2)$ a deviation pays $0 < 2$. For the mixed one, $4p = 2 + p$ gives $p = \frac{2}{3}$ and $4q = 2 + q$ gives $q = \frac{2}{3}$, so the mixed equilibrium is $(\frac{2}{3}, \frac{1}{3}), (\frac{2}{3}, \frac{1}{3})$.

Practice 30: A larger saddle search

Problem. Solve $A = \begin{bmatrix} 8 & 6 & 7 & 9 \\ 5 & 6 & 4 & 3 \\ 7 & 6 & 8 & 5 \end{bmatrix}$. *Solution.* Row minima are 6, 3, 5, so the maximin is 6 (row 1). Column maxima are 8, 6, 8, 9, so the minimax is 6 (column 2). They agree, and the entry $a_{12} = 6$ is a saddle: smallest in its row, largest in its column. The value is 6, with Player I playing row 1 and Player II column 2; no mixing is needed.

Practice 31: Dominance down to a saddle

Problem. Solve the zero-sum game $A = \begin{bmatrix} 3 & 2 & 4 \\ 2 & 1 & 3 \\ 5 & 4 & 6 \end{bmatrix}$. *Solution.* Row 3 beats both other rows entrywise ($5 > 3 > 2$ in column 1, and likewise in every column), so Player I will only ever play row 3; delete rows 1 and 2. Against row 3 the payoffs are $(5, 4, 6)$, and Player II, minimizing, takes the smallest, column 2 at 4. The entry $a_{32} = 4$ is a saddle (largest in its column, smallest in the surviving row), so $V = 4$ with pure strategies row 3 and column 2.

Practice 32: A 2×2 by the formula

Problem. Solve $A = \begin{bmatrix} 0 & 3 \\ 4 & 1 \end{bmatrix}$. *Solution.* Row minima 0, 1 give maximin 1; column maxima 4, 3 give minimax 3. No saddle. With $D = 0 + 1 - 3 - 4 = -6$, the formula gives $V = (0 - 12)/(-6) = 2$, $\mathbf{p}^* = (\frac{1}{2}, \frac{1}{2})$, and $\mathbf{q}^* = (\frac{1}{3}, \frac{2}{3})$. Against $\mathbf{q}^*$ both rows pay 2.

Practice 33: A Shapley value with a useless player

Problem. Three players have singleton worths 0; only $v(\{1, 2\}) = 3$ and $v(\{1, 2, 3\}) = 3$ are positive (every other coalition is worth 0). Find the Shapley value. *Solution.* Players 1 and 2 create all the value and are symmetric; player 3 adds nothing to any coalition and is a null player. Averaging marginal contributions over the six orders gives $\varphi = (\frac{3}{2}, \frac{3}{2}, 0)$: the two productive players split the 3, and the null player gets nothing, exactly as the null-player and symmetry axioms demand.

Practice 34: Nash bargaining on a tilted frontier

Problem. Two players bargain over the frontier $2u + v = 12$ with disagreement point $d = (1, 2)$. Find the Nash bargaining solution. *Solution.* Maximize the gain product $(u - 1)(v - 2)$ along $v = 12 - 2u$, that is $(u - 1)(10 - 2u)$, whose derivative $12 - 4u$ vanishes at $u = 3$, giving $v = 6$. The solution is $(3, 6)$. The equal-surplus shortcut does not apply here, because the frontier is not at slope -1 ; the product must be maximized directly.

Practice 35: A core that is a single point

Problem. For singleton worths 0, pair worths $v(\{i, j\}) = 10$, and $v(N) = 15$, find the core. *Solution.* Each pair constraint $x_i + x_j \geq 10$ caps the excluded player at $x_k \leq 5$. With the total fixed at 15 and each coordinate at most 5, all three equal 5. The core is the single point $(5, 5, 5)$. The balancedness check is tight: the three pairs at weight $\frac{1}{2}$ produce $\frac{1}{2}(10 + 10 + 10) = 15 = v(N)$, so the game is balanced and its core is the lone point that just satisfies every coalition.

B.4 Examination methods at a glance

Every examination problem is an instance of a method developed in the text. To solve a zero-sum matrix game, follow Procedure 4.19: test for a saddle, reduce by dominance, then use the 2×2 formula, the graphical method, the inversion formula, or the linear program. To solve an impartial game, compute Grundy values by mex and combine by nim-sum (Theorem 2.15). To solve a recursive game, substitute the value and solve the fixed-point equation (Procedure 7.1). To find equilibria of a general-sum game, use best responses and the indifference conditions (Chapter 8). To divide a cooperative surplus, compute the Shapley value by (10.1) or the nucleolus by minimizing excesses (Chapter 10).

APPENDIX C

Computational Companion

Four Python notebooks accompanied the four problem sets, and two more collect the remaining algorithms from the lectures, turning the theorems into working code. This appendix describes what each of the six computes and gives the load-bearing routines. The code is written for clarity, using only `numpy` and the standard library, with a linear-programming routine standing in for the simplex method on the larger games.

C.1 The Nim and Grundy toolkit

The first notebook implements the combinatorial theory of Chapter 2. The heart is the nim-sum and the winning-move rule from the proof of Bouton's theorem (Theorem 2.7).

```
from functools import reduce

def nim_sum(piles):
    "Bitwise XOR of all pile sizes."
    return reduce(lambda a, b: a ^ b, piles, 0)

def nim_player(piles):
    "Return all winning first moves (pile index, new size), or [] if losing."
    S = nim_sum(piles)
    if S == 0:
        return [] # P-position: no winning move
    moves = []
    for i, c in enumerate(piles):
        target = c ^ S # the equalizing size
        if target < c: # a legal (decreasing) move
            moves.append((i, target))
    return moves
```

For (1, 2, 3, 4, 5) the nim-sum is 1 and the routine returns the three winning moves of Example 2.9's companion. The Grundy machinery for general impartial games uses the minimal excludant and a memoized recurrence.

```
def mex(values):
    "Smallest nonnegative integer not in the set."
    s, m = set(values), 0
    while m in s:
        m += 1
    return m

def grundy_subtraction(n, S, memo):
    "Grundy value of a one-pile subtraction game with move set S."
    if n in memo:
        return memo[n]
    memo[n] = mex(grundy_subtraction(n - s, S, memo) for s in S if s <= n)
    return memo[n]
```

The take-and-break games of Section 2.6 add splitting moves, whose value is the nim-sum of the parts. For Kayles, a move removes one or two chips and may split the pile:

```
def grundy_kayles(n, memo):
    if n in memo: return memo[n]
    opts = set()
    for take in (1, 2): # remove 1 or 2 chips
        if take <= n:
            opts.add(grundy_kayles(n - take, memo))
            for left in range(1, n - take): # then split the remainder
                opts.add(grundy_kayles(left, memo)
                        ^ grundy_kayles(n - take - left, memo))
    memo[n] = mex(opts)
    return memo[n]
```

Tabulating to 100 confirms that Kayles is eventually periodic with period 12, Dawson's Chess with period 34, and that Grundy's game shows no period in range. A directed-acyclic-graph evaluator (for the graph problem of [Chapter 11](#)) memoizes $g(x) = \text{mex}\{g(y) : y \in F(x)\}$ over a dictionary of followers.

C.2 The universal zero-sum solver

The second notebook solves any two-person zero-sum game by the seven methods of [Procedure 4.19](#), trying the cheapest first and falling through to the linear program. The saddle test and the 2×2 closed form are immediate:

```
import numpy as np

def saddle_point(A):
    "Return (value, i, j) if a pure saddle exists, else None."
    row_min = A.min(axis=1); col_max = A.max(axis=0)
    if row_min.max() == col_max.min():
        i = int(row_min.argmax()); j = int(col_max.argmin())
        return A[i, j], i, j
    return None

def solve_2x2(A):
    "Mixed solution of a 2x2 game with no saddle."
    a, b, c, d = A[0,0], A[0,1], A[1,0], A[1,1]
    D = a - b - c + d
    p = (d - c) / D # P1 prob of row 1
    q = (d - b) / D # P2 prob of col 1
    v = (a*d - b*c) / D
    return v, np.array([p, 1-p]), np.array([q, 1-q])
```

Larger games reduce by dominance, then by the principle of indifference (solving the linear system that equalizes the active strategies), by the inversion formula $V = 1/(\mathbf{1}^\top A^{-1} \mathbf{1})$ for non-degenerate square games, by symmetry reduction for invariant games, or by the linear program

$$\min \mathbf{1}^\top \mathbf{x} \text{ s.t. } A^\top \mathbf{x} \geq \mathbf{1}, \mathbf{x} \geq \mathbf{0},$$

after shifting A to be positive ([Section 6.4](#)). The solver was checked against an independent reference on a hundred random matrices of sizes up to 100×100 and agreed on every well-posed instance.

C.3 Recursive game solvers

The third notebook solves the recursive games of [Chapter 7](#). A numerical solver memoizes the value of each state, assembling a 2×2 subgame from the values of the smaller states and solving it by the saddle test and `solve_2x2`. A symbolic solver, built on `sympy`, handles a self-referential cell by setting the value v equal to the value of the matrix containing v , then solving the resulting polynomial fixed-point equation and selecting the root that respects the security bounds. It reproduces the nested value $\frac{9}{7}$, the mutually recursive trio's $3 - \sqrt{5}$, and the coin-toss limit $4 - 2\sqrt{3}$ of [Chapter 7](#).

C.4 Fixed points and the cooperative solver

The fourth notebook accompanies Chapters 8 and 10. An approximate fixed-point finder for a continuous map $f : B^2 \rightarrow B^2$ discretizes the square, evaluates the displacement $f(z) - z$ on the grid, and refines where its components change sign; on a contraction toward the center it returns the center. A transferable-utility solver forms the combined matrix $C = A + B$, solves it with the zero-sum solver to find the cooperative value, and splits the surplus by the threat value $\delta = V(A - B)$ as in Theorem 9.5. The bankruptcy routine builds the characteristic function $v(S) = \max\{0, E - \sum_{i \notin S} c_i\}$, computes the Shapley value by averaging marginal contributions over coalitions, and finds the nucleolus by repeatedly minimizing the largest excess. For the claims (100, 300, 500) and estate 500 it returns the Shapley value (50, 150, 300) and the nucleolus (50, 150, 300), which coincide for this estate; rerun at a scarcer estate (for example $E = 200$) and the two diverge, as in Appendix B.2.

```
from itertools import combinations
from math import factorial

def bankruptcy_v(S, E, claims):
    "Worth of coalition S in the bankruptcy game."
    outside = sum(claims[i] for i in range(len(claims)) if i not in S)
    return max(0.0, E - outside)

def shapley(n, v):
    "Shapley value from the marginal-contribution formula."
    phi = [0.0] * n
    players = range(n)
    for i in players:
        others = [j for j in players if j != i]
        for k in range(len(others) + 1):
            for S in combinations(others, k):
                S = set(S)
                w = factorial(k) * factorial(n - k - 1) / factorial(n)
                phi[i] += w * (v(S | {i}) - v(S))
    return phi
```

C.5 Backward induction and fictitious play

A fifth notebook covers two algorithms from the non-cooperative chapters. Backward induction solves a perfect-information game tree by folding from the leaves: each decision node returns the child outcome that maximizes the mover's own payoff, exactly the procedure of Chapter 3.

```
def solve(node):
    "Fold a perfect-information game tree to its subgame perfect payoff."
    if node[0] == "leaf":
        return node[1] # a payoff vector
    _, player, children = node # ("node", whose turn, [subtrees])
    outcomes = [solve(child) for child in children]
    return max(outcomes, key=lambda payoff: payoff[player])
```

The second routine is fictitious play for a zero-sum game: each player repeatedly best-responds to the running frequency of the opponent's past moves. Robinson's theorem guarantees the empirical mixes converge to the optimal strategies and the average payoff to the value, though only at rate $O(1/\sqrt{t})$, so thousands of rounds buy a couple of digits.

```
def fictitious_play(A, rounds=10000):
    "Empirical optimal strategies and value of a zero-sum game (row maximizes)."
    m, n = len(A), len(A[0])
    rc, cc = [0]*m, [0]*n # cumulative play counts
    for _ in range(rounds):
        i = max(range(m), key=lambda i: sum(A[i][j]*cc[j] for j in range(n)))
        j = min(range(n), key=lambda j: sum(A[i][j]*rc[i] for i in range(m)))
        rc[i] += 1
        cc[j] += 1
```

```

p = [x/rounds for x in rc]
q = [x/rounds for x in cc]
val = sum(p[i]*A[i][j]*q[j] for i in range(m) for j in range(n))
return p, q, val

```

On Matching Pennies the empirical mixes settle near $(\frac{1}{2}, \frac{1}{2})$ and the value near 0, matching the minimax solution of Chapter 4; the same loop on Rock–Paper–Scissors converges to the uniform mix $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, cycling toward it as Robinson’s proof predicts.

C.6 Combinatorial and cooperative solvers

A sixth notebook gathers small solvers for the two ends of the book. For combinatorial games, a Wythoff generator and a misère-Nim oracle:

```

def wythoff(k):
    "The k-th cold (P-) position of Wythoff's game."
    phi = (1 + 5**0.5) / 2
    a = int(k * phi)
    return (a, a + k) # = (floor(k*phi), floor(k*phi^2))

def misere_nim_winner(heaps):
    "Who wins Nim under the misere rule (last to move loses)."
```

```

from functools import reduce
nim = reduce(lambda x, y: x ^ y, heaps, 0)
if all(h <= 1 for h in heaps): # all singletons: parity decides
    return "second" if sum(heaps) % 2 == 1 else "first"
return "first" if nim != 0 else "second" # otherwise the normal-play rule

```

For cooperative games, a Bondareva–Shapley balancedness test that decides, for three players, whether the core is nonempty:

```

def core_nonempty_3(v):
    "True iff the 3-player game v (a dict over coalitions) has a nonempty core."
    fs = frozenset
    s = [v[fs({i})] for i in range(3)]
    p12, p13, p23 = v[fs({0,1})], v[fs({0,2})], v[fs({1,2})]
    N = v[fs({0,1,2})]
    partitions = [s[0]+s[1]+s[2], s[0]+p23, s[1]+p13, s[2]+p12]
    half_pairs = 0.5 * (p12 + p13 + p23) # the three pairs, weight 1/2 each
    return max(partitions + [half_pairs]) <= N

```

The first returns the golden-ratio pairs $(0,0), (1,2), (3,5), \dots$; the second confirms that $(1,1,1)$ loses under misère while $(2,3,5)$ wins; the third reports a nonempty core for the game of Example 10.9 and an empty one for its flattened cousin, exactly as Theorem 10.8 predicts.

Across the six notebooks the pattern is the same: a theorem from the text becomes a short, exact routine, and the routine confirms the worked examples to the last digit. The code is the proof that the theory computes.

APPENDIX D

Solutions to Selected Exercises

Solutions to a representative selection of the end-of-chapter exercises follow, grouped by chapter. Each solution *restates the problem* before solving it, so this appendix can be read on its own, and each works the steps in full. The remaining exercises are answered by the same methods. Problem-set and examination exercises are solved, also in full, in Chapter 11 and Appendix B.

Chapter 1

Golden Balls. *Problem.* Two contestants face a jackpot of $J = 100,000$ dollars. Each secretly chooses SPLIT or STEAL. If both split, they share the jackpot equally; if one steals and the other splits, the stealer takes all and the splitter gets nothing; if both steal, both get nothing. Write the bimatrix and find the rational outcome. *Solution.* The cells are $(\text{SPLIT}, \text{SPLIT}) = (50,000, 50,000)$, $(\text{SPLIT}, \text{STEAL}) = (0, 100,000)$, $(\text{STEAL}, \text{SPLIT}) = (100,000, 0)$, $(\text{STEAL}, \text{STEAL}) = (0, 0)$. Compare Player I's two actions against each choice of Player II. If Player II splits, Player I earns 100,000 by stealing versus 50,000 by splitting, so stealing is strictly better; if Player II steals, Player I earns 0 either way, so stealing ties. Thus STEAL weakly (not strictly) dominates SPLIT, and by symmetry the same holds for Player II, giving the outcome $(\text{STEAL}, \text{STEAL})$ paying $(0, 0)$. The tie when the opponent steals is the slack that lets real contestants sometimes gamble on SPLIT.

The zero-sum test. *Problem.* Show that a bimatrix game (A, B) is zero-sum exactly when $B = -A$, and explain why a single matrix then suffices to describe it. *Solution.* By definition the game is zero-sum when $u_1(s) + u_2(s) = 0$ for every profile, that is $a_{ij} + b_{ij} = 0$ for all i, j , equivalently $b_{ij} = -a_{ij}$, i.e. $B = -A$. Player II then maximizes $b_{ij} = -a_{ij}$, which is the same as minimizing a_{ij} . So all the information lives in A : Player I maximizes it and Player II minimizes it, and we write the game as the one matrix A .

Dominated strategies and iterated removal. *Problem.* In each game find all strictly and weakly dominated strategies, then apply iterated removal of strictly dominated strategies:

(a)		L	R	(b)		L	C	R
	T	$(2, 1)$	$(0, 0)$		T	$(4, 3)$	$(2, 7)$	$(0, 4)$
	B	$(1, 1)$	$(1, 2)$		M	$(5, 5)$	$(5, -1)$	$(-4, -2)$
					B	$(3, 5)$	$(1, 5)$	$(-1, 6)$

Solution. (a) Compare Player I's rows: against L row T pays $2 > 1$, against R row B pays $1 > 0$, so neither row dominates the other. Player II's columns: against T column L pays $1 > 0$, against B column R pays $2 > 1$, so neither column dominates. Nothing is even weakly dominated, iterated removal changes nothing, and the game keeps both pure equilibria $(T, L) = (2, 1)$ and $(B, R) = (1, 2)$. Not every game simplifies. (b) Here removal runs all the way down. Row T beats row B in every column ($4 > 3$, $2 > 1$, $0 > -1$), so T strictly dominates B ; delete B . Among the surviving rows T, M , Player II's column C beats R ($7 > 4$ against T , $-1 > -2$ against M), so delete R . With columns L, C left, Player I's row M beats T ($5 > 4$, $5 > 2$), so delete T . In the lone row M , Player II prefers L ($5 > -1$), so delete C . One cell survives, $(M, L) = (5, 5)$, the unique outcome left by iterated strict dominance.

Pure equilibria by hand. *Problem.* Using the best-response marking method, count the pure Nash equilibria of game (a) above and of the Prisoner's Dilemma. *Solution.* In game (a), mark Player I's best row in each column (T under L , B under R) and Player II's best column in each row (L across T , R across B). The doubly marked cells are (T, L) and (B, R) , so there are *two* pure equilibria. In the Prisoner's Dilemma each player has a strictly dominant strategy to defect, so only $(\text{DEFECT}, \text{DEFECT})$ carries both marks: exactly *one* pure equilibrium, and it is the inefficient one.

Proportional division against the Talmud. *Problem.* Three creditors are owed $c = (100, 200, 300)$, totalling 600. For estates $E \in \{100, 200, 300\}$ compute the proportional award $E \cdot c_i / 600$ and compare it with the Talmud's divisions. *Solution.* Proportional division rescales the claims to the size of the estate:

E	creditor 1	2	3	Talmud
100	$16\frac{2}{3}$	$33\frac{1}{3}$	50	$(33\frac{1}{3}, 33\frac{1}{3}, 33\frac{1}{3})$
200	$33\frac{1}{3}$	$66\frac{2}{3}$	100	(50, 75, 75)
300	50	100	150	(50, 100, 150)

The rules disagree at $E = 100$, where the Talmud divides equally, and at $E = 200$, but they coincide at $E = 300$, both awarding (50, 100, 150). That estate is special because 300 is exactly half the total claim: at the half-way point the contested-garment principle behind the Talmud table reduces to proportional division. The full reconciliation, through the nucleolus, is [Example 10.21](#).

The tragedy of the commons. *Problem.* Each of n herders chooses how many cattle to graze on a shared pasture; more total grazing lowers the value of every animal. Argue that the Nash equilibrium overgrazes relative to the value-maximizing total. *Solution.* Make the model concrete: let herder i graze g_i animals, write $G = \sum_j g_j$ for the total herd, and let each animal be worth $a - G$ once the herd reaches G . Herder i earns $g_i(a - G)$. Maximizing over g_i with the others held fixed gives the condition $a - 2g_i - \sum_{j \neq i} g_j = 0$, and by symmetry every herder grazes $g^* = a/(n+1)$, for a total $G_{NE} = na/(n+1)$. The value-maximizing total instead maximizes $G(a - G)$, giving $G_{opt} = a/2$. Since $na/(n+1) > a/2$ for every $n \geq 2$, the equilibrium overgrazes, and the gap widens with n : as the pasture is shared more thinly, $G_{NE} \rightarrow a$ and the common value collapses toward zero. The cause is an externality. Each herder reaps the full private gain of one more animal but bears only a fraction of the damage it does to the pasture, the Prisoner's Dilemma scaled to n players.

Randomization in the goblet duel. *Problem.* Westley poisons goblet A with probability w (and B otherwise); Vizzini, unable to observe the choice, drinks from A with probability d . Write Vizzini's probability of drinking poison, and find the w that makes his survival the same whatever d he picks. *Solution.* Vizzini drinks poison precisely when he drinks the poisoned goblet, so

$$\Pr(\text{poison}) = wd + (1-w)(1-d) = (1-w) + (2w-1)d.$$

The dependence on d sits in the coefficient $2w-1$. Setting $w = \frac{1}{2}$ kills it, leaving $\Pr(\text{poison}) = \frac{1}{2}$ whatever d Vizzini chooses. By poisoning each goblet with equal probability, Westley makes Vizzini's reasoning worthless: every choice, and every counter-bluff about where the poison "really" is, returns the same even odds. Randomizing on the unobserved move neutralizes an opponent who is trying to out-think it.

Two equilibria in a coordination game. *Problem.* The coordination game $\begin{bmatrix} (2,2) & (0,0) \\ (0,0) & (1,1) \end{bmatrix}$ has each player choosing A or B . Find the pure Nash equilibria, say which the players prefer, and what this reveals about prediction. *Solution.* Marking best responses, both diagonal cells are equilibria: at (A, A) a deviation pays $0 < 2$, and at (B, B) a deviation pays $0 < 1$, so neither player moves from either. There are two pure equilibria, $(A, A) = (2, 2)$ and $(B, B) = (1, 1)$. Both players strictly prefer (A, A) , which Pareto-dominates (B, B) , yet the equilibrium concept endorses both: nothing in the definition tells the players to coordinate on the better one, and if each merely expects the other to play B , then (B, B) is self-confirming. Equilibrium pins down stability without fixing which stable point gets reached; among equilibria needs something more, a focal point, a convention, or communication.

Chapter 2

Nim positions. *Problem.* For each Nim position compute the nim-sum, classify it as P or N, and give a winning first move if one exists: (a) (3, 5, 7); (b) (1, 4, 9, 16); (c) (13, 17, 26). *Solution.* (a) (3, 5, 7): in binary $011 \oplus 101 \oplus 111$; column parities give $(001)_2 = 1 \neq 0$, an N-position. The leftmost 1 of $S = 1$ is the units bit, which every (odd) pile carries, so each pile yields a winning move: $7 \oplus 1 = 6$, $5 \oplus 1 = 4$, or $3 \oplus 1 = 2$. Reducing $7 \rightarrow 6$ leaves $3 \oplus 5 \oplus 6 = 0$. (b) (1, 4, 9, 16): aligning bits, $S = 1 \oplus 4 \oplus 9 \oplus 16 = (11100)_2 = 16 + 8 + 4 = 28 \neq 0$, an N-position. The leftmost 1 of 28 is the 16s bit, carried only by the pile 16, whose target is $16 \oplus 28 = (10000)_2 \oplus (11100)_2 = (01100)_2 = 12$; so the winning move removes 4 from the heap of 16, after which $1 \oplus 4 \oplus 9 \oplus 12 = 0$. (c) (13, 17, 26): $S = 6 \neq 0$, an N-position; the 4s bit (the leftmost of $6 = (110)_2$) is carried by 13, and $13 \oplus 6 = 11$, so reduce $13 \rightarrow 11$, leaving $11 \oplus 17 \oplus 26 = 0$.

Wythoff openings. *Problem.* In Wythoff's game (two piles; remove any amount from one pile, or an equal amount from both) show by hand that (1, 2) and (3, 5) are P-positions. *Solution.* The position (1, 2) is a P-position: every move either changes one pile (leaving an N-position such as (0, 2), (1, 1), (1, 0)) or takes equally from both to reach

$(0, 1)$, all of which the opponent wins from. Likewise $(3, 5)$ is a P-position. These are the first nonzero entries of the golden-ratio sequence $(\lfloor k\phi \rfloor, \lfloor k\phi^2 \rfloor)$.

Misère Nim. *Problem.* Using the misère theory (Theorem 2.22, where the player taking the last chip *loses*), decide the outcome of the Nim positions $(1, 1, 1)$, $(1, 1, 1, 1)$, $(2, 2)$, and $(1, 2, 3)$, and give a winning move for each N-position. *Solution.* Apply Theorem 2.22. $(1, 1, 1)$: all piles are ≤ 1 and the number of nonempty piles is 3 (odd), so it is a P-position; the player to move loses. $(1, 1, 1, 1)$: all ≤ 1 but the count 4 is even, so it is an N-position; the winning move removes one whole pile to reach $(1, 1, 1)$, an odd count, handing the opponent the loss. $(2, 2)$: a pile exceeds 1 and the nim-sum is $2 \oplus 2 = 0$, so it is a P-position. $(1, 2, 3)$: a pile exceeds 1 and $1 \oplus 2 \oplus 3 = 0$, so it too is a P-position. In all three P-positions the player to move loses under optimal play.

Three heaps $(7, 9, 12)$. *Problem.* Three Nim heaps hold 7, 9, and 12 chips, last chip wins. Compute the nim-sum in binary, classify the position, and give every winning move, verifying that it leaves nim-sum zero. *Solution.* Write each heap in four bits and add the columns modulo 2:

$$\begin{array}{rcccc} 7 = & 0 & 1 & 1 & 1 \\ 9 = & 1 & 0 & 0 & 1 \\ 12 = & 1 & 1 & 0 & 0 \\ \hline \oplus = & 0 & 0 & 1 & 0 \end{array} \quad \Longrightarrow \quad 7 \oplus 9 \oplus 12 = (0010)_2 = 2.$$

The nim-sum is $2 \neq 0$, so this is an **N-position** and the player to move wins. To find a winning move, the rule says target each heap h at $h \oplus 2$ and keep only the moves that *shrink* the heap. Heap 7: $7 \oplus 2 = (0111)_2 \oplus (0010)_2 = (0101)_2 = 5 < 7$, a legal move (remove 2). Heap 9: $9 \oplus 2 = (1001)_2 \oplus (0010)_2 = (1011)_2 = 11 > 9$, not legal. Heap 12: $12 \oplus 2 = (1100)_2 \oplus (0010)_2 = (1110)_2 = 14 > 12$, not legal. So the *only* winning move is $7 \rightarrow 5$ (remove 2 chips from the first heap). Check: $5 \oplus 9 \oplus 12 = (0101) \oplus (1001) \oplus (1100) = (0000)_2 = 0$, a P-position handed to the opponent. The leftmost (and here only) set bit of the nim-sum sits in the 2s column, which only heap 7 carries with a 1, which is exactly why just one heap offers a winning move.

Subtraction game $S = \{1, 2, 3\}$. *Problem.* A heap is played with $S = \{1, 2, 3\}$, last chip wins. (a) Compute $g(0), \dots, g(8)$. (b) Find $g(10), g(15), g(21)$. (c) Classify the disjunctive sum of heaps 10, 15, 21 and say whether to move first. *Solution.* (a) The Grundy value is $g(n) = \text{mex}\{g(n-s) : s \in S, n-s \geq 0\}$, where mex is the smallest nonnegative integer not in the set. Build it up:

n	$\{g(n-1), g(n-2), g(n-3)\}$	$g(n) = \text{mex}$
0	$\emptyset$	0
1	$\{g(0)\} = \{0\}$	1
2	$\{g(1), g(0)\} = \{1, 0\}$	2
3	$\{g(2), g(1), g(0)\} = \{2, 1, 0\}$	3
4	$\{g(3), g(2), g(1)\} = \{3, 2, 1\}$	0
5	$\{g(4), g(3), g(2)\} = \{0, 3, 2\}$	1
6	$\{g(5), g(4), g(3)\} = \{1, 0, 3\}$	2
7	$\{g(6), g(5), g(4)\} = \{2, 1, 0\}$	3
8	$\{g(7), g(6), g(5)\} = \{3, 2, 1\}$	0

The values cycle $0, 1, 2, 3, 0, 1, 2, 3, \dots$, so $g(n) = n \bmod 4$. This makes sense: from a heap you can move to any of the next three smaller residues, so the mex always lands on the missing fourth residue. (b) $g(10) = 10 \bmod 4 = 2$, $g(15) = 15 \bmod 4 = 3$, $g(21) = 21 \bmod 4 = 1$. (c) By Sprague–Grundy the combined Grundy value is the nim-sum of the parts: $2 \oplus 3 \oplus 1 = (10)_2 \oplus (11)_2 \oplus (01)_2 = (00)_2 = 0$. The sum is a **P-position**, so you want to move *second*: whoever moves first from $(10, 15, 21)$ loses to correct play. (A winning reply to any opponent move restores nim-sum 0, e.g. mirror within the residue classes.)

The take-away game with set $\{1, \dots, 6\}$. *Problem.* A single pile holds n chips; a move removes 1 to 6 chips and the player taking the last chip wins. Identify all P-positions, give the winning move at $n = 31$, and say whether you want to move first at $n = 28$. *Solution.* A position is a P-position exactly when $n \equiv 0 \pmod{7}$. From a multiple of 7 any removal of 1 through 6 lands on a non-multiple, while from a non-multiple the mover can always take $n \bmod 7$ chips to reach the multiple of 7 just below, so the multiples of 7 are precisely the positions that lose under best play. At $n = 31 = 4 \cdot 7 + 3$, remove 3 chips to reach 28, a P-position, and win. At $n = 28 = 4 \cdot 7$ the position is already a P-position, so the mover loses: prefer to move *second*.

Lasker's Nim, the first nine Grundy values. *Problem.* In Lasker's Nim a move either removes any number of chips from one pile, as in Nim, or splits one pile of size at least 2 into two nonempty piles. Compute $g(0)$ through $g(8)$ and confirm the values $0, 1, 2, 4, 3, 5, 6, 8, 7$. *Solution.* Writing g_i for $g(i)$, the value $g(n)$ is the mex of the removal

followers $g_0, \dots, g_{n-1}$ together with the split nim-sums $g_a \oplus g_b$ over all $a + b = n$ with $a, b \geq 1$. The removals reproduce the sequence built so far; the table records the split nim-sums and the resulting value, starting from $g_0 = 0$ and $g_1 = 1$:

n	split nim-sums $g_a \oplus g_b$	$g(n)$
2	$g_1 \oplus g_1 = 0$	2
3	$g_1 \oplus g_2 = 3$	4
4	$g_1 \oplus g_3, g_2 \oplus g_2 = 5, 0$	3
5	$g_1 \oplus g_4, g_2 \oplus g_3 = 2, 6$	5
6	$g_1 \oplus g_5, g_2 \oplus g_4, g_3 \oplus g_3 = 4, 1, 0$	6
7	$g_1 \oplus g_6, g_2 \oplus g_5, g_3 \oplus g_4 = 7, 7, 7$	8
8	$g_2 \oplus g_6, g_3 \oplus g_5, g_1 \oplus g_7, g_4 \oplus g_4 = 4, 1, 9, 0$	7

For instance at $n = 8$ the removals supply $g_0, \dots, g_7 = 0, 1, 2, 4, 3, 5, 6, 8$ and the splits add 4, 1, 9, 0, so the reachable set is $\{0, 1, 2, 3, 4, 5, 6, 8, 9\}$ and its mex is 7. The sequence 0, 1, 2, 4, 3, 5, 6, 8, 7 matches, and its pattern is $g(n) = n$ for $n \equiv 1, 2 \pmod{4}$, $g(n) = n + 1$ for $n \equiv 3$, and $g(n) = n - 1$ for $n \equiv 0$ (with $g(0) = 0$), the closed form for Lasker's Nim.

Chapter 3

Backward induction. *Problem.* Player I chooses A or B ; Player II then responds, choosing C or D after A and E or F after B . Payoffs (u_1, u_2) are $A \rightarrow C : (2, 1)$, $A \rightarrow D : (0, 0)$, $B \rightarrow E : (3, 0)$, $B \rightarrow F : (1, 2)$. Find the subgame perfect equilibrium and its outcome. *Solution.* Fold the tree from the leaves. At Player II's node after A , she compares her own payoffs $C : 1$ against $D : 0$ and plays C , so that node is worth $(2, 1)$. After B she compares $E : 0$ against $F : 2$ and plays F , worth $(1, 2)$. Player I now compares $A \rightarrow (2, 1)$, paying him 2, against $B \rightarrow (1, 2)$, paying him 1, and chooses A . The SPNE is $(A, (C \text{ after } A, F \text{ after } B))$ with outcome $(2, 1)$.

Nim as a tree. *Problem.* Draw the game tree for two-pile Nim from $(1, 2)$ and solve it by backward induction, confirming agreement with Bouton's theorem. *Solution.* The nim-sum is $1 \oplus 2 = 3 \neq 0$, so $(1, 2)$ is an N-position. The winning first move equalizes the piles to $(1, 1)$ (remove one chip from the pile of 2); from $(1, 1)$ any move leaves a single chip, which the opponent must take, and then the first player takes the last. Backward induction over the (small) tree confirms the move to $(1, 1)$ wins, matching Bouton's rule that the winning move zeroes the nim-sum, since $1 \oplus 1 = 0$.

From a tree to its normal form. *Problem.* Player I chooses A or B ; after A Player II chooses C or D , and after B Player II chooses E or F , with payoffs (u_1, u_2) given by $A \rightarrow C : (2, 1)$, $A \rightarrow D : (0, 0)$, $B \rightarrow E : (3, 0)$, $B \rightarrow F : (1, 2)$. Write the normal form, find all pure equilibria, say which is subgame perfect, and name the threat that sustains the other. *Solution.* A strategy for Player II must prescribe a reply at *both* of her nodes, so she has four: CE, CF, DE, DF , where the first letter answers A and the second answers B . Player I has A and B . Reading the resulting leaf into each cell,

	CE	CF	DE	DF
A	(2, 1)	(2, 1)	(0, 0)	(0, 0)
B	(3, 0)	(1, 2)	(3, 0)	(1, 2)

Marking best responses (Player I down each column, Player II across each row) leaves two doubly marked cells: $(A, CF) = (2, 1)$ and $(B, DF) = (1, 2)$. Backward induction selects CF for Player II, since after A she prefers C ($1 > 0$) and after B she prefers F ($2 > 0$), and then Player I takes A ($2 > 1$). So (A, CF) is the subgame perfect equilibrium. The other equilibrium (B, DF) survives only on the non-credible threat "I will play D after A ." Playing D would cost Player II (0 rather than 1), so she would never carry it out, but the threat sits off the path that (B, DF) induces and ordinary Nash equilibrium never tests off-path play. Subgame perfection tests it and rejects the equilibrium.

Counting pure strategies. *Problem.* Player I moves first with three actions; after each, Player II moves with two actions; after each of those, Player I moves again with two actions. How many pure strategies does each player have? *Solution.* A pure strategy fixes an action at *every* one of that player's decision nodes. Player II has one node after each of Player I's three openings, three nodes of two actions each, for $2^3 = 8$ strategies. Player I acts at the root (three actions) and again at each of the $3 \times 2 = 6$ third-level nodes (two actions each), for $3 \cdot 2^6 = 192$ strategies. The strategy count explodes long before the tree looks large, which is why the normal form of an extensive game is usually far bigger than the tree behind it.

The ultimatum game. *Problem.* Player I offers Player II an amount $x \in \{0, 0.25, 0.5, 0.75, 1\}$ out of one dollar and keeps $1 - x$; Player II then accepts, giving $(1 - x, x)$, or rejects, giving $(0, 0)$. Find the subgame perfect equilibrium

and compare it with observed play. *Solution.* Solve the responder's move first. For any $x > 0$, accepting pays $x > 0$ and rejecting pays 0, so Player II accepts every positive offer (and is indifferent at $x = 0$). Anticipating this, Player I offers the smallest positive amount sure to be accepted, $x = 0.25$, and keeps 0.75; the subgame perfect outcome is $(0.75, 0.25)$. Real play departs sharply. Proposers usually offer near an even split, and responders reject offers below roughly a fifth, paying real money to punish a stingy proposer. The gap is the one the Centipede game also exposes: backward induction assumes each player maximizes their own monetary payoff and that this is common knowledge, while people also value fairness and will spend to enforce it.

Zermelo's theorem and checkers. *Problem.* Explain why Zermelo's theorem guarantees that exactly one of win, lose, or draw holds for checkers under perfect play, yet does not say which, while Bouton's theorem settles every Nim position in closed form. *Solution.* Zermelo's theorem is an existence result. Checkers is a finite game of perfect information with no chance, so backward induction from its terminal positions assigns every position a definite value, and the opening is therefore a forced win for one side or a forced draw. The argument is non-constructive: it shows the value exists without computing it, because folding the actual game tree of checkers is astronomically large (the game was settled, as a draw, only by a decades-long computation). Nim differs in kind. Bouton's theorem gives the value of every position by a formula, the nim-sum, so no search is needed at all. The contrast is existence against construction: Zermelo guarantees an answer, Bouton hands you one. Nim's special structure, a disjunctive sum governed by binary exclusive-or, is what collapses the search into arithmetic.

Chapter 4

Maximin, minimax, and saddles. *Problem.* For each matrix find the maximin and minimax, decide whether a saddle point exists, and if so give the value and optimal pure strategies: (a) $\begin{bmatrix} 3 & 1 \\ 0 & 2 \end{bmatrix}$; (b) $\begin{bmatrix} -2 & 2 & -1 \\ 5 & 3 & 1 \\ 2 & 0 & 1 \end{bmatrix}$; (c) $\begin{bmatrix} 1 & 2 & 3 & 4 & 5 \\ 3 & 4 & 5 & 6 & 7 \end{bmatrix}$. *Solution.* (a) Row minima $\min(3, 1) = 1$ and $\min(0, 2) = 0$, so maximin 1; column maxima $\max(3, 0) = 3$ and $\max(1, 2) = 2$, so minimax 2. Since $1 \neq 2$ there is no saddle (the game needs mixing). (b) Row minima $-2, 1, 0$, so maximin 1 (row 2); column maxima $5, 3, 1$, so minimax 1 (column 3). They agree, so the entry $a_{23} = 1$ is a saddle: $V = 1$, Player I plays row 2, Player II plays column 3. (c) Row minima 1 and 3, so maximin 3; column maxima 3, 4, 5, 6, 7, so minimax 3. The entry $a_{21} = 3$ is a saddle (row 2 dominates row 1, so the maximizer never leaves it): $V = 3$, with row 2 and column 1.

2×2 solutions. *Problem.* Solve each game: check for a saddle first, and otherwise apply the 2×2 formula $p^* = \frac{d-c}{D}$, $q^* = \frac{d-b}{D}$, $V = \frac{ad-bc}{D}$ with $D = a - b - c + d$: (a) $\begin{bmatrix} 0 & -10 \\ 1 & 2 \end{bmatrix}$; (b) $\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$; (c) $\begin{bmatrix} 2 & -1 \\ -1 & 1 \end{bmatrix}$. *Solution.* (a) Row minima $\min(0, -10) = -10$, $\min(1, 2) = 1$, so maximin 1; column maxima $\max(0, 1) = 1$, $\max(-10, 2) = 2$, so minimax 1. They agree, so $a_{21} = 1$ is a saddle, $V = 1$, with the pure pair (row 2, column 1). (b) Matching Pennies has maximin $-1 \neq \text{minimax } 1$, no saddle; with $a = 1, b = -1, c = -1, d = 1$ and $D = 1 + 1 + 1 + 1 = 4$, the formula gives $p^* = \frac{1-(-1)}{4} = \frac{1}{2}$, $q^* = \frac{1}{2}$, and $V = \frac{1-1}{4} = 0$, a fair game (it is also symmetric under swapping the players). (c) Maximin $\max(-1, -1) = -1 \neq \text{minimax } \min(2, 1) = 1$, no saddle; with $a = 2, b = -1, c = -1, d = 1$ and $D = 2 + 1 + 1 + 1 = 5$, $p^* = \frac{1-(-1)}{5} = \frac{2}{5}$, $q^* = \frac{2}{5}$, $V = \frac{2 \cdot 1 - (-1)(-1)}{5} = \frac{1}{5}$, so $p^* = q^* = (\frac{2}{5}, \frac{3}{5})$.

The value shift. *Problem.* Show that adding a constant c to every entry of a payoff matrix raises the value by c and leaves the optimal strategies unchanged. *Solution.* Adding c to every entry replaces A by $A + c\mathbf{1}\mathbf{1}^\top$. For any strategies p, q , $p^\top(A + c\mathbf{1}\mathbf{1}^\top)q = p^\top Aq + c(p^\top \mathbf{1})(\mathbf{1}^\top q) = p^\top Aq + c$, since $p^\top \mathbf{1} = \mathbf{1}^\top q = 1$. Every payoff shifts by the same constant c , so the maximin and minimax both rise by c (the value rises by c) while the strategies that attain them are unchanged. This is the shift that makes a value positive before forming the linear program of Chapter 6.

A game that needs no mixing. *Problem.* Solve $A = \begin{bmatrix} 6 & 5 \\ 2 & 3 \end{bmatrix}$; decide whether a saddle point exists, give the value and optimal strategies, and say why randomizing does not help. *Solution.* Tabulate the row minima and column maxima:

	c_1	c_2	row min
r_1	6	5	<u>5</u>
r_2	2	3	<u>2</u>
col max	6	5	

The maximin is $\max\{5, 2\} = 5$ (row 1) and the minimax is $\min\{6, 5\} = 5$ (column 2). They are equal, so the entry $a_{12} = 5$ is a saddle point: it is the smallest in its row and the largest in its column. The value is $V = 5$, with optimal pure strategies r_1 for Player I and c_2 for Player II. Randomizing cannot help either side: from r_1 , Player II already holds Player I to 5 by playing c_2 , and Player I already guarantees 5 by playing r_1 , so the pure pair meets the $\underline{V} = \bar{V}$ test of Proposition 4.5 and no mixture beats it.

A 2×2 that does need mixing. *Problem.* Solve $A = \begin{bmatrix} 2 & 5 \\ 6 & 3 \end{bmatrix}$: rule out a saddle, apply the closed form, and verify the value two ways. *Solution.* Row minima are 2 and 3, so the maximin is 3; column maxima are 6 and 5, so the minimax is 5. Since $3 \neq 5$ there is no saddle and the solution is mixed. With $a = 2, b = 5, c = 6, d = 3$ the denominator is $D = a - b - c + d = 2 - 5 - 6 + 3 = -6$, so

$$p^* = \frac{d - c}{D} = \frac{3 - 6}{-6} = \frac{1}{2}, \quad q^* = \frac{d - b}{D} = \frac{3 - 5}{-6} = \frac{1}{3}, \quad V = \frac{ad - bc}{D} = \frac{2 \cdot 3 - 5 \cdot 6}{-6} = \frac{-24}{-6} = 4.$$

Player I plays $\mathbf{p}^* = (\frac{1}{2}, \frac{1}{2})$, Player II plays $\mathbf{q}^* = (\frac{1}{3}, \frac{2}{3})$, and $V = 4$. *Check via Player I's mix against each column:* column 1 pays $\frac{1}{2}(2) + \frac{1}{2}(6) = 4$, column 2 pays $\frac{1}{2}(5) + \frac{1}{2}(3) = 4$. Both equal 4, so $\mathbf{p}^*$ guarantees 4 no matter what Player II does. *Check via Player II's mix against each row:* row 1 pays $\frac{1}{3}(2) + \frac{2}{3}(5) = \frac{12}{3} = 4$, row 2 pays $\frac{1}{3}(6) + \frac{2}{3}(3) = \frac{12}{3} = 4$. Both equal 4, confirming $V = 4$ and the indifference that marks an optimum.

Graphical solution of a 2×3 . *Problem.* Solve $A = \begin{bmatrix} 0 & 2 & 5 \\ 4 & 3 & 1 \end{bmatrix}$ graphically. *Solution.* Let Player I play row 1 with probability p and row 2 with $1 - p$. Against each of Player II's columns the expected payoff to Player I is a line in p :

$$U_{c_1}(p) = 0p + 4(1 - p) = 4 - 4p, \quad U_{c_2}(p) = 2p + 3(1 - p) = 3 - p, \quad U_{c_3}(p) = 5p + 1(1 - p) = 1 + 4p.$$

Player II, the minimizer, will hold Player I to the *lowest* of the three lines at each p ; Player I picks the p that maximizes that lower envelope. At $p = 0$ the lowest is $U_{c_3} = 1$; at $p = 1$ the lowest is $U_{c_1} = 0$. The envelope therefore rises then falls, and its peak is where the two *descending-into* lines cross. The rising line $U_{c_3} = 1 + 4p$ meets the falling line $U_{c_1} = 4 - 4p$ at

$$1 + 4p = 4 - 4p \implies 8p = 3 \implies p^* = \frac{3}{8}, \quad V = 1 + 4 \cdot \frac{3}{8} = \frac{5}{2}.$$

At $p^* = \frac{3}{8}$ the middle line gives $U_{c_2} = 3 - \frac{3}{8} = \frac{21}{8} = 2.625 > \frac{5}{2}$, so column 2 is *slack* (above the envelope) and Player II never uses it: $q_2^* = 0$. Player II mixes only c_1 and c_3 , which is the 2×2 game $\begin{bmatrix} 0 & 5 \\ 4 & 1 \end{bmatrix}$ with $D = 0 - 5 - 4 + 1 = -8$, giving $q_1^* = \frac{d - b}{D} = \frac{1 - 5}{-8} = \frac{1}{2}$. Hence $\mathbf{p}^* = (\frac{3}{8}, \frac{5}{8})$, $\mathbf{q}^* = (\frac{1}{2}, 0, \frac{1}{2})$, and $V = \frac{5}{2}$.

Dominance down to a pure value. *Problem.* Solve $A = \begin{bmatrix} 3 & 2 & 4 \\ 2 & 1 & 3 \\ 5 & 4 & 6 \end{bmatrix}$ by dominance. *Solution.* Compare row 3 with the other rows entry by entry. Row 3 versus row 1: $5 \geq 3$, $4 \geq 2$, $6 \geq 4$, all true, so row 3 dominates row 1. Row 3 versus row 2: $5 \geq 2$, $4 \geq 1$, $6 \geq 3$, all true, so row 3 dominates row 2 as well. Player I (the maximizer) will never play a dominated row, so both fall away in red:

$$\begin{bmatrix} \textcolor{red}{3} & \textcolor{red}{2} & \textcolor{red}{4} \\ \textcolor{red}{2} & \textcolor{red}{1} & \textcolor{red}{3} \\ 5 & 4 & 6 \end{bmatrix} \xrightarrow{\text{keep only } r_3} \begin{bmatrix} 5 & 4 & 6 \end{bmatrix}.$$

Now Player I always plays r_3 , and Player II (the minimizer) reads across the single remaining row and picks its smallest entry, $\min\{5, 4, 6\} = 4$, in column 2. The reduced game has the obvious saddle at (r_3, c_2) , so $V = 4$ with optimal pure strategies r_3 and c_2 . The original matrix *did* have a saddle all along, hidden by the dominated rows: $a_{32} = 4$ is the minimum of its row and the maximum of its column. Dominance simply made it visible.

Reducing to a 2×2 by dominance. *Problem.* Reduce each game by dominance to a 2×2 and solve: (a) $\begin{bmatrix} 2 & 0 & 4 \\ 1 & 2 & 3 \\ 4 & 1 & 2 \end{bmatrix}$,

(b) $\begin{bmatrix} 0 & 4 & 6 \\ 5 & 7 & 4 \\ 9 & 6 & 3 \end{bmatrix}$; for (b) exhibit the mixed column that dominates the deleted one. *Solution.* (a) Player II minimizes, so a column sitting entrywise below another is deletable. Column 2 $(0, 2, 1)$ is $\leq$ column 3 $(4, 3, 2)$ in every row, so delete column 3, then row 3 $(4, 1)$ beats row 1 $(2, 0)$ entrywise, so delete row 1:

$$\begin{bmatrix} 2 & 0 & 4 \\ 1 & 2 & 3 \\ 4 & 1 & 2 \end{bmatrix} \rightarrow \begin{bmatrix} 2 & 0 \\ 1 & 2 \\ 4 & 1 \end{bmatrix} \rightarrow \begin{bmatrix} 1 & 2 \\ 4 & 1 \end{bmatrix}.$$

The surviving 2×2 (rows 2,3 against columns 1,2) has no saddle, and the formula gives $V = \frac{1 \cdot 1 - 2 \cdot 4}{1 + 1 - 2 - 4} = \frac{-7}{-4} = \frac{7}{4}$, with $\mathbf{p}^* = (\frac{3}{4}, \frac{1}{4})$ on rows 2,3 and $\mathbf{q}^* = (\frac{1}{4}, \frac{3}{4})$ on columns 1,2. Restoring the deleted strategies, $\mathbf{p}^* = (0, \frac{3}{4}, \frac{1}{4})$, $\mathbf{q}^* = (\frac{1}{4}, \frac{3}{4}, 0)$, $V = \frac{7}{4}$.

(b) No single column dominates another, but a *mixture* does. The even blend of columns 1 and 3, $\frac{1}{2}(0, 5, 9) + \frac{1}{2}(6, 4, 3) = (3, \frac{9}{2}, 6)$, is $\leq$ column 2 $(4, 7, 6)$ in every row, so column 2 is dominated and deleted:

$$\begin{bmatrix} 0 & 4 & 6 \\ 5 & 7 & 4 \\ 9 & 6 & 3 \end{bmatrix} \rightarrow \begin{bmatrix} 0 & 6 \\ 5 & 4 \\ 9 & 3 \end{bmatrix}.$$

This 3×2 has no row dominance, so solve it graphically: against Player II's weight q on column 1, the three rows pay $6 - 6q$, $4 + q$, and $3 + 6q$. Player II minimizes the upper envelope, whose lowest point is where the falling line $6 - 6q$ meets the steep rising line $3 + 6q$, at $q = \frac{1}{4}$, both equal to $\frac{9}{2}$; the middle row pays $\frac{17}{4} < \frac{9}{2}$ there and is slack. Rows 1,3 are active, and equalizing them sets Player I's weight on each to $\frac{1}{2}$. Restoring the deletions, $\mathbf{p}^* = (\frac{1}{2}, 0, \frac{1}{2})$, $\mathbf{q}^* = (\frac{1}{4}, 0, \frac{3}{4})$, $V = \frac{9}{2}$.

Chapter 5

Rock–Paper–Scissors. *Problem.* Solve the zero-sum game with matrix $\begin{bmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{bmatrix}$. *Solution.* The matrix is skew-symmetric ($A^\top = -A$, with zero diagonal), so by Theorem 5.7 the value is $V = 0$ and $\mathbf{p}^* = \mathbf{q}^*$. Setting each row equal to 0 against $\mathbf{q}^* = (q_1, q_2, q_3)$ gives $-q_2 + q_3 = 0$, $q_1 - q_3 = 0$, $-q_1 + q_2 = 0$, so $q_1 = q_2 = q_3$; with $\sum q_i = 1$ this is $\mathbf{q}^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Hence $\mathbf{p}^* = \mathbf{q}^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, $V = 0$.

Degenerate 2×2 . *Problem.* What does the inversion formula give for $\begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}$, and which assumption fails? *Solution.* Here $\det A = 1 \cdot 1 - 1 \cdot 1 = 0$, so A is singular and the inversion formula $V = 1/(\mathbf{1}^\top A^{-1} \mathbf{1})$ is undefined. The game itself is trivial: every cell pays 1, so $V = 1$ and every strategy is optimal. The failed assumption is that the optimal solution is unique and fully mixed with A invertible; with all entries equal, the support is not pinned down.

A 2×2 by indifference. *Problem.* Solve $A = \begin{bmatrix} 3 & 1 \\ 2 & 5 \end{bmatrix}$ by the principle of indifference and verify the value two ways. *Solution.* First check for a saddle: row minima are $\min(3, 1) = 1$ and $\min(2, 5) = 2$, so $\underline{V} = 2$; column maxima are $\max(3, 2) = 3$ and $\max(1, 5) = 5$, so $\overline{V} = 3$. Since $2 \neq 3$ there is no saddle and the game is mixed. Let Player II play column 1 with probability q . Player I's rows then earn

$$\text{row 1: } 3q + 1(1 - q) = 1 + 2q, \quad \text{row 2: } 2q + 5(1 - q) = 5 - 3q.$$

Indifference $1 + 2q = 5 - 3q$ gives $5q = 4$, so $q = \frac{4}{5}$ and the value is $1 + 2 \cdot \frac{4}{5} = \frac{13}{5}$. Now let Player I play row 1 with probability p ; Player II's columns cost

$$\text{col 1: } 3p + 2(1 - p) = 2 + p, \quad \text{col 2: } 1p + 5(1 - p) = 5 - 4p,$$

and $2 + p = 5 - 4p$ gives $5p = 3$, so $p = \frac{3}{5}$. Thus $\mathbf{p}^* = (\frac{3}{5}, \frac{2}{5})$, $\mathbf{q}^* = (\frac{4}{5}, \frac{1}{5})$, $V = \frac{13}{5}$. Check, Player I's mix vs. each column: col 1 pays $\frac{3}{5} \cdot 3 + \frac{2}{5} \cdot 2 = \frac{9+4}{5} = \frac{13}{5}$; col 2 pays $\frac{3}{5} \cdot 1 + \frac{2}{5} \cdot 5 = \frac{3+10}{5} = \frac{13}{5}$. Check, Player II's mix vs. each row: row 1 pays $\frac{4}{5} \cdot 3 + \frac{1}{5} \cdot 1 = \frac{13}{5}$; row 2 pays $\frac{4}{5} \cdot 2 + \frac{1}{5} \cdot 5 = \frac{13}{5}$. Both mixes lock in $\frac{13}{5}$, confirming the value. (The same answer comes from the closed form: $D = 3 - 1 - 2 + 5 = 5$, $p^* = (d - c)/D = 3/5$, $q^* = (d - b)/D = 4/5$, $V = (ad - bc)/D = 13/5$.)

Solving a 3×3 by the inversion formula. *Problem.* Solve $A = \begin{bmatrix} 1 & 2 & -1 \\ 2 & -1 & 4 \\ -1 & 4 & 3 \end{bmatrix}$ by the inversion formula, assuming every strategy is active. *Solution.* The recipe solves $A\mathbf{x} = \mathbf{1}$, then sets $V = 1/(\mathbf{1}^\top \mathbf{x})$ and $\mathbf{q}^* = \mathbf{x}/(\mathbf{1}^\top \mathbf{x})$. The system

$$x_1 + 2x_2 - x_3 = 1, \quad 2x_1 - x_2 + 4x_3 = 1, \quad -x_1 + 4x_2 + 3x_3 = 1$$

solves to $\mathbf{x} = \frac{1}{23}(11, 7, 2)$ with $\mathbf{1}^\top \mathbf{x} = \frac{20}{23}$, so

$$V = \frac{23}{20}, \quad \mathbf{q}^* = \left(\frac{11}{20}, \frac{7}{20}, \frac{1}{10}\right).$$

Since A is symmetric, Player I's optimal mix is the same, $\mathbf{p}^* = \mathbf{q}^*$. The answer checks: each coordinate of $A\mathbf{q}^*$ equals $\frac{23}{20}$, the indifference condition $A\mathbf{q}^* = V\mathbf{1}$, and all three probabilities are positive, so the full-support assumption was justified.

Two games by the principle of indifference. *Problem.* Solve by indifference: (a) $\begin{bmatrix} 1 & 2 & -1 \\ 0 & 1 & -2 \end{bmatrix}$, (b) $\begin{bmatrix} 1 & 2 & -1 \\ 2 & -1 & 4 \\ -1 & 4 & 3 \end{bmatrix}$. *Solution.* (a) There is no saddle (maximin 0, minimax 1). Equalize the columns by Player I's mix $(p, 1 - p)$: $0p + 1(1 - p) = 1p + (-2)(1 - p)$, that is $1 - p = 3p - 2$, so $p = \frac{3}{4}$ and Player I plays $(\frac{3}{4}, \frac{1}{4})$. The mirror computation (equalizing the rows) gives Player II the same $(\frac{3}{4}, \frac{1}{4})$, and the value is $0 \cdot \frac{3}{4} + 1 \cdot \frac{1}{4} = \frac{1}{4}$. (b) The matrix is symmetric, so the two indifference systems coincide and $\mathbf{p}^* = \mathbf{q}^*$. Solving $A\mathbf{q} = V\mathbf{1}$ at full support gives $\mathbf{q}^* = (\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$, and each row then pays 1: row 1 gives $\frac{1}{4} + 1 - \frac{1}{4}$, row 2 gives $\frac{1}{2} - \frac{1}{2} + 1$, row 3 gives $-\frac{1}{4} + 2 - \frac{3}{4}$, all equal to 1. So $\mathbf{p}^* = \mathbf{q}^* = (\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$, $V = 1$.

Symmetric Colonel Blotto over two posts. *Problem.* Both commanders have three units to spread over two posts; a post is won by whoever sends more units there, and the payoff is posts won minus posts lost. Write the

matrix, use the post-reversal invariance, and decide whether the game is fair. *Solution.* Each commander chooses one of four splits, $(3, 0), (2, 1), (1, 2), (0, 3)$. Whatever the pair of splits, a commander who wins one post has sent fewer units to the other and loses it, or ties both, so every cell of the 4×4 payoff matrix is 0:

$$A = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

The game is trivially symmetric ($A^\top = -A$, both being zero), so by Theorem 5.7 the value is 0 and the game is fair; every strategy is optimal, and the post-reversal symmetry $(a, b) \mapsto (b, a)$ pairing the splits adds no constraint, since nothing distinguishes the two posts. Equal forces over two posts give neither side an edge. Blotto only comes alive when the forces are unequal or the posts more numerous.

Chapter 6

Corners of an LP. *Problem.* Maximize $3y_1 + 2y_2$ subject to $y_1 + y_2 \leq 4$, $y_1 + 3y_2 \leq 6$, $y_1, y_2 \geq 0$, by listing the vertices and evaluating the objective at each. *Solution.* The feasible region is bounded by the two axes and the two slanted lines. Its vertices are the origin $(0, 0)$; $(4, 0)$ where $y_1 + y_2 = 4$ meets the y_1 -axis; $(0, 2)$ where $y_1 + 3y_2 = 6$ meets the y_2 -axis; and the crossing of the two slanted lines, found by subtracting $y_1 + y_2 = 4$ from $y_1 + 3y_2 = 6$ to get $2y_2 = 2$, so $y_2 = 1$, $y_1 = 3$, the point $(3, 1)$. Evaluating $3y_1 + 2y_2$: $(0, 0)$: 0, $(4, 0)$: 12, $(3, 1)$: 11, $(0, 2)$: 4. The maximum is 12, attained at $(4, 0)$.

A game by its LP. *Problem.* For the game $\begin{bmatrix} 3 & 1 \\ 1 & 2 \end{bmatrix}$, write Player I's and Player II's linear programs, solve, and check against the 2×2 formula. *Solution.* With $a = 3, b = 1, c = 1, d = 2$ and $D = a - b - c + d = 3$, the closed form gives $V = \frac{ad-bc}{D} = \frac{6-1}{3} = \frac{5}{3}$, $p^* = \frac{d-c}{D} = \frac{1}{3}$, and $q^* = \frac{d-b}{D} = \frac{1}{3}$, so $\mathbf{p}^* = \mathbf{q}^* = (\frac{1}{3}, \frac{2}{3})$. The two LPs are $\min \mathbf{1}^\top \mathbf{x}$ s.t. $A^\top \mathbf{x} \geq \mathbf{1}$ and $\max \mathbf{1}^\top \mathbf{y}$ s.t. $A\mathbf{y} \leq \mathbf{1}$; both attain optimal objective $\frac{1}{V} = \frac{3}{5}$, and rescaling $\mathbf{p}^* = V\mathbf{x}^*$, $\mathbf{q}^* = V\mathbf{y}^*$ recovers the same strategies, confirming the agreement guaranteed by strong duality.

The shift trick. *Problem.* The game $A = \begin{bmatrix} -1 & 2 \\ 3 & -2 \end{bmatrix}$ has entries of both signs. Add a constant to make all entries positive, solve the shifted game, and recover the value of the original; confirm the optimal strategies are unchanged. *Solution.* Add 3 to every entry to get $A' = \begin{bmatrix} 2 & 5 \\ 6 & 1 \end{bmatrix}$, all positive. It has no saddle (maximin 2, minimax 5), and with $a = 2, b = 5, c = 6, d = 1$, $D = 2 - 5 - 6 + 1 = -8$, so $V' = \frac{ad-bc}{D} = \frac{2-30}{-8} = \frac{7}{2}$, $p^* = \frac{d-c}{D} = \frac{1-6}{-8} = \frac{5}{8}$, $q^* = \frac{d-b}{D} = \frac{1-5}{-8} = \frac{1}{2}$. Since adding a constant raises the value by that constant and leaves strategies fixed, the original value is $V = V' - 3 = \frac{7}{2} - 3 = \frac{1}{2}$, with the same $\mathbf{p}^* = (\frac{5}{8}, \frac{3}{8})$, $\mathbf{q}^* = (\frac{1}{2}, \frac{1}{2})$. (A direct solve of A confirms it: $D = -1 - 2 - 3 - 2 = -8$ and $V = \frac{(-1)(-2)-2 \cdot 3}{-8} = \frac{-4}{-8} = \frac{1}{2}$.)

Complementary slackness in a solved game. *Problem.* For the game $A = \begin{bmatrix} 2 & 1 \\ 1 & 3 \end{bmatrix}$ with value $V = \frac{5}{3}$ and optimal $\mathbf{p}^* = \mathbf{q}^* = (\frac{2}{3}, \frac{1}{3})$, verify complementary slackness directly. *Solution.* Both of Player I's pure rows earn exactly the value against $\mathbf{q}^*$:

$$\text{row 1: } 2 \cdot \frac{2}{3} + 1 \cdot \frac{1}{3} = \frac{5}{3}, \quad \text{row 2: } 1 \cdot \frac{2}{3} + 3 \cdot \frac{1}{3} = \frac{5}{3}.$$

Neither row beats the value and both meet it, which is the principle of indifference: a strategy played with positive probability must earn the value. In the language of the linear program, each row constraint of Player II's program is tight, so the matching primal weight p_i^* is free to be positive, and indeed both are. The same calculation on the columns shows both of Player II's columns cost exactly $\frac{5}{3}$ against $\mathbf{p}^*$. Complementary slackness is the indifference principle read through duality: positive weight forces a tight constraint, and a slack constraint forces zero weight.

A dual bound by weak duality. *Problem.* State the dual of "maximize $3y_1 + 2y_2$ subject to $y_1 + y_2 \leq 4$, $y_1 + 3y_2 \leq 6$, $y \geq 0$," and use a feasible dual point to bound the primal optimum. *Solution.* The dual assigns a price $x_i \geq 0$ to each primal constraint and reads

$$\min 4x_1 + 6x_2 \quad \text{subject to} \quad x_1 + x_2 \geq 3, \quad x_1 + 3x_2 \geq 2, \quad \mathbf{x} \geq \mathbf{0}.$$

The point $\mathbf{x} = (3, 0)$ is dual-feasible ($3 \geq 3$ and $3 \geq 2$) with objective $4 \cdot 3 = 12$, so by weak duality every primal-feasible $\mathbf{y}$ obeys $3y_1 + 2y_2 \leq 12$. The bound is met: the primal vertex $\mathbf{y} = (4, 0)$ is feasible and scores 12, so 12 is the common optimum, pinned from above by one dual point before any vertex search.

Why duality is minimax. *Problem.* Explain why strong duality for linear programs is equivalent to the equality $\underline{V} = \bar{V}$ of the minimax theorem, and say which direction is weak duality. *Solution.* Casting a matrix game as the dual pair of Section 6.4 sends Player I's guarantee $\underline{V}$ to the optimum of one program and Player II's ceiling $\bar{V}$ to the

optimum of its dual. Weak duality, that every feasible dual value bounds every feasible primal value, becomes the easy inequality $\underline{V} \leq \overline{V}$: whatever Player I can guarantee is at most what Player II can hold them to. That direction is immediate. The substance is strong duality, that the two optima are *equal*, which becomes $\underline{V} = \overline{V} = V$, the minimax theorem. A game has a value for exactly the reason a feasible, bounded linear program has no duality gap.

Solving an LP by checking corners. *Problem.* Maximize $3y_1 + 2y_2$ subject to $y_1 + y_2 \leq 4$, $y_1 + 3y_2 \leq 6$, $\mathbf{y} \geq \mathbf{0}$. List the vertices and the objective at each. *Solution.* The feasible region is a quadrilateral; evaluating the objective at its corners,

vertex	$3y_1 + 2y_2$
(0, 0)	0
(4, 0)	12
(3, 1)	11
(0, 2)	4

where (3, 1) solves $y_1 + y_2 = 4$ and $y_1 + 3y_2 = 6$ together. A linear objective attains its maximum at a vertex, and the largest value is 12 at (4, 0). (The companion dual point in the weak-duality solution above also scores 12, confirming optimality.)

A simplex walk. *Problem.* Apply the simplex method to $\max y_1 + y_2$ subject to $2y_1 + y_2 \leq 1$, $y_1 + 3y_2 \leq 1$, $\mathbf{y} \geq \mathbf{0}$, starting at $\mathbf{y} = \mathbf{0}$, and confirm the optimum $\mathbf{y}^* = (\frac{2}{5}, \frac{1}{5})$. *Solution.* Add slacks s_1, s_2 and start at the origin with basis $\{s_1, s_2\}$ and objective 0. Both y_1 and y_2 would improve it. *Pivot 1:* enter y_1 ; the ratio test $\min(\frac{1}{2}, \frac{1}{1})$ selects the first constraint, so y_1 replaces s_1 , moving to $(\frac{1}{2}, 0)$ with objective $\frac{1}{2}$. The reduced cost of y_2 is still negative. *Pivot 2:* enter y_2 ; the ratio test now selects the second constraint, y_2 replaces s_2 , and the objective rises to $\frac{3}{5}$. All reduced costs are then nonnegative, so the vertex is optimal: $\mathbf{y}^* = (\frac{2}{5}, \frac{1}{5})$ with value $\frac{3}{5}$. As this is Player II's program for the game of value $\frac{5}{3}$, the strategy is $\mathbf{q}^* = V\mathbf{y}^* = \frac{5}{3}(\frac{2}{5}, \frac{1}{5}) = (\frac{2}{3}, \frac{1}{3})$.

Chapter 7

A recursive value. *Problem.* Solve the role-switching recursive game whose payoff matrix to Player I is $\begin{bmatrix} \text{repeat, roles switched} & 2 \\ 3 & 0 \end{bmatrix}$, where the top-left outcome returns play to the same game with the players exchanged. *Solution.* The role switch replaces the top-left cell by $-v$, giving the effective matrix $\begin{bmatrix} -v & 2 \\ 3 & 0 \end{bmatrix}$, which has no saddle for $v \geq 0$ (maximin 0, minimax 2). The 2×2 formula gives $v = \frac{(-v)(0) - 2 \cdot 3}{-v - 2 - 3 + 0} = \frac{-6}{-v-5} = \frac{6}{v+5}$, so $v(v+5) = 6$, that is $v^2 + 5v - 6 = (v+6)(v-1) = 0$, with roots $v = 1$ and $v = -6$. Player I can guarantee 0 by the bottom row (3, 0), so the security bound $v \geq 0$ excludes $v = -6$ and the value is $v = 1$.

Inspection. *Problem.* An inspector with limited inspections faces a thief over n time slots, with stage matrix $M_n = \begin{bmatrix} 1 & 0 \\ 0 & G_{n-1} \end{bmatrix}$ and $G_1 = 1$. Find the value G_n for all n . *Solution.* The stage matrix is $M_n = \begin{bmatrix} 1 & 0 \\ 0 & G_{n-1} \end{bmatrix}$ with $G_1 = 1$. Each has no saddle, so the 2×2 formula gives $G_n = \frac{1 \cdot G_{n-1} - 0}{1 - 0 - 0 + G_{n-1}} = \frac{G_{n-1}}{1 + G_{n-1}}$. Thus $G_2 = \frac{1}{1+1} = \frac{1}{2}$ and $G_3 = \frac{1/2}{1+1/2} = \frac{1}{3}$, and by induction $G_{n-1} = 1/(n-1)$ gives $G_n = \frac{1/(n-1)}{1+1/(n-1)} = \frac{1}{n}$, confirming $V_N = 1/N$.

A continuum of self-consistent values. *Problem.* Explain why a recursive game can have a whole interval of self-consistent values while an ordinary finite matrix game has a single value. *Solution.* A finite matrix game has one value V by the minimax theorem, the unique saddle value of $\max_p \min_q$. A recursive game replaces some payoff cells with “play again,” so its value v must satisfy a fixed-point equation $v = V(M(v))$, where the matrix $M(v)$ carries v itself in the recursive cells. When that equation collapses to an identity over a range, as it does when a recursive cell leaves the reduced equation reading $v = v$, every number in the range solves it, and each is genuinely self-consistent: a strategy promising v is sustained by a continuation that also pays v . The multiplicity is not a flaw in the analysis but a feature of self-reference. Selecting a single value then needs an extra principle, such as the security level each player can actually guarantee, which picks one point out of the interval.

Chapter 8

Battle of the Sexes. *Problem.* For $(A, B) = \begin{bmatrix} (2,1) & (0,0) \\ (0,0) & (1,2) \end{bmatrix}$ (rows/columns F, R), confirm the two pure equilibria and derive the mixed equilibrium and its payoffs. *Solution.* The pure equilibria are the coordinated cells (F, F) , paying (2, 1), and (R, R) , paying (1, 2), since at each neither player gains by deviating alone. For the mixed equilibrium let Player I play F with probability p ; Player II is indifferent when her F payoff $1 \cdot p + 0(1-p) = p$ equals her R payoff $0 \cdot p + 2(1-p) = 2 - 2p$, i.e. $p = 2 - 2p$, so $p = \frac{2}{3}$. Let Player II play F with probability q ; Player I is indifferent

when $2q = 1 - q$, so $q = \frac{1}{3}$. The mixed payoffs are $u_1 = 2q = \frac{2}{3}$ and $u_2 = p = \frac{2}{3}$, below either pure equilibrium, the standard tension of coordination under conflicting preferences.

Coordination game $\begin{bmatrix} (2,2) & (0,0) \\ (0,0) & (1,1) \end{bmatrix}$. The two pure equilibria are the diagonal cells (2,2) and (1,1). For the mixed one, let Player I play the top row (option A) with probability p . Player II is indifferent when her column A payoff $2p + 0(1 - p) = 2p$ equals her column B payoff $0 \cdot p + 1(1 - p) = 1 - p$, that is $2p = 1 - p$, giving $p = \frac{1}{3}$; by symmetry $q = \frac{1}{3}$. The mixed equilibrium pays each player $2 \cdot \frac{1}{3} = \frac{2}{3}$, worse than either pure equilibrium. Three equilibria coexist, and which is reached is a question of coordination rather than of the equilibrium concept, since all three are stable against unilateral deviation.

Safety levels. *Problem.* For $(A, B) = \begin{bmatrix} (1,2) & (3,1) \\ (2,0) & (0,3) \end{bmatrix}$, compute both safety levels $V_I = V(A)$ and $V_{II} = V(B^\top)$ and the safety strategies that achieve them. *Solution.* Player I's matrix is $A = \begin{bmatrix} 1 & 3 \\ 2 & 0 \end{bmatrix}$. Its row minima are 1 and 0 (maximin 1) and column maxima 2 and 3 (minimax 2); no saddle, so with $D = 1 - 3 - 2 + 0 = -4$, $V_I = \frac{1 \cdot 0 - 3 \cdot 2}{-4} = \frac{-6}{-4} = \frac{3}{2}$ and $p^* = \frac{0 - 2}{-4} = \frac{1}{2}$, giving the safety mix $(\frac{1}{2}, \frac{1}{2})$. Player II's matrix transposed is $B^\top = \begin{bmatrix} 2 & 0 \\ 1 & 3 \end{bmatrix}$ (Player II is its row player); with $D = 2 - 0 - 1 + 3 = 4$, $V_{II} = \frac{2 \cdot 3 - 0 \cdot 1}{4} = \frac{3}{2}$ and its row probability $\frac{3 - 1}{4} = \frac{1}{2}$, so Player II's safety mix is also $(\frac{1}{2}, \frac{1}{2})$. Both players can guarantee $\frac{3}{2}$.

One-dimensional Brouwer. *Problem.* Prove that every continuous $f : [0, 1] \rightarrow [0, 1]$ has a fixed point, and give a continuous map on $(0, 1]$ with none, identifying the hypothesis that fails. *Solution.* Let $g(x) = f(x) - x$. Since f maps into $[0, 1]$, $g(0) = f(0) \geq 0$ and $g(1) = f(1) - 1 \leq 0$. As g is continuous, the intermediate value theorem gives an x^* with $g(x^*) = 0$, that is $f(x^*) = x^*$. On the half-open interval $(0, 1]$ the map $f(x) = x/2$ is continuous but its only fixed point is $0 \notin (0, 1]$, so it has none in the set; the failed hypothesis of Brouwer's theorem is compactness, since $(0, 1]$ is not closed.

All equilibria of a coordination game. *Problem.* For $(A, B) = \begin{bmatrix} (3,3) & (0,2) \\ (2,1) & (5,5) \end{bmatrix}$, find every Nash equilibrium and check that the mixed-equilibrium payoffs equal the two safety levels. *Solution.* The diagonal cells are pure equilibria: at (3,3) a deviation pays $2 < 3$ for each player, and at (5,5) deviations pay $2 < 5$ and $1 < 5$. For the mixed equilibrium, let Player I play row 1 with probability p and Player II play column 1 with probability q . Player II is indifferent when $1 + 2p = 5 - 3p$, giving $p = \frac{4}{5}$; Player I is indifferent when $3q = 5 - 3q$, giving $q = \frac{5}{6}$. The mixed equilibrium $((\frac{4}{5}, \frac{1}{5}), (\frac{5}{6}, \frac{1}{6}))$ pays Player I $3q = \frac{5}{2}$ and Player II $1 + 2p = \frac{13}{5}$. The safety levels are the values of the one-player matrices, $V(A) = V[\begin{smallmatrix} 3 & 0 \\ 2 & 5 \end{smallmatrix}] = \frac{15}{6} = \frac{5}{2}$ and $V(B^\top) = V[\begin{smallmatrix} 3 & 1 \\ 2 & 5 \end{smallmatrix}] = \frac{13}{5}$, exactly the equilibrium payoffs. The match is forced: at a fully mixed equilibrium each player's mix makes the opponent indifferent across all pure replies, so that mix already pins the opponent to one number whatever they do. The equilibrium strategy *is* the safety strategy, and the equilibrium payoff is the safety level.

Nash's map on a 2×2 game. *Problem.* Write Nash's map (8.1) explicitly for Player I in a 2×2 game, and check that a profile is a fixed point exactly when neither player has a profitable pure deviation. *Solution.* Fix Player II's mix and let Player I's two rows earn expected payoffs r_1 and r_2 , so a mix $(p, 1 - p)$ earns $u = pr_1 + (1 - p)r_2$. The gains are

$$\varphi_1 = \max\{0, r_1 - u\} = \max\{0, (1 - p)(r_1 - r_2)\}, \quad \varphi_2 = \max\{0, r_2 - u\} = \max\{0, p(r_2 - r_1)\},$$

and the map sends $p \mapsto (p + \varphi_1)/(1 + \varphi_1 + \varphi_2)$. At a fixed point this reduces to $p(\varphi_1 + \varphi_2) = \varphi_1$. If $r_1 > r_2$, then $\varphi_2 = 0$ and $\varphi_1 = (1 - p)(r_1 - r_2)$, so the equation forces $\varphi_1 = 0$, that is $p = 1$: Player I already plays the better row and no deviation helps. If $r_1 = r_2$, both gains vanish for every p , and again no deviation helps. The case $r_2 > r_1$ is symmetric and forces $p = 0$. So Player I sits at a fixed point exactly when all her weight rests on a best pure row, which says no pure deviation pays. The same holds for Player II, so a joint fixed point is precisely a profile in which neither player gains by a pure deviation, a Nash equilibrium.

Chapter 9

A transferable-utility division. *Problem.* For the TU game $(A, B) = \begin{bmatrix} (1,5) & (2,2) & (0,1) \\ (4,2) & (1,0) & (2,1) \\ (5,0) & (2,3) & (0,0) \end{bmatrix}$, find the cooperative value σ , the threat value $\delta = V(A - B)$, and the TU solution φ . *Solution.* The cooperative value is the largest cell total $a_{ij} + b_{ij}$; scanning, the totals are 6, 4, 1; 6, 1, 3; 5, 5, 0, so $\sigma = 6$ (attained at cell (1, 1), payoff (1, 5), or cell (2, 1), payoff (4, 2)). The threat game is

$$A - B = \begin{bmatrix} 1-5 & 2-2 & 0-1 \\ 4-2 & 1-0 & 2-1 \\ 5-0 & 2-3 & 0-0 \end{bmatrix} = \begin{bmatrix} -4 & 0 & -1 \\ 2 & 1 & 1 \\ 5 & -1 & 0 \end{bmatrix}.$$

Reduce by dominance: row 2 = (2, 1, 1) dominates row 1 = (-4, 0, -1) entrywise, so drop row 1; in what remains, column 2 = (1, -1) is entrywise below both column 1 = (2, 5) and column 3 = (1, 0), so Player II keeps only column 2, leaving the single column (1, -1) on which Player I takes the larger, 1. Equivalently the entry $a_{22} = 1$ of $A - B$ is a saddle (the minimum 1 of its row and the maximum 1 of its column), so $\delta = V(A - B) = 1$. The TU solution is

$$\varphi = \left(\frac{\sigma + \delta}{2}, \frac{\sigma - \delta}{2} \right) = \left(\frac{6+1}{2}, \frac{6-1}{2} \right) = \left(\frac{7}{2}, \frac{5}{2} \right).$$

If the players realize σ at the cooperative cell (2, 1), paying (4, 2), Player I transfers $4 - \frac{7}{2} = \frac{1}{2}$ to Player II to reach φ .

Transferable Prisoner's Dilemma. *Problem.* For the cooperative Prisoner's Dilemma $(A, B) = \begin{bmatrix} (1,1) & (4,0) \\ (2,4) & (3,3) \end{bmatrix}$, find the transferable-utility solution. *Solution.* Cooperation maximizes the total, $\sigma = \max_{i,j}(a_{ij} + b_{ij}) = 6$ at the cell (3, 3). The threat game is $A - B = \begin{bmatrix} 0 & 4 \\ -2 & 0 \end{bmatrix}$, whose top-left entry 0 is the minimum of its row and the maximum of its column (maximin 0 = minimax 0), a saddle, so $\delta = V(A - B) = 0$. The TU value is $\varphi = \left(\frac{\sigma + \delta}{2}, \frac{\sigma - \delta}{2} \right) = (3, 3)$, requiring no side payment, and it recovers the efficient outcome that non-cooperative play (Example 8.9) destroyed.

Bargaining. *Problem.* On the feasible triangle with vertices (0, 0), (3, 0), (0, 1) and disagreement point (0, 0), find the Nash bargaining solution, and say how raising Player I's disagreement payoff moves it. *Solution.* The Pareto frontier is the segment from (3, 0) to (0, 1), the line $v = -\frac{1}{3}u + 1$. Maximize the Nash product uv on it. The tangency condition equates the level curve's slope $-v/u$ to the frontier's slope $-\frac{1}{3}$, giving $\bar{u} = 3\bar{v}$; substituting into the frontier, $\bar{v} = -\frac{1}{3}(3\bar{v}) + 1 = 1 - \bar{v}$, so $2\bar{v} = 1$, $\bar{v} = \frac{1}{2}$ and $\bar{u} = \frac{3}{2}$. Raising Player I's disagreement payoff u^* above 0 enlarges the gain term $(u - u^*)$ less favorably for him only where the frontier is steep, and the maximizer shifts toward Player I, awarding him more.

A merger split. *Problem.* Two divisions pool budgets for a project worth \$12M, so the frontier is $u_1 + u_2 = 12$; apart they deliver \$5M and \$1M, so $d = (5, 1)$. Find the Nash bargaining split. *Solution.* Maximize the gain product $N = (u_1 - 5)(u_2 - 1)$ along $u_2 = 12 - u_1$:

$$N(u_1) = (u_1 - 5)(11 - u_1), \quad N'(u_1) = (11 - u_1) - (u_1 - 5) = 16 - 2u_1 = 0 \Rightarrow u_1 = 8,$$

so $u_2 = 4$ and the split is (8, 4) with product (3)(3) = 9. The “floor plus half the surplus” rule agrees: the surplus is $s = 12 - 5 - 1 = 6$, and $u_1 = 5 + \frac{s}{2} = 5 + 3 = 8$, $u_2 = 1 + \frac{s}{2} = 1 + 3 = 4$. Division 1's larger stand-alone value, rather than any difference in bargaining skill, is what earns it the larger share.

The Nash bargaining axioms. *Problem.* State the four axioms that characterize the Nash bargaining solution, and check that the solution of the triangle problem (frontier $v = 1 - u/3$, disagreement (0, 0)) is unchanged when Player I's payoffs are doubled. *Solution.* Nash's solution is the unique rule obeying (1) *Pareto efficiency*, the outcome sits on the upper-right frontier; (2) *symmetry*, a problem unchanged by swapping the players splits equally; (3) *invariance to affine rescaling*, applying a positive affine map to one player's utility rescales that player's payoff and leaves the rest untouched; and (4) *independence of irrelevant alternatives*, deleting unchosen options does not move the solution. On the triangle the solution maximizes the product uv along $v = 1 - u/3$, that is $u(1 - u/3)$, whose derivative $1 - 2u/3$ vanishes at $u = \frac{3}{2}$, giving the split $(\frac{3}{2}, \frac{1}{2})$. Doubling Player I's payoffs sends $u \mapsto u' = 2u$, so the frontier becomes $v = 1 - u'/6$ and the product $u'(1 - u'/6)$ peaks at $u' = 3$, $v = \frac{1}{2}$. Back in the original units $u = u'/2 = \frac{3}{2}$: the real split is exactly as before. The recorded number scales with the units, the agreement does not, which is precisely axiom (3).

The exchange rate in a two-currency NTU game. *Problem.* Explain why the exchange rate λ cannot be chosen freely, and describe how the NTU split $(\sigma + \frac{\delta}{2\lambda}, \sigma - \frac{\delta}{2\lambda})$ depends on it. *Solution.* The rate λ converts player 2's currency into player 1's so that “transferable” utility means the same thing on both scales. As λ grows, the transfer $\delta/(2\lambda)$ shrinks and the split slides toward the equal division (σ, σ) ; as $\lambda \rightarrow 0$ the transfer blows up and the split spreads far

apart. A freely chosen λ could therefore name almost any division, which is exactly why it cannot be arbitrary. It is pinned by feasibility: the correct λ is the one at which the bargaining solution's tangency direction matches the slope of the NTU frontier, so the side payment the formula calls for is actually achievable inside the feasible set. At that rate, and only that rate, the transferable-utility computation lands on the NTU frontier instead of outside it.

Chapter 10

A three-player Shapley value by arrival orders. *Problem.* Compute the Shapley value of the game with $v(\{i\}) = 0$ for each i , $v(\{1, 2\}) = v(\{1, 3\}) = 0$, $v(\{2, 3\}) = 100$, and $v(N) = 200$, and compare it with the nucleolus $(50, 75, 75)$. *Solution.* Average each player's marginal contribution over the six arrival orders. The marginal triples $(\Delta_1, \Delta_2, \Delta_3)$ are

order	Δ_1	Δ_2	Δ_3
123	0	0	200
132	0	200	0
213	0	0	200
231	100	0	100
312	0	200	0
321	100	100	0

For instance, in order 231 player 3 joins $\{2\}$ and adds $v(\{2, 3\}) - v(\{2\}) = 100$, then player 1 completes the grand coalition and adds $v(N) - v(\{2, 3\}) = 100$. Summing each column and dividing by six,

$$\varphi = \left(\frac{200}{6}, \frac{500}{6}, \frac{500}{6} \right) = \left(\frac{100}{3}, \frac{250}{3}, \frac{250}{3} \right) \approx (33.3, 83.3, 83.3),$$

which sums to $200 = v(N)$. The Shapley value rewards players 2 and 3 for the valuable pair they form, splitting the surplus evenly between them, and pays player 1 only what they add to the grand coalition. The nucleolus $(50, 75, 75)$ is more even, because it protects player 1 by minimizing the largest coalition's complaint rather than averaging marginal contributions.

The Shapley value of a unanimity game. *Problem.* For the unanimity game W_T , worth 1 to every coalition that contains T and 0 otherwise, prove from the axioms that $\varphi_i(W_T) = 1/|T|$ for $i \in T$ and 0 for $i \notin T$, naming the axiom that forces each part. *Solution.* A player $i \notin T$ is a *null player*: adjoining them to any coalition S leaves $W_T(S \cup \{i\}) = W_T(S)$, since i is never needed to contain T . The null-player axiom gives $\varphi_i(W_T) = 0$ for every $i \notin T$. Any two players $i, j \in T$ are *symmetric*: for a coalition S holding neither, $W_T(S \cup \{i\}) = W_T(S \cup \{j\}) = 0$, because the other member of T is still absent, so the two additions are interchangeable. The symmetry axiom makes their values equal, so all $|T|$ members of T receive the same amount. Efficiency fixes that amount: the values sum to $W_T(N) = 1$ (as $T \subseteq N$), and $|T|$ equal shares of 1 are each $1/|T|$. Null forces the zeros, symmetry the equality, efficiency the common value $1/|T|$.

Shapley in the core. *Problem.* For the three-player game of Example 10.5 ($v(\{1\}) = 1, v(\{2\}) = 0, v(\{3\}) = 1, v(\{1, 2\}) = 4, v(\{1, 3\}) = 3, v(\{2, 3\}) = 5, v(N) = 8$), verify that the Shapley value lies in the core, and exhibit another core point. *Solution.* The Shapley value is $\varphi = (\frac{7}{3}, \frac{17}{6}, \frac{17}{6})$. Check the core inequalities: $x_1 = \frac{7}{3} \leq 3$, $x_2 = \frac{17}{6} \leq 5$, $x_3 = \frac{17}{6} \leq 4$, and all pair-sums exceed their worths, so the Shapley value lies in the core. A different core point is, for instance, $(2, 3, 3)$, which also satisfies every constraint and sums to 8.

ClosedAI bankruptcy. With claims $(100, 300, 500)$ and $E = 500$, the worths are $v(\{3\}) = 100, v(\{1, 3\}) = 200, v(\{2, 3\}) = 400, v(N) = 500$, others 0. Averaging marginal contributions gives $\varphi = (50, 150, 300)$. The nucleolus balances the opposing complaints of $\{1\}$ (excess $-x_1$) and $\{2, 3\}$ (excess $x_1 - 100$) at $x_1 = 50$, then $\{2\}$ against $\{1, 3\}$ at $x_2 = 150, x_3 = 300$, so $\mathcal{N} = (50, 150, 300)$ as well: the two solutions coincide at this estate. Lowering the estate to $E = 200$ leaves only $v(\{2, 3\}) = 100$ and $v(N) = 200$, where $\varphi = (\frac{100}{3}, \frac{250}{3}, \frac{250}{3})$ but $\mathcal{N} = (50, 75, 75)$, so a scarcer estate makes the average-contribution and equalize-complaints principles diverge.

An empty core. Take $v(\{i\}) = 0$ for each player, $v(S) = 1$ for every two-player coalition, and $v(N) = 1$. A core allocation would need $x_1 + x_2 + x_3 = 1$ together with $x_1 + x_2 \geq 1, x_1 + x_3 \geq 1$, and $x_2 + x_3 \geq 1$. Adding the three pair inequalities gives $2(x_1 + x_2 + x_3) \geq 3$, that is $x_1 + x_2 + x_3 \geq \frac{3}{2}$; but efficiency fixes the sum at $1 < \frac{3}{2}$, a contradiction. The core is empty: any proposed split of the single unit leaves some pair able to claim more than it receives, so no allocation is stable. (This is the "majority" game, where any two players can seize everything.)

The Talmud at $E = 300$. *Problem.* For claims $(100, 200, 300)$ and estate $E = 300$, build the bankruptcy game $v(S) = \max\{0, E - \sum_{i \notin S} c_i\}$ and verify its nucleolus is $(50, 100, 150)$, the third row of the Talmud table. *Solution.* A

coalition is worth what the estate leaves after the outsiders are paid. Computing, $v(\{1, 3\}) = \max\{0, 300 - 200\} = 100$, $v(\{2, 3\}) = \max\{0, 300 - 100\} = 200$, $v(N) = 300$, and every other proper coalition is 0. Use efficiency $x_1 + x_2 + x_3 = 300$ and drive down the largest excess. The two strongest opposing complaints are $\{1\}$, with excess $-x_1$, and $\{2, 3\}$, with excess $200 - (x_2 + x_3) = 200 - (300 - x_1) = x_1 - 100$; balancing, $-x_1 = x_1 - 100$ gives $x_1 = 50$ and a top excess of -50 . With $x_1 = 50$ fixed, the next pair to balance is $\{2\}$ (excess $-x_2$) against $\{1, 3\}$ (excess $100 - (50 + x_3) = x_2 - 200$, using $x_2 + x_3 = 250$); balancing, $-x_2 = x_2 - 200$ gives $x_2 = 100$, hence $x_3 = 150$. The nucleolus is $(50, 100, 150)$, matching the Talmud.

A three-player Shapley value. *Problem.* For $v(\{1, 2\}) = 4$, $v(\{1, 3\}) = 6$, $v(\{2, 3\}) = 8$, $v(N) = 12$, all singletons 0, compute the Shapley value and check it lies in the core. *Solution.* Average each player's marginal contribution over the $3! = 6$ arrival orders. In order 1, 2, 3, player 3's marginal is $v(N) - v(\{1, 2\}) = 12 - 4 = 8$, and so on:

order	1	2	3
1, 2, 3	0	4	8
1, 3, 2	0	6	6
2, 1, 3	4	0	8
2, 3, 1	4	0	8
3, 1, 2	6	6	0
3, 2, 1	4	8	0
sum	18	24	30
$\div 6$	3	4	5

The Shapley value is $\varphi = (3, 4, 5)$, which sums to $12 = v(N)$ (efficiency). It lies in the core: singletons get $3, 4, 5 \geq 0$; the pairs get $\varphi_1 + \varphi_2 = 7 \geq 4$, $\varphi_1 + \varphi_3 = 8 \geq 6$, and $\varphi_2 + \varphi_3 = 9 \geq 8$, so no coalition can object. Here the fair allocation is also stable, which need not happen in general (Section 10.7).

APPENDIX E

Glossary

A concise reference to the central terms, gathered in one place. The index points to where each is developed in full.

Backward induction The procedure that solves a finite game of perfect information by working from the leaves to the root, replacing each subgame by the payoff its last mover would secure (Procedure 3.6). It produces a subgame perfect equilibrium.

Balanced collection A family of coalitions with nonnegative weights that sum to one across every player. A game is *balanced* when no balanced collection can manufacture more worth than the grand coalition; by the Bondareva–Shapley theorem this holds exactly when the core is nonempty.

Banzhaf index A measure of voting power that counts the *swings* of each player, the coalitions they turn from losing to winning by joining, normalized to sum to one. It is a companion to the Shapley–Shubik index and can rank power differently.

Bayesian game A game of incomplete information in which Nature first draws each player’s private *type*, Harsanyi’s device for analyzing uncertainty about payoffs.

Behavioral strategy A rule that randomizes independently at each information set of an extensive-form game. Under perfect recall, Kuhn’s theorem makes behavioral and mixed strategies equivalent.

Best response A strategy that maximizes a player’s payoff against a fixed profile of the others. A Nash equilibrium is a profile in which every strategy is a best response to the rest.

Bimatrix game A two-player game given by a pair of payoff matrices (A, B) , one for each player. It is zero-sum exactly when $B = -A$.

Brouwer’s fixed-point theorem Every continuous self-map of a nonempty compact convex set has a fixed point, the topological fact behind Nash’s existence proof.

Characteristic function The function v assigning to each coalition S the worth $v(S)$ it can guarantee on its own. It is the data of a coalitional game.

Coalition A subset of the players who may coordinate. The whole player set N is the *grand coalition*.

Combinatorial game A two-player game of perfect information with no chance and no hidden state, in which players alternate moves. Under normal play the last player to move wins.

Common knowledge of rationality The assumption that all players are rational, all know it, all know that all know it, and so on without end, the condition backward induction quietly requires.

Complementary slackness The duality condition that a positive primal variable forces its dual constraint to bind, and a slack constraint forces its variable to zero. For matrix games it is the principle of indifference: a strategy used with positive weight earns exactly the value.

Core The set of imputations no coalition can improve upon, every coalition receiving at least its worth. It may be empty (Definition 10.4).

Correlated equilibrium A joint distribution over action profiles such that, when a referee privately recommends each player their part of a draw, no one wishes to deviate. Every Nash equilibrium is one, but correlation can reach payoffs none of them attains.

Disjunctive sum The game in which a move is a legal move in exactly one component. By the Sprague–Grundy theorem its Grundy value is the nim-sum of the components’ values.

Dominance Strategy s *strictly dominates* s' if it pays more against every opponent choice, and *weakly dominates* if it pays at least as much always and strictly more somewhere. A rational player never uses a strictly dominated strategy.

Dominant strategy A strategy that is a best response to every choice of the others. A player who has one needs no beliefs about the rest.

- Duality** The pairing of a linear program with a second program over its constraints. Weak duality bounds one by the other; strong duality makes their optima equal, which for matrix games is the minimax theorem.
- Evolutionarily stable strategy** A strategy that, once adopted by a whole population, resists invasion by any small group of mutants. Every such strategy is a Nash equilibrium with an added margin of stability.
- Extensive form** The representation of a game as a tree of moves, with decision nodes, chance nodes, information sets, and payoffs at the leaves.
- Fictitious play** A learning process in which each player best-responds to the empirical frequency of the opponent's past moves. In zero-sum games it converges to the value (Robinson's theorem).
- Focal point** An equilibrium the players select by shared expectation or salience rather than by the payoffs alone, Schelling's resolution of a coordination problem.
- Folk theorem** The result that an infinitely repeated game has subgame perfect equilibria realizing essentially any feasible, individually rational payoff once the players are patient enough.
- Grundy value** The value $g(x) = \text{mex}\{g(y) : y \text{ a follower of } x\}$ of a position in an impartial game. A position is a P-position exactly when its Grundy value is zero.
- Imputation** A payment vector that is efficient (it distributes the grand coalition's worth) and individually rational (no player gets less than they could secure alone).
- Information set** A collection of decision nodes a player cannot tell apart when moving, the device that encodes imperfect information in the extensive form.
- Kakutani's fixed-point theorem** The set-valued generalization of Brouwer's: a convex-valued, closed-graph correspondence on a compact convex set has a fixed point, the tool of Nash's original proof.
- Linear program** The problem of maximizing or minimizing a linear objective subject to linear inequality constraints. Its optimum, when finite, is attained at a vertex of the feasible polytope.
- Maximin and minimax** Player I's guaranteed floor $\underline{V} = \max_p \min_q p^\top Aq$ and Player II's ceiling $\bar{V} = \min_q \max_p p^\top Aq$. Always $\underline{V} \leq \bar{V}$, with equality the content of the minimax theorem.
- Mex** The *minimum excludant*, the smallest nonnegative integer not in a given set. It drives the Grundy-value recurrence.
- Minimax theorem** Von Neumann's theorem that every finite two-person zero-sum game has a value and optimal mixed strategies, so $\underline{V} = \bar{V} = V$.
- Mixed strategy** A probability distribution over a player's pure strategies. Allowing mixing is what makes the minimax and Nash existence theorems true.
- Nash bargaining solution** The split of a cooperative surplus that maximizes the product of the players' gains over a disagreement point, the unique rule satisfying efficiency, symmetry, scale invariance, and independence of irrelevant alternatives.
- Nash equilibrium** A strategy profile in which no player can gain by deviating alone. Every finite game has one in mixed strategies (Nash's theorem).
- Nim and nim-sum** Nim is the game of removing chips from heaps, last to move winning. The *nim-sum* is the bitwise exclusive-or of the heap sizes; a position is lost for the mover exactly when it is zero (Bouton's theorem).
- Normal form** The representation of a game as a table of strategies and payoffs, collapsing the timing of an extensive game into a single simultaneous choice.
- Nucleolus** The single imputation that lexicographically minimizes the sorted vector of coalition complaints (excesses). It always lies in the core when the core is nonempty.
- P-position and N-position** A *P-position* is lost for the Player about to move (the *Previous* player wins); an *N-position* is won for the player about to move (the *Next* player wins).
- Pareto efficiency** An outcome from which no player can be made better off without making another worse off. Nash equilibria need not be efficient, as the Prisoner's Dilemma shows.
- Payoff function** The map from strategy profiles to each player's utility, the numbers that encode preferences over outcomes.
- Perfect information** The property that every player, when moving, knows the full history of play, so the tree has no nontrivial information sets.
- Price of anarchy** The ratio of the social cost of the worst equilibrium to the cost of the optimal coordinated outcome, a measure of the inefficiency of self-interest.
- Pure strategy** A complete deterministic plan of action, one entry of the strategy set over which a mixed strategy randomizes.

Recursive game A game some of whose outcomes are to play the game again, so its value satisfies a fixed-point equation solved by substitution.

Repeated game A game that plays a fixed stage game over many periods, where strategies may condition on history. Patience can sustain cooperation impossible in one shot.

Risk dominance A criterion that selects the equilibrium which is the best response against an opponent equally likely to play either way. In the Stag Hunt it favors the safe hunt over the efficient one.

Saddle point An entry of a payoff matrix that is the smallest in its row and the largest in its column. When one exists the game has a value in pure strategies and no one needs to randomize.

Safety level The payoff a player can guarantee unilaterally, $V(A)$ for Player I, regardless of what the others do. It is the disagreement point from which cooperative bargaining starts.

Shapley value The unique division of a coalitional game's worth satisfying efficiency, symmetry, the null-player rule, and additivity. It pays each player their average marginal contribution over all orders of arrival (Equation (10.1)).

Simplex method Dantzig's algorithm that solves a linear program by walking from vertex to vertex of the feasible polytope, improving the objective until no neighbor is better.

Sprague–Grundy theorem Every impartial game under normal play is equivalent to a single Nim heap, whose size is the position's Grundy value, and sums combine by nim-sum.

Stochastic game Shapley's generalization in which play moves among states by random transitions that the players' choices influence, with values obeying Bellman fixed-point equations.

Strategy A complete plan specifying a player's action in every situation that can arise; pure when deterministic, mixed when randomized.

Subgame perfect equilibrium A strategy profile that induces a Nash equilibrium in every subgame, ruling out non-credible threats. Backward induction finds one in any finite perfect-information game.

Threat, non-credible An off-path action a player would not actually carry out, which a Nash equilibrium may rely on but subgame perfection rejects.

Transferable utility The assumption that players can make binding side payments in a common currency, so a coalition's worth is a single number to be divided freely.

Value of a game The common number V that Player I can guarantee and Player II can hold them to in a zero-sum game, granted by the minimax theorem.

Wythoff's game A two-heap game allowing removal from one heap or an equal amount from both. Its cold positions are the golden-ratio Beatty pairs $(\lfloor k\varphi \rfloor, \lfloor k\varphi^2 \rfloor)$.

Zero-sum game A game in which the players' payoffs sum to zero, so one's gain is the other's loss. It is described by a single matrix that Player I maximizes and Player II minimizes.

APPENDIX F

A Timeline of Game Theory

The subject grew in bursts, often from the meeting of a concrete puzzle and a new piece of mathematics. The milestones below trace the ideas this book develops, from the first minimax calculation to the machines that now play games better than people.

- 1713** In a letter about the card game *le Her*, James Waldegrave writes down a mixed-strategy minimax solution, the earliest known. The idea then lies dormant for two centuries.
- 1838** Augustin Cournot models two firms choosing output, solving for what we now call a Nash equilibrium by intersecting reaction curves.
- 1913** Ernst Zermelo proves that in chess one of the players has a non-losing strategy, the first theorem about a game of perfect information.
- 1921** Émile Borel begins a series of notes on symmetric games, mixed strategies, and the minimax question, conjecturing that small games need not have a value.
- 1928** John von Neumann proves the minimax theorem: every finite two-person zero-sum game has a value and optimal mixed strategies. Game theory has its first deep theorem.
- 1935** Roland Sprague, followed independently by Patrick Grundy in 1939, shows every impartial game is equivalent to a single Nim heap, completing the theory Charles Bouton began with Nim in 1901.
- 1944** Von Neumann and Oskar Morgenstern publish *Theory of Games and Economic Behavior*, founding the field as a discipline and introducing cooperative games and utility theory.
- 1947** George Dantzig invents the simplex method; von Neumann, in conversation, conjectures that linear programming duality and the minimax theorem are the same statement.
- 1950** John Nash defines his equilibrium and proves it always exists. The same year, Merrill Flood and Melvin Dresher run the first Prisoner's Dilemma experiment at RAND, and Albert Tucker gives the game its name and story.
- 1951** Nash's full paper appears in the *Annals of Mathematics*. George Brown proposes fictitious play, and Julia Robinson proves it converges in zero-sum games.
- 1953** Lloyd Shapley introduces both the Shapley value and stochastic games; Harold Kuhn proves that behavioral and mixed strategies coincide under perfect recall and formalizes extensive games; Donald Gillies defines the core.
- 1965** Reinhard Selten refines Nash equilibrium with subgame perfection, ruling out non-credible threats in extensive games.
- 1967** John Harsanyi shows how to analyze games of incomplete information by letting nature choose each player's type, opening the theory of Bayesian games.
- 1969** David Schmeidler defines the nucleolus, the single division that makes the most aggrieved coalition as content as possible.
- 1973** John Maynard Smith and George Price introduce the evolutionarily stable strategy, carrying game theory into evolutionary biology.
- 1974** Robert Aumann defines correlated equilibrium, showing that a shared signal can beat independent randomization.
- 1975** Selten sharpens perfection again with the trembling-hand idea, asking that an equilibrium survive small mistakes.
- 1984** Robert Axelrod publishes the results of his repeated Prisoner's Dilemma tournaments, where the simple TIT-FOR-TAT rule prevails and cooperation emerges among strangers.
- 1994** Nash, Harsanyi, and Selten share the Nobel Memorial Prize in Economics for the theory of non-cooperative equilibrium.

- 2005** Aumann and Thomas Schelling share the prize for conflict and cooperation through the lens of game theory.
- 2007** Jonathan Schaeffer's team announces that checkers is solved: with perfect play it is a draw, the largest game settled by Zermelo's guarantee made constructive.
- 2012** Alvin Roth and Shapley share the prize for the theory of stable allocations and market design, the cooperative theory turned to practice.
- 2016** A program combining tree search with learned evaluation defeats a world-class human at Go, a perfect-information game long thought to resist the methods of [Chapter 3](#).

APPENDIX G

Further Reading

This book is a first course. The literature it draws on is deep, and a reader who wants to go further has excellent company. The notes below point to the texts that shaped the subject and the ones that teach it best today.

The founding texts

Everything starts with von Neumann and Morgenstern's *Theory of Games and Economic Behavior* (1944), the book that turned scattered results into a discipline. It is still worth opening, if only to see how much was there at the beginning. Von Neumann's 1928 minimax paper and Nash's two short notes from 1950 and 1951 are short enough to read in an afternoon and repay the effort: the whole of two-person theory grows from them.

Undergraduate introductions

Thomas Ferguson's *Game Theory*, the lecture notes this course follows, is freely available and covers combinatorial, zero-sum, and cooperative games with unusual care for worked technique. Philip Straffin's *Game Theory and Strategy* is gentler still, rich in applications to voting and fair division. Martin Osborne's *An Introduction to Game Theory* is the standard modern text for the non-cooperative side, careful about the difference between a model and its solution.

Graduate and comprehensive treatments

Osborne and Rubinstein's *A Course in Game Theory* is the concise classic, spare and exact. Fudenberg and Tirole's *Game Theory* is the graduate standard for the dynamic and incomplete-information theory this book only touches, including the full Folk Theorem. Roger Myerson's *Game Theory: Analysis of Conflict*, by one of the founders of mechanism design, is the deepest single-author account. For breadth, Maschler, Solan, and Zamir's *Game Theory* is the most complete recent reference, and Karlin and Peres's *Game Theory, Alive* brings a probabilist's eye and a wealth of modern examples.

Combinatorial games

The theory of Chapter 2 has its own literature. Berlekamp, Conway, and Guy's *Winning Ways for Your Mathematical Plays* is the encyclopedic and playful source, and Conway's *On Numbers and Games* builds the astonishing theory of game values and surreal numbers from the ground up. Both reward the patient reader with mathematics found nowhere else.

Where the subject is going

Game theory now reaches into computer science and machine learning. The complexity of finding equilibria, the design of auctions and matching markets, and the self-play algorithms that have mastered chess, Go, and poker all rest on the foundations in this book. A reader who has understood the four theorems here, Sprague–Grundy, minimax, Nash, and Shapley, is ready for any of it.

APPENDIX H

A Catalog of Canonical Games

The small games that recur throughout the subject, collected for quick reference. Each is solved in full where it first appears; the cross-reference points there.

Two-person zero-sum

Matching Pennies $\begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$ Value 0; both players mix $(\frac{1}{2}, \frac{1}{2})$. No saddle, pure conflict (Chapter 4).

Rock–Paper–Scissors $\begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix}$ Value 0; both mix $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. The symmetric three-cycle (Exercise 5.4).

A game with a saddle $\begin{pmatrix} 5 & 4 \\ 6 & 3 \end{pmatrix}$ Value 4 at the saddle (R_1, C_2) ; pure strategies, no mixing needed.

A game without a saddle $\begin{pmatrix} 2 & 5 \\ 6 & 3 \end{pmatrix}$ Value 4; Player I mixes $(\frac{1}{2}, \frac{1}{2})$ and Player II mixes $(\frac{1}{3}, \frac{2}{3})$.

Two-person general-sum

Each cell is (u_1, u_2) .

Prisoner's Dilemma $\begin{pmatrix} (3,3) & (0,5) \\ (5,0) & (1,1) \end{pmatrix}$ Unique Nash equilibrium $(\text{DEFECT}, \text{DEFECT}) = (1, 1)$, dominated by the unreachable $(3, 3)$ (Example 1.8).

Battle of the Sexes $\begin{pmatrix} (2,1) & (0,0) \\ (0,0) & (1,2) \end{pmatrix}$ Two pure equilibria on the diagonal, plus the mixed $((\frac{2}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{2}{3}))$. Coordination with conflict (Example 8.8).

Chicken (Hawk–Dove) $\begin{pmatrix} (0,0) & (-1,1) \\ (1,-1) & (-10,-10) \end{pmatrix}$ Two pure equilibria off the diagonal, plus a mixed one driving straight with probability $\frac{1}{10}$. Anti-coordination (Example 8.10).

Stag Hunt $\begin{pmatrix} (3,3) & (0,2) \\ (2,0) & (2,2) \end{pmatrix}$ Two pure equilibria on the diagonal: the payoff-dominant $(3, 3)$ and the risk-dominant $(2, 2)$, plus a mixed one hunting the stag with probability $\frac{2}{3}$.

Pure coordination $\begin{pmatrix} (1,1) & (0,0) \\ (0,0) & (1,1) \end{pmatrix}$ Two equivalent pure equilibria and a mixed one; the players need only agree.

Combinatorial and cooperative

Outside the matrix world a handful of structures recur. **Nim**: heaps of chips, last to move wins; the mover loses exactly when the nim-sum of the heap sizes is zero. **Wythoff's game**: two heaps with an extra diagonal move, whose cold positions are the golden-ratio Beatty pairs. **The subtraction game** $\{1, \dots, k\}$: the mover loses exactly when the pile size is a multiple of $k + 1$. **Bankruptcy**: claims against too small an estate, divided by the contested-garment rule that is both the Talmud's and the nucleolus. **Weighted voting** $[q; w_1, \dots, w_n]$: power measured by the Shapley–Shubik or Banzhaf index, which need not track the weights.

Common Pitfalls

The same handful of mistakes trip up almost everyone learning the subject. Recognizing them is half of avoiding them.

Calling the maximin the value too early The lower value $\underline{V}$ is the value only when it equals the upper value $\bar{V}$, which is to say when a saddle exists. When the two differ, mixing strictly raises what Player I can guarantee, and the maximin understates the value.

Using the 2×2 formula when a saddle exists The closed form assumes no saddle; applied to a game that has one it returns a meaningless “mixed” answer with a negative probability or a value outside the entries. Always test for a saddle first, and only mix if none is found.

Adding nim-sums in decimal The nim-sum is bitwise exclusive-or rather than ordinary addition. For heaps 1, 2, 3 the ordinary sum is 6, but the nim-sum is $1 \oplus 2 \oplus 3 = 0$, so the mover loses there even though the heaps look ample.

Trusting the order of elimination for weak dominance Iterated removal of *strictly* dominated strategies is order-independent and safe. For *weak* dominance the surviving set can depend on the order, and a weakly dominated strategy can still tie as a best response, so it may not be discardable.

Misreading the mixed-equilibrium payoff At a fully mixed equilibrium each player is *indifferent* among their active strategies. The mix is chosen to make the *opponent* indifferent, so a player’s equilibrium payoff is fixed by the opponent’s constraint rather than by maximizing anything of their own.

Expecting the Shapley value to be an equal split The Shapley value averages marginal contributions and rewards indispensability. It coincides with the equal division only when the players are symmetric; in general the pivotal player earns more.

Assuming the core is always nonempty The core can be empty, by the Bondareva–Shapley theorem. Three players who each profit by pairing up but cannot all agree at once are the smallest instance, and no imputation survives every coalition’s objection.

Confusing subgame perfection with Nash equilibrium Every subgame perfect equilibrium is a Nash equilibrium, but the converse fails. A Nash equilibrium can be propped up by a non-credible threat that the player would never carry out; backward induction tests the threat off the equilibrium path and discards it.

Folding the wrong player’s payoff in backward induction At each decision node the mover maximizes *their own* payoff coordinate rather than the first player’s. Folding the wrong coordinate up the tree gives a confidently wrong answer.

Forgetting expectations at a chance node When Nature moves, the value of that node is the *expected* payoff, weighted by the move probabilities. Treating a chance node like a decision node, or omitting the weights, breaks the induction.

Reversing P- and N-positions A *P-position* is a loss for the player about to move (the *Previous* player wins); an *N-position* is a win for the mover (the *Next* player wins). Many sign errors trace back to flipping this.

Choosing the wrong root of a recursive game The fixed-point equation $v = V(M(v))$ is often a quadratic with two roots. Only one is the value: the one the player can actually guarantee, selected by the security bound. The other is spurious.

Thinking the shift trick changes the strategies Adding a constant c to every entry of a zero-sum matrix raises the value by exactly c and leaves the optimal mixed strategies unchanged. The trick makes a matrix positive for the linear program without altering the solution.

Summing Grundy values instead of nim-summing them The Sprague–Grundy theorem combines a disjunctive sum by the *nim-sum* of the components’ Grundy values rather than their ordinary sum or their maximum. A heap-5 Nim pile plus a Grundy-3 game behaves like Nim on $5 \oplus 3 = 6$.

Reporting a strategy that is not a distribution A mixed strategy's weights must be nonnegative and sum to one. An answer with a negative weight signals a wrong support (often a missed saddle or an inactive strategy forced into the solution), and is a cue to re-examine which strategies are really used.

Bibliography

- [1] Robert J. Aumann and Michael Maschler. Game theoretic analysis of a bankruptcy problem from the talmud. *Journal of Economic Theory*, 36(2):195–213, 1985. doi: 10.1016/0022-0531(85)90102-4.
- [2] Rudolf Avenhaus, Bernhard von Stengel, and Shmuel Zamir. Inspection games. In Robert J. Aumann and Sergiu Hart, editors, *Handbook of Game Theory with Economic Applications, Volume 3*, pages 1947–1987. Elsevier, 2002.
- [3] Robert Axelrod. *The Evolution of Cooperation*. Basic Books, 1984.
- [4] Elwyn R. Berlekamp, John H. Conway, and Richard K. Guy. *Winning Ways for Your Mathematical Plays*. A K Peters, Natick, MA, 2 edition, 2001. Four volumes, published 2001–2004. Vol. 1 ISBN 978-1-56881-130-7; Vol. 2 ISBN 978-1-56881-142-0; Vol. 3 ISBN 978-1-56881-143-7; Vol. 4 ISBN 978-1-56881-144-4.
- [5] Émile Borel. La théorie du jeu et les équations intégrales à noyau symétrique gauche. *Comptes Rendus de l'Académie des Sciences*, 173:1304–1308, 1921.
- [6] Charles L. Bouton. Nim, a game with a complete mathematical theory. *Annals of Mathematics*, 3(1/4):35–39, 1901. doi: 10.2307/1967631.
- [7] Michael Bowling, Neil Burch, Michael Johanson, and Oskari Tammelin. Heads-up limit hold'em poker is solved. *Science*, 347(6218):145–149, 2015.
- [8] Dietrich Braess. Über ein paradoxon aus der verkehrsplanung. *Unternehmensforschung*, 12(1):258–268, 1968.
- [9] L. E. J. Brouwer. Über abbildung von mannigfaltigkeiten. *Mathematische Annalen*, 71(1):97–115, 1911. doi: 10.1007/BF01456931.
- [10] George W. Brown. Iterative solution of games by fictitious play. In Tjalling C. Koopmans, editor, *Activity Analysis of Production and Allocation*, pages 374–376. Wiley, New York, 1951.
- [11] Noam Brown and Tuomas Sandholm. Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374):418–424, 2018.
- [12] Noam Brown and Tuomas Sandholm. Superhuman AI for multiplayer poker. *Science*, 365(6456):885–890, 2019.
- [13] John H. Conway. *On Numbers and Games*. Academic Press, 1976.
- [14] George B. Dantzig. *Linear Programming and Extensions*. Princeton University Press, Princeton, NJ, 1963. ISBN 978-0-691-08000-0.
- [15] George B. Dantzig. Reminiscences about the origins of linear programming. *Operations Research Letters*, 1(2): 43–48, 1982.
- [16] Thomas S. Ferguson. *A Course in Game Theory*. World Scientific, Singapore, 2020. ISBN 978-981-3227-34-7. Based on the author's UCLA game theory lecture notes. Author page: <https://www.math.ucla.edu/~tom/>.
- [17] Merrill M. Flood. Some experimental games. *Management Science*, 5(1):5–26, 1958.
- [18] James W. Friedman. A non-cooperative equilibrium for supergames. *Review of Economic Studies*, 38(1):1–12, 1971.
- [19] Drew Fudenberg and Eric Maskin. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54(3):533–554, 1986.
- [20] Drew Fudenberg and Jean Tirole. *Game Theory*. The MIT Press, Cambridge, MA, 1991. ISBN 978-0-262-06141-4.
- [21] David Gale and Lloyd S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.

- [22] David Gale, Harold W. Kuhn, and Albert W. Tucker. Linear programming and the theory of games. In Tjalling C. Koopmans, editor, *Activity Analysis of Production and Allocation*, pages 317–329. Wiley, 1951.
- [23] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems 27 (NIPS 2014)*, pages 2672–2680, 2014.
- [24] Patrick M. Grundy. Mathematics and games. *Eureka*, 2:6–8, 1939.
- [25] Ehud Kalai and Meir Smorodinsky. Other solutions to Nash’s bargaining problem. *Econometrica*, 43(3):513–518, 1975.
- [26] Anna R. Karlin and Yuval Peres. *Game Theory, Alive*. American Mathematical Society, Providence, RI, 2017. ISBN 978-1-4704-1982-0.
- [27] Victor Klee and George J. Minty. How good is the simplex algorithm? In Oved Shisha, editor, *Inequalities III*, pages 159–175. Academic Press, 1972.
- [28] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In Christoph Meinel and Sophie Tison, editors, *STACS 99: 16th Annual Symposium on Theoretical Aspects of Computer Science*, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413, Berlin, Heidelberg, 1999. Springer. doi: 10.1007/3-540-49116-3_38.
- [29] Harold W. Kuhn. Extensive games and the problem of information. In Harold W. Kuhn and Albert W. Tucker, editors, *Contributions to the Theory of Games, Volume II*, number 28 in *Annals of Mathematics Studies*, pages 193–216. Princeton University Press, Princeton, NJ, 1953.
- [30] S. C. Littlechild and G. Owen. A simple expression for the Shapley value in a special case. *Management Science*, 20(3):370–372, 1973.
- [31] Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems 30 (NeurIPS)*, pages 4765–4774, 2017.
- [32] Michael Maschler, Eilon Solan, and Shmuel Zamir. *Game Theory*. Cambridge University Press, Cambridge, 2 edition, 2020. ISBN 978-1-108-49345-1.
- [33] Richard D. McKelvey and Thomas R. Palfrey. An experimental study of the centipede game. *Econometrica*, 60(4):803–836, 1992.
- [34] Jean-François Mertens and Abraham Neyman. Stochastic games. *International Journal of Game Theory*, 10(2): 53–66, 1981.
- [35] Roger B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, MA, 1991. ISBN 978-0-674-34115-9.
- [36] John Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951. doi: 10.2307/1969529.
- [37] John F. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950. doi: 10.2307/1907266.
- [38] John F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1): 48–49, 1950. doi: 10.1073/pnas.36.1.48.
- [39] Martin J. Osborne. *An Introduction to Game Theory*. Oxford University Press, New York, 2004. ISBN 978-0-19-512895-6.
- [40] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. The MIT Press, Cambridge, MA, 1994. ISBN 978-0-262-65040-3.
- [41] Ignacio Palacios-Huerta. Professionals play minimax. *The Review of Economic Studies*, 70(2):395–415, 2003.
- [42] Christos H. Papadimitriou. Algorithms, games, and the internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 749–753, 2001.
- [43] Robert Powell. Defending against terrorist attacks with limited resources. *American Political Science Review*, 101(3):527–541, 2007.
- [44] Brian Roberson. The colonel blotto game. *Economic Theory*, 29(1):1–24, 2006. doi: 10.1007/s00199-005-0071-5.

- [45] Julia Robinson. An iterative method of solving a game. *Annals of Mathematics*, 54(2):296–301, 1951. doi: 10.2307/1969530.
- [46] Robert W. Rosenthal. Games of perfect information, predatory pricing and the chain-store paradox. *Journal of Economic Theory*, 25(1):92–100, 1981. doi: 10.1016/0022-0531(81)90018-1.
- [47] Alvin E. Roth, Tayfun Sönmez, and M. Utku Ünver. Kidney exchange. *The Quarterly Journal of Economics*, 119(2):457–488, 2004.
- [48] Tim Roughgarden and Éva Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.
- [49] Jonathan Schaeffer, Neil Burch, Yngvi Björnsson, Akihiro Kishimoto, Martin Müller, Robert Lake, Paul Lu, and Steve Sutphen. Checkers is solved. *Science*, 317(5844):1518–1522, 2007.
- [50] David Schmeidler. The nucleolus of a characteristic function game. *SIAM Journal on Applied Mathematics*, 17(6):1163–1170, 1969. doi: 10.1137/0117107.
- [51] Reinhard Selten. Spieltheoretische behandlung eines oligopolmodells mit nachfrageträgheit. *Zeitschrift für die gesamte Staatswissenschaft*, 121:301–324 and 667–689, 1965. Published in two parts in volume 121: pp. 301–324 and pp. 667–689.
- [52] Reinhard Selten. Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4(1):25–55, 1975. doi: 10.1007/BF01766400.
- [53] Claude E. Shannon. Programming a computer for playing chess. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 41(314):256–275, 1950.
- [54] Lloyd S. Shapley. Stochastic games. *Proceedings of the National Academy of Sciences*, 39(10):1095–1100, 1953. doi: 10.1073/pnas.39.10.1095.
- [55] Lloyd S. Shapley. A value for n-person games. In Harold W. Kuhn and Albert W. Tucker, editors, *Contributions to the Theory of Games, Volume II*, number 28 in Annals of Mathematics Studies, pages 307–317. Princeton University Press, Princeton, NJ, 1953.
- [56] Lloyd S. Shapley. On balanced sets and cores. *Naval Research Logistics Quarterly*, 14(4):453–460, 1967.
- [57] Lloyd S. Shapley and Martin Shubik. A method for evaluating the distribution of power in a committee system. *American Political Science Review*, 48(3):787–792, 1954.
- [58] David Silver, Aja Huang, Chris J. Maddison, et al. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- [59] David Silver, Thomas Hubert, Julian Schrittwieser, et al. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419):1140–1144, 2018.
- [60] Roland P. Sprague. Über mathematische kampfspiele. *Tôhoku Mathematical Journal*, 41:438–444, 1935. Tohoku Math. J. (first series), vol. 41 (1935–1936).
- [61] George J. Stigler. The cost of subsistence. *Journal of Farm Economics*, 27(2):303–314, 1945.
- [62] Jean Ville. Sur la théorie générale des jeux où intervient l’habileté des joueurs. In Émile Borel, editor, *Traité du calcul des probabilités et de ses applications*, volume IV, pages 105–113. Gauthier-Villars, 1938.
- [63] John von Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100(1):295–320, 1928. doi: 10.1007/BF01448847.
- [64] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1944. 60th Anniversary Commemorative Edition, Princeton University Press, 2004, ISBN 978-0-691-13061-3, with an introduction by Harold W. Kuhn and an afterword by Ariel Rubinstein.
- [65] Jonathan Weinstein. On a mistranslation of a mistake about minimax. Working paper, Washington University in St. Louis, 2023.
- [66] David J. Wu. Accelerating self-play learning in Go. *arXiv preprint arXiv:1902.10565*, 2019.
- [67] W. A. Wythoff. A modification of the game of Nim. *Nieuw Archief voor Wiskunde*, 7:199–202, 1907.
- [68] Ernst Zermelo. Über eine anwendung der mengenlehre auf die theorie des schachspiels. In *Proceedings of the Fifth International Congress of Mathematicians*, volume 2, pages 501–504. Cambridge University Press, 1913.

Index

- backward induction, 28
- balanced collection, 115
- balanced game, 115
- Banzhaf index, 120
- basic feasible solution, 71
- Beatty sequence, 22
- behavioral strategy, 35
- best response, 9, 88
- bimatrix game, 6, 87
- Bondareva–Shapley theorem, 115
- Bouton’s theorem, 15
- Brouwer’s fixed-point theorem, 94

- Centipede game, 35
- characteristic function, 113
- coalitional game, 113
- combinatorial game, 13
- common knowledge of rationality, 35
- complementary slackness, 35
- cooperative game, 103
- cooperative strategy, 104
- cooperative value, 104
- core, 114
 - nonemptiness, 115
- correlated equilibrium, 97

- disagreement point, 105
- disjunctive sum, 18
- dominance
 - in matrix games, 46
 - strict, 8
 - weak, 8
- dominant strategy, 8
- dual, 68
- duality
 - linear programming, 68

- equilibrium
 - correlated, 97
 - Nash, 9
 - payoff-dominant, 10
 - risk-dominant, 10
 - strategic, 88
 - subgame perfect, 29
- excess, 121
- exchange rate
 - Shapley, 110
- exclusive-or, 15
- extensive-form game, 27

- Farkas’s lemma, 69
- feasible region, 67

- feasible set
 - NTU, 104
 - TU, 104
- fixed point, 94
 - recursive game, 75
- folk theorem, 81
- follower, 13

- game
 - bimatrix, 6
 - cooperative, 10
 - general-sum, 10, 87
 - impartial, 10
 - n -person, 88
 - Nim, 14
 - non-cooperative, 10
 - strategic-form, 6
 - zero-sum, 10
- game tree, 27
- Golden Balls, 8
- golden ratio, 22
- grand coalition, 113
- grim trigger, 80
- Grundy value, 17

- impartial game, 10, 13
- imperfect information, 10, 35
- imputation, 114
- information set, 27
- inspection game, 79
- invariance, 60

- Kuhn’s theorem, 36

- lexicographic order, 121
- linear program, 67
- lower value, 42

- marginal contribution, 117
- maximin, 42
- mex, 16
- minimax, 42
- minimax theorem, 44
- minimum excludant, 16
- mixed strategy, 7
- multilinear extension, 88

- N-position, 13
- Nash bargaining solution, 106
- Nash equilibrium, 9, 88
- Nash product, 106
- Nash’s existence theorem, 94

- Nim, 14
- nim-sum, 15
- number, 17
- node
 - decision, 27
 - terminal, 27
- non-transferable utility, 103
- normal play, 13
- nucleolus, 121

- objective function, 67

- P-position, 13
- Pareto optimal, 104
- payoff function, 6
- payoff matrix, 41
- perfect information, 10, 27
- perfect recall, 35
- pivot, 71
- Player I, 41
- Player II, 41
- position, 13
- power index
 - Banzhaf, 120
- primal, 68
- principle of indifference, 55
- Prisoner's Dilemma, 8

- recursive game, 75
- repeated game, 79
 - folk theorem, 80

- saddle point, 42
- safety level, 87
- Shapley value, 117
 - axioms, 117
- simplex, 7
- simplex method, 71
- skew-symmetric matrix, 59
- Sprague–Grundy function, 17
- Sprague–Grundy theorem, 18
- Stag Hunt, 10
- strategic form, 6
- strategy
 - extensive form, 28
 - mixed, 7
 - pure, 6
- strong duality, 69
- subgame, 28
- subgame perfect equilibrium, 29
- substitution principle, 75
- sum of games, 18
- super-additivity, 113
- support, 7
- symmetric game, 59

- threat strategy, 105
- transferable utility, 103

- upper value, 42
- value of a game, 44

- weak duality, 69
- Wythoff's game, 22

- Zermelo's theorem, 33
- zero-sum game, 10, 41