

Game Theory, Alive

Anna R. Karlin and Yuval Peres

Draft January 20, 2013

Please send comments and corrections to

`karlin@cs.washington.edu` and `peres@microsoft.com`

We are grateful to Alan Hammond, Yun Long, Gábor Pete, and Peter Ralph for scribing early drafts of this book from lectures by Yuval Peres. These drafts were edited by Liat Kessler, Asaf Nachmias, Yelena Shvets, Sara Robinson and David Wilson; Yelena also drew many of the figures. We also thank Ranjit Samra for the lemon figure, Barry Sinervo for the Lizard picture, and Davis Sheperd for additional figures.

Sourav Chatterjee, Elchanan Mossel, Asaf Nachmias, and Shobhana Stoyanov taught from drafts of the book and provided valuable suggestions. Thanks also to Varsha Dani, Kieran Kishore, Itamar Landau, Eric Lei, Mallory Monasterio, Andrea McCool, Katia Nepom, Colleen Ross, Zhuohong Shen, Davis Sheperd, Stephanie Somersille, and Sithparran Vanniasegaram for comments and corrections.

The support of the NSF VIGRE grant to the Department of Statistics at the University of California, Berkeley, and NSF grants DMS-0244479, DMS-0104073, and CCF-1016509 is acknowledged.

Contents

1 Introduction	<i>page</i> 1
2 Two-person zero-sum games	6
2.1 Examples	6
2.2 Definitions	8
2.3 Saddle points and Nash equilibria	10
2.4 Simplifying and solving zero-sum games	11
2.4.1 The technique of domination	12
2.4.2 Summary of Domination	13
2.4.3 The use of symmetry	14
2.4.4 Series and Parallel Game Combinations	17
2.5 Games on graphs	18
2.5.1 Maximum Matchings	18
2.5.2 Hide-and-see games	20
2.5.3 Weighted hide-and-see games	22
2.5.4 The bomber and battleship game	25
2.6 Von Neumann's minimax theorem	26
2.7 Linear Programming and the Minimax Theorem	31
2.7.1 Linear Programming Basics	32
2.7.2 Linear Programming Duality	33
2.7.3 Duality, more formally	34
2.7.4 The proof of the duality theorem	35
2.7.5 An interpretation of a primal/dual pair	38
2.8 Zero-Sum Games With Infinite Action Spaces*	40
2.9 Solved Exercises	47
2.9.1 Another Betting Game	47
3 General-sum games	49
3.1 Some examples	49

<i>Contents</i>	iii
3.2 Nash equilibria	52
3.3 General-sum games with more than two players	56
3.4 Games with Infinite Strategy Spaces	60
3.5 Evolutionary game theory	62
3.5.1 Hawks and Doves	62
3.5.2 Evolutionarily stable strategies	65
3.6 Potential games	68
3.7 Correlated equilibria	71
3.8 The proof of Nash's theorem	75
3.9 Fixed-point theorems*	77
3.9.1 Easier fixed-point theorems	77
3.9.2 Sperner's lemma	80
3.9.3 Brouwer's fixed-point theorem	83
4 Signaling and asymmetric games	94
4.1 Signaling and asymmetric information	95
4.1.1 Examples of signaling (and not)	96
4.1.2 The collapsing used car market	98
4.2 Some further examples	99
5 Social choice	102
5.1 Voting and Ranking Mechanisms	102
5.1.1 Definitions	103
5.1.2 Instant runoff elections	105
5.1.3 Borda count	106
5.1.4 Dictatorship	107
5.2 Arrow's impossibility theorem	107
5.3 Strategy-proof Voting	109
6 Auctions and Mechanism Design	114
6.1 Auctions	114
6.2 Single Item Auctions	114
6.2.1 Profit in single-item auctions	116
6.2.2 More profit?	117
6.2.3 Exercise: Vickrey with Reserve Price	118
6.3 The general mechanism design problem	118
6.4 Social Welfare Maximization	121
6.5 Win/Lose Mechanism Design	125
6.6 Profit Maximization in Win/Lose Settings	126
6.6.1 Profit maximization in digital goods auctions	127
6.6.2 Profit Extraction	128
6.6.3 A profit-making digital goods auction	129

6.6.4	With priors: the Bayesian setting	131
6.6.5	Revenue Equivalence	134
6.6.6	The revelation principle	136
6.6.7	Optimal Auctions	137
7	Stable matching	140
7.1	Introduction	140
7.2	Algorithms for finding stable matchings	141
7.3	Properties of stable matchings	142
7.4	A special preference order case	143
8	Coalitions and Shapley value	146
8.1	The Shapley value and the glove market	146
8.2	The Shapley value	150
8.3	Two more examples	152
9	Interactive Protocols	155
9.1	Keeping the meteorologist honest	155
9.2	Secret sharing	158
9.2.1	A simple secret sharing method	159
9.2.2	Polynomial method	160
9.3	Private computation	162
9.4	Cake cutting	163
9.5	Zero-knowledge proofs	164
9.6	Remote coin tossing	166
10	Combinatorial games	168
10.1	Impartial games	169
10.1.1	Nim and Bouton's solution	175
10.1.2	Other impartial games	178
10.1.3	Impartial games and the Sprague-Grundy theorem	186
10.2	Partisan games	192
10.2.1	The game of Hex	195
10.2.2	Topology and Hex: a path of arrows*	196
10.2.3	Hex and Y	198
10.2.4	More general boards*	200
10.2.5	Other partisan games played on graphs	201
10.3	Brouwer's fixed-point theorem via Hex	206
11	Random-turn and auctioned-turn games	212
11.1	Random-turn games defined	212
11.2	Random-turn selection games	213
11.2.1	Hex	213

<i>Contents</i>	v
11.2.2 Bridg-It	214
11.2.3 Surround	215
11.2.4 Full-board Tic-Tac-Toe	215
11.2.5 Recursive majority	215
11.2.6 Team captains	216
11.3 Optimal strategy for random-turn selection games	217
11.4 Win-or-lose selection games	219
11.4.1 Length of play for random-turn Recursive Majority	220
11.5 Richman games	221
11.6 Additional notes on random-turn Hex	224
11.6.1 Odds of winning on large boards under biased play.	224
11.7 Random-turn Bridg-It	225
<i>Bibliography</i>	227
<i>Index</i>	231

1

Introduction

In this course on game theory, we will be studying a range of mathematical models of conflict and cooperation between two or more agents. We begin with an outline of the content of this course.

We begin with the classic **two-person zero-sum games**. In such games, both players move simultaneously, and depending on their actions, they each get a certain payoff. What makes these games “zero-sum” is that each player benefits only at the expense of the other. We will show how to find optimal strategies for each player in such games. These strategies will typically turn out to be a randomized choice of the available options.

For example, in **Penalty Kicks**, a soccer/football-inspired zero-sum game, one player, the penalty-taker, chooses to kick the ball either to the left or to the right of the other player, the goal-keeper. At the same instant as the kick, the goal-keeper guesses whether to dive left or right.




Fig. 1.1. The game of Penalty Kicks.

The goal-keeper has a chance of saving the goal if he dives in the same direction as the kick. The penalty-taker, being left-footed, has a greater likelihood of success if he kicks left. The probabilities that the penalty kick scores are displayed in the table below:

		goal-keeper	
		L	R
penalty-taker	L	0.8	1
	R	1	0.5

For this set of scoring probabilities, the optimal strategy for the penalty-taker is to kick left with probability $5/7$ and kick right with probability $2/7$ — then regardless of what the goal-keeper does, the probability of scoring is $6/7$. Similarly, the optimal strategy for the goal-keeper is to dive left with probability $5/7$ and dive right with probability $2/7$.

In **general-sum games**, the topic of Chapter 3, we no longer have optimal strategies. Nevertheless, there is still a notion of a “rational choice” for the players. A **Nash equilibrium** is a set of strategies, one for each player, with the property that no player can gain by unilaterally changing his strategy. It turns out that every general-sum game has at least one Nash equilibrium. The proof of this fact requires an important geometric tool, the **Brouwer fixed-point theorem**.

One interesting class of general-sum games, important in computer science, is that of **congestion games**. In a congestion game, there are two drivers, I and II, who must navigate as quickly as possible through a congested network of roads. Driver I must travel from city B to city D , and driver II, from city A to city C .




Fig. 1.2. A congestion game. Shown here are the commute times for the four roads connecting four cities. For each road, the first number is the commute time when only one driver uses the road, the second number is the commute time when two drivers use the road.

The travel time for using a road is less when the road is less congested. In the ordered pair (t_1, t_2) attached to each road in the diagram below, t_1 represents the travel time when only one driver uses the road, and t_2 represents the travel time when the road is shared. For example, if drivers I and II both use road AB , with I traveling from A to B and II from B to A , then each must wait 5 units of time. If only one driver uses the road, then it takes only 3 units of time.

A development of the last twenty years is the application of general-sum game theory to evolutionary biology. In economic applications, it is often assumed that the agents are acting “rationally,” which can be a hazardous assumption in many economic applications. In some biological applications, however, Nash equilibria arise as stable points of evolutionary systems composed of agents who are “just doing their own thing.” There is no need for a notion of rationality.

Chapter 4 considers games with asymmetric information and signaling. If one player has some information that another does not, that may be to his advantage. But if he plays differently, might he give away what he knows, thereby removing this advantage?

The topic of Chapter 8 is **cooperative game theory**, in which players form coalitions to work toward a common goal. As an example, suppose that three people are selling their wares in a market. Two are each selling a single, left-handed glove, while the third is selling a right-handed one. A wealthy tourist enters the store in dire need of a pair of gloves. She refuses to deal with the glove-bearers individually, so that it becomes their job to form coalitions to make a sale of a left- and right-handed glove to her. The third player has an advantage, because his commodity is in scarcer supply. This means that he should be able to obtain a higher fraction of the payment that the tourist makes than either of the other players. However, if he holds out for too high a fraction of the earnings, the other players may agree between them to refuse to deal with him at all, blocking any sale, and thereby risking his earnings. Finding a solution for such a game involves a mathematical concept known as the **Shapley value**.

Another major topic within game theory concerns the *design* of markets or schemes (which are themselves games) that achieve desirable outcomes in equilibrium. This is called **mechanism design**. Chapter 5 considers **social choice**, settings in which we wish to design a mechanism that aggregates the preferences of a collection of individuals in some socially desirable way. The most basic example is the design of **voting schemes**. Unfortunately, the most important result here, Arrow’s Impossibility Theorem, is negative. It states, more or less, that if there is an election with more than two candidates, then no matter which system one chooses to use for voting, there is trouble ahead: at least one desirable property that we might wish for the election will be violated.

Chapter 6 shows how introducing payments into the mechanism design problem can alleviate some of the difficulties presented in Chapter 5. One of the most important results here is the famous VCG mechanism which shows how to use payments to design a mechanism that maximizes **social**

welfare, the total happiness of society while incentivizing the participants in the mechanism to report their private information truthfully. The simplest example of this is a sealed-bid auction for a single item. In this setting, there is always a temptation for bidders to bid less than their true value for an item. But suppose the goal of the auction designer is to ensure that the item ends up in the hands of the bidder that values it the most. Bandwidth auctions conducted by governments are an example of a setting where this is the goal. If bidders are not incentivized to report their value for the item truthfully, then there is no guarantee that the auction designer's goal will be achieved. An elegant solution to this problem is to conduct a **second-price auction**, in which the item is sold to the bidder that bid highest, but that bidder only pays the bid of the second highest bidder. This turns out to incentivize bidders to bid truthfully.

Another problem in the realm of social choice is the **stable matching** problem, the topic of Chapter 7. Suppose that there are n men and n women, each man has a sorted list of the women he prefers, and each woman has a sorted list of the men that she prefers. A matching between them is stable if there is no man and woman who both prefer one another to their partners in the matching. Gale and Shapley showed that there always is a stable matching, and showed how to find one. Stable matchings generalize to stable assignments, and these are found by centralized clearinghouses for markets, such as the National Resident Matching Program which each year matches about 20,000 new doctors to residency programs at hospitals.

Chapter 9 studies a variety of other types of mechanism design problems. An example is the problem of fairly sharing a resource. Consider the problem of a pizza with several different toppings, each distributed over portions of the pizza. The game has two or more players, each of whom prefers certain toppings. If there are just two players, there is a well-known mechanism for dividing the pizza: One splits it into two sections, and the other chooses which section he would like to take. Under this system, each player is at least as happy with what he receives as he would be with the other player's share. What if there are three or more players? We will study this question, as well as an interesting variant of it.

Finally, we turn to **combinatorial games**, in which two players take turns making moves until a winning position for one of the players is reached. The solution concept for this type of game is a **winning strategy** — a collection of moves for one of the players, one for each possible situation, that guarantees his victory.

A classic example of a combinatorial game is **Nim**. In Nim, there are several piles of chips, and the players take turns choosing a pile and removing

one or more chips from it. The goal for each player is to take the last chip. We will describe a winning strategy for Nim and show that a large class of combinatorial games are essentially similar to it.

Chess and Go are examples of popular combinatorial games that are famously difficult to analyze. We will restrict our attention to simpler examples, such as the game of **Hex**, which was invented by Danish mathematician, Piet Hein, and independently by the famous game theorist John Nash, while he was a graduate student at Princeton. Hex is played on a rhombus shaped board tiled with small hexagons (see Figure 1.3). Two players, Blue and Yellow, alternate coloring in hexagons in their assigned color, blue or yellow, one hexagon per turn. The goal for Blue is to produce a blue chain crossing between his two sides of the board. The goal for Yellow is to produce a yellow chain connecting the other two sides.




Fig. 1.3. The board for the game of Hex.

As we will see, it is possible to prove that the player who moves first can always win. Finding the winning strategy, however, remains an unsolved problem, except when the size of the board is small.

In an interesting variant of the game, the players, instead of alternating turns, toss a coin to determine who moves next. In this case, we are able to give an explicit description of the optimal strategies of the players. Such **random-turn combinatorial games** are the subject of Chapter 11.

Game theory and mechanism design remain an active area of research, and our goal is whet the reader's appetite by introducing some of its many facets.

2

Two-person zero-sum games

We begin with the theory of **two-person zero-sum games**, developed in a seminal paper by John von Neumann and Oskar Morgenstern. Two-person zero-sum games are perfectly competitive games, in which one player's loss is the other player's gain. The central theorem for two-person zero-sum games is that even if each player's strategy is known to the other, there is an amount that one player can guarantee as her expected gain, and the other, as his maximum expected loss. This amount is known as the **value** of the game.

2.1 Examples

Consider the following game:

Example 2.1.1 (*Pick-a-Hand*, a betting game). There are two players, Chooser (player I), and Hider (player II). Hider has two gold coins in his back pocket. At the beginning of a turn, he† puts his hands behind his back and either takes out one coin and holds it in his left hand (strategy $L1$), or takes out both and holds them in his right hand (strategy $R2$). Chooser picks a hand and wins any coins the hider has hidden there. She may get nothing (if the hand is empty), or she might win one coin, or two.

The following matrix summarizes the payoffs to Chooser in each of the cases.

		Hider	
		$L1$	$R2$
Chooser	L	1	0
	R	0	2

† In all two-person games, we adopt the convention that player I is female and player II is male.

How should Hider and Chooser play? Imagine that they are conservative and want to optimize for the worst case scenario. Hider can guarantee himself a loss of at most 1 by selecting action L1 (whereas if he selects R2, he has the potential to lose 2). Chooser cannot guarantee herself any positive gain since if she selects L, in the worst case, Hider selects R2, whereas if she selects R, in the worst case, Hider selects L1.

Now consider expanding the possibilities available to the players by incorporating randomness. Suppose that Hider selects L1 with probability y_1 and R2 with probability $y_2 = 1 - y_1$. Hider's expected loss is y_1 if Chooser plays L, and $2(1 - y_1)$ if Chooser plays R. Thus Hider's worst-case expected loss is $\max(y_1, 2(1 - y_1))$. To minimize this, Hider will choose $y_1 = 2/3$, guaranteeing himself an expected loss of at most $2/3$. See Figure 2.1.




Fig. 2.1. The left side of the figure shows the worst-case expected gain of Chooser as a function of x_1 , the probability with which she plays L. The right side of the figure shows the worst-case expected loss of Hider as a function of y_1 , the probability with which he plays L1.

Similarly, suppose that Chooser selects L with probability x_1 and R with probability $x_2 = 1 - x_1$. Then Chooser's worst-case expected gain is $\min(x_1, 2(1 - x_1))$. To maximize this, she will choose $x_1 = 2/3$, guaranteeing herself an expected gain of at least $2/3$.

This example illustrates a striking general feature of zero-sum games. *With randomness, conservative play is optimal:* Since Hider can guarantee himself an expected loss of at most $2/3$, Chooser cannot do better than the strategy that guarantees her an expected gain of $2/3$, and vice versa.

Notice that without some extra incentive, it is not in Hider's interest to play *Pick-a-hand* because he can only lose by playing. To be enticed into joining the game, Hider will need to be paid at least $2/3$.

Exercise 2.1.2 (Another Betting Game). Consider the betting game with the following payoff matrix:

		player II	
		L	R
player I	T	0	2
	B	5	1

Draw graphs for this game analogous to those shown in Figure 2.1. This exercise is solved in §Section 2.9.

2.2 Definitions

A two-person zero-sum game can be represented by an $m \times n$ payoff **pay-off matrix** $A = (a_{ij})$, whose rows are indexed by the m possible actions of player I, and whose columns are indexed by the n possible actions of player II. Player I selects an action i and player II selects an action j , each unaware of the other's selection. Their selections are then revealed and player II pays player I the amount a_{ij} .

It is elementary to verify that

$$\max_i \min_j a_{ij} \leq \min_j \max_i a_{ij} \quad (2.1)$$

since player I can guarantee gaining the left hand side and player II can guarantee not losing more than the right hand side. (For a formal proof, see Lemma 2.6.3.) Unfortunately, as in Example 2.1.1, without randomness, the inequality is usually strict. With randomness, the situation is more promising.

A strategy in which each action is selected with some probability is a **mixed strategy**. A mixed strategy for player I is determined by a vector $(x_1, \dots, x_m)^T$ where x_i represents the probability of playing action i . The set of mixed strategies for player I is denoted by

$$\Delta_m = \left\{ \mathbf{x} \in \mathbb{R}^m : x_i \geq 0, \sum_{i=1}^m x_i = 1 \right\}.$$

Similarly, the set of mixed strategies for player II is denoted by

$$\Delta_n = \left\{ \mathbf{y} \in \mathbb{R}^n : y_j \geq 0, \sum_{j=1}^n y_j = 1 \right\}.$$

A mixed strategy in which a particular action is played with probability 1 is called a **pure strategy**. Observe that in this vector notation, pure

strategies are represented by the standard basis vectors, though we often identify the pure strategy \mathbf{e}_i with the corresponding action i .

If player I employs strategy \mathbf{x} , she can guarantee herself an expected gain of

$$\min_{\mathbf{y} \in \Delta_m} \mathbf{x}^T \mathbf{A} \mathbf{y} = \min_j (\mathbf{x}^T \mathbf{A})_j \quad (2.2)$$

(Write $\mathbf{z} = \mathbf{x}^T \mathbf{A}$. Then $\mathbf{z}^T \mathbf{y} = \sum_j z_j y_j$ is a weighted average of the z_j 's for $\mathbf{y} \in \Delta_m$, so $\min_{\mathbf{y} \in \Delta_m} \mathbf{z}^T \mathbf{y} = \min_j z_j$.)

A conservative player will choose \mathbf{x} to maximize (2.2).

Definition 2.2.1. A mixed strategy $\mathbf{x}^* \in \Delta_m$ is a **safety strategy for player I** if the maximum over $\mathbf{x} \in \Delta_m$ of the function

$$\mathbf{x} \mapsto \min_{\mathbf{y} \in \Delta_n} \mathbf{x}^T \mathbf{A} \mathbf{y}$$

is attained at \mathbf{x}^* . The value of this function at \mathbf{x}^* is the **safety value for player I**. Similarly, a mixed strategy $\mathbf{y}^* \in \Delta_n$ is a **safety strategy for player II** if the minimum over $\mathbf{y} \in \Delta_n$ of the function

$$\mathbf{y} \mapsto \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^T \mathbf{A} \mathbf{y}$$

is attained at \mathbf{y}^* . The value of this function at \mathbf{y}^* is the **safety value for player II**.

Remark. For the existence of safety strategies see Lemma 2.6.3.

Safety strategies might appear conservative, but the following celebrated theorem shows that the two players' safety values coincide.

Theorem 2.2.2. von Neumann's Minimax Theorem For any finite two-person zero-sum game, there is a number V , called the value of the game, satisfying

$$\max_{\mathbf{x} \in \Delta_m} \min_{\mathbf{y} \in \Delta_n} \mathbf{x}^T \mathbf{A} \mathbf{y} = V = \min_{\mathbf{y} \in \Delta_n} \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^T \mathbf{A} \mathbf{y}. \quad (2.3)$$

We will prove the minimax theorem in §Section 2.6.

Remarks:

- (i) It is easy to check that the left hand side of equation (2.3) is upper

bounded by the right hand side, i.e.

$$\max_{\mathbf{x} \in \Delta_m} \min_{\mathbf{y} \in \Delta_n} \mathbf{x}^T \mathbf{A} \mathbf{y} \leq \min_{\mathbf{y} \in \Delta_n} \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^T \mathbf{A} \mathbf{y}. \quad (2.4)$$

(See the argument for equation 2.1 and Lemma 2.6.3). The magic of zero-sum games is that, in mixed strategies, this inequality becomes an equality.

- (ii) If \mathbf{x}^* is a safety strategy for player I and \mathbf{y}^* is a safety strategy for player II, then it follows from Theorem 2.2.2 that:

$$\min_{\mathbf{y} \in \Delta_n} (\mathbf{x}^*)^T \mathbf{A} \mathbf{y} = V = \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^T \mathbf{A} \mathbf{y}^*. \quad (2.5)$$

In words, this means that the mixed strategy \mathbf{x}^* yields player I an expected gain of at least V , no matter how II plays, and the mixed strategy \mathbf{y}^* yields player II an expected loss of at most V , no matter how I plays. Therefore, from now on, we will refer to the safety strategies as **optimal strategies**.

2.3 Saddle points and Nash equilibria

A notion of great importance in game theory is the notion of **Nash equilibrium**. In this section, we introduce this notion and show that a pair of strategies in a zero-sum game is optimal if and only if they are in Nash equilibrium.

Example 2.3.1. In the following game, if both players are playing action 1, then neither has an incentive to switch. Note this game has value 1.

		player II	
		action 1	action 2
player I	action 1	1	2
	action 2	0	-1

Such an entry, which is the largest in its column and the smallest in its row, is called a saddle point.

Definition 2.3.2. A **saddle point**[†] of a payoff matrix A is a pair (i^*, j^*) such that

$$\max_i a_{ij^*} = a_{i^*j^*} = \min_j a_{i^*j} \quad (2.6)$$

[†] The term saddle point comes from the continuous setting where a function $f(x, y)$ of two variables has a point (x^*, y^*) at which locally $\max_x f(x, y^*) = f(x^*, y^*) = \min_y f(x^*, y)$. Thus, the surface resembles a saddle that curves up the the y direction and curves down in the x direction.

A saddle point is also known as a **pure Nash equilibrium**. More generally we have the following definition:

Definition 2.3.3. A pair of strategies $(\mathbf{x}^*, \mathbf{y}^*)$ is a **Nash equilibrium** in a zero-sum game with payoff matrix A if

$$\min_{\mathbf{y} \in \Delta_n} (\mathbf{x}^*)^T A \mathbf{y} = (\mathbf{x}^*)^T A \mathbf{y}^* = \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^T A \mathbf{y}^*. \quad (2.7)$$

Thus, even if player I when selecting her strategy *knows* player II is playing strategy \mathbf{y}^* , she has no incentive to switch from \mathbf{x}^* to a different strategy \mathbf{x} (and similarly for player II).

Remark. If $\mathbf{x}^* = \mathbf{e}_{i^*}$ and $\mathbf{y}^* = \mathbf{e}_{j^*}$, then by Equation (2.2), this definition coincides with Definition 2.3.2.

Theorem 2.3.4. A pair of strategies $(\mathbf{x}^*, \mathbf{y}^*)$ is optimal if and only if $(\mathbf{x}^*, \mathbf{y}^*)$ is a Nash equilibrium.

Proof. Suppose that \mathbf{x}^* is an optimal strategy for player I and \mathbf{y}^* is an optimal strategy for player II. Then, by Theorem 2.2.2 and Definition 2.2.1, we have

$$(\mathbf{x}^*)^T A \mathbf{y}^* \geq \min_{\mathbf{y} \in \Delta_n} (\mathbf{x}^*)^T A \mathbf{y} = V = \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^T A \mathbf{y}^* \geq (\mathbf{x}^*)^T A \mathbf{y}^*,$$

and thus all these inequalities are equalities and 2.7 holds. For the other direction, observe that for any pair of vectors \mathbf{x}^* and \mathbf{y}^*

$$\min_{\mathbf{y} \in \Delta_n} (\mathbf{x}^*)^T A \mathbf{y} \leq \max_{\mathbf{x} \in \Delta_m} \min_{\mathbf{y} \in \Delta_n} \mathbf{x}^T A \mathbf{y} \leq \min_{\mathbf{y} \in \Delta_n} \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^T A \mathbf{y} \leq \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^T A \mathbf{y}^*.$$

If (2.7) holds, then all these inequalities become equalities and thus \mathbf{x}^* and \mathbf{y}^* are optimal. \square

Remark. It follows that if (i^*, j^*) is a saddle point, then i^* is an optimal strategy for player I and j^* is an optimal strategy for player II.

2.4 Simplifying and solving zero-sum games

In this section, we will discuss techniques that help us understand zero-sum games and solve them (that is, find their value and determine optimal strategies for the two players).

2.4.1 The technique of domination

Domination is a technique for reducing the size of a game’s payoff matrix, enabling it to be more easily analyzed. Consider the following example.

Example 2.4.1 (*Plus One*). Each player chooses a number from $\{1, 2, \dots, n\}$ and writes it down; then the players compare the two numbers. If the numbers differ by one, the player with the higher number wins \$1 from the other player. If the players’ choices differ by two or more, the player with the higher number pays \$2 to the other player. In the event of a tie, no money changes hands.

The payoff matrix for the game is:

		player II							
		1	2	3	4	5	6	...	n
player I	1	0	-1	2	2	2	2	...	2
	2	1	0	-1	2	2	2	...	2
	3	-2	1	0	-1	2	2	...	2
	4	-2	-2	1	0	-1	2	...	2
	5	-2	-2	-2	1	0	-1	2	2
	6	-2	-2	-2	-2	1	0	2	2
	\vdots	\vdots	\vdots			-2	-2	\ddots	\vdots
	$n - 1$	-2	-2	...					0
n	-2	-2	...					1	0

In this payoff matrix, every entry in row 4 is at most the corresponding entry in row 1. Thus player I has no incentive to play 4 since it is *dominated* by row 1. In fact, rows 4 through n are all dominated by row 1, and hence player I can ignore those strategies.

By symmetry, we see that player II need never play any of strategies 4 through n . Thus, in *Plus One* we can search for optimal strategies in the reduced payoff matrix:

		player II		
		1	2	3
player I	1	0	-1	2
	2	1	0	-1
	3	-2	1	0

To analyze the reduced game, let $\mathbf{x}^T = (x_1, x_2, x_3)$ be player I’s mixed strategy. For \mathbf{x} to be optimal, each component of

$$\mathbf{x}^T A = (x_2 - 2x_3, -x_1 + x_3, 2x_1 - x_2) \tag{2.8}$$

must be at least the value of the game. In this game, there is complete symmetry between the players. This implies that the payoff matrix is **anti-symmetric**: the game matrix is square and $a_{ij} = -a_{ji}$ for every i and j .

Claim 2.4.2. *If the payoff matrix of a zero-sum game is anti-symmetric, then the game has value 0.*

Proof. If V is the safety value for player I, then by symmetry, $-V$ is the safety value for II, and since these coincide, $V = 0$. \square

We conclude that for any optimal strategy \mathbf{x} in *Plus One*

$$\begin{aligned} x_2 - 2x_3 &\geq 0 \\ -x_1 + x_3 &\geq 0 \\ 2x_1 - x_2 &\geq 0, \end{aligned}$$

Thus $x_2 \geq 2x_3$, $x_3 \geq x_1$, and $2x_1 \geq x_2$. If one of these inequalities was strict, then adding the first, twice the second and the third, we could deduce $x_2 > x_2$, so in fact each of them must be an equality. Solving the resulting system, with the constraint $x_1 + x_2 + x_3 = 1$, we find that the optimal strategy for each player is $(1/4, 1/2, 1/4)$.

2.4.2 Summary of Domination

We say a row ℓ of a two-person zero-sum game dominates row i if $a_{\ell j} \geq a_{ij}$ for all j . When row i is dominated, then there is no loss to player I if she never plays it. More generally, we say that subset I of rows dominates row i if there is a convex combination β_ℓ , for $\ell \in I$ (i.e. $\beta_\ell \geq 0$ for all $\ell \in I$ and $\sum_{\ell \in I} \beta_\ell = 1$) such that for every j

$$\sum_{\ell \in I} \beta_\ell a_{\ell j} \geq a_{ij}. \quad (2.9)$$

In this situation, there is no loss to player I in ignoring row i .

Analogously for columns, we say that subset J of columns dominates column j if there is a convex combination β_ℓ , for $\ell \in J$ such that

$$\sum_{\ell \in J} \beta_\ell a_{i\ell} \leq a_{ij}$$

for every i . In this situation, there is no loss to player II in ignoring column j .

Exercise 2.4.3. Prove that if equation (2.9) holds, then player I can safely ignore row i .

Solution: Consider any mixed strategy \mathbf{x} for player I, and use it to construct a new strategy \mathbf{z} in which $z_i = 0$, $z_\ell = x_\ell + \beta_\ell x_i$, for $\ell \in I$, and $z_k = x_k$ for $k \notin I \cup \{i\}$. Then, against II's j -th strategy:

$$(\mathbf{z}^T A - \mathbf{x}^T A)_j = \sum_{\ell \in I} (x_\ell + \beta_\ell x_i - x_\ell) a_{\ell j} - x_i a_{ij} \geq 0.$$

2.4.3 The use of symmetry

Another way to simplify the analysis of a game is via the technique of **symmetry**. We illustrate a symmetry argument in the following example:

Example 2.4.4 (*Submarine Salvo*).




Fig. 2.2.

A submarine is located on two adjacent squares of a three-by-three grid. A bomber (player I), who cannot see the submerged craft, hovers overhead and drops a bomb on one of the nine squares. She wins \$1 if she hits the submarine and \$0 if she misses it. There are nine pure strategies for the bomber and twelve for the submarine, so the payoff matrix for the game is quite large. Symmetry arguments can simplify the analysis.

There are three types of moves that the bomber can make: She can drop a bomb in the center, in the middle of one of the sides, or in a corner. Similarly, there are two types of positions that the submarine can assume: taking up the center square, or taking up a corner square.

It is intuitive (and true) that both players have optimal strategies that assign equal probability to actions of the same type. To see this, observe that in *Submarine Salvo* a 90 degree rotation describes a permutation π of the possible submarine positions and a permutation σ of the possible bomber actions. Clearly π^4 (rotating by 90 degrees four times) is the identity and so is σ^4 . For any bomber strategy \mathbf{x} , let $\pi\mathbf{x}$ be the rotated row strategy. (Formally $(\pi\mathbf{x})_i = x_{\pi(i)}$). Clearly, the probability that the bomber will hit the submarine if they play $\pi\mathbf{x}$ and $\sigma\mathbf{y}$ is the same as it is when they play \mathbf{x} and \mathbf{y} , and therefore

$$\min_{\mathbf{y}} \mathbf{x}^T A \mathbf{y} = \min_{\mathbf{y}} (\pi\mathbf{x})^T A \mathbf{y}.$$

Thus, if v is the value of the game and \mathbf{x} is optimal, then $\pi^k\mathbf{x}$ is also optimal for all k .

Fix any submarine strategy \mathbf{y} . Then $\pi^k\mathbf{x}$ gains at least v against \mathbf{y} , hence so does

$$\mathbf{x}^* = \frac{1}{4}(\mathbf{x} + \pi\mathbf{x} + \pi^2\mathbf{x} + \pi^3\mathbf{x}).$$

Therefore \mathbf{x}^* is an optimal rotation-invariant strategy.

Using these equivalences, we may write down a more manageable payoff matrix:

		submarine	
		center	corner
bomber	corner	0	1/4
	midside	1/4	1/4
	middle	1	0

Note that the values for the new payoff matrix are different from those in the standard payoff matrix. They incorporate the fact that when the bomber and submarine are both playing *corner* there is only a one-in-four chance that there will be a hit. In fact, the pure strategy of *corner* for the bomber in this reduced game corresponds to the mixed strategy of bombing each corner with probability 1/4 in the original game. Similar reasoning applies to each of the pure strategies in the reduced game.

We can use domination to simplify the matrix even further. This is because for the bomber, the strategy *midside* dominates that of *corner* (because the sub, when touching a corner, must also be touching a midside). This observation reduces the matrix to:

		submarine	
		center	corner
bomber	midside	1/4	1/4
	middle	1	0

Now note that for the submarine, *corner* dominates *center*, and thus we obtain the reduced matrix:

		corner
		corner
bomber	midside	1/4
	middle	0

The bomber picks the better alternative — technically, another application of domination — and picks *midside* over *middle*. The value of the game is 1/4; the bomber’s optimal strategy is to hit one of the four mid-sides with probability 1/4 each, and the optimal submarine strategy is to hide with probability 1/8 each in one of the eight possible pairs of adjacent squares that exclude the center.

The symmetry argument is generalized in the following theorem:

Theorem 2.4.5. *Suppose that π and σ are permutations of $\{1, \dots, m\}$ and $\{1, \dots, n\}$ respectively such that*

$$a_{\pi(i)\sigma(j)} = a_{ij} \tag{2.10}$$

for all i and j . Then there exist optimal strategies \mathbf{x}^* and \mathbf{y}^* such that $x_i^* = x_{\pi(i)}^*$ for all i and $y_j^* = y_{\sigma(j)}^*$ for all j .

Proof. First, observe that there is an ℓ such that π^ℓ is the identity permutation (since there must be $k > r$ with $\pi^k = \pi^r$, in which case $\ell = k - r$.)

Let $(\pi\mathbf{x})_i = x_{\pi(i)}$ and $(\sigma\mathbf{y})_j = y_{\sigma(j)}$.

Let $\Psi(\mathbf{x}) = \min_{\mathbf{y}} \mathbf{x}^T A \mathbf{y}$. Since $(\pi\mathbf{x})^T A (\sigma\mathbf{y}) = \mathbf{x}^T A \mathbf{y}$, we have $\Psi(\mathbf{x}) = \Psi(\pi\mathbf{x})$ for all $\mathbf{x} \in \Delta_m$. Therefore, for all $\mathbf{y} \in \Delta_n$

$$\left(\frac{1}{\ell} \sum_{k=0}^{\ell-1} \pi^k \mathbf{x} \right)^T A \mathbf{y} \geq \frac{1}{\ell} \sum_{k=0}^{\ell-1} \Psi(\pi^k \mathbf{x}) = \Psi(\mathbf{x}).$$

Thus, if \mathbf{x} is optimal, so is $\mathbf{x}^* = \frac{1}{\ell} \sum_{k=0}^{\ell-1} \pi^k \mathbf{x}$. Clearly $\pi \mathbf{x}^* = \mathbf{x}^*$. □

Remark. It is perhaps surprising that in *Submarine Salvo* there also exist optimal strategies that do not assign equal probability to all actions of the same type. (See exercise 2.18.)

2.4.4 Series and Parallel Game Combinations

In this section, we will analyze two ways that zero-sum games can be combined: in series and in parallel.

Definition 2.4.6. Given two zero-sum games G_1 and G_2 with values v_1 and v_2 , their **series sum-game** corresponds to playing G_1 and then G_2 . The series sum-game has value $v_1 + v_2$. In a **parallel sum-game**, each player chooses either G_1 or G_2 to play. If each picks the same game, then it is that game which is played. If they differ, then no game is played, and the payoff is zero.

We may write a big payoff matrix for the parallel sum-game, in which player I's strategies are the union of her strategies in G_1 and her strategies in G_2 as follows:

		player II	
		pure strategies of G_1	pure strategies of G_2
player I	pure strategies of G_1	G_1	0
	pure strategies of G_2	0	G_2

In this payoff matrix, we have abused notation and written G_1 and G_2 inside the matrix to denote the payoff matrix of G_1 and G_2 respectively. If the two players play G_1 and G_2 optimally, the payoff matrix is effectively:

		player II	
		play in G_1	play in G_2
player I	play in G_1	v_1	0
	play in G_2	0	v_2

Thus to find optimal strategies, the players just need to determine with what probability they should play G_1 and with what probability they should play G_2 . If both payoffs v_1 and v_2 are positive, the optimal strategy for each player consists of playing G_1 with probability $v_2/(v_1 + v_2)$, and G_2 with probability $v_1/(v_1 + v_2)$. Assuming both v_1 and v_2 are positive, the expected payoff of the parallel sum-game is

$$\frac{v_1 v_2}{v_1 + v_2} = \frac{1}{1/v_1 + 1/v_2}.$$

For those familiar with electrical networks, it is interesting to observe that the rules for computing the value of parallel or series games in terms of the values of the component games are precisely the same as the the rules for computing the effective resistance of a pair of resistors in series or in parallel. We will explore some games that exploit this connection in Chapter (??).

2.5 Games on graphs

2.5.1 Maximum Matchings

Given a set of boys B and a set of girls G , draw an edge between a boy and a girl if they know each other. The resulting graph is called a **bipartite graph** since there are two disjoint sets of nodes, and all edges go between them. Bipartite graphs are ubiquitous. For instance, there is a natural bipartite graph where one set of nodes represents workers, the other set represents jobs, and an edge from worker w to job j means that worker w can perform job j . Other examples involve customers and suppliers, and students and colleges.

A **matching** in a bipartite graph is a collection of disjoint edges, e.g. a set of boy-girl pairs that know each other, where every individual occurs in at most one pair. (See figure 2.3.)

Suppose $|B| \leq |G|$. Then clearly there cannot be a matching that includes more than $|B|$ edges. Under what condition is there a matching of this size, i.e. a matching in which every boy is matched to a girl he knows?




Fig. 2.3. On the left is a bipartite graph where an edge between a boy and a girl means that they know each other. The edges in a matching are shown in purple in the figure on the right..

An obvious necessary condition, known as **Hall's condition**, is that each subset B' of the boys collectively knows enough girls, at least $|B'|$ of them. What Hall's theorem says is that this condition is not only necessary, but sufficient.

Theorem 2.5.1 (Hall's marriage theorem). *Suppose that B is a finite set of boys and G is a finite set of girls. For any particular boy $b \in B$, let $f(b)$ denote the set of girls that b knows. For a subset $B' \subseteq B$ of the boys, let $f(B')$ denote the set of girls that boys in B' collectively know, i.e., $f(B') = \cup_{b \in B'} f(b)$. There is a matching of size $|B|$ if and only if Hall's condition holds: every subset $B' \subseteq B$ satisfies $|f(B')| \geq |B'|$.*

Proof. We need only prove that Hall's condition is sufficient, which we do by induction on the number of boys.

The base case when $|B| = 1$ is straightforward. For the induction step, we consider two cases.

Case 1: $|f(B')| > |B'|$ for each nonempty $B' \subsetneq B$. Then we can just match an arbitrary boy b to any girl he knows. The set of remaining boys and girls still satisfy Hall's condition, so by the inductive hypothesis, we can match them up. (Of course this approach does not work for the example in Figure 2.3: there are three sets of boys B' for which $|f(B')| = |B'|$, and indeed, if the third boy is paired with the first girl, there is no way to match the remaining boys and girls.)

Case 2: There is a nonempty $B' \subsetneq B$ for which $|f(B')| = |B'|$. By the induction hypothesis, there is a matching of size $|B'|$ between B' and $f(B')$. Once we show that Hall's condition holds for the bipartite graph between $B \setminus B'$ and $G \setminus f(B')$, another application of the inductive hypothesis yields the theorem.

Suppose Hall's condition fails, i.e., there is a set A of boys disjoint from B' such that the set $S = f(A) \setminus f(B')$ of girls they know outside $f(B')$ has $|S| < |A|$. Then

$$|f(A \cup B')| = |S \cup f(B')| < |A| + |B'|$$

violating Hall's condition for the full graph, a contradiction. □

A useful way to represent a bipartite graph whose edges go between vertex sets I and J is via its **adjacency matrix** H . This is a 0/1 matrix where the rows correspond to vertices in I , the columns to vertices in J , and $h_{ij} = 1$ if and only if there is an edge between i and j . Conversely, any 0/1 matrix is the adjacency matrix of a bipartite graph. A set of pairs $\mathcal{S} \subset I \times J$ is a **matching** for the adjacency matrix H if $h_{ij} = 1$ for all $(i, j) \in \mathcal{S}$ and no two elements of \mathcal{S} are in the same row or column. This corresponds to a matching between I and J in the graph represented by H .

For example, the following matrix is the adjacency matrix for the bipartite graph shown in Figure 2.3, with the edges corresponding to the matching in bold red. (Rows represent boys from left to right and columns represent girls from left to right.)

$$\begin{pmatrix} \mathbf{1} & 1 & 0 & 0 \\ 0 & \mathbf{1} & 0 & 0 \\ 1 & 0 & 0 & \mathbf{1} \\ 0 & 0 & \mathbf{1} & 0 \end{pmatrix}$$

2.5.2 *Hide-and-seek games*

Example 2.5.2 (*Hide-and-Seek*). A robber, player II, hides in one of a set of safehouses located at certain street/avenue intersections in Manhattan. A cop, player I, chooses one of the avenues or streets to travel along. The cop wins a unit payoff if she travels on a road that intersects the robber's location.




Fig. 2.4. The figure shows an example scenario for the *Hide-and-Seek* game. In this example, the robber chooses to hide at the safehouse at the intersection of 2nd St. and 4th Ave., and the cop chooses to travel along 1st St. Thus, the payoff to the cop is -1.

We represent this situation with a 0/1 matrix H where rows represent streets, columns represent avenues, and $h_{ij} = 1$ if there is a safehouse at the intersection of street i and avenue j , and $h_{ij} = 0$ otherwise. The following is the matrix H corresponding to the scenario shown in Figure 2.4:

$$\begin{pmatrix} 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 & 0 \end{pmatrix}$$

Given a scenario represented by a 0/1 matrix H , the cop's strategy corresponds to choosing a row or column of this matrix and the robber's strategy corresponds to picking a 1 in the matrix.

Clearly, the cop can restrict attention to roads that contain safehouses; a natural strategy for her is to find a smallest set of roads that contain all safehouses, and choose one of these at random. Formally, a **line-cover** of the matrix H is a set of lines (rows and columns) that cover all nonzero entries of H . The proposed cop strategy is to fix a minimum-sized line cover \mathcal{C} and choose one of the lines in \mathcal{C} uniformly at random. This guarantees the cop an expected gain of at least $1/|\mathcal{C}|$ against any robber strategy.

Next we consider robber strategies. A bad strategy would be to choose from among a set of safehouses that all lie on the same road. The “opposite” of that is to find a maximum-sized set \mathcal{M} of safehouses, where no two lie on the same road, and choose one of these uniformly at random. This guarantees that the cop’s expected gain is at most $1/|\mathcal{M}|$.

Observing that the set \mathcal{M} is a *matching* in the matrix H , the following lemma implies that $|\mathcal{C}| = |\mathcal{M}|$. This means that the proposed pair of strategies is a Nash equilibrium and thus, by Theorem 2.3.4, jointly optimal for *Hide-and-Seek*.

Lemma 2.5.3 (König’s lemma). *Given an m by n 0/1 matrix H , the size of the maximum matching is equal to the size of the minimum line-cover.*

Proof. Suppose the maximum matching has size k and the minimum line-cover \mathcal{C} has size ℓ . At least one member of each pair in the matching has to be in \mathcal{C} and therefore $k \leq \ell$.

For the other direction, we use Hall’s Theorem. Suppose that there are r rows and c columns in the \mathcal{C} , so $r + c = \ell$. We claim that there is a matching M of size r in the submatrix defined by the rows in \mathcal{C} and the columns outside \mathcal{C} , and a matching M' of size c in the submatrix defined by the rows outside \mathcal{C} and the columns in \mathcal{C} . If so, since M and M' are disjoint, there is a matching of size at least ℓ , and hence $\ell \leq k$, completing the proof.

Suppose there is no matching of size r in the submatrix defined by the rows in \mathcal{C} and the columns outside \mathcal{C} . View the rows in \mathcal{C} as boys, the columns outside \mathcal{C} as girls and a 1 in entry (i, j) as indicating that boy i and girl j know each other. Then applying Hall’s theorem, we conclude that there is a subset S of rows in \mathcal{C} who collectively know fewer than $|S|$ columns outside \mathcal{C} . But then if we replace S in \mathcal{C} by the uncovered columns that know them, we will reduce the size of the line-cover, contradicting our assumption that it was minimum. A similar argument shows that there is a matching of size c in the submatrix defined by the rows outside \mathcal{C} and the columns in \mathcal{C} .

□

2.5.3 Weighted hide-and-peek games

Example 2.5.4 (Generalized Hide-and-Seek). We generalize the previous game by allowing H to be a nonnegative matrix. The nonzero entries still correspond to safehouses, but the value $h_{ij} > 0$ represents the payoff to the cop if the robber hides at location (i, j) and the cop chooses row i or column j . (E.g., certain safehouses might be safer than others, and h_{ij} could represent the probability the cop actually catches the robber if she chooses either i or j when he is hiding at (i, j) .)

Suppose that both players have n strategies and consider the following class of player II strategies: Player II first chooses a fixed permutation π of the set $\{1, \dots, n\}$ and then hides at location (i, π_i) with a probability p_i that he chooses. For example, if $n = 5$, and the fixed permutation π is 3, 1, 4, 2, 5, then the following matrix gives the probability of player II hiding in different places:

$$\begin{array}{ccccc} 0 & 0 & p_1 & 0 & 0 \\ p_2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & p_3 & 0 \\ 0 & p_4 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & p_5 \end{array}$$

Given a permutation π , the optimal choice for p_i is $p_i = d_{i,\pi_i}/D_\pi$, where

$$d_{ij} = \frac{1}{h_{ij}},$$

if $h_{ij} > 0$, and 0 otherwise. and

$$D_\pi = \sum_{i=1}^n d_{i,\pi_i}.$$

This choice equalizes the expected payments to player I no matter what row or column she chooses. To see this, observe that if player I selects row i , she obtains an expected payoff of $p_i h_{i,\pi(i)} = 1/D_\pi$, whereas if she chooses column j , she obtains an expected payoff of $p_j h_{\pi^{-1}(j),j} = 1/D_\pi$. Thus, if player II is going to use this type of strategy, the right permutation to pick is one that maximizes D_π . We will in fact show that doing this is an optimal strategy, not just in this restricted class of strategies, but in general.

To find an optimal strategy for player I, we need an analogue of König's lemma. In this context, a *covering* of the matrix $D = (d_{ij})_{n \times n}$ will be a pair of vectors $\mathbf{u} = (u_1, \dots, u_n)$ and $\mathbf{w} = (w_1, \dots, w_n)$, with non-negative components, such that $u_i + w_j \geq d_{ij}$ for each pair (i, j) . The analogue of König's lemma is:

Lemma 2.5.5. Consider a minimum covering $(\mathbf{u}^*, \mathbf{w}^*)$ of $D = (d_{ij})_{n \times n}$ (i.e., one for which $\sum_{i=1}^n (u_i + w_i)$ is minimum). Then

$$\sum_{i=1}^n (u_i^* + w_i^*) = \max_{\pi} D_{\pi}. \quad (2.11)$$

Remark. Note that a minimum covering exists, because the continuous map

$$(\mathbf{u}, \mathbf{w}) \mapsto \sum_{i=1}^n (u_i + w_i),$$

defined on the closed and bounded set

$$\{(\mathbf{u}, \mathbf{w}) : 0 \leq u_i, w_i \leq M, \text{ and } u_i + w_j \geq d_{ij}\},$$

where $M = \max_{i,j} d_{ij}$, does indeed attain its infimum. Note also that we may assume that $\min_i u_i^* > 0$.

Proof. We first show that $\sum_{i=1}^n (u_i^* + w_i^*) \geq \max_{\pi} D_{\pi}$. This is straightforward, since for any π , $u_i^* + w_{\pi_i}^* \geq d_{i,\pi_i}$. Summing over i yields the inequality.

Showing the other inequality is harder; we will use Hall's marriage theorem (Theorem 2.5.1). To this end, we need a definition of "knowing": We say that row i knows column j if

$$u_i^* + w_j^* = d_{ij}.$$

We first show that every subset of k rows know at least k columns. For contradiction, suppose that the k rows i_1, \dots, i_k know between them only $\ell < k$ columns j_1, \dots, j_{ℓ} . We claim that this contradicts the minimality of $(\mathbf{u}^*, \mathbf{w}^*)$.

To see this, define $\tilde{\mathbf{u}}$ from \mathbf{u}^* by reducing u_i^* on these k rows by a small amount $\varepsilon > 0$ leaving the other rows unchanged, in such a way that all \tilde{u}_i 's remain positive, and we do not violate the constraints that $\tilde{u}_i + \tilde{w}_j \geq d_{ij}$ for any $j \notin \{j_1, \dots, j_{\ell}\}$. (Thus, we must have $0 < \varepsilon \leq \min_i u_i^*$ and $\varepsilon \leq \min \{u_i^* + w_j^* - d_{ij} : (i, j) \text{ such that } u_i^* + w_j^* > d_{ij}\}$.)

Similarly, define $\tilde{\mathbf{w}}$ from \mathbf{w}^* by adding ε to the ℓ columns known by the k rows. Leave the other columns unchanged. That is

$$\tilde{w}_{j_i} = w_{j_i}^* + \varepsilon \text{ for } i \in \{1, \dots, \ell\}.$$

Clearly, by construction, $(\tilde{\mathbf{u}}, \tilde{\mathbf{w}})$ is a covering of the matrix.

Moreover, the covering $(\tilde{\mathbf{u}}, \tilde{\mathbf{w}})$ has a strictly smaller sum of components than does $(\mathbf{u}^*, \mathbf{w}^*)$, contradicting the fact that this latter covering is minimum

Thus, Hall's condition holds, and there is a perfect matching between rows

and columns that “know” each other. This is a permutation π^* such that, for each i :

$$u_i^* + w_{\pi_i^*}^* = d_{i,\pi_i^*},$$

from which it follows that

$$\sum_{i=1}^n u_i^* + \sum_{i=1}^n w_i^* = D_{\pi^*} \leq \max_{\pi} D_{\pi}.$$

This proves that $\sum_{i=1}^n (u_i^* + w_i^*) \leq \max_{\pi} D_{\pi}$, and completes the proof of the theorem. \square

The lemma and the proof give us a pair of optimal strategies for the players. Player I chooses row i with probability u_i^*/D_{π^*} , and column j with probability w_j^*/D_{π^*} . Against this strategy, if player II chooses some (i, j) , then the payoff will be

$$\frac{(u_i^* + v_j^*)}{D_{\pi^*}} h_{ij} \geq \frac{d_{ij} h_{ij}}{D_{\pi^*}} = \frac{1}{D_{\pi^*}}.$$

We deduce that the permutation strategy for player II described before the lemma is indeed optimal.

Example 2.5.6. Consider the *Generalized Hide-and-Seek* game with probabilities given by the following matrix:

$$\begin{bmatrix} 1 & 1/2 \\ 1/3 & 1/5 \end{bmatrix}.$$

This means that the matrix D is equal to

$$\begin{bmatrix} 1 & 2 \\ 3 & 5 \end{bmatrix}.$$

To determine a minimum cover of the matrix D , consider first a cover that has all of its mass on the rows: $\mathbf{u} = (2, 5)$ and $\mathbf{v} = (0, 0)$. Note that rows 1 and 2 know only column 2, according to the definition of “knowing” introduced in the analysis of this game. Modifying the vectors \mathbf{u} and \mathbf{v} according to the rule given in this analysis, we obtain updated vectors, $\mathbf{u} = (1, 4)$ and $\mathbf{v} = (0, 1)$, whose sum is 6, equal to the expression $\max_{\pi} D_{\pi}$ (obtained by choosing the identity permutation).

Thus, an optimal strategy for the robber is to hide at location $(1, 1)$ with probability $1/6$ and location $(2, 2)$ with probability $5/6$. An optimal strategy for the cop is to choose avenue (row) 1 with probability $1/6$, avenue 2 with probability $2/3$ and street 2 with probability $1/6$. The value of the game is $1/6$.

2.5.4 The bomber and battleship game

Example 2.5.7 (*Bomber and Battleship*). In this family of games, a battleship is initially located at the origin in \mathbb{Z} . At each time step in $\{0, 1, \dots\}$, the ship moves either left or right to a new site where it remains until the next time step. The bomber (player I), who can see the current location of the battleship (player II), drops one bomb at some time j over some site in \mathbb{Z} . The bomb arrives at time $j + 2$, and destroys the battleship if it hits it. (The battleship cannot see the bomber or its bomb in time to change course.) For the game G_n , the bomber has enough fuel to drop its bomb at any time $j \in \{0, 1, \dots, n\}$. What is the value of the game?

Exercise 2.5.8. (i) Show that the value of G_0 is $1/3$. (ii) Show that the value of G_1 is also $1/3$. (iii) Show that the value of G_2 is greater than $1/3$.




Fig. 2.5. The bomber drops its bomb where it hopes the battleship will be two time units later. The battleship does not see the bomb coming, and randomizes its path to avoid the bomb. (The length of each arrow is 2.)

Consider the following strategy for G_n . On the first move, go left with probability a and right with probability $1 - a$. From then on, at each step turn with probability of $1 - a$, and keep going with probability of a .

We choose a to optimize the probability of evasion for the battleship. Its probabilities of arrival at sites -2 , 0 , or 2 at time 2 are a^2 , $1 - a$ and $a(1 - a)$. We have to choose a so that $\max\{a^2, 1 - a\}$ is minimal. This value is achieved when $a^2 = 1 - a$, whose solution in $(0, 1)$ is given by $a = 2/(1 + \sqrt{5})$. Since at any time j that the bomber chooses to drop a bomb, the battleship's position two time steps later has the same distribution, the payoff for the

bomber against this strategy is at most $1 - a$. Thus, $v(G_n)$ is at most $1 - a$ for each n . While this strategy is not optimal for any G_n , it has the merit of converging to optimal play, as $n \rightarrow \infty$. See the notes for a discussion of the result.

2.6 Von Neumann's minimax theorem

We now prove the von Neumann Minimax Theorem. The proof will rely on a basic theorem from convex geometry.

Definition 2.6.1. A set $K \subseteq \mathbb{R}^d$ is **convex** if, for any two points $\mathbf{a}, \mathbf{b} \in K$, the line segment that connects them,

$$\{p\mathbf{a} + (1 - p)\mathbf{b} : p \in [0, 1]\},$$

also lies in K .

Theorem 2.6.2 (The Separating Hyperplane Theorem). *Suppose that $K \subseteq \mathbb{R}^d$ is closed and convex. If $\mathbf{0} \notin K$, then there exists $\mathbf{z} \in \mathbb{R}^d$ and $c \in \mathbb{R}$ such that*

$$0 < c < \mathbf{z}^T \mathbf{v}$$

for all $\mathbf{v} \in K$.

Here $\mathbf{0}$ denotes the vector of all 0's, and $\mathbf{z}^T \mathbf{v}$ is the usual dot product $\sum_i z_i v_i$. The theorem says that there is a **hyperplane** (a line in two dimensions, a plane in three dimensions, or, more generally, an affine \mathbb{R}^{d-1} -subspace in \mathbb{R}^d) that separates $\mathbf{0}$ from K . In particular, on any continuous path from $\mathbf{0}$ to K , there is some point that lies on this hyperplane. The separating hyperplane is given by $\{\mathbf{x} \in \mathbb{R}^d : \mathbf{z}^T \mathbf{x} = c\}$. The point $\mathbf{0}$ lies in the half-space $\{\mathbf{x} \in \mathbb{R}^d : \mathbf{z}^T \mathbf{x} < c\}$, while the convex body K lies in the complementary half-space $\{\mathbf{x} \in \mathbb{R}^d : \mathbf{z}^T \mathbf{x} > c\}$.

Recall first that the **(Euclidean) norm of \mathbf{v}** is the (Euclidean) distance between $\mathbf{0}$ and \mathbf{v} , and is denoted by $\|\mathbf{v}\|$. Thus $\|\mathbf{v}\| = \sqrt{\mathbf{v}^T \mathbf{v}}$. A subset of a metric space is **closed** if it contains all its limit points, and **bounded** if it is contained inside a ball of some finite radius R . In what follows, the metric is the Euclidean metric.

Proof of Theorem 2.6.2. If we pick R so that the ball of radius R centered at $\mathbf{0}$ intersects K , the function $\mathbf{w} \mapsto \|\mathbf{w}\|$, considered as a map from $K \cap \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\| \leq R\}$ to $[0, \infty)$, is continuous, with a domain that is nonempty,




Fig. 2.6. Hyperplane separating the closed convex body K from $\mathbf{0}$.

closed and bounded (see Figure 2.7). Thus the map attains its infimum at some point \mathbf{z} in K . For this $\mathbf{z} \in K$ we have

$$\|\mathbf{z}\| = \inf_{\mathbf{w} \in K} \|\mathbf{w}\|.$$




Fig. 2.7. Intersecting K with a ball to get a nonempty closed bounded domain.

Let $\mathbf{v} \in K$. Because K is convex, for any $\varepsilon \in (0, 1)$, we have that $\varepsilon\mathbf{v} + (1 - \varepsilon)\mathbf{z} = \mathbf{z} - \varepsilon(\mathbf{z} - \mathbf{v}) \in K$. Since \mathbf{z} has the minimum norm of any point in K ,

$$\|\mathbf{z}\|^2 \leq \|\mathbf{z} - \varepsilon(\mathbf{z} - \mathbf{v})\|^2.$$

Multiplying this out, we get

$$\|\mathbf{z}\|^2 \leq \|\mathbf{z}\|^2 - 2\varepsilon\mathbf{z}^T(\mathbf{z} - \mathbf{v}) + \varepsilon^2\|\mathbf{z} - \mathbf{v}\|^2$$

Cancelling $\|\mathbf{z}\|^2$ and rearranging terms we get

$$2\varepsilon\mathbf{z}^T(\mathbf{z} - \mathbf{v}) \leq \varepsilon^2\|\mathbf{z} - \mathbf{v}\|^2$$

or

$$\mathbf{z}^T(\mathbf{z} - \mathbf{v}) \leq \frac{\varepsilon}{2}\|\mathbf{z} - \mathbf{v}\|^2.$$

Letting ε approach 0, we find

$$\mathbf{z}^T(\mathbf{z} - \mathbf{v}) \leq 0 \tag{2.12}$$

which means that

$$\|\mathbf{z}\|^2 \leq \mathbf{z}^T\mathbf{v}.$$

Since $z \in K$ and $\mathbf{0} \notin K$, the norm $\|\mathbf{z}\| > 0$. Choosing $c = \frac{1}{2}\|\mathbf{z}\|^2$, we get $0 < c < \mathbf{z}^T\mathbf{v}$ for each $\mathbf{v} \in K$. \square

We will also need the following simple lemma:

Lemma 2.6.3. *Let X and Y be closed and bounded sets in \mathbb{R}^d . Let $f : X \times Y \rightarrow \mathbb{R}$ be continuous. Then*

$$\max_{\mathbf{x} \in X} \min_{\mathbf{y} \in Y} f(\mathbf{x}, \mathbf{y}) \leq \min_{\mathbf{y} \in Y} \max_{\mathbf{x} \in X} f(\mathbf{x}, \mathbf{y}).$$

Proof. We first prove the lemma for the case where X and Y are finite sets. Let $(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \in X \times Y$. Clearly we have $f(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \leq \max_{\mathbf{x} \in X} f(\mathbf{x}, \tilde{\mathbf{y}})$ and $\min_{\mathbf{y} \in Y} f(\tilde{\mathbf{x}}, \mathbf{y}) \leq f(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$, which gives us

$$\min_{\mathbf{y} \in Y} f(\tilde{\mathbf{x}}, \mathbf{y}) \leq \max_{\mathbf{x} \in X} f(\mathbf{x}, \tilde{\mathbf{y}}).$$

Because the inequality holds for any $\tilde{\mathbf{x}} \in X$, it holds for $\max_{\mathbf{x} \in X}$ of the quantity on the left. Similarly, because the inequality holds for all $\tilde{\mathbf{y}} \in Y$, it must hold for the $\min_{\mathbf{y} \in Y}$ of the quantity on the right. We thus have:

$$\max_{\mathbf{x} \in X} \min_{\mathbf{y} \in Y} f(\mathbf{x}, \mathbf{y}) \leq \min_{\mathbf{y} \in Y} \max_{\mathbf{x} \in X} f(\mathbf{x}, \mathbf{y}).$$

To prove the lemma in the general case, we just need to verify the existence of the relevant maxima and minima. Since continuous functions achieve their

minimum on compact sets, $g(\mathbf{x}) = \min_{\mathbf{y} \in Y} f(\mathbf{x}, \mathbf{y})$ is well-defined. The continuity of f and compactness of $X \times Y$ imply that f is uniformly continuous on $X \times Y$. In particular,

$$\forall \epsilon \exists \delta : |\mathbf{x}_1 - \mathbf{x}_2| < \delta \implies |f(\mathbf{x}_1, \mathbf{y}) - f(\mathbf{x}_2, \mathbf{y})| \leq \epsilon$$

and hence $|g(\mathbf{x}_1) - g(\mathbf{x}_2)| \leq \epsilon$. Thus, $g : X \rightarrow \mathbb{R}$ is continuous and $\max_{\mathbf{x} \in X} g(\mathbf{x})$ exists. □

We can now prove:

Theorem 2.6.4 (Von Neumann's Minimax Theorem). *Let A be an $m \times n$ payoff matrix, and let $\Delta_m = \{\mathbf{x} \in \mathbb{R}^m : \mathbf{x} \geq \mathbf{0}, \sum_i x_i = 1\}$ and $\Delta_n = \{\mathbf{y} \in \mathbb{R}^n : \mathbf{y} \geq \mathbf{0}, \sum_j y_j = 1\}$. Then*

$$\max_{\mathbf{x} \in \Delta_m} \min_{\mathbf{y} \in \Delta_n} \mathbf{x}^T A \mathbf{y} = \min_{\mathbf{y} \in \Delta_n} \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^T A \mathbf{y}.$$

*This quantity is called the **value** of the two-person zero-sum game with payoff matrix A .*

Proof. The inequality

$$\max_{\mathbf{x} \in \Delta_m} \min_{\mathbf{y} \in \Delta_n} \mathbf{x}^T A \mathbf{y} \leq \min_{\mathbf{y} \in \Delta_n} \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^T A \mathbf{y}$$

follows immediately from the Lemma 2.6.3 because $f(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T A \mathbf{y}$ is a continuous function in both variables and $\Delta_m \subset \mathbb{R}^m$, $\Delta_n \subset \mathbb{R}^n$ are closed and bounded.

For the other inequality, suppose towards a contradiction that

$$\max_{\mathbf{x} \in \Delta_m} \min_{\mathbf{y} \in \Delta_n} \mathbf{x}^T A \mathbf{y} < \lambda < \min_{\mathbf{y} \in \Delta_n} \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^T A \mathbf{y}.$$

Define a new game with payoff matrix \hat{A} given by $\hat{a}_{i,j} = a_{i,j} - \lambda$. For this new game, since each payoff in the matrix is reduced by λ , the expected payoffs for every pair of mixed strategies are also reduced by λ and hence:

$$\max_{\mathbf{x} \in \Delta_m} \min_{\mathbf{y} \in \Delta_n} \mathbf{x}^T \hat{A} \mathbf{y} < 0 < \min_{\mathbf{y} \in \Delta_n} \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^T \hat{A} \mathbf{y}. \quad (2.13)$$

Each mixed strategy $\mathbf{y} \in \Delta_n$ for player II yields a gain vector $\hat{A} \mathbf{y} \in \mathbb{R}^m$. Let K denote the set of all vectors which dominate the gain vectors $\hat{A} \mathbf{y}$, that is,

$$K = \left\{ \hat{A} \mathbf{y} + \mathbf{v} : \mathbf{y} \in \Delta_n, \mathbf{v} \in \mathbb{R}^m, \mathbf{v} \geq \mathbf{0} \right\}.$$

See Figure 2.8.




Fig. 2.8. The figure shows the set K for an example where \hat{A} has two rows. Here $\hat{A}_{(i)}$ represents the i th row of \hat{A} .

It is easy to see that K is convex and closed: this follows immediately from the fact that Δ_n , the set of probability vectors corresponding to mixed strategies \mathbf{y} for player II, is closed, bounded and convex, and the set $\{\mathbf{v} \in \mathbb{R}^m, \mathbf{v} \geq \mathbf{0}\}$ is closed and convex. Also, K cannot contain the $\mathbf{0}$ vector, because if $\mathbf{0}$ were in K , there would be some mixed strategy $\mathbf{y} \in \Delta_n$ such that $\hat{A}\mathbf{y} \leq \mathbf{0}$, whence for any $\mathbf{x} \in \Delta_m$ we have $\mathbf{x}^T \hat{A}\mathbf{y} \leq 0$, contradicting the right-hand side of (2.13).

Thus K satisfies the conditions of the separating hyperplane theorem (Theorem 2.6.2), which gives us $\mathbf{z} \in \mathbb{R}^m$ and $c > 0$ such that $\mathbf{z}^T \mathbf{w} > c > 0$ for all $\mathbf{w} \in K$. That is,

$$\mathbf{z}^T (\hat{A}\mathbf{y} + \mathbf{v}) > c > 0 \text{ for all } \mathbf{y} \in \Delta_n \text{ and } \mathbf{v} \geq \mathbf{0}. \quad (2.14)$$

We claim also that $\mathbf{z} \geq \mathbf{0}$. If not, say $z_j < 0$ for some j , then for $\mathbf{v} \in \mathbb{R}^m$ with v_j sufficiently large and $v_i = 0$ for all $i \neq j$, we would have $\mathbf{z}^T (\hat{A}\mathbf{y} + \mathbf{v}) = \mathbf{z}^T \hat{A}\mathbf{y} + z_j v_j < 0$ for some $\mathbf{y} \in \Delta_n$ which would contradict (2.14).

The same condition (2.14) shows that not all of the z_i 's can be zero. Thus $s = \sum_{i=1}^m z_i$ is strictly positive, so that $\tilde{\mathbf{x}} = \frac{1}{s}(z_1, \dots, z_m)^T = \mathbf{z}/s \in \Delta_m$, with $\tilde{\mathbf{x}}^T \hat{A}\mathbf{y} > c/s > 0$ for all $\mathbf{y} \in \Delta_n$.

In other words, $\tilde{\mathbf{x}}$ is a mixed strategy for player I that gives a positive expected payoff against any mixed strategy of player II. This contradicts the left hand inequality of (2.13). \square

Note that the above proof merely shows that the value always exists; it doesn't give a way of finding it. In fact, there are efficient algorithms for finding the value and the optimal strategies in a 2-person zero-sum game and we discuss those in the next section.

2.7 Linear Programming and the Minimax Theorem

Suppose that we want to determine if player I in a two-person zero-sum game with m by n payoff matrix $A = (a_{ij})$ can guarantee an expected gain of at least v . It suffices for her to find a mixed strategy \mathbf{x} which guarantees her an expected gain of at least v for each possible pure strategy j player II might play. These conditions are captured by the following system of inequalities:

$$x_1 a_{1j} + x_2 a_{2j} + \dots + x_m a_{mj} \geq v \text{ for } 1 \leq j \leq n.$$

In matrix-vector notation, this system of inequalities becomes:

$$\mathbf{x}^T A \geq v \mathbf{e}^T,$$

where \mathbf{e} is an all-1's vector. (Its length will be clear from context.)

Thus, to maximize her guaranteed expected gain, player I should

$$\begin{aligned} & \text{maximize } v \\ \text{subject to } & \mathbf{x}^T A \geq v \mathbf{e}^T \\ & \sum_{1 \leq i \leq m} x_i = 1 \\ & x_i \geq 0 \text{ for all } 1 \leq i \leq m. \end{aligned} \tag{2.15}$$

This is an example of a **linear programming problem**. Linear programming is the process of minimizing or maximizing a linear function of a finite set of real-valued variables, subject to linear equality and inequality constraints on those variables. In the linear program 2.15, the variables are v and x_1, \dots, x_m .

The problem of finding the optimal strategy for player II can similarly be formulated as a linear program:

$$\begin{aligned} & \text{minimize } v \\ \text{subject to } & A \mathbf{y} \leq v \mathbf{e} \\ & \sum_{1 \leq j \leq n} y_j = 1 \\ & y_j \geq 0 \text{ for all } 1 \leq j \leq n. \end{aligned} \tag{2.16}$$

As many fundamental problems can be formulated as linear programs, this is a tremendously important class of problems. Conveniently, there are well-known efficient (polynomial time) algorithms for solving linear programs (see notes) and, thus, we can use these algorithms to solve for optimal

strategies in large zero-sum games. In the rest of the chapter, we give a brief introduction to the theory of linear programming.

2.7.1 Linear Programming Basics

Example 2.7.1. (The protein problem). Consider the dilemma faced by a student-athlete interested in maximizing her protein consumption, while consuming no more than 5 units of fat per day and spending no more than \$6 a day. She considers two alternatives: steak, which costs \$4 per pound, and contains 2 units of protein and 1 unit of fat per pound; and peanut butter, which costs \$1 per pound and contains 1 unit of protein and 2 units of fat per pound.

Let x_1 be the number of pounds of steak she buys per day, and let x_2 be the number of pounds of peanut butter she buys per day. Then her goal is to

$$\begin{aligned} & \max 2x_1 + x_2 \\ & \text{subject to} \quad 4x_1 + x_2 \leq 6 \\ & \quad \quad \quad x_1 + 2x_2 \leq 5 \\ & \quad \quad \quad x_1, x_2 \geq 0 \end{aligned} \tag{2.17}$$

The feasible region for the LP and its optimal solution are shown in Figure 2.9.




Fig. 2.9. This figure shows the feasible region for LP 2.17 and illustrates its solution. The arrow from the origin on the right is perpendicular to all the lines $2x_1 + x_2 = c$ for any c .

The **objective function** of a linear program is the linear function being optimized, in this case $2x_1 + x_2$. The **feasible set** of a linear program is the

set of **feasible** vectors (x_1, x_2) that satisfy the constraints of the program, in this case, all nonnegative vectors (x_1, x_2) such that $4x_1 + x_2 \leq 6$ and $x_1 + 2x_2 \leq 5$.

The left hand side of Figure 2.9 shows this set. A linear program is said to be **feasible** if the feasible set is non-empty. The question then becomes: which point in this feasible set maximizes $2x_1 + x_2$? In this example, this point is $(x_1, x_2) = (1, 2)$, and at this point $2x_1 + x_2 = 4$. Thus, the optimal solution to the linear program is 4.

2.7.2 Linear Programming Duality

The minimax theorem that we proved earlier shows that for any zero-sum game, the two linear programs (2.15) and (2.16) have the same optimal value V^* . This is a special case of the most important theorem of linear programming, **the duality theorem**.

To motivate this theorem, let's consider the LP from the previous section more analytically. The first constraint of (2.17) immediately implies that the objective function is upper bounded by 6 on the feasible set. Doubling the second constraint gives a worse bound of 10. But combining them we can do better.

Multiplying the first constraint by $y_1 \geq 0$, the second by $y_2 \geq 0$, and adding the results yields

$$y_1(4x_1 + x_2) + y_2(x_1 + 2x_2) \leq 6y_1 + 5y_2 \quad (2.18)$$

The left hand side of equation (2.18) dominates the objective function $2x_1 + x_2$ as long as

$$\begin{aligned} 4y_1 + y_2 &\geq 2 \\ y_1 + 2y_2 &\geq 1 \\ y_1, y_2 &\geq 0 \end{aligned} \quad (2.19)$$

So for any (y_1, y_2) that satisfy inequalities in (2.19), we have $2x_1 + x_2 \leq 6y_1 + 5y_2$ for all feasible (x_1, x_2) . The best upper bound we can obtain this way on the optimal value of (2.17) is the solution to the linear program

$$\min 6y_1 + 5y_2 \text{ subject to (2.19)}. \quad (2.20)$$

This minimization problem is called the **dual** of LP (2.17). Observing that $(y_1, y_2) = (3/7, 2/7)$ is feasible for LP (2.20) with objective value 4, we can conclude that $(x_1, x_2) = (1, 2)$, which attains objective value 4 for the original problem, must be optimal.

2.7.3 Duality, more formally

Consider a maximization linear program in so-called standard form[†]. We will call such a linear program the **primal LP (P)**:

$$\left. \begin{array}{l} \max \quad \mathbf{c}^T \mathbf{x} \\ \text{subject to} \\ \mathbf{Ax} \leq \mathbf{b} \\ \mathbf{x} \geq 0 \end{array} \right\} \quad (\mathbf{P})$$

where $A \in \mathbb{R}^{m \times n}$, $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{c} \in \mathbb{R}^n$, and $\mathbf{b} \in \mathbb{R}^m$. We say the primal LP is **feasible** if the feasible set $\mathcal{F}(P) = \{\mathbf{x} \mid \mathbf{Ax} \leq \mathbf{b}, \mathbf{x} \geq 0\}$ is nonempty.

As in the example at the beginning of this section, if $\mathbf{y} \geq 0 \in \mathbb{R}^m$ satisfies $\mathbf{y}^T A \geq \mathbf{c}^T$, then

$$\forall \mathbf{x} \in \mathcal{F}(P), \quad \mathbf{y}^T \mathbf{b} \geq \mathbf{y}^T \mathbf{Ax} \geq \mathbf{c}^T \mathbf{x}. \quad (2.21)$$

This motivates the general definition of the **dual LP**:

$$\left. \begin{array}{l} \min \quad \mathbf{b}^T \mathbf{y} \\ \text{such that} \\ \mathbf{y}^T A \geq \mathbf{c}^T \\ \mathbf{y} \geq 0 \end{array} \right\} \quad (\mathbf{D})$$

where $\mathbf{y} \in \mathbb{R}^m$. As with the primal LP, we say the dual LP is feasible if the set $\mathcal{F}(D) = \{\mathbf{y} \mid \mathbf{y}^T A \geq \mathbf{c}^T; \mathbf{y} \geq 0\}$ is nonempty.

It is easy to check that the dual of the dual LP is the primal LP.[†]

Theorem 2.7.2 (The Duality Theorem of Linear Programming). *Suppose $A \in \mathbb{R}^{m \times n}$, $\mathbf{x}, \mathbf{c} \in \mathbb{R}^n$, and $\mathbf{y}, \mathbf{b} \in \mathbb{R}^m$. Suppose $\mathcal{F}(P)$ and $\mathcal{F}(D)$ are nonempty. Then:*

- $\mathbf{b}^T \mathbf{y} \geq \mathbf{c}^T \mathbf{x}$ for all $\mathbf{x} \in \mathcal{F}(P)$ and $\mathbf{y} \in \mathcal{F}(D)$. (This is called *weak duality*.)
- (P) has an optimal solution \mathbf{x}^* , (D) has an optimal solution \mathbf{y}^* and $\mathbf{c}^T \mathbf{x}^* = \mathbf{b}^T \mathbf{y}^*$.

[†] It is a simple exercise to convert from non-standard form (such as a game LP) to standard form. For example, an equality constraint such as $a_1 x_1 + a_2 x_2 + \dots + a_n x_n = b$ can be converted to two inequalities: $a_1 x_1 + a_2 x_2 + \dots + a_n x_n \geq b$ and $a_1 x_1 + a_2 x_2 + \dots + a_n x_n \leq b$. A \geq inequality can be converted to a \leq inequality and vice versa by multiplying by -1. A variable x that is not constrained to be nonnegative, can be replaced by the difference $x' - x''$ of two nonnegative variables, and so on.

[†] A standard form minimization LP can be converted to a maximization LP (and vice versa) by observing that minimizing $\mathbf{b}^T \mathbf{y}$ is the same as maximizing $-\mathbf{b}^T \mathbf{y}$, and \geq inequalities can be converted to \leq inequalities by multiplying the inequality by -1.

Remark. The proof of the duality theorem is similar to the proof of the minimax theorem. This is not accidental; see the chapter notes.

Corollary 2.7.3 (Complementary Slackness). *Let \mathbf{x}^* be feasible for (P) and let \mathbf{y}^* be feasible for (D). Then the following two statements are equivalent:*

- (i) \mathbf{x}^* is optimal for (P) and \mathbf{y}^* is optimal for (D).
- (ii) For all i , $\sum_{1 \leq j \leq n} a_{ij}x_j^* < b_i$ if and only if $y_i^* = 0$, and for all j , $c_j < \sum_{1 \leq i \leq m} y_i^* a_{ij}$ if and only if $x_j^* = 0$.

Proof. We have

$$\sum_j c_j x_j^* \leq \sum_j x_j^* \sum_i y_i^* a_{ij} = \sum_i y_i^* \sum_j a_{ij} x_j^* \leq \sum_i b_i y_i^*. \quad (2.22)$$

Optimality of \mathbf{x}^* and \mathbf{y}^* implies that both of the above inequalities are equalities. Moreover by feasibility, for each j we have $c_j x_j^* \leq x_j^* \sum_i y_i^* a_{ij}$, and for each i we have $y_i^* \sum_j a_{ij} x_j^* \leq b_i y_i^*$. Thus equality holds in (2.22) if and only if (ii) holds. \square

2.7.4 The proof of the duality theorem

Weak duality follows from (2.21). We complete the proof of the duality theorem in two steps. First, we will use the separating hyperplane theorem to show that $\sup_{\mathbf{x} \in \mathcal{F}(P)} \mathbf{c}^T \mathbf{x} = \inf_{\mathbf{y} \in \mathcal{F}(D)} \mathbf{b}^T \mathbf{y}$, and then we will show that the sup and inf above are attained. For the first step, it will be convenient to establish the following “alternative” theorem known as **Farkas’ Lemma**, from which the proof of duality will follow.

Lemma 2.7.4 (Farkas’ Lemma – 2 versions). *Let $A \in R^{m \times n}$ and $\mathbf{x} \in R^n$. Then*

- (i) *Exactly one of the following holds:*
 - (a) *There exists a $\mathbf{x} \in R^n$ such that $A\mathbf{x} = \mathbf{b}$ and $\mathbf{x} \geq 0$; or*
 - (b) *there exists a $\mathbf{y} \in R^m$ such that $\mathbf{y}^T A \geq 0$ and $\mathbf{y}^T \mathbf{b} < 0$.*
- (ii) *Exactly one of the following holds:*
 - (a) *There exists a $\mathbf{x} \in R^n$ such that $A\mathbf{x} \leq \mathbf{b}$ and $\mathbf{x} \geq 0$; or*
 - (b) *there exists a $\mathbf{y} \in R^m$ such that $\mathbf{y}^T A \geq 0$, $\mathbf{y}^T \mathbf{b} < 0$ and $\mathbf{y} \geq 0$.*

Proof. Proof of Part (i): (See Figure 2.10 for a visualization of Part (i).) We first show by contradiction that (a) and (b) can’t hold simultaneously: Suppose that \mathbf{x} satisfies (a) and \mathbf{y} satisfies (b). Then

$$0 > \mathbf{y}^T \mathbf{b} = \mathbf{y}^T A \mathbf{x} \geq 0,$$




Fig. 2.10. The figure illustrates the two cases (i)a and (i)b of the Farkas' Lemma

a contradiction.

We next show that if (a) is infeasible, then (b) is feasible: Let $S = \{Ax \mid x \geq 0\}$. It is easy to check that S is closed and convex. In addition, $b \notin S$, since (a) is infeasible. Therefore, by the separating hyperplane theorem, there is a hyperplane that separates b from S , i.e. $y^T b < a$ and $y^T z \geq a$ for all $z \in S$. Since 0 is in S , $a \leq 0$ and therefore $y^T b < 0$. Moreover, all entries of $y^T A$ are nonnegative. If not, say the k th entry is negative, by taking x_k arbitrarily large and $x_i = 0$ for $i \neq k$, the inequality $y^T Ax \geq a$ would be violated for some $x \geq 0$. Thus, it must be that $y^T A \geq 0$.

Proof of Part (ii): We apply part (i) to an equivalent pair of systems. The existence of an $x \in R^n$ such that $Ax \leq b$ and $x \geq 0$ is equivalent to the existence of an $x \geq 0 \in R^n$ and $v \geq 0 \in R^m$ such that

$$Ax + Iv = b$$

where I is the m by m identity matrix. Applying part 1 to this system means that either it is feasible or there is a $y \in R^m$ such that

$$\begin{aligned} y^T A &\geq 0 \\ Iy &\geq 0 \\ y^T b &< 0, \end{aligned}$$

which is precisely equivalent to (b). □

Corollary 2.7.5. Under the assumptions of Theorem 2.7.2

$$\sup_{x \in \mathcal{F}(P)} c^T x = \inf_{y \in \mathcal{F}(D)} b^T y.$$

Proof. Suppose that $\sup_{\mathbf{x} \in \mathcal{F}(P)} \mathbf{c}^T \mathbf{x} < \gamma$. Then $\{\mathbf{A}\mathbf{x} \leq \mathbf{b}; -\mathbf{c}^T \mathbf{x} \leq -\gamma; \mathbf{x} \geq 0\}$ is infeasible, and therefore by the second part of the Farkas lemma there is $(\mathbf{y}, \lambda) \geq 0$ in \mathbb{R}^{m+1} such that $\mathbf{y}^T \mathbf{A} \geq 0$, $-\lambda \mathbf{c}^T \geq 0$ and $\mathbf{y}^T \mathbf{b} - \lambda \gamma < 0$. Since there is an $\mathbf{x} \in \mathcal{F}(P)$, we have $\mathbf{y}^T \mathbf{b} \geq \mathbf{y}^T \mathbf{A}\mathbf{x} \geq 0$ and therefore $\lambda > 0$. We conclude that \mathbf{y}/λ is feasible for (D) with objective value less than γ . \square

To complete the proof of the duality theorem, we need to show that the sup and inf in Corollary 2.7.5 are achieved. This will follow from the next theorem.

Theorem 2.7.6. *Let $A \in \mathbb{R}^{m \times n}$ and $\mathbf{x} \in \mathbb{R}^n$.*

- (i) *Let $\mathcal{F}(P_-) = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x} \geq 0 \text{ and } \mathbf{A}\mathbf{x} = \mathbf{b}\}$. If $\mathcal{F}(P_-) \neq \emptyset$ and $\sup\{\mathbf{c}^T \mathbf{x} | \mathbf{x} \in \mathcal{F}(P_-)\} < \infty$, then this sup is achieved.*
- (ii) *If $\mathcal{F}(P) \neq \emptyset$ and $\sup\{\mathbf{c}^T \mathbf{x} | \mathbf{x} \in \mathcal{F}(P)\} < \infty$, then this sup is achieved.*

The proof of (i) will show that the sup is attained at one of a distinguished, finite set of points in $\mathcal{F}(P_-)$ known as extreme points or vertices.

- Definition 2.7.7.**
- (i) Let S be a convex set. A point $\mathbf{x} \in S$ is an **extreme point** of S if whenever $\mathbf{x} = \alpha \mathbf{u} + (1 - \alpha) \mathbf{v}$ with $\mathbf{u}, \mathbf{v} \in S$ and $0 < \alpha < 1$, we must have $\mathbf{x} = \mathbf{u} = \mathbf{v}$.
 - (ii) If S is the feasible set of a linear program, then S is convex; an extreme point of S is called a **vertex**.

We will need the following lemma.

Lemma 2.7.8. *Let $\mathbf{x} \in \mathcal{F}(P_-)$. Then \mathbf{x} is a vertex of $\mathcal{F}(P_-)$ if and only if the columns $\{A^{(j)} \mid x_j > 0\}$ are linearly independent.*

Proof. Suppose \mathbf{x} is not extreme, i.e., $\mathbf{x} = \alpha \mathbf{v} + (1 - \alpha) \mathbf{w}$, where $\mathbf{v} \neq \mathbf{w}$, $0 < \alpha < 1$, and $\mathbf{v}, \mathbf{w} \in \mathcal{F}(P_-)$. Thus, $A(\mathbf{v} - \mathbf{w}) = 0$, and $\mathbf{v} - \mathbf{w} \neq \mathbf{0}$. Observe that $v_j = w_j = 0$ for all $j \notin S$, where $S = \{j \mid x_j > 0\}$; otherwise, one of w_j or v_j is negative. We conclude that the columns $\{A^{(j)} \mid x_j > 0\}$ are linearly dependent.

For the other direction, suppose that the vectors $\{A^{(j)} \mid x_j > 0\}$ are linearly dependent. Then there is $\mathbf{w} \neq \mathbf{0}$ such that $A\mathbf{w} = 0$ and $w_j = 0$ for all $j \notin S$. Then for ϵ sufficiently small $\mathbf{x} \pm \epsilon \mathbf{w} \in \mathcal{F}(P_-)$ and therefore \mathbf{x} is not extreme. \square

Lemma 2.7.9. *For any $\mathbf{x} \in \mathcal{F}(P_-)$, there is a vertex $\tilde{\mathbf{x}} \in \mathcal{F}(P_-)$ with $\mathbf{c}^T \tilde{\mathbf{x}} \geq \mathbf{c}^T \mathbf{x}$.*

Proof. We show that if \mathbf{x} is not a vertex, then there is $\mathbf{x}' \in \mathcal{F}(P_{=})$ with a strictly larger number of zero entries than \mathbf{x} such that $\mathbf{c}^T \mathbf{x}' \geq \mathbf{c}^T \mathbf{x}$. This step can be applied only a finite number of times before we reach a vertex that satisfies the conditions of the lemma.

Let $S = \{j | x_j > 0\}$. If \mathbf{x} is not a vertex, then the columns $\{A^{(j)} | j \in S\}$ are linearly dependent and there is a vector $\boldsymbol{\lambda} \neq \mathbf{0}$ such that $\sum_j \lambda_j A^{(j)} = A\boldsymbol{\lambda} = \mathbf{0}$ and $\lambda_j = 0$ for $j \notin S$.

Without loss of generality, $\mathbf{c}^T \boldsymbol{\lambda} \geq 0$ (if not, negate $\boldsymbol{\lambda}$). Consider the vector $\hat{\mathbf{x}}(t) = \mathbf{x} + t\boldsymbol{\lambda}$. For $t \geq 0$, we have $\mathbf{c}^T \hat{\mathbf{x}}(t) \geq \mathbf{c}^T \mathbf{x}$ and $A\hat{\mathbf{x}}(t) = \mathbf{b}$. For t sufficiently small, $\hat{\mathbf{x}}(t)$ is also nonnegative and thus feasible.

If there is $j \in S$ such that $\lambda_j < 0$, then there is a positive t such that $\hat{\mathbf{x}}(t)$ is feasible with strictly more zeros than \mathbf{x} , so we can take $\mathbf{x}' = \hat{\mathbf{x}}(t)$.

The same conclusion holds if $\lambda_j \geq 0$ for all j and $\mathbf{c}^T \boldsymbol{\lambda} = 0$; simply negate $\boldsymbol{\lambda}$ and apply the previous argument.

To complete the argument, we show that the previous two cases are exhaustive: if $\lambda_j \geq 0$ for all j and $\mathbf{c}^T \boldsymbol{\lambda} > 0$, then $\hat{\mathbf{x}}(t) \geq 0$ for all $t \geq 0$ and $\lim_{t \rightarrow \infty} \mathbf{c}^T \hat{\mathbf{x}}(t) = \infty$, contradicting the assumption that the objective value is bounded on $\mathcal{F}(P_{=})$. □

Proof of Theorem 2.7.6:

Part (i): Lemma 2.7.9 shows that if the linear program

$$\text{maximize } \mathbf{c}^T \mathbf{x} \text{ subject to } \mathbf{x} \in \mathcal{F}(P_{=})$$

is feasible and bounded, then for every feasible solution, there is a vertex with at least that objective value. Thus, we can search for the optimum of the linear program by considering only vertices of $\mathcal{F}(P_{=})$. Since there are only finitely many, the optimum is achieved.

Part (ii): We apply the reduction from part (ii) of the Farkas' Lemma to show that linear program (P) is equivalent to a program of the type considered in part (i) with a matrix $(A; I)$ in place of A . □

2.7.5 An interpretation of a primal/dual pair

Consider an advertiser about to purchase advertising space in a set of n newspapers, and suppose that c_j is the price of placing an ad in newspaper j . The advertiser is targeting m different kinds of users, for example, based on geographic location, interests, age and gender, and wants to ensure that, on average, b_i users of type i will see the ad over the course of each month.

Denote by a_{ij} the number of type i users expected to see each ad in newspaper j . The advertiser is deciding how many ads to place in each newspaper per month in order to meet his various demographic targets at minimum cost. To this end, the advertiser solves the following linear program, where x_j is the number of ad slots from newspaper j that she will purchase.

$$\begin{aligned} & \min \sum_{1 \leq j \leq n} c_j x_j \\ \text{subject to} & \sum_{1 \leq j \leq n} a_{ij} x_j \geq b_i \quad \text{for all } 1 \leq i \leq m \\ & x_1, x_2, \dots, x_n \geq 0. \end{aligned} \quad (2.23)$$

The dual program is:

$$\begin{aligned} & \min \sum_{1 \leq i \leq m} b_i y_i \\ \text{subject to} & \sum_{1 \leq i \leq m} y_i a_{ij} \leq c_j \quad \text{for all } 1 \leq j \leq n \\ & y_1, y_2, \dots, y_m \geq 0. \end{aligned} \quad (2.24)$$

This dual program has a nice interpretation: Consider an advertising exchange that matches advertisers with display ad slots. The exchange needs to determine y_i , how much to charge the advertiser for each impression (displayed ad) shown to a user of type i . Observing that $y_i a_{ij}$ is the expected cost of reaching the same number of type i users online that would be reached by placing a single ad in newspaper j , we see that if the prices y_i are set so that $\sum_{1 \leq i \leq m} y_i a_{ij} \leq c_j$, then the advertiser can switch from advertising in newspaper j to advertising online, reaching the same combination of user types without increasing her cost. If the advertiser switches entirely from advertising in newspapers to advertising online, the exchange's revenue will be

$$\sum_{1 \leq i \leq m} b_i y_i.$$

The duality theorem says that the exchange can price the impressions so as to satisfy (2.24) and incentivize the advertiser to switch while still ensuring that its revenue $\sum_i b_i y_i$ matches the total revenue of the newspapers.

Moreover, Theorem 2.7.3 implies that whenever inequality (2.23) is not tight, say $\sum_{1 \leq j \leq n} a_{ij} x_j > b_i$ for user type i , in the optimal solution of the dual, $y_i = 0$. In other words, if the optimal combination of ads

the advertiser buys from the newspapers results in the advertisement being shown to more users of type i than necessary, then in the optimal pricing for the exchange, impressions shown to users of type i will be provided to the advertiser for free. In other words, the exchange concentrates its fixed total charges on the user types which correspond to tight constraints in the primal. Thus, the advertiser can switch to advertising exclusively on the exchange without paying more, and without sacrificing any of the “bonus” advertising the newspapers were providing.

(The fact that some impressions are free may seem counterintuitive, but it is a consequence of the assumption that the exchange maximizes revenue from this advertiser. In reality, the exchange would maximize profit, and these goals are equivalent only when the cost of production is zero.)

Finally, the other consequence of Theorem 2.7.3 is that if $x_j > 0$, i.e., some ads were purchased from newspaper j , then the corresponding dual constraint must be tight, i.e., $\sum_{1 \leq i \leq m} y_i a_{ij} = c_j$.

2.8 Zero-Sum Games With Infinite Action Spaces*

Theorem 2.8.1. *Consider a zero-sum game in which the players' action spaces are $[0, 1]$ and the payoff $A(x, y)$ when player I chooses action x and player II chooses action y is continuous on $[0, 1]^2$. Let $\Delta = \Delta_{[0,1]}$ be the space of probability distributions on $[0, 1]$. Then*

$$\max_{F \in \Delta} \min_{G \in \Delta} \int \int A(x, y) dF(x) dG(y) = \min_{G \in \Delta} \max_{F \in \Delta} \int \int A(x, y) dF(x) dG(y) \quad (2.25)$$

Proof. If there is a matrix (a_{ij}) for which

$$A(x, y) = a_{\lceil nx \rceil, \lceil ny \rceil} \quad (2.26)$$

then (2.25) reduces to the finite case. If A is continuous, there are functions A_0 and A_1 of the form (2.26) so that $A_0 \leq A \leq A_1$ and $|A_1 - A_0| \leq \epsilon$. This implies (2.25) with infs and sups in place of min and max. The existence of the maxima and minima follows from compactness of $\Delta_{[0,1]}$ as in the proof of Lemma 2.6.3. \square

Next, we show how a theorem in geometry due to Berge[] can be deduced from the minimax theorem.

Theorem 2.8.2. *Let $S_1, \dots, S_n \subset \mathbb{R}^\ell$ compact, convex sets such that every subset of $n - 1$ of them intersects and $S = \cup_{i=1}^n S_i$ is convex. Then $S = \cap_{i=1}^n S_i \neq \emptyset$.*

We prove the theorem by considering the following zero-sum game G : Player I chooses $i \in [n]$, and player II chooses $z \in S$. The payoff to player I is the distance $d(z, S_i)$ from z to S_i .

Lemma 2.8.3. *The game G has a value with a mixed optimal strategy for player I and a pure optimal strategy for player II.*

Proof. For each positive integer k , let $S_i(k) = S_i \cap 2^{-k}\mathbb{Z}^\ell$, and let $S(k) = \cup_{i=1}^n S_i(k)$.

Define a sequence of games G_k in which Player I chooses $i \in [n]$, and player II chooses $z \in S(k)$, where the payoff to player I is $d(z, S_i(k))$. Since G_k is a finite game, it has a value v_k , and each player has an optimal strategy, say $\mathbf{x}^{(k)}$ for player I and $\mathbf{y}^{(k)}$ for player II. Thus, for all $s \in S(k)$, we have $\sum_i x_i^{(k)} d(s, S_i(k)) \geq v_k$ and for all $i \in [n]$, we have $\sum_{s \in S(k)} y_s^{(k)} d(s, S_i(k)) \leq v_k$. The v_k 's are decreasing and bounded so they converge to a limit, say v . We now claim that

$$\sup_{\mathbf{x} \in \Delta_n} \inf_{s \in S} \sum_i x_i d(s, S_i) = v = \inf_{s \in S} \sup_{\mathbf{x} \in \Delta_n} \sum_i x_i d(s, S_i). \quad (2.27)$$

From Lemma 2.6.3 we know that the left-hand side of equation (2.27) is at most the right-hand side. We now show that the left-hand side is greater than or equal to the right-hand side. We have

$$\forall s' \in S(k) \quad \sum_i x_i^{(k)} d(s', S_i(k)) \geq v_k \geq v,$$

and thus

$$\forall s \in S \quad \sum_i x_i^{(k)} d(s, S_i) \geq v - 2\ell 2^{-k}.$$

This proves that the left-hand side of (2.27) is at least v .

Also, since $\sum_{s \in S(k)} y_s^{(k)} d(s, S_i(k)) \leq v_k$ for any i , we have

$$\sum_{s \in S(k)} y_s^{(k)} d(s, S_i) \leq v_k + \ell 2^{-k}.$$

Let $z_k = \sum_{s \in S(k)} y_s^{(k)} s$. Then by Exercise 2.8.6 and Jensen's Inequality (Exercise 2.8.5),

$$d(z_k, S_i) \leq \sum_{s \in S(k)} y_s^{(k)} d(s, S_i) \leq v_k + \ell 2^{-k}.$$

Hence

$$\forall x \in \Delta_n \quad \sum_i x_i d(z_k, S_i) \leq v_k + \ell 2^{-k}.$$

This proves that the right-hand side of (2.27) is at most v . \square

Definition 2.8.4. Let D be a convex set in \mathbb{R}^ℓ . Then $f : D \rightarrow \mathbb{R}$ is a convex function if for any two points z and w in D , and $0 \leq \alpha \leq 1$,

$$f(\alpha z + (1 - \alpha)w) \leq \alpha f(z) + (1 - \alpha)f(w).$$

Exercise 2.8.5 (Jensen's Inequality for Finite Sets). Let $f : D \rightarrow \mathbb{R}$ be convex. Let $z_1, \dots, z_m \in D$, and $\alpha \in \Delta_m$. Show that

$$f\left(\sum_i \alpha_i z_i\right) \leq \sum_i \alpha_i f(z_i).$$

Exercise 2.8.6. Let S be a convex set in \mathbb{R}^ℓ . Show that the function $f(z) = d(z, S)$ is convex on \mathbb{R}^ℓ .

Proof of Theorem 2.8.2: Let \mathbf{x} be I's optimal strategy and let $z \in S$ be player II's optimal strategy in the game G of Lemma 2.8.3. We will show that $v = 0$. If so, we have $d(z, S_i) = 0$ for all i , and thus $z \in \bigcap_{i=1}^n S_i$, completing the proof.

Suppose that $\sum_i x_i d(z, S_i) = v > 0$. We have that $z \in S_j$ for some j , and thus $d(z, S_j) = 0$. Since $d(z, S_i) \leq v$ for all i , it must be that $x_j = 0$. But then, since there is a point $w \in \bigcap_{i \neq j} S_i$, we have $\sum_i x_i d(w, S_i) = 0$, contradicting the assumption that $v > 0$. \square

Exercise 2.8.7. Two players each choose a positive integer. The player that chose the lower number pays \$1 to the player who chose the higher number (with no payment in case of a tie). Show that this game has no Nash equilibrium. Show that the safety values for players I and II are -1 and 1 respectively.

Exercises

- 2.1 Show that all saddle points in a zero-sum game (assuming there is at least one) result in the same payoff to player I.
- 2.2 Find the value of the following zero-sum game. Find some optimal strategies for each of the players.

		player II			
player I	8	3	4	1	
	4	7	1	6	
	0	3	8	5	

- 2.3 Find the value of the zero-sum game given by the following payoff matrix, and determine optimal strategies for both players.

$$\begin{pmatrix} 0 & 9 & 1 & 1 \\ 5 & 0 & 6 & 7 \\ 2 & 4 & 3 & 3 \end{pmatrix}$$

- 2.4 Find the value of the zero-sum game given by the following payoff matrix and determine *all* optimal strategies for both players.

$$\begin{pmatrix} 3 & 0 \\ 0 & 3 \\ 2 & 2 \end{pmatrix}$$

- 2.5 Define a zero-sum game in which one player's optimal strategy is pure and the other player's optimal strategy is mixed.
- 2.6 Prove that the value of any antisymmetric zero-sum game is zero.
- 2.7 Player II is moving an important item in one of three cars, labeled 1, 2, and 3. Player I will drop a bomb on one of the cars of his choosing. He has no chance of destroying the item if he bombs the wrong car. If he chooses the right car, then his probability of destroying the item depends on that car. The probabilities for cars 1, 2, and 3 are equal to $3/4$, $1/4$, and $1/2$.

Write the 3×3 payoff matrix for the game, and find some optimal winning strategies for each of the players.

- 2.8 Verify the following two facts: Every strategy that has positive probability of being played in an optimal strategy for one of the players results in the same expected payoff against an optimal opponent (one playing an optimal strategy). A strategy that is not played in an optimal strategy can't have higher expected payoff than a strategy that *is* played against an optimal opponent.
- 2.9 Let \mathbf{x} and \mathbf{y} be mixed strategies for the two players in a zero-sum. Prove that this pair of strategies is optimal if and only if there is a

number V such that:

$$V = \sum_j a_{ij}y_j \text{ for every } i \text{ such that } x_i > 0.$$

$$V \leq \sum_j a_{ij}y_j \text{ for every } i \text{ such that } x_i = 0.$$

$$V = \sum_i x_i a_{ij} \text{ for every } j \text{ such that } y_j > 0$$

$$V \geq \sum_i x_i a_{ij} \text{ for every } j \text{ such that } y_j = 0.$$

- 2.10 Using the result of the previous exercise, give an exponential time algorithm to solve an n by m two-person zero-sum game. Hint: Consider each possibility for which subset S of player I strategies have $x_i > 0$ and which subset of player II strategies T have $y_j > 0$.
- 2.11 Consider a two-person zero-sum game in which there are two maps, π_1 , a permutation (a relabelling) of the possible moves of player I, and π_2 a permutation of the possible moves of player II, for which the payoffs a_{ij} satisfy

$$a_{\pi_1(i), \pi_2(j)} = a_{ij}.$$

Prove that there is an optimal mixed strategy for player I that gives equal probability to $\pi_1(i)$ and i for each i and that there is an optimal mixed strategy for player II that gives equal probability to the moves $\pi_2(j)$ and j for each j .

- 2.12 Recall the bomber and battleship game from section 2.5.4. Set up the payoff matrix and find the value of the game G_2 .
- 2.13 Consider the following two-person zero-sum game. Both players simultaneously call out one of the numbers $\{2, 3\}$. Player 1 wins if the sum of the numbers called is odd and player 2 wins if their sum is even. The loser pays the winner the product of the two numbers called (in dollars). Find the payoff matrix, the value of the game, and an optimal strategy for each player.
- 2.14 There are two roads that leave city A and head towards city B . One goes there directly. The other branches into two new roads, each of which arrives in city B . A traveler and a troll each choose paths from city A to city B . The traveler will pay the troll a toll equal to

the number of common roads that they traverse. Set up the payoff matrix, find the value of the game, and find some optimal mixed strategies.

- 2.15 Company I opens one restaurant and company II opens two. Each company decides in which of three locations each of its restaurants will be opened. The three locations are on the line, at Central and at Left and Right, with the distance between Left and Central, and between Central and Right, equal to half a mile. A customer is located at an unknown location according to a uniform random variable within one mile each way of Central (so that he is within one mile of Central, and has an even probability of appearing in any part of this two-mile stretch). He walks to whichever of Left, Central, or Right is the nearest, and then into one of the restaurants there, chosen uniformly at random. The payoff to company I is the probability that the customer visits a company I restaurant.

Solve the game: that is, find its value, and some optimal mixed strategies for the companies.

- 2.16 Bob has a concession at Yankee Stadium. He can sell 500 umbrellas at \$10 each if it rains. (The umbrellas cost him \$5 each.) If it shines, he can sell only 100 umbrellas at \$10 each and 1000 sunglasses at \$5 each. (The sunglasses cost him \$2 each.) He has \$2500 to invest in one day, but everything that isn't sold is trampled by the fans and is a total loss.

This is a game against nature. Nature has two strategies: rain and shine. Bob also has two strategies: buy for rain or buy for shine.

Find the optimal strategy for Bob assuming that the probability for rain is 50%.

- 2.17 **The number picking game.** Two players I and II pick a positive integer each. If the two numbers are the same, no money changes hands. If the players' choices differ by 1 the player with the lower number pays \$1 to the opponent. If the difference is at least 2 the player with the higher number pays \$2 to the opponent. Find the value of this zero-sum game and determine optimal strategies for both players. (Hint: use domination.)
- 2.18 Show that in *Submarine Salvo* the submarine has an optimal strat-

egy hwere all choices containing a corner and a clockwise adjacent site are excluded. PICTURE??

- 2.19 A zebra has four possible locations to cross the Zambezi river, call them $a, b, c,$ and $d,$ arranged from north to south. A crocodile can wait (undetected) at one of these locations. If the zebra and the crocodile choose the same location, the payoff to the crocodile (that is, the chance it will catch the zebra) is 1. The payoff to the crocodile is $1/2$ if they choose adjacent locations, and 0 in the remaining cases, when the locations chosen are distinct and non-adjacent.
- (a) Write the payoff matrix for this zero-sum game in normal form.
 - (b) Can you reduce this game to a 2×2 game?
 - (c) Find the value of the game (to the crocodile) and optimal strategies for both.

For the following two exercises, see the definition of effective resistance in the notes at the end of the chapter.

- 2.20 **FIX FIX FIX** The troll-and-traveler game can be played on an arbitrary (not necessarily series-parallel) network with two distinguished points A and B . On general networks, we get a similarly elegant solution for the game defined as follows: If the troll and the traveler traverse an edge in the same direction, the traveler pays the cost of the road to the troll, whereas if they traverse a road in opposite directions, then the troll pays the cost of the road to the traveler. The value of the game turns out to be the effective resistance between A and B . **PROBLEM 1:** the simple non series-parallel network. **PROBLEM 2:** the general case.
- 2.21 **A recursive zero-sum game.** An inspector can inspect a facility on just one occasion, on one of the days $1, \dots, n$. The worker at the facility can cheat or be honest on any given day. The payoff to the inspector is 1 if he inspects while the worker is cheating. The payoff is -1 if the worker cheats and is not caught. The payoff is also -1 if the inspector inspects but the worker did not cheat, and there is at least one day left. This leads to the following matrices Γ_n for the game with n days: the matrix Γ_1 is shown on the left, and the matrix Γ_n is shown on the right.

		worker	
		cheat	honest
inspector	inspect	1	0
	wait	-1	0

		worker	
		cheat	honest
inspector	inspect	1	-1
	wait	-1	Γ_{n-1}

Find the optimal strategies and the value of Γ_n .

- 2.22 Prove that every k -regular bipartite graph has a perfect matching.
- 2.23 FIX Prove that every bistochastic n by n matrix is a convex combination of permutation matrices.
- 2.24 • Prove that if set $G \subseteq \mathbb{R}^d$ is compact and $H \subseteq \mathbb{R}^d$ is closed, then $G + H$ is closed. (This fact is used in the proof of the minimax theorem to show that the set K is closed.)
 Proof: $x_n + y_n \rightarrow z$. x_n from G , $x_{n_k} \rightarrow x \in G$ $y_{n_k} \rightarrow z - x$ implies $z - x \in H$.
- Find $F_1, F_2 \subset \mathbb{R}^2$ closed such that $F_1 - F_2$ is not closed.
 Solution: $F_1 = \{xy \geq 1\}$, $F_2 = \{x = 0\}$, $F_1 + F_2 = \{x > 0\}$.
- 2.25 Prove that linear programs (2.15) and (2.16) are dual to each other.

2.9 Solved Exercises

2.9.1 Another Betting Game

Consider the betting game with the following payoff matrix:

		player II	
		L	R
player I	T	0	2
	B	5	1

Draw graphs for this game analogous to those shown in Figure 2.1.

Solution:

Suppose player I plays T with probability x_1 and B with probability $1 - x_1$, and player II plays L with probability y_1 and R with probability $1 - y_1$. (We note that in this game, there is no saddle point.)

Reasoning from player I's perspective, her expected gain is $2(1 - y_2)$ for playing the pure strategy T , and $4y_2 + 1$ for playing the pure strategy B . Thus, if she knows y_2 , she will pick the strategy corresponding to the maximum of $2(1 - y_2)$ and $4y_2 + 1$. Player II can choose $y_2 = 1/6$ so as to minimize this maximum, and the expected amount player II will pay player I is $5/3$. This is the player II strategy that minimizes his worst-case loss. See Figure 2.11 for an illustration.

From player II's perspective, his expected loss is $5(1 - x_1)$ if he plays the pure strategy L and $1 + x_1$ if he plays the pure strategy R , and he will




Fig. 2.11. The left side of the figure shows the worst-case expected gain of player I as a function of her mixed strategy (where she plays T with probability x_1 and B with probability $1 - x_1$). This worst case expected gain is maximized when she plays T with probability $2/3$ and B with probability $1/3$. The right side of the figure shows the worst-case expected loss of player II as a function of his mixed strategy (where he plays L with probability y_1 and R with probability $1 - y_1$). The worst case expected loss is minimized when he plays L with probability $1/6$ and R with probability $5/6$.

aim to minimize this expected payout. In order to maximize this minimum, player I will choose $x_1 = 2/3$, which again yields an expected gain of $5/3$.

3

General-sum games

We now turn to the theory of **general-sum games**. Such a game is given by two matrices A and B , whose entries give the payoffs to the two players for each pair of pure strategies that they might play. Usually there is no joint optimal strategy for the players, but the notion of Nash equilibrium remains relevant. These equilibria give the strategies that “rational” players might choose. However, there are often several Nash equilibria, and in choosing one of them, some degree of cooperation between the players may be desirable. Moreover, a pair of strategies based on cooperation might be better for both players than any of the Nash equilibria. We begin with two examples.

3.1 Some examples

Example 3.1.1 (The prisoner’s dilemma). Two suspects are held and questioned by police who ask each of them to confess. The charge is serious, but the police don’t have enough evidence to convict. Separately, each suspect is offered the following plea deal. If he confesses and the other prisoner remains silent, the confessor goes free, and his confession is used to sentence the other prisoner to ten years in prison. If both confess, they will both spend eight years in prison. If both remain silent, the sentence is one year to each for the minor crime that can be proved without additional evidence. The following matrix summarizes the payoffs, where negative numbers represent years in prison.

		prisoner II	
		silent	confess
prisoner I	silent	(-1, -1)	(-10, 0)
	confess	(0, -10)	(-8, -8)




Fig. 3.1. Two prisoners considering whether to confess or remain silent.

If the players are playing this game once, the payoff a player secures by confessing is always greater than the payoff a player will get by remaining silent, no matter how the other player behaves. However, if both follow this reasoning, then both will confess and each of them will be worse off than they would have been had they both remained silent. Unfortunately, to achieve this latter, mutually preferable outcome, each player must suppress his or her natural desire to act selfishly. As we know from real life, this is nontrivial!

The same phenomenon occurs even if the players were to play this same game a fixed number of times. This can be shown by a backwards induction argument. However, as we shall see in §Section ??, if the game is played repeatedly, but ends at a random time, the mutually preferable solution may arise even with selfish play.

Example 3.1.2 (Investing in communication infrastructure). Two firms are interested in setting up infrastructure that will enable them to communicate with each other. Each of the firms decides independently whether to buy high bandwidth equipment (H) or low bandwidth equipment (L). High bandwidth equipment is more expensive than low bandwidth equipment, but more than pays for itself in communication quality as long as both firms employ it. Low bandwidth equipment yields a payoff of 1 to the firm employing it regardless of the equipment employed by the other firm. This leads to the following payoff matrix:

		Firm II	
		high (H)	low (L)
Firm I	high (H)	(2, 2)	(0, 1)
	low (L)	(1, 0)	(1, 1)

What are good strategies for the firms? We begin by considering safety strategies. L is the unique safety strategy for each player, and results in a payoff of 1 to each player. The strategy pair (L, L) is also a pure Nash equilibrium, since given the choice of low bandwidth by the other firm, neither firm has an incentive to switch to high bandwidth. There is another pure Nash equilibrium in this game, (H, H), which yields both players a payoff of 2. Finally, there is a mixed Nash equilibrium in this game, in which both players choose each action with probability 1/2. This also results in an expected payoff of 1 to both players.

This example illustrates one of the new phenomena that arise in general sum games: multiplicity of equilibria with different expected payoffs to the players.

Example 3.1.3 (Driver and parking inspector game). Player I is choosing between parking in a convenient but illegal parking spot (payoff 10 if she's not caught), and parking in a legal but inconvenient spot (payoff 0). If she parks illegally and is caught, she will pay a hefty fine (payoff -90). Player II, the inspector representing the city, needs to decide whether to check for illegal parking. There is a small cost (payoff -1) to inspecting. However, there is a greater cost to the city if player I has parked illegally since that can disrupt traffic (payoff -10). This cost is partially mitigated if the inspector catches the offender (payoff -6).

The resulting payoff matrix is the following:

		Inspector	
		Don't Inspect	Inspect
Driver	Legal	(0, 0)	(0, -1)
	Illegal	(10, -10)	(-90, -6)

In this game, the safety strategy for the driver is to park legally (guaranteeing her a payoff of 0), and the safety strategy for the inspector is to inspect (guaranteeing him/the city a payoff of -6). However, the strategy pair (legal, inspect) is *not* a Nash equilibrium. Indeed, knowing the driver is parking legally, the inspector's best response is not to inspect. It is easy

to check that this game has no Nash equilibrium in which either player uses a pure strategy.

There is, however, a mixed Nash equilibrium. Suppose the strategy pair $(x, 1-x)$ for the driver and $(y, 1-y)$ for the inspector are a Nash equilibrium. Since $0 < y < 1$, both possible actions of the inspector yield the same payoff and thus $-10(1-x) = -x - 6(1-x)$. Similarly, $0 = 10y - 90(1-y)$. These equations yield $x = 0.8$ (the driver parks legally with probability 0.8 and obtains an expected payoff of 0) and $y = 0.9$ (the inspector inspects with probability 0.1 and obtains an expected payoff of -2).

3.2 Nash equilibria

A two-person general-sum game can be represented by a pair of $m \times n$ **payoff matrices** $A = (a_{ij})$ and $B = (b_{ij})$, whose rows are indexed by the m possible actions of player I, and whose columns are indexed by the n possible actions of player II. (In the examples, we represent the payoffs by an $m \times n$ matrix of pairs (a_{ij}, b_{ij}) .) Player I selects an action i and player II selects an action j , each unaware of the other's selection. Their selections are then revealed and player I receives a payoff of a_{ij} and player II a payoff of b_{ij} .

A **mixed strategy** for player I is determined by a vector $(x_1, \dots, x_m)^T$ where x_i represents the probability that player I plays action i and a mixed strategy for player II is determined by a vector $(y_1, \dots, y_n)^T$ where y_j is the probability that player II plays action j . A mixed strategy in which a particular action is played with probability 1 is called a **pure strategy**.

Definition 3.2.1 (Nash equilibrium). A pair of mixed strategy vectors $(\mathbf{x}^*, \mathbf{y}^*)$ with $\mathbf{x}^* \in \Delta_m$ (where $\Delta_m = \{\mathbf{x} \in \mathbb{R}^m : x_i \geq 0, \sum_{i=1}^m x_i = 1\}$), and $\mathbf{y}^* \in \Delta_n$ (where $\Delta_n = \{\mathbf{y} \in \mathbb{R}^n : y_j \geq 0, \sum_{j=1}^n y_j = 1\}$) is a **Nash equilibrium** if no player gains by unilaterally deviating from it. That is,

$$(\mathbf{x}^*)^T A \mathbf{y}^* \geq \mathbf{x}^T A \mathbf{y}^*$$

for all $\mathbf{x} \in \Delta_m$ and

$$(\mathbf{x}^*)^T B \mathbf{y}^* \geq (\mathbf{x}^*)^T B \mathbf{y}$$

for all $\mathbf{y} \in \Delta_n$.

The game is called **symmetric** if $m = n$ and $a_{i,j} = b_{j,i}$ for all $i, j \in \{1, 2, \dots, n\}$. A pair (\mathbf{x}, \mathbf{y}) of strategies is called **symmetric** if $x_i = y_i$ for all $i = 1, \dots, n$.

We will see that there always exists a Nash equilibrium; however, there can be many of them, and they may yield different payoffs to the players.

Thus, Nash equilibria do not have the predictive power in general sum games that safety strategies have in zero-sum games. We discuss in the notes to what extent Nash equilibria are a reasonable model for rational behavior.

Example 3.2.2 (Cheetahs and antelopes). Consider a simple model, where two cheetahs are giving chase to two antelopes, one large and one small. Each cheetah has two possible strategies: chase the large cheetah (L) or chase the small cheetah (S). The cheetahs will catch any antelope they choose, but if they choose the same one, they must share the spoils. Otherwise, the catch is unshared. The large antelope is worth ℓ and the small one is worth s . Here is the payoff matrix:

		cheetah II	
		L	S
cheetah I	L	$(\ell/2, \ell/2)$	(ℓ, s)
	S	(s, ℓ)	$(s/2, s/2)$




Fig. 3.2. Cheetahs deciding whether to chase the large or the small antelope.

If the larger antelope is worth at least twice as much as the smaller ($\ell \geq 2s$), then strategy L dominates strategy S . Hence each cheetah should just chase the larger antelope. If $s < \ell < 2s$, then there are two pure Nash equilibria, (L, S) and (S, L) . These pay off quite well for both cheetahs — but how would two healthy cheetahs agree which should chase the smaller antelope? Therefore it makes sense to look for symmetric mixed equilibria.

If the first cheetah chases the large antelope with probability x , then the expected payoff to the second cheetah by chasing the larger antelope is

$$L(x) = \frac{\ell}{2}x + (1 - x)\ell,$$

and the expected payoff arising from chasing the smaller antelope is

$$S(x) = xs + (1 - x)\frac{s}{2}.$$

These expected payoffs are equal when

$$x = \frac{2\ell - s}{\ell + s}.$$

For any other value of x , the second cheetah would prefer either the pure strategy L or the pure strategy S , and then the first cheetah would do better by simply playing pure strategy S or pure strategy L . But if both cheetahs chase the large antelope with probability

$$x^* = \frac{2\ell - s}{\ell + s},$$

then neither one has an incentive to deviate from this strategy, so this a Nash equilibrium, in fact a symmetric Nash equilibrium.

ADD A GRAPH OF $L(p)$ and $S(p)$

There is a fascinating connection between symmetric mixed Nash equilibria in games such as this and equilibria in biological populations. Consider a population of cheetahs, and suppose a fraction x of them are greedy (i.e., play strategy L). Each time a cheetah plays this game, he plays it against a random cheetah in the population. Then a greedy cheetah obtains an expected payoff of $L(x)$, whereas a non-greedy cheetah obtains an expected payoff of $S(x)$. If $x > x^*$, then $S(x) > L(x)$ and non-greedy cheetahs have an advantage over greedy cheetahs. On the other hand, if $x < x^*$, greedy cheetahs have an advantage. Altogether, the population seems to be pushed by evolution towards the symmetric mixed Nash equilibrium $(x^*, 1 - x^*)$. Indeed, such phenomena have been observed in real biological systems. The related notion of an **evolutionarily stable strategy** is formalized in section 3.5.

Example 3.2.3 (The game of chicken). Two drivers speed head-on toward each other and a collision is bound to occur unless one of them chickens out at the last minute. If both chicken out, everything is OK (we'll say that in this case, they both get a payoff of 1). If one chickens out and the other does not, then it is a great success for the player with iron nerves (payoff = 2) and a great disgrace for the chicken (payoff = -1). If both players have iron nerves, disaster strikes (both incur a large penalty M).

		player II	
		Chicken (C)	Drive (D)
player I	Chicken (C)	(1, 1)	(-1, 2)
	Drive (D)	(2, -1)	(-M, -M)




Fig. 3.3. The game of chicken.

Let's determine its Nash equilibria. First, we see that there are two pure Nash equilibria (C, D) and (D, C) : if one player knows with certainty that the other will drive on (resp. chicken out), that player is better off chickening out (resp. driving on).

To determine the mixed equilibria, suppose that player I plays C with probability x and D with probability $1 - x$. This presents player II with expected payoffs of $x \times 1 + (1 - x) \times (-1) = 2x - 1$ if she plays C , and $x \times 2 + (1 - x) \times (-M) = (M + 2)x - M$ if she plays D . We seek an equilibrium where player II has positive probability on each of C and D . Thus,

$$2x - 1 = (M + 2)x - M.$$

That is, $x = 1 - 1/M$. The payoff for player II is $2x - 1$, which equals $1 - 2/M$.

Remarks:

- (i) Notice that even though both payoff matrices decrease as M increases, the equilibrium payoffs become larger. This contrasts with the situation in zero sum games where decreasing a player's payoff matrix can only lower her expected payoff in equilibrium.
- (ii) The payoff for a player is lower in the symmetric Nash equilibrium than it is in the pure equilibrium where that player plays D and the other plays C. One way for a player to ensure that the higher payoff asymmetric Nash equilibrium is reached is to irrevocably commit to the strategy D, for example, by ripping out the steering wheel and throwing it out of the car. In this way, it becomes impossible for him to chicken out, and if the other player sees this and believes her eyes, then she has no other choice but to chicken out.

In a number of games, making this kind of *binding commitment* pushes the game into a pure Nash equilibrium, and the nature of that equilibrium strongly depends on who managed to commit first. Here, the payoff for the player who did not make the commitment is lower than the payoff in the unique mixed Nash equilibrium, while in some games it is higher (e.g., see Battle of the Sexes in §Section 3.7).

- (iii) An amusing real-life example of commitments arises in a certain narrow two-way street in Jerusalem. Only one car at a time can pass. If two cars headed in opposite directions meet in the street, the driver that can signal to the opponent that he “has time for a face-off” will be able to force the other to back out. Some drivers carry a newspaper with them which they can strategically pull out to signal that they are not in any particular rush.

3.3 General-sum games with more than two players

We now consider general sum games with more than two players and generalize the notion of Nash equilibrium to this setting. Each player i has a set S_i of pure strategies. We are given payoff or *utility* functions $u_i : S_1 \times S_2 \times \cdots \times S_k \rightarrow \mathbb{R}$, for each player i , where $i \in \{1, \dots, k\}$. If player j plays strategy $s_j \in S_j$ for each $j \in \{1, \dots, k\}$, then player i has a payoff or utility of $u_i(s_1, \dots, s_k)$.

Example 3.3.1 (An ecology game). Three firms will either pollute a lake in the following year, or purify it. They pay 1 unit to purify, but it is free to pollute. If two or more pollute, then the water in the lake is useless, and each firm must pay 3 units to obtain the water that they need from

elsewhere. If at most one firm pollutes, then the water is usable, and the firms incur no further costs.

Assuming that firm III purifies, the cost matrix (cost=-payoff) is:

		firm II	
		purify	pollute
firm I	purify	(1,1,1)	(1,0,1)
	pollute	(0,1,1)	(3,3,3+1)

If firm III pollutes, then it is:

		firm II	
		purify	pollute
firm I	purify	(1,1,0)	(3+1,3,3)
	pollute	(3,3+1,3)	(3,3,3)




Fig. 3.4.

To discuss the game, we generalize the notion of Nash equilibrium to games with more players.

Definition 3.3.2. For a vector $\mathbf{s} = (s_1, \dots, s_n)$, we use \mathbf{s}_{-i} to denote the vector obtained by excluding s_i , i.e.,

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

We interchangeably refer to the full vector (s_1, \dots, s_n) as either \mathbf{s} or, slightly abusing notation, (s_i, \mathbf{s}_{-i}) .

Definition 3.3.3. A **pure Nash equilibrium** in a k -player game is a sequence of pure strategies

$$(s_1^*, \dots, s_k^*) \in S_1 \times \dots \times S_k$$

such that for each player $j \in \{1, \dots, k\}$ and each $s_j \in S_j$, we have

$$u_j(s_j^*, \mathbf{s}_{-j}^*) \geq u_j(s_j, \mathbf{s}_{-j}^*).$$

In other words, for each player j , his selected strategy s_j^* is a best response to the selected strategies \mathbf{s}_{-j}^* of the other players.

A **mixed Nash equilibrium** is a sequence of k mixed strategies, with $\mathbf{x}_i^* \in \Delta_{|S_i|}$ the mixed strategy of player i , such that for each player $j \in \{1, \dots, k\}$ and each probability vector $\mathbf{x}_j \in \Delta_{|S_j|}$, we have

$$\bar{u}_j(\mathbf{x}_j^*, \mathbf{x}_{-j}^*) \geq \bar{u}_j(\mathbf{x}_j, \mathbf{x}_{-j}^*).$$

Here,

$$\bar{u}_j(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k) := \sum_{s_1 \in S_1, \dots, s_k \in S_k} \mathbf{x}_1(s_1) \dots \mathbf{x}_k(s_k) u_j(s_1, \dots, s_k),$$

where $\mathbf{x}_i(s)$ is the probability with which player i plays pure strategy s in the mixed strategy \mathbf{x}_i .

Definition 3.3.4. A game is **symmetric** if the players strategies and payoffs are identical, up to relabelling, i.e., for every $i_0, j_0 \in \{1, \dots, k\}$, there is a permutation π of the set $\{1, \dots, k\}$ such that $\pi(i_0) = j_0$ and

$$u_{\pi(i)}(\ell_{\pi(1)}, \dots, \ell_{\pi(k)}) = u_i(\ell_1, \dots, \ell_k).$$

(For this definition to make sense, we require that the strategy sets of the players coincide.)

We will prove the following result in §Section 3.8.

Theorem 3.3.5 (Nash's theorem). *Every finite general sum game has a Nash equilibrium. Moreover, in a symmetric game, there is a symmetric Nash equilibrium.*

For determining Nash equilibria in (small) games, the following lemma (which we have already applied several times for 2-player games) is useful.

Lemma 3.3.6. *Consider a k -player game with \mathbf{x}_i the mixed strategy of*

player i . For each i , let $T_i = \{s \in S_i \mid \mathbf{x}_i(s) > 0\}$. Then $(\mathbf{x}_1, \dots, \mathbf{x}_k)$ is a Nash equilibrium if and only if for each i , there is a constant c_i such that[†]

$$\forall s_i \in T_i \quad u_i(s_i, \mathbf{x}_{-i}) = c_i$$

and

$$\forall s_i \notin T_i \quad u_i(s_i, \mathbf{x}_{-i}) \leq c_i.$$

Exercise:

- Prove Lemma 3.3.6.
- Use Lemma 3.3.6 to derive an exponential time algorithm for finding a Nash equilibrium in two-player general sum games using linear programming (Section 2.7).

Returning to the ecology game, it is easy to check that the pure equilibria consist of all three firms polluting, or one of the three firms polluting, and the remaining two purifying.

Next we consider mixed strategies. Suppose that player i 's strategy is $\mathbf{x}_i = (p_i, 1 - p_i)$ (i.e. i purifies with probability p_i). It follows from lemma 3.3.6 these strategies are a Nash equilibrium with $0 < p_i < 1$ if and only if:

$$u_i(\text{purify}, \mathbf{x}_{-i}) = u_i(\text{pollute}, \mathbf{x}_{-i}).$$

Thus, if player 1 plays a mixed strategy, then

$$\begin{aligned} p_2 p_3 + p_2(1 - p_3) + p_3(1 - p_2) + 4(1 - p_2)(1 - p_3) \\ = 3p_2(1 - p_3) + 3p_3(1 - p_2) + 3(1 - p_2)(1 - p_3), \end{aligned}$$

or, equivalently,

$$1 = 3(p_2 + p_3 - 2p_2 p_3). \quad (3.1)$$

Similarly, if player 2 plays a mixed strategy, then

$$1 = 3(p_1 + p_3 - 2p_1 p_3), \quad (3.2)$$

and if player 3 plays a mixed strategy, then

$$1 = 3(p_1 + p_2 - 2p_1 p_2). \quad (3.3)$$

Subtracting (3.2) from (3.3), we get $0 = 3(p_2 - p_3)(1 - 2p_1)$. This means that if all three firms use mixed strategies, then either $p_2 = p_3$ or $p_1 = 1/2$. In the first case ($p_2 = p_3$), equation (3.1) becomes quadratic in p_2 , with two

[†] The notation (s_i, \mathbf{x}_{-i}) is an abbreviation where we identify the pure strategy s_i with the probability vector $\mathbf{1}_{s_i}$ that assigns s_i probability 1.

solutions $p_2 = p_3 = (3 \pm \sqrt{3})/6$, both in $(0, 1)$. Substituting these solutions into the first equation, yields $p_1 = p_2 = p_3$, resulting in two symmetric mixed equilibria. If, instead of $p_2 = p_3$, we let $p_1 = 1/2$, then the first equation becomes $1 = 3/2$, which is nonsense. This means that there is no asymmetric equilibrium with at least two mixed strategies. It is easy to check that there is no equilibrium with two pure and one mixed strategy. Thus we have found all Nash equilibria: one symmetric and three asymmetric pure equilibria, and two symmetric mixed ones.

3.4 Games with Infinite Strategy Spaces

In some cases, an agent's strategy space S_i is unbounded.

Example 3.4.1 (Tragedy of the commons). Consider a set of k players that each want to send information along a shared channel of maximum capacity 1. Each player decides how much information to send along the channel, measured as a fraction of the capacity. Ideally, a player would like to send as much information as possible. The problem is that the quality of the channel degrades as a larger and larger fraction of it is utilized, and if it is over-utilized, no information gets through. In this setting, each agent's strategy space $S_i = [0, 1]$. The utility function of each player i is

$$u_i(s_i, \mathbf{s}_{-i}) = s_i \left(1 - \sum_{j \neq i} s_j \right),$$

if $\sum_j s_j \leq 1$ and 0 otherwise.

We check that there is a pure Nash equilibrium in this game. Fix a player i and suppose that the other player's select strategies \mathbf{s}_{-i} . Then player i 's best response consists of choosing that $s_i \in [0, 1]$ so that $s_i(1 - \sum_{j \neq i} s_j)$ is maximized, which occurs at

$$s_i = \left(1 - \sum_{j \neq i} s_j \right) / 2. \quad (3.4)$$

To be in Nash equilibrium, (3.4) must hold for all i . The unique solution to this system of equations has $s_i = 1/(k+1)$ for all i .

This is a "tragedy" because the resulting sum of utilities is

$$\sum_{1 \leq i \leq k} u_i(s_i, \mathbf{s}_{-i}) = \frac{k}{(k+1)^2} = O\left(\frac{1}{k}\right).$$

However, if the players acted globally, rather than optimizing just their own

utility, and chose, for example $s_i = 1/2k$, then each player would have utility approximately $1/4k$ (instead of $1/(k+1)^2$), and the sum of utilities would be constant.

Example 3.4.2 (A pricing game). Consider a setting with two sellers selling the same product and three buyers each interested in buying one unit of the product. Seller I can be assured that buyer A will buy the product from her, and seller II can be assured that buyer C will buy the product from him. However, the two sellers compete to sell the product to buyer B. The strategy space for each of the sellers is their choice of price in $[0, 1]$. (We assume neither buyer is willing to spend more than 1 on the product.) Buyer B will buy from the seller with the lower priced offer, unless their prices are the same, in which case he buys from seller I.

Thus, if seller I sets her price at p_1 and seller II sets his price at p_2 , with $p_1 \leq p_2$, then seller I's utility is $2p_1$, and seller II's utility is p_2 , whereas if $p_1 > p_2$, then seller I's utility is p_1 , and seller II's utility is $2p_2$.

In this game, there is no pure Nash equilibrium. To see this, suppose that seller II chooses a price $x_2 > 1/2$. Then seller I's best response is to choose $x_1 = x_2$. But then x_2 is no longer a best response to x_1 . If $x_2 = 1/2$, then player I's best response is either $x_1 = 1/2$ or $x_1 = 1$, but in either case, $x_2 = 1/2$ is not a best response. Finally, we observe that seller II will never set $x_2 < 1/2$, since this ensures a payoff less than 1, whereas a payoff of 1 is always achievable.

There is, however, a symmetric mixed Nash equilibrium. Any pure strategy with $x < 1/2$ is dominated by the strategy $x = 1$, and thus we can restrict attention to mixed strategies supported on $[1/2, 1]$. Suppose that both sellers choose their prices X and Y from distributions F and G supported on all of $[1/2, 1]$. Then the expected payoff to seller II for any price y he might choose is $yF(y) + 2y(1 - F(y)) = y(2 - F(y))$, which must be equal for all y in $[1/2, 1]$. This holds when $F(x) = 2 - 1/x$ in $[1/2, 1]$ (corresponding to density $f(x) = 1/x^2$ on that interval). Setting $G = F$ yields a Nash equilibrium. Note that the continuous distributions ensure the chance of a tie is zero.

Exercise 3.4.3. Consider the pricing game with two sellers and one buyer who buys at the lower price, however prices are required to be strictly positive. Thus, if the prices selected are x and y then payoffs will be $(x, 0)$ if $x \leq y$ and $(0, y)$ if $x > y$. Show that for any $c > 0$, there is a mixed Nash equilibrium that yields expected payoff c for both players.

3.5 Evolutionary game theory

Evolutionary biology is based on the principle that the genetic makeup of an organism determines many of its behaviors and characteristics. These behaviors and characteristics in turn determine how successful that organism is in life and, therefore, at reproducing. Thus, genes that give rise to behaviors and characteristics that promote reproduction tend to increase in frequency in the population.

One major factor in the reproductive success of an organism is how it interacts with other organisms, and this is where evolutionary game theory comes in. Think of these interactions as a series of encounters between random organisms in a population. An organisms' genes determine how it behaves in each of these encounters, and depending on what happens in these encounters, each participant obtains a certain reward. The greater the reward, the greater the reproductive success that organism has.

We model each encounter between two organisms as a game. The type of an organism, which determines how they behave in the game, corresponds to a pure strategy. The rewards from the encounter as a function of the types are the payoffs, and, finally, the population frequencies of each type of organism correspond to mixed strategies in the game. This is because we think of the encounter or game as transpiring between two random members of the overall population.

One of the fundamental questions we then ask is: what population frequencies are stable? The answer we will consider in this section is the notion of an **evolutionary stable strategy (ESS)**. We will see that every ESS in a game is a symmetric mixed Nash equilibrium, but not vice versa.

We begin with an example, a variant of our old friend, the game of Chicken:

3.5.1 *Hawks and Doves*

The game described in Figure 3.5 is a simple model for two behaviors — one bellicose, the other pacifistic — within the population of a single species.

This game has the following payoff matrix:

		player II	
		<i>H</i>	<i>D</i>
player I	<i>H</i>	$(\frac{v}{2} - c, \frac{v}{2} - c)$	$(v, 0)$
	<i>D</i>	$(0, v)$	$(\frac{v}{2}, \frac{v}{2})$

Now imagine a large population, each of whose members are hardwired genetically either as hawks or as doves, and assume that the payoffs in

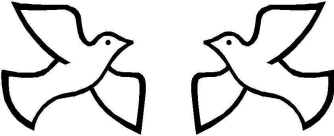


$v/2$		$v/2$
v		0
$v/2 - c$		$v/2 - c$

Fig. 3.5. Two players play this game, for a prize of value $v > 0$. They confront each other, and each chooses (simultaneously) to fight or to flee; these two strategies are called the “hawk” (H) and the “dove” (D) strategies, respectively. If they both choose to fight (two hawks), then each pays a cost c to fight, and the winner (either is equally likely) takes the prize. If a hawk faces a dove, the dove flees, and the hawk takes the prize. If two doves meet, they split the prize equally.

the game translate directly into reproductive success, so that those who do better at this game have more offspring. We will argue that if $(x, 1 - x)$ is a symmetric Nash equilibrium in this game, then these will also be equilibrium proportions in the population.

Let’s see what the Nash equilibria are. If $c < \frac{v}{2}$, the game is a version of Prisoner’s Dilemma and (H, H) is the only equilibrium. When $c > \frac{v}{2}$, there are two pure Nash equilibria: (H, D) and (D, H) ; and since the game is symmetric, there is a symmetric mixed Nash equilibrium. Suppose I plays H with probability x . To be a Nash equilibrium, we need the payoffs for player II to play H and D to be equal:

$$(L) \quad x\left(\frac{v}{2} - c\right) + (1 - x)v = (1 - x)\frac{v}{2} \quad (R). \quad (3.5)$$

For this to be true, we need $x = \frac{v}{2c}$, which by the assumption, is less than one. By symmetry, player II will do the same thing.

Population Dynamics for Hawks and Doves: Now suppose we have

the following dynamics in the population: throughout their lives, random members of the population pair off and play Hawks and Doves; at the end of each generation, members reproduce in numbers proportional to their winnings. Let x denote the fraction of hawks in the population. If the population is large, then by the Law of Large Numbers, the total payoff accumulated by the hawks in the population, properly normalized, will be the expected payoff of a hawk playing against an opponent whose mixed strategy is to play H with probability x and D with probability $(1 - x)$ — and so also will go the proportion of hawks and doves in the next generation.

If $x < \frac{v}{2c}$, then in equation (3.5), (L) $>$ (R) — the expected payoff for a hawk is greater than that for a dove, and so in the next generation, x , the fraction of hawks, will increase.

On the other hand, if $x > \frac{v}{2c}$, then (L) $<$ (R) — the expected payoff for a dove is higher than that of a hawk, and so, in the next generation, x will decrease.

Example 3.5.1 (Sex Ratios). Evolutionary stability can be used to explain sex ratios in nature. In mostly monogamous species, it seems natural that the birth rate of males and females should be roughly equal. But what about sea lions, in which a single male gathers a large harem of females, while many males never reproduce? Game theory helps explain why reproducing at a 1:1 ratio remains stable. To illustrate this, consider the following highly simplified model. Suppose that each harem consists of one male and ten females. If M is the number of males in the population and F the number of females, then the number of “lucky” males, that is, males with a harem, is $M_L = \min(M, F/10)$. Suppose also that each mating pair has b offspring on average. A random male has a harem with probability M_L/M , and if he does, he has $10b$ offspring on average. Thus, the expected number of offspring a random male has is $\mathbb{E}[C_m] = 10bM_L/M = b \min(10, F/M)$. On the other hand, the number of females that belong to a harem is $F_L = \min(F, 10M)$, and thus the expected number of offspring a female has is $\mathbb{E}[C_f] = bF_L/F = b \min(1, 10M/F)$.

If $M < F$, then $\mathbb{E}[C_m] > \mathbb{E}[C_f]$, and individuals with a higher propensity to have male offspring than females will tend to have more grandchildren, resulting in a higher proportion of genes in the population with a propensity for male offspring. In other words, the relative birthrate of males increases. On the other hand, if $M > F$, then $\mathbb{E}[C_m] < \mathbb{E}[C_f]$, and the relative birthrate of females increases. (Of course, when $M = F$, we have $\mathbb{E}[C_m] = \mathbb{E}[C_f]$, and the sex ratio is stable.)

3.5.2 Evolutionarily stable strategies

Consider a symmetric, two-player game with n pure strategies each, and payoff matrices A and B for players I and II, with $A_{i,j} = B_{j,i}$.

We take the point of view that a symmetric mixed strategy in this game corresponds to the proportions of each type within the population.

To motivate the formalism, suppose a population with strategy \mathbf{x} is invaded by a small population of mutants of type \mathbf{z} (that is, playing strategy \mathbf{z}), so the new composition is $\varepsilon\mathbf{z} + (1-\varepsilon)\mathbf{x}$, where ε is small. The new payoffs will be:

$$\varepsilon\mathbf{x}^T A\mathbf{z} + (1-\varepsilon)\mathbf{x}^T A\mathbf{x} \quad (\text{for } \mathbf{x}'\text{s}) \quad (3.6)$$

$$\varepsilon\mathbf{z}^T A\mathbf{z} + (1-\varepsilon)\mathbf{z}^T A\mathbf{x} \quad (\text{for } \mathbf{z}'\text{s}). \quad (3.7)$$

The criteria for \mathbf{x} to be an evolutionary stable strategy will imply that, for small enough ε , the average payoff for \mathbf{x} 's will be strictly greater than that for \mathbf{z} 's, so the invaders will disappear. Formally:

Definition 3.5.2. A mixed strategy \mathbf{x} in Δ_n is an **evolutionarily stable strategy (ESS)** if for any pure “mutant” strategy \mathbf{z} :

- (i) $\mathbf{z}^T A\mathbf{x} \leq \mathbf{x}^T A\mathbf{x}$.
- (ii) if $\mathbf{z}^T A\mathbf{x} = \mathbf{x}^T A\mathbf{x}$, then $\mathbf{z}^T A\mathbf{z} < \mathbf{x}^T A\mathbf{z}$.

Observe that criterion (i) is equivalent to saying that \mathbf{x} is a Nash equilibrium. Thus, if \mathbf{x} is a Nash equilibrium, criterion (i) holds with equality for any \mathbf{z} in the support of \mathbf{x} .

Example 3.5.3 (Hawks and Doves). We will verify that the mixed Nash equilibrium $\mathbf{x} = (\frac{v}{2c}, 1 - \frac{v}{2c})$ (i.e., H is played with probability $\frac{v}{2c}$) is an ESS when $c > \frac{v}{2}$. First, we observe that both pure strategies satisfy constraint (i) with equality, so we check (ii).

- If $\mathbf{z} = (1, 0)$ (“H”) then $\mathbf{z}^T A\mathbf{z} = \frac{v}{2} - c$, which is strictly less than $\mathbf{x}^T A\mathbf{z} = x(\frac{v}{2} - c) + (1-x)0$.
- If $\mathbf{z} = (0, 1)$ (“D”) then $\mathbf{z}^T A\mathbf{z} = \frac{v}{2} < \mathbf{x}^T A\mathbf{z} = xv + (1-x)\frac{v}{2}$.

Thus, the mixed Nash equilibrium for Hawks and Doves is an ESS.

Example 3.5.4 (Rock-Paper-Scissors). The unique Nash equilibrium in Rock-Paper-Scissors, $\mathbf{x} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, is **not** evolutionarily stable. This is because the payoff of \mathbf{x} against any strategy is 0, and the payoff of any pure strategy against itself is also 0, and thus, the expected payoff of \mathbf{x} and

\mathbf{z} will be equal. This means that under appropriate notions of population dynamics, cycling will occur: a population with many Rocks will be taken over by Paper, which in turn will be invaded (bloodily, no doubt) by Scissors, and so forth. These dynamics have been observed in actual populations of organisms — in particular, in a California lizard.

The side-blotched lizard *Uta stansburiana* has three distinct types of male: orange-throats, blue-throats and yellow-striped. The orange-throats are violently aggressive, keep large harems of females and defend large territories. The blue-throats are less aggressive, keep smaller harems and defend small territories. The yellow-striped are very docile and look like receptive females. They do not defend territory or keep harems. Instead, they sneak into another male's territory and secretly copulate with the females. In 1996, B. Sinervo and C. M. Lively published the first article in *Nature* describing the regular succession in the frequencies of different types of males from generation to generation [SL96].

The researchers observed a six-year cycle which started with a domination by the orange-throats. Eventually, the orange-throats have amassed territories and harems large enough so they could no longer be guarded effectively against the sneaky yellow-striped males, who were able to secure a majority of copulations and produce the largest number of offspring. When the yellow-striped have become very common, however, the males of the blue-throated variety got an edge, since they could detect and ward off the yellow-striped, as the blue-throats have smaller territories and fewer females to monitor. So a period when the blue-throats became dominant followed. However, the vigorous orange-throats do comparatively well against blue-throats, since they can challenge them and acquire their harems and territories, thus propagating themselves. In this manner, the population frequencies eventually returned to the original ones, and the cycle began anew.

Example 3.5.5 (Unstable mixed Nash equilibrium). In this game,

		player II	
		A	B
player I	A	(10, 10)	(0, 0)
	B	(0, 0)	(5, 5)

both pure strategies (A, A) and (B, B) are evolutionarily stable, while the mixed Nash equilibrium is not.




Fig. 3.6. The three types of male of the lizard *Uta stansburiana*. Picture courtesy of Barry Sinervo; see <http://bio.research.ucsc.edu/~barrylab>.

Notice that although (B, B) is evolutionarily stable, if a sufficiently large population of A 's invades, then the “stable” population will in fact shift to being entirely composed of A 's. Specifically, if after the A 's invade the new composition is ε fraction A 's and $1 - \varepsilon$ fraction B 's, then using (3.6), the payoffs for each type are

$$\begin{aligned} (1 - \varepsilon)5 & \quad (\text{for } B\text{s}) \\ \varepsilon 10 & \quad (\text{for } A\text{s}). \end{aligned}$$

Thus if $\varepsilon > 1/3$, the payoffs of the A s will be higher and they will “take over”.

Exercise 3.5.6 (Mixed population invasion). Consider the following game:

		player II		
		A	B	C
player I	A	(0, 0)	(6, 2)	(-1, -1)
	B	(2, 6)	(0, 0)	(3, 9)
	C	(-1, -1)	(9, 3)	(0, 0)

Find two mixed Nash equilibria, one supported on $\{A, B\}$, the other supported on $\{B, C\}$. Show they are both ESS, but the $\{A, B\}$ equilibrium is not stable when invaded by an arbitrarily small population composed of half B 's and half C 's.

3.6 Potential games

Consider a set of k players repeatedly playing a finite game. Suppose that in each round, some player who can improve his payoff chooses a best response to the actions of the other players and switches to that action, while the other players repeat their action from the previous round. There are two possibilities for what can happen. The first is that eventually nobody has an improving move, in which case, the set of strategies being played is a Nash equilibrium. The second possibility is that the process cycles.

The most natural way to prove that a Nash equilibrium is reached is to construct a potential function $\psi(\cdot)$ mapping strategy profiles to \mathbb{R} with the property that each time a player improves his payoff, the potential function value increases by a positive amount. Since these improvements cannot continue indefinitely, the process must reach a pure Nash equilibrium. A game for which there exists such a potential function is called a **potential game**.

Formally, consider k -player games, in which player j 's strategy space is the finite set S_j . Let $u_i(s_1, s_2, \dots, s_k)$ denote the payoff to player i when player j plays strategy s_j for each j . In a potential game, there is a function $\psi : S_1 \times \dots \times S_k \rightarrow \mathbb{R}$ such that for each i , $s_i, \tilde{s}_i \in S_i$ and $\mathbf{s}_{-i} \in S_{-i}$

$$u_i(\tilde{s}_i, \mathbf{s}_{-i}) - u_i(s_i, \mathbf{s}_{-i}) = \psi(\tilde{s}_i, \mathbf{s}_{-i}) - \psi(s_i, \mathbf{s}_{-i}). \quad (3.8)$$

We call the function ψ the **potential function** associated with the game.

Claim 3.6.1. *Every potential game has a Nash equilibrium in pure strategies.*

Proof. The set $S_1 \times \dots \times S_k$ is finite so there exists some \mathbf{s} that maximizes $\psi(\mathbf{s})$. Note that for this \mathbf{s} the expression on the right hand side in Equation (3.8) is at most zero for any $i \in \{1, \dots, k\}$ and any choice of \tilde{s}_i . This implies that \mathbf{s} is a Nash equilibrium. \square

Example 3.6.2 (A congestion game). There is a road network with R roads and k drivers, where the j th driver wishes to drive from point s_j to point t_j . Each driver, say the j -th, chooses a path γ_j from s_j to t_j and incurs a cost or latency due to the congestion on the path selected.

This cost is determined as follows. Suppose that the choices made by the k drivers are $\gamma = (\gamma_1, \gamma_2, \dots, \gamma_k)$. This determines the total number of drivers on each road, say r , as

$$n_r(\gamma) = \left| \left\{ j \in \{1, \dots, k\} : \text{driver } j \text{ uses road } r \text{ when he drives on path } \gamma_j \right\} \right|.$$

In addition, there is a real valued cost function C_r for each road such that $C_r(n)$ is the cost incurred by any driver using road r when the total number of drivers using road r is n . The total cost $Cost_i(\gamma)$ experienced by a driver is the sum of the costs on each road the driver uses, i.e. for driver i it is

$$Cost_i(\gamma) = \sum_{r \in \gamma_i} C_r(n_r(\gamma)).$$

(Note that the utility of the driver is then $u_i(\gamma) = -Cost_i(\gamma)$.)

Claim 3.6.3. *The function ψ defined on strategy tuples $\gamma = (\gamma_1, \dots, \gamma_k)$ as*

$$\psi(\gamma) = - \sum_{r=1}^R \sum_{\ell=1}^{n_r(\gamma)} C_r(\ell).$$

is a potential function for this congestion game.

To get some intuition for this potential function and why it satisfies Equation (3.8), imagine adding the players one at a time, and looking at the cost each player incurs at the moment he's added. The sum of these quantities is the potential function value. If we remove the last player, say on path P , and add him back in, say on path P' , then the change in potential is equal to the change in the cost he incurs when he switches from P to P' . Since the potential function value doesn't depend on the order in which the players are added, any player can be viewed as the last player.

Formally, the observation in the previous paragraph is that, for any i ,

$$\psi(\gamma) = \psi(\gamma_{-i}) - Cost_i(\gamma).$$

Thus,

$$\begin{aligned} \psi(\gamma'_i, \gamma_{-i}) - \psi(\gamma_i, \gamma_{-i}) &= -Cost_i(\gamma'_i, \gamma_{-i}) + Cost_i(\gamma_i, \gamma_{-i}) \\ &= u_i(\gamma'_i, \gamma_{-i}) - u_i(\gamma_i, \gamma_{-i}). \end{aligned}$$

An example illustrating this argument is shown in Figure ??.

Example 3.6.4 (Graph Coloring Game). Consider an arbitrary undirected graph $G = (V, E)$. In this game, each vertex $\{v_1, \dots, v_n\} \in V$ is a

player, and their action consists of choosing a color from the set $[n]$. We represent vertex i 's color choice by $s_i \in [n]$ for each i , and, for any color c , define

$n_c(\mathbf{s}) =$ number of vertices with color c when players color according to \mathbf{s} .

The payoff of a vertex v_j (with color s_j) is then

$$u_j(\mathbf{s}) = \begin{cases} n_{s_j}(\mathbf{s}) & \text{if no neighbor of } v_j \text{ has the same color as } v_j \\ 0 & \text{otherwise.} \end{cases}$$

Consider a series of moves in which one player at a time makes a best response move. Then as soon as every player who has an improving move to make has done so, the graph will be **properly colored**, that is, no neighbors will have the same color. This is because a node's payoff is positive if it doesn't share its color with any neighbor and 0 otherwise. Moreover, once the graph is properly colored, it will never become improperly colored by a best response move. Thus, we can restrict attention to strategy profiles \mathbf{s} in which the graph is properly colored.

Lemma 3.6.5. *The graph coloring game has a pure Nash equilibrium.*

Proof. We claim that, restricted to proper colorings, this game is a potential game with potential function

$$\psi(\mathbf{s}) = \sum_{c=1}^n \sum_{\ell=1}^{n_c(\mathbf{s})} \ell,$$

i.e., that for any i , (s_i, \mathbf{s}_{-i}) and $(\tilde{s}_i, \mathbf{s}_{-i})$ that are proper colorings,

$$u_i(\tilde{s}_i, \mathbf{s}_{-i}) - u_i(s_i, \mathbf{s}_{-i}) = \psi(\tilde{s}_i, \mathbf{s}_{-i}) - \psi(s_i, \mathbf{s}_{-i}).$$

This follows from the same line of reasoning as the congestion game example: $\psi(\mathbf{s})$ is obtained by coloring the nodes one at a time, adding in the payoff of the new node relative to the nodes that have already been colored. Thus, for any \mathbf{s} that is a proper coloring, and any player i ,

$$\psi(\mathbf{s}) = \psi(\mathbf{s}_{-i}) + u_i(\mathbf{s}) = \psi(\mathbf{s}_{-i}) + n_{s_i}(\mathbf{s}).$$

The rest of the argument follows as in the previous example. \square

Corollary 3.6.6. *Let $\chi(G)$ be the **chromatic number** of the graph G , that is, the minimum number of colors in any proper coloring of G . Then the graph coloring game has a pure Nash equilibrium with $\chi(G)$ colors.*

Proof. Suppose that \mathbf{s} is the coloring corresponding to $\chi(G)$. Then in a series of single-player best response moves, no player will ever introduce an additional color, and the coloring will remain proper always. In addition, since the game is a potential game, the series of moves will end in a pure Nash equilibrium. Thus, this Nash equilibrium will have $\chi(G)$ colors. \square

3.7 Correlated equilibria

Example 3.7.1 (The battle of the sexes). The wife wants to head to the opera, but the husband yearns instead to spend an evening watching baseball. Neither is satisfied by an evening without the other. In numbers, player I being the wife and II the husband, here is the scenario:

		husband	
		opera	baseball
wife	opera	(4,1)	(0,0)
	baseball	(0,0)	(1,4)

How do we expect a rational couple to work this dilemma out?

In this game there are two pure Nash equilibria: both go the opera or both watch baseball. There is also a mixed Nash equilibrium which yields each player an expected payoff of $4/5$ (when the wife plays $(4/5, 1/5)$ and the husband plays $(1/5, 4/5)$). This mixed equilibrium hardly seems rational: the payoff a player gets is lower than what they would obtain by agreeing to go along with what their spouse wants. How might this couple decide between the two pure Nash equilibria?

One way to do this would be to pick a joint action based on a flip of a single coin. For example, the two players could agree that if the coin lands heads then both go to the opera, otherwise both watch baseball. Observe that even after the coin toss, neither player has an incentive to unilaterally deviate from the agreement.

This idea was introduced in 1974 by Aumann ([Aum87]) and is called a **correlated equilibrium**. To motivate the formal definition, observe that a mixed strategy pair in a two-player general-sum game with action spaces $[m]$ and $[n]$ can be described by a random pair of actions, \mathcal{R} with distribution $\mathbf{x} \in \Delta_m$, and \mathcal{C} with distribution $\mathbf{y} \in \Delta_n$, picked independently by player I and II. Thus,

$$\mathbb{P}[\mathcal{R} = i, \mathcal{C} = j] = x_i y_j.$$

It follows from Lemma 3.3.6 that \mathbf{x}, \mathbf{y} is a Nash equilibrium if and only if

$$\mathbb{P}[\mathcal{R} = i] > 0 \implies \mathbb{E}[a_{i,\mathcal{C}}] \geq \mathbb{E}[a_{\ell,\mathcal{C}}]$$

for all i and ℓ in $[n]$, and

$$\mathbb{P}[\mathcal{C} = j] > 0 \implies \mathbb{E}[b_{\mathcal{R},j}] \geq \mathbb{E}[b_{\mathcal{R},k}].$$

for all j and k in $[m]$.

Definition 3.7.2. A **correlated strategy pair** is a pair of random actions $(\mathcal{R}, \mathcal{C})$ with an arbitrary joint distribution

$$z_{ij} = \mathbb{P}[\mathcal{R} = i, \mathcal{C} = j].$$

The next definition formalizes the idea that, in a correlated equilibrium, if player I knows that the players' actions $(\mathcal{R}, \mathcal{C})$ are picked according to the joint distribution \mathbf{z} and player I is informed only that $\mathcal{R} = i$, then she has no incentive to switch to some other action ℓ .

Definition 3.7.3. A correlated strategy pair in a two-player game with payoff matrices A and B is a **correlated equilibrium** if

$$\mathbb{P}[\mathcal{R} = i] > 0 \implies \mathbb{E}[a_{i,\mathcal{C}} \mid \mathcal{R} = i] \geq \mathbb{E}[a_{\ell,\mathcal{C}} \mid \mathcal{R} = i] \quad (3.9)$$

for all i and ℓ in $[n]$, and

$$\mathbb{P}[\mathcal{C} = j] > 0 \implies \mathbb{E}[b_{\mathcal{R},j} \mid \mathcal{C} = j] \geq \mathbb{E}[b_{\mathcal{R},k} \mid \mathcal{C} = j].$$

for all j and k in $[m]$.

Remark. In terms of the distribution \mathbf{z} , the inequality in condition (3.9) is

$$\sum_j \left(\frac{z_{ij}}{\sum_k z_{ik}} \right) a_{ij} \geq \sum_j \left(\frac{z_{ij}}{\sum_k z_{ik}} \right) a_{\ell j}.$$

Thus, \mathbf{z} is a correlated equilibrium iff for all i and ℓ ,

$$\sum_j z_{ij} a_{ij} \geq \sum_j z_{ij} a_{\ell j},$$

and for all j and k ,

$$\sum_i z_{ij} a_{ij} \geq \sum_i z_{ij} a_{ik}.$$

The next example illustrates a more sophisticated correlated equilibrium that is not simply a mixture of Nash equilibria.

Example 3.7.4. A Game of Chicken:

		player II	
		Chicken (C)	Drive (D)
player I	Chicken (C)	(6, 6)	(2, 7)
	Drive (D)	(7, 2)	(0, 0)

In this game, (C, D) and (D, C) are Nash equilibria with payoffs of $(2, 7)$ and $(7, 2)$ respectively. There is also a mixed Nash equilibrium in which each player plays C with probability $2/3$ and D with probability $1/3$ resulting in an expected payoff of $4\frac{2}{3}$.

The following probability distribution \mathbf{z} is a correlated equilibrium which results in an expected payoff of $4\frac{1}{2}$ to each player, worse than the mixed Nash equilibrium.

		player II	
		Chicken (C)	Drive (D)
player I	Chicken (C)	0	$1/2$
	Drive (D)	$1/2$	0

A more interesting correlated equilibrium that yields a payoff outside the convex hull of the Nash equilibrium payoffs is the following:

		player II	
		Chicken (C)	Drive (D)
player I	Chicken (C)	$1/3$	$1/3$
	Drive (D)	$1/3$	0

For this correlated equilibrium, it is crucial that the row player only know \mathcal{R} and the column player only know \mathcal{C} . Otherwise, in the case that the outcome is (C, C) , both players would have an incentive to deviate (unilaterally).

Thus, to implement a correlated equilibrium, an external mediator is typically needed. Here, the external mediator chooses the strategy pair according to this distribution $((C, D), (D, C), (C, C))$ with probability $\frac{1}{3}$ each, and then discloses to each player which strategy he or she should use (but not the strategy of the opponent). At this point, the players are free to follow or to reject the suggested strategy. We claim that it is in their best interest to follow the mediator's suggestion, and thus this distribution is a correlated equilibrium.

To see this, suppose the mediator tells player I to play D . In this case, she knows that player II was told to play C and player I does best by complying to collect the payoff of 7. She has no incentive to deviate.

On the other hand, if the mediator tells her to play C , she is uncertain about what player II is told, but conditioned on what she is told, she knows

that (C, C) and (C, D) are equally likely. If she follows the mediator's suggestion and plays C , her payoff will be of $6 \times \frac{1}{2} + 2 \times \frac{1}{2} = 4$, while her expected payoff from switching is $7 \times \frac{1}{2} = 3.5$, so the player is better off following the suggestion.

We emphasize that the mixed strategies used ($z_{1,1} = z_{1,2} = z_{2,1} = 1/3$ and $z_{1,4} = 0$) in the correlated equilibrium are dependent, so this is not a Nash equilibrium. Moreover, the expected payoff to player I when both follow the suggestion is $2 \times \frac{1}{3} + 6 \times \frac{1}{3} + 7 \times \frac{1}{3} = 5$. This is better than they could do by following an uncorrelated (or regular) mixed Nash equilibrium.

Surprisingly, finding a correlated equilibrium in large scale problems is computationally easier than finding a Nash equilibrium. In fact, there are no computationally efficient algorithms known for finding Nash equilibria, even in two player games. However, correlated equilibria computation reduces to linear programming (see Exercise ??).

Exercise 3.7.5. Occasionally, two parties resolve a dispute (pick a “winner”) by playing a variant of Rock-Paper-Scissors. In this version, the parties are penalized if there is a delay before a winner is declared; a delay occurs when both players choose the same strategy. The resulting payoff matrix is the following:

		Player II		
		Rock	Paper	Scissors
Player I	Rock	(-1, -1)	(0, 1)	(1, 0)
	Paper	(1, 0)	(-1, -1)	(0, 1)
	Scissors	(0, 1)	(1, 0)	(-1, -1)

Show that this game has a unique Nash equilibrium that is fully mixed, and results in expected payoffs of 0 to both players. Then show that the following probability distribution is a correlated equilibrium in which the players obtain expected payoffs of $1/2$.

		Player II		
		Rock	Paper	Scissors
Player I	Rock	0	1/6	1/6
	Paper	1/6	0	1/6
	Scissors	1/6	1/6	0

3.8 The proof of Nash's theorem

Recall Nash's theorem:

Theorem 3.8.1. *For any general-sum game with $k \geq 2$ players, there exists at least one Nash equilibrium.*

To prove this theorem, we use the following theorem that will be proved in the next section.

Theorem 3.8.2. [*Brouwer's fixed-point theorem*] *If $K \subseteq \mathbb{R}^d$ is closed, convex and bounded, and $T : K \rightarrow K$ is continuous, then there exists $\mathbf{x} \in K$ such that $T(\mathbf{x}) = \mathbf{x}$.*

Proof of Nash's theorem using Brouwer's theorem. Suppose that there are two players and the game is specified by payoff matrices $A_{m \times n}$ and $B_{m \times n}$ for players I and II. Let $K = \Delta_m \times \Delta_n$. We will define a continuous map $T : K \rightarrow K$ that takes a pair of strategies (\mathbf{x}, \mathbf{y}) to a new pair $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ with the following properties:

- (i) $\hat{\mathbf{x}}$ is a better response to \mathbf{y} than \mathbf{x} is, if there is one; otherwise $\hat{\mathbf{x}} = \mathbf{x}$.
- (ii) $\hat{\mathbf{y}}$ is a better response to \mathbf{x} than \mathbf{y} is, if there is one; otherwise, $\hat{\mathbf{y}} = \mathbf{y}$.

A fixed point of T will then be a Nash equilibrium.

Define c_i to be the gain player I obtains by switching from strategy \mathbf{x} to pure strategy i , when playing against \mathbf{y} if this gain is positive, and zero otherwise. Formally, for $\mathbf{x} \in \Delta_m$

$$c_i = c_i(\mathbf{x}, \mathbf{y}) = \max \{ A_{(i)} \mathbf{y} - \mathbf{x}^T A \mathbf{y}, 0 \},$$

where A_i denotes the i^{th} row of the matrix A . Define $\hat{\mathbf{x}} \in \Delta_m$ by

$$\hat{x}_i = \frac{x_i + c_i}{1 + \sum_{k=1}^m c_k},$$

i.e., the weight of each action for player I is increased according to its performance against the mixed strategy \mathbf{y} .

Similarly, define d_j to be the gain player II obtains by switching from strategy \mathbf{y} to pure strategy j when playing against \mathbf{x} , if positive. Formally,

$$d_j = d_j(\mathbf{x}, \mathbf{y}) = \max \{ \mathbf{x}^T B^{(j)} - \mathbf{x}^T B \mathbf{y}, 0 \},$$

where $B^{(j)}$ denotes the j^{th} column of B , and define $\hat{\mathbf{y}} \in \Delta_n$ by

$$\hat{y}_j = \frac{y_j + d_j}{1 + \sum_{k=1}^n d_k}.$$

Finally, define $T(\mathbf{x}, \mathbf{y}) = (\hat{\mathbf{x}}, \hat{\mathbf{y}})$.

We now prove that property (i) holds for this mapping. If $c_i = 0$ for all i , (i.e. $\mathbf{x}^T \mathbf{A} \mathbf{y} \geq A_i \mathbf{y}$), then $\hat{\mathbf{x}} = \mathbf{x}$ is a best response to \mathbf{y} . Otherwise, if there is a better response to \mathbf{y} than \mathbf{x} , then there must be some $c_\ell > 0$. We need to show that

$$\sum_{i=1}^m \hat{x}_i A_i \mathbf{y} > \mathbf{x}^T \mathbf{A} \mathbf{y}. \quad (3.10)$$

Multiplying both sides by $1 + S$ where $S = \sum_k c_k$, this is equivalent to

$$\sum_{i=1}^m (x_i + c_i) A_i \mathbf{y} > (1 + S) \mathbf{x}^T \mathbf{A} \mathbf{y},$$

which holds since

$$\sum_{i=1}^m \frac{c_i}{S} A_i \mathbf{y} > \mathbf{x}^T \mathbf{A} \mathbf{y}.$$

Similarly property (ii) is satisfied.

Finally, we observe that K is convex, closed and bounded, and that T is continuous, since c_i and d_j are. Thus, an application of Brouwer's theorem shows that there exists $(\mathbf{x}, \mathbf{y}) \in K$ for which $T(\mathbf{x}, \mathbf{y}) = (\mathbf{x}, \mathbf{y})$; by properties (i) and (ii), (\mathbf{x}, \mathbf{y}) is a Nash equilibrium.

For $k > 2$ players, we define for each player j and pure strategy ℓ of that player, the quantity $c_\ell^{(j)}$ which is the gain player j gets by switching from their current strategy $\mathbf{x}^{(j)}$ to pure strategy ℓ , if positive, given the current strategies of all the other players. The rest of the argument follows as before. \square

We also stated that in a symmetric game, there is always a **symmetric Nash equilibrium**. This also follows from the above proof, by noting that the map T , defined from the k -fold product $\Delta_n \times \cdots \times \Delta_n$ to itself, can be restricted to the diagonal

$$D = \{(\mathbf{x}, \dots, \mathbf{x}) \in \Delta_n^k : \mathbf{x} \in \Delta_n\}.$$

The image of D under T is again in D , because, in a symmetric game, $c_i^{(1)}(\mathbf{x}, \dots, \mathbf{x}) = \cdots = c_i^{(k)}(\mathbf{x}, \dots, \mathbf{x})$ for all pure strategies i and $\mathbf{x} \in \Delta_n$. Then, Brouwer's fixed-point theorem gives us a fixed point within D , which is a symmetric Nash equilibrium.

3.9 Fixed-point theorems*

Brouwer's theorem is straightforward in one dimension $d = 1$. Given $T : [a, b] \rightarrow [a, b]$, define $f(x) = T(x) - x$. Clearly, $f(a) \geq 0$, while $f(b) \leq 0$. By the intermediate value theorem, there is $x \in [a, b]$ for which $f(x) = 0$, so $T(x) = x$.

In higher dimensions, Brouwer's theorem is rather subtle; in particular, there is no generally applicable recipe to find or approximate a fixed point, and there may be many fixed points. Thus, before we turn to a proof of Theorem 3.8.2, we discuss some easier fixed point theorems, where iteration of the mapping from any starting point converges to the fixed point.

3.9.1 Easier fixed-point theorems

Banach's fixed-point theorem applies when the mapping T contracts distances, as in the following figure.




Fig. 3.7. Under the transformation T a square is mapped to a smaller square, rotated with respect to the original. When iterated repeatedly, the map produces a sequence of nested squares. If we were to continue this process indefinitely, a single point (fixed by T) would emerge.

Theorem 3.9.1 (Banach's fixed-point theorem). *Let K be a complete metric space. Suppose that $T : K \rightarrow K$ satisfies $d(T\mathbf{x}, T\mathbf{y}) \leq \lambda d(\mathbf{x}, \mathbf{y})$ for all $\mathbf{x}, \mathbf{y} \in K$, with $0 < \lambda < 1$ fixed. Then T has a unique fixed point $\mathbf{z} \in K$. Moreover, for any $\mathbf{x} \in K$, we have*

$$d(T^n \mathbf{x}, \mathbf{z}) \leq \frac{d(\mathbf{x}, T\mathbf{x}) \lambda^n}{1 - \lambda}.$$

Remark. Recall that a metric space is **complete** if each Cauchy sequence therein converges to a point in the space. For example, any closed subset of \mathbb{R}^n endowed with the Euclidean metric is complete. See [Rud64] for a discussion of general metric spaces.

Proof. Uniqueness of the fixed point: if $T\mathbf{x} = \mathbf{x}$ and $T\mathbf{y} = \mathbf{y}$, then

$$d(\mathbf{x}, \mathbf{y}) = d(T\mathbf{x}, T\mathbf{y}) \leq \lambda d(\mathbf{x}, \mathbf{y}).$$

Thus, $d(\mathbf{x}, \mathbf{y}) = 0$, so $\mathbf{x} = \mathbf{y}$.

As for existence, given any $\mathbf{x} \in K$, we define $\mathbf{x}_n = T\mathbf{x}_{n-1}$ for each $n \geq 1$, setting $\mathbf{x}_0 = \mathbf{x}$. Set $a = d(\mathbf{x}_0, \mathbf{x}_1)$, and note that $d(\mathbf{x}_n, \mathbf{x}_{n+1}) \leq \lambda^n a$. If $k > n$, then by triangle inequality,

$$\begin{aligned} d(\mathbf{x}_n, \mathbf{x}_k) &\leq d(\mathbf{x}_n, \mathbf{x}_{n+1}) + \cdots + d(\mathbf{x}_{k-1}, \mathbf{x}_k) \\ &\leq a(\lambda^n + \cdots + \lambda^{k-1}) \leq \frac{a\lambda^n}{1-\lambda}. \end{aligned} \quad (3.11)$$

This implies that $\{\mathbf{x}_n : n \in \mathbb{N}\}$ is a Cauchy sequence. The metric space K is complete, whence $\mathbf{x}_n \rightarrow \mathbf{z}$ as $n \rightarrow \infty$. Note that

$$d(\mathbf{z}, T\mathbf{z}) \leq d(\mathbf{z}, \mathbf{x}_n) + d(\mathbf{x}_n, \mathbf{x}_{n+1}) + d(\mathbf{x}_{n+1}, T\mathbf{z}) \leq (1+\lambda)d(\mathbf{z}, \mathbf{x}_n) + \lambda^n a \rightarrow 0$$

as $n \rightarrow \infty$. Hence, $d(T\mathbf{z}, \mathbf{z}) = 0$, and $T\mathbf{z} = \mathbf{z}$.

Thus, letting $k \rightarrow \infty$ in (3.11) yields

$$d(T^n \mathbf{x}, \mathbf{z}) = d(\mathbf{x}_n, \mathbf{z}) \leq \frac{a\lambda^n}{1-\lambda}.$$

□

It is not sufficient, however, for distances to decrease in order for there to be a fixed point, as the following example shows.

Example 3.9.2 (A map that decreases distances but has no fixed points). Consider the map $T : \mathbb{R} \rightarrow \mathbb{R}$ given by

$$T(x) = x + \frac{1}{1 + \exp(x)}.$$

Note that, if $x < y$, then

$$T(x) - x = \frac{1}{1 + \exp(x)} > \frac{1}{1 + \exp(y)} = T(y) - y,$$

implying that $T(y) - T(x) < y - x$. Note also that

$$T'(x) = 1 - \frac{\exp(x)}{(1 + \exp(x))^2} > 0,$$

so that $T(y) - T(x) > 0$. Thus, T decreases distances, but it has no fixed points. This is not a counterexample to Banach's fixed-point theorem, however, because there does not exist any $\lambda \in (0, 1)$ for which $|T(x) - T(y)| < \lambda|x - y|$ for all $x, y \in \mathbb{R}$.

This requirement can sometimes be relaxed, in particular for compact metric spaces.

Remark. Recall that a metric space is **compact** if each sequence therein has a subsequence that converges to a point in the space. A subset of the Euclidean space \mathbb{R}^d is compact if and only if it is closed and bounded. See [Rud64].

Theorem 3.9.3 (Compact fixed-point theorem). *If K is a compact metric space and $T : K \rightarrow K$ satisfies $d(T(\mathbf{x}), T(\mathbf{y})) < d(\mathbf{x}, \mathbf{y})$ for all $\mathbf{x} \neq \mathbf{y} \in K$, then T has a fixed point $\mathbf{z} \in K$. Moreover, for any $x \in K$, we have $T^n(x) \rightarrow z$.*

Proof. Let $f : K \rightarrow \mathbb{R}$ be given by $f(\mathbf{x}) = d(\mathbf{x}, T\mathbf{x})$. We first show that f is continuous. By triangle inequality we have:

$$d(\mathbf{x}, T\mathbf{x}) \leq d(\mathbf{x}, \mathbf{y}) + d(\mathbf{y}, T\mathbf{y}) + d(T\mathbf{y}, T\mathbf{x}),$$

so

$$f(\mathbf{x}) - f(\mathbf{y}) \leq d(\mathbf{x}, \mathbf{y}) + d(T\mathbf{y}, T\mathbf{x}) \leq 2d(\mathbf{x}, \mathbf{y}).$$

By symmetry, we also have: $f(\mathbf{y}) - f(\mathbf{x}) \leq 2d(\mathbf{x}, \mathbf{y})$ and hence

$$|f(\mathbf{x}) - f(\mathbf{y})| \leq 2d(\mathbf{x}, \mathbf{y}),$$

which implies that f is continuous.

Since f is a continuous function and K is compact, there exists $\mathbf{z} \in K$ such that

$$f(\mathbf{z}) = \min_{\mathbf{x} \in K} f(\mathbf{x}). \quad (3.12)$$

If $T\mathbf{z} \neq \mathbf{z}$, then $f(T\mathbf{z}) = d(T\mathbf{z}, T^2\mathbf{z}) < d(\mathbf{z}, T\mathbf{z}) = f(\mathbf{z})$, and we have a contradiction to the minimizing property (3.12) of \mathbf{z} . This implies that $T\mathbf{z} = \mathbf{z}$.

Finally, we observe that iteration converges from any starting point x . Let $x_n = T^n x$, and suppose that x_n does not converge to z . Then for some $\epsilon > 0$, the set $S = \{n | d(x_n, z) \geq \epsilon\}$ is infinite. Let $\{n_k\} \subset S$ be an increasing sequence such that $y_k = x_{n_k} \rightarrow y \neq z$. Now

$$d(Ty_k, z) \rightarrow d(Ty, z) < d(y, z). \quad (3.13)$$

But $T^{n_{k+1}-n_k-1}(Ty_k) = y_{k+1}$, so

$$d(Ty_k, z) \geq d(y_{k+1}, z) \rightarrow d(y, z)$$

contradicting (3.13). □

Exercise 3.9.4. Show that the convergence in the compact fixed point theorem can be arbitrarily slow by showing that for any decreasing sequence

$\{a_n\} \downarrow_{n \geq 0} 0$, there is a distance decreasing $T : [0, a_0] \rightarrow [0, a_0]$ such that $T(0) = 0$ and $d(T^n a_0, 0) \geq a_n$.

3.9.2 Sperner's lemma

In this section, we state and prove a combinatorial lemma that is key to proving Brouwer's fixed-point theorem.

Definition 3.9.5 (Simplex). An n -simplex $\Delta(v_0, v_1, \dots, v_n)$ is the convex hull of a set of $n + 1$ points $v_0, v_1, \dots, v_n \in \mathbb{R}^d$ that are affinely independent, i.e. the n vectors $v_i - v_0$, for $1 \leq i \leq n$, are linearly independent.

Definition 3.9.6 (Face). A k -face of an n -simplex $\Delta(v_0, v_1, \dots, v_n)$ is the convex hull of any $k + 1$ of the points v_0, v_1, \dots, v_n .

Exercise 3.9.7. (1) Show that $n + 1$ points $v_0, v_1, \dots, v_n \in \mathbb{R}^d$ are affinely independent if and only if for every non-zero vector $(\alpha_0, \dots, \alpha_n)$ for which $\sum_{0 \leq i \leq n} \alpha_i = 0$, it must be that $\sum_{0 \leq i \leq n} \alpha_i v_i \neq 0$. Thus, affine independence is a symmetric notion.

(2) Show that a k -face of an n -simplex is a k -simplex.

Definition 3.9.8 (Subdivision of a simplex). A subdivision of a simplex $\Delta(v_0, v_1, \dots, v_n)$ is a collection Γ of n -simplices such that any two simplices in Γ are disjoint or their intersection is a face of both.

Remark. Call an $n - 1$ -face of $\Delta_1 \in \Gamma$ an **outer face** if it lies on an $n - 1$ -face of $\Delta(v_0, v_1, \dots, v_n)$; otherwise, call it an **inner face**. It follows from the definition of subdivision that each inner face of $\Delta_1 \in \Gamma$ is an $n - 1$ -face of exactly one other simplex in Γ . Moreover, if F is an $n - 1$ -face of $\Delta(v_0, v_1, \dots, v_n)$ then

$$\Gamma(F) := \{\Delta_1 \cap F\}_{\Delta_1 \in \Gamma}$$

is a subdivision of F .

Lemma 3.9.9. For any simplex $\Delta(v_0, v_1, \dots, v_n)$ and $\epsilon > 0$, there is a subdivision Γ such that all simplices in Γ have diameter less than ϵ .

Proof. We will use the **barycentric subdivision**

$$\Gamma_1 = \{\Delta_\pi \mid \pi \text{ a permutation of } \{0, \dots, n\}\},$$

where

$$\Delta_\pi = \left\{ \sum_{0 \leq i \leq n} \alpha_i v_i \mid \alpha_{\pi(0)} \geq \dots \geq \alpha_{\pi(n)} \geq 0 \text{ and } \sum_{0 \leq i \leq n} \alpha_i = 1 \right\}.$$

Let $w_i = v_{\pi(i)}$. Then the vertices of Δ_π are

$$w_0, \frac{w_0 + w_1}{2}, \frac{w_0 + w_1 + w_2}{3}, \dots, \frac{1}{n+1} \sum_{i=0}^n w_i.$$

Finally, the diameter of each simplex in Γ_1 is the maximum distance between any two vertices of the sub-simplex.

See Exercise 3.9.10 below for the verification that this subdivision has the required intersection property.

$$\begin{aligned} \left| \frac{1}{k} \sum_{i=0}^{k-1} w_i - \frac{1}{r} \sum_{j=0}^{r-1} w_j \right| &= \frac{1}{kr} \left| \sum_{i=0}^{k-1} \sum_{j=0}^{r-1} (w_i - w_j) \right| \\ &\leq \frac{kr - r}{kr} D \\ &= \left(\frac{k-1}{k} \right) D, \end{aligned}$$

where D is the diameter of $\Delta(v_0, \dots, v_n)$.

Iterating the barycentric subdivision m times yields a subdivision Γ_m in which the maximum diameter of any simplex is at most $\left(\frac{n}{n+1}\right)^m D$. \square


Exercise 3.9.10. (1) Verify that Δ_π has one outer face determined by the equation $\alpha_{\pi(n)} = 0$ and n inner faces determined by the equations $\alpha_{\pi(k)} = \alpha_{\pi(k+1)}$ for $0 \leq k \leq n-1$. (2) Verify that Γ_1 is indeed a subdivision. (3) Verify that for any $n-1$ -face F of $\Delta(v_0, v_1, \dots, v_n)$, the subdivision $\Gamma_1(F)$ is the barycentric subdivision of F .

Definition 3.9.11 (Proper Labeling of a Simplex). A labeling ℓ of the vertices of an n -simplex $\Delta(v_0, v_1, \dots, v_n)$ is proper if $\ell(v_0), \ell(v_1), \dots, \ell(v_n)$ are all different.

Definition 3.9.12 (Sperner Labeling of a Subdivision). A Sperner Labeling ℓ of the vertices in a subdivision Γ of an n -simplex $\Delta(v_0, v_1, \dots, v_n)$ is a labeling in which

- $\Delta(v_0, v_1, \dots, v_n)$ is properly labeled,
- All vertices in Γ are assigned labels in $\{\ell(v_0), \ell(v_1), \dots, \ell(v_n)\}$, and
- The labeling restricted to each face of $\Delta(v_0, \dots, v_n)$ is a Sperner labeling there.

Lemma 3.9.13 (Sperner). Let ℓ be a Sperner labeling of the vertices in Γ ,

Fig. 3.8. Sperner's lemma when $d = 2$.

where Γ is a subdivision of the n -simplex $\Delta(v_0, v_1, \dots, v_n)$. Then the number of properly labeled simplices in Γ is odd.

Proof. We prove the lemma by induction on n . For $n = 1$, this is obvious: In a string of bits that starts with 0 (the label of v_0) and ends with 1 (the label of v_1), the number of bit flips is odd.

For $n = 2$, the simplex Δ is a triangle, and the subdivision is a triangulation of the triangle. We think of the three labels as colors: red (R), blue (B) and yellow (Y). (See Figure 3.8.) We say a 1-face, i.e. an edge, in this triangulation is *good* if its two endpoints are colored red and blue. By the inductive hypothesis, on the red/blue side of Δ , there are an odd number of good edges.

Construct a graph with a node for each triangle in the subdivision (call these inner nodes), and a node for each good edge on the red/blue side of Δ (call these outer nodes). Two inner nodes are adjacent if the corresponding triangles share a red/blue edge. An outer node and an inner node are adjacent if the corresponding outer red/blue edge is one of the sides of the inner triangle. Observe that each outer node has degree 1, and each inner node either has degree 2, if the corresponding triangle has vertex labels RBB or RRB, degree 1, if it is properly labeled RGB, and degree 0 otherwise. Thus, the graph consists of a collection of isolated nodes, paths, and cycles. Since there are an odd number of outer nodes, an odd number of them are endpoints of a path whose other endpoint is an inner node, i.e. properly labeled.

The proof of the previous paragraph can be generalized to higher dimensions. (See exercise ??) Here we give a slightly different proof based on

direct counting. For $n \geq 2$, consider a Sperner labeling of Γ . Call an $(n-1)$ face *good* if its vertex labels are $\ell(v_0), \dots, \ell(v_{n-1})$.

Let $g = \sharp$ of good inner faces; let $g_\partial = \sharp$ of good outer faces on $\Delta(v_0, \dots, v_{n-1})$, and let $N_j = \sharp$ of simplices in Γ with labels $\{\ell(v_i)\}_{i=0}^{n-1}$ and $\ell(v_j)$. Counting pairs

(simplex in Γ , good face of that simplex),

by Remark 3.9.2 we obtain

$$2 \sum_{j=0}^{n-1} N_j + N_n = 2g + g_\partial.$$

Since g_∂ is odd by the inductive hypothesis, so is N_n . \square

Exercise 3.9.14. Extend the proof above for $n = 2$ to give an alternative proof of the induction step.

Solution: By hypothesis, each $n-1$ -face of $\Delta(v_0, \dots, v_{n-1})$ has an odd number of properly labeled simplices. Call an $(n-1)$ face *good* if its vertex labels are $\ell(v_0), \dots, \ell(v_{n-1})$. Define a graph, with a node for each simplex in Γ (call these inner nodes), and a node for each good outer face (call these outer nodes). Two nodes in the graph are adjacent if the corresponding simplices share a good face. Observe that every outer node has degree 1, and each inner node either has degree 2 (if the corresponding simplex has vertices with labels $\ell(v_0), \ell(v_1), \dots, \ell(v_{n-1}), \ell(v_i)$ for some i with $0 \leq i \leq n-1$), degree 1 (if the corresponding simplex is properly labeled), or degree 0. Thus, the graph consists of a collection of cycles and paths, where the endpoints of the paths are either outer nodes or properly labeled inner nodes. Since the number of degree one nodes is even, and the number of outer nodes is odd, the number of properly labeled simplices in Γ must be odd.

3.9.3 Brouwer's fixed-point theorem

Definition 3.9.15. A set $S \subseteq \mathbb{R}^d$ has the **fixed point property** (abbreviated **f.p.p.**) if for any continuous function $T : S \rightarrow S$, there exists $\mathbf{x} \in S$ such that $T(\mathbf{x}) = \mathbf{x}$.

Brouwer's Theorem asserts that every closed, bounded, convex set $K \subset \mathbb{R}^d$ has the f.p.p. Each of the hypotheses on K in the theorem is needed, as the following examples show:

- (i) $K = \mathbb{R}$ (closed, convex, not bounded) with $T(x) = x + 1$.
- (ii) $K = (0, 1)$ (bounded, convex, not closed) with $T(x) = x/2$.

- (iii) $K = \{x \in \mathbb{R} : |x| \in [1, 2]\}$ (bounded, closed, not convex) with $T(x) = -x$.

Theorem 3.9.16. *Brouwer's fixed-point theorem for the simplex*
 The standard n -simplex $\Delta = \{\mathbf{x} \mid \sum_{i=0}^n x_i = 1, \forall i \ x_i \geq 0\}$ has the fixed point property.

Proof. Let Γ be a subdivision of Δ with maximal diameter ϵ and let $T(\mathbf{x}) = (T_0(\mathbf{x}), \dots, T_n(\mathbf{x}))$. For any vertex \mathbf{x} of Γ , let

$$\ell(\mathbf{x}) = \min\{i : T_i(\mathbf{x}) < x_i\}.$$

(Note that since $\sum_{i=0}^n x_i = 1$ and $\sum_{i=0}^n T_i(\mathbf{x}) = 1$, if there is no i with $T_i(\mathbf{x}) < x_i$, then \mathbf{x} is a fixed point.)

By Sperner's Lemma, there is a properly labeled simplex Δ_1 in Γ , and this can already be used to produce an approximate fixed point of T ; see the remark below.

To get a fixed point, find, for each k , a simplex with vertices $\{z^i(k)\}_{i=0}^n$ in Δ and diameter at most $\frac{1}{k}$ satisfying

$$T_i(z^i(k)) < z_i^i(k) \text{ for all } i \in [0, n]. \quad (3.14)$$

Find a convergent subsequence $\mathbf{z}^0(k_j) \rightarrow \mathbf{z}$ and observe that $\mathbf{z}^i(k_j) \rightarrow \mathbf{z}$ for all i . Thus, $T_i(\mathbf{z}) \leq z_i$ for all i , so $T(\mathbf{z}) = \mathbf{z}$. □

Remark. Let Δ_1 be a properly labeled simplex of diameter at most ϵ as in the proof above. Denote by $\mathbf{z}^0, \mathbf{z}^1, \dots, \mathbf{z}^n$ the vertices of Δ_1 , where $\ell(\mathbf{z}^i) = i$.

Then

$$T_i(\mathbf{z}^0) \leq T_i(\mathbf{z}^i) + \omega(\epsilon) < z_i^i + \omega(\epsilon) \leq z_i^0 + \epsilon + \omega(\epsilon),$$

where $\omega(\epsilon) = \max_{|\mathbf{x}-\mathbf{y}| \leq \epsilon} |T(\mathbf{x}) - T(\mathbf{y})|$. On the other hand,

$$T_i(\mathbf{z}^0) = 1 - \sum_{j \neq i} T_j(\mathbf{z}^0) \geq 1 - \sum_{j \neq i} (z_j^0 + \epsilon + \omega(\epsilon)) = z_i^0 - n(\epsilon + \omega(\epsilon)).$$

Thus,

$$|T(\mathbf{z}^0) - \mathbf{z}^0| \leq n(n+1)(\epsilon + \omega(\epsilon)),$$

so \mathbf{z}^0 is an approximate fixed point.

Definition 3.9.17. Let $S \subseteq \mathbb{R}^d$ and $\tilde{S} \subseteq \mathbb{R}^n$. A **homeomorphism** $h : S \rightarrow \tilde{S}$ is a one-to-one continuous map with a continuous inverse.

Definition 3.9.18. Let $S \subseteq A \subseteq \mathbb{R}^d$. A **retraction** $g : A \rightarrow S$ is a continuous map where g restricted to S is the identity map.

Lemma 3.9.19. Let $A \subseteq S \subseteq \mathbb{R}^d$ and $\tilde{S} \subseteq \mathbb{R}^n$.

(i) If S has the f.p.p. and $h : S \rightarrow \tilde{S}$ is a homeomorphism, then \tilde{S} has the f.p.p.

(ii) If $g : A \rightarrow S$ is a retraction and A has the f.p.p., then S has the f.p.p.

Proof. (i): Given $T : \tilde{S} \rightarrow \tilde{S}$ continuous, let $\mathbf{x} \in S$ be a fixed point of $h^{-1} \circ T \circ h : S \rightarrow S$. Then $h(\mathbf{x})$ is a fixed point of T .

(ii): Given $T : S \rightarrow S$, any fixed point of $T \circ g : A \rightarrow S$ is a fixed point of T . \square

Lemma 3.9.20. For $K \subset \mathbb{R}^d$ closed and convex, the nearest point map $\Psi : \mathbb{R}^d \rightarrow K$ where

$$\|\mathbf{x} - \Psi(\mathbf{x})\| = d(\mathbf{x}, K) := \min_{y \in K} \|\mathbf{x} - \mathbf{y}\|$$

is uniquely defined and continuous.

Proof. For uniqueness, suppose that $\|\mathbf{x} - \mathbf{y}\| = \|\mathbf{x} - \mathbf{z}\| = d(\mathbf{x}, K)$ with $\mathbf{y}, \mathbf{z} \in K$. Assume by translation that $\mathbf{x} = 0$. Then

$$d(0, K)^2 + \frac{\|\mathbf{y} - \mathbf{z}\|^2}{2} \leq \frac{\|\mathbf{y} + \mathbf{z}\|^2}{2} + \frac{\|\mathbf{y} - \mathbf{z}\|^2}{2} = \frac{\|\mathbf{y}\|^2 + \|\mathbf{z}\|^2}{2} = d(0, K)^2,$$

so $\mathbf{y} = \mathbf{z}$.

First proof of continuity: Suppose $\mathbf{x}_k \rightarrow \mathbf{x}$, but $\mathbf{y}_k := \Psi(\mathbf{x}_k) \not\rightarrow \Psi(\mathbf{x})$. Then

$$\{k : \|\mathbf{y}_k - \Psi(\mathbf{x})\| \geq \epsilon\}$$

is infinite for some $\epsilon > 0$. Passing to a subsequence, we have $\mathbf{y}_{k(j)} \rightarrow \mathbf{y} \in K$ with $\|\mathbf{y} - \Psi(\mathbf{x})\| \geq \epsilon$. Finally,

$$\|\mathbf{x} - \mathbf{y}\| = \lim_j \|\mathbf{x}_{k(j)} - \mathbf{y}_{k(j)}\| = \lim_j d(\mathbf{x}_{k(j)}, K) = d(\mathbf{x}, K),$$

contradicting the uniqueness of $\Psi(\mathbf{x})$.

Second proof of continuity: Let $\Psi(\mathbf{x}) = \mathbf{y}$ and $\Psi(\mathbf{x} + \mathbf{u}) = \mathbf{y} + \mathbf{v}$. PICTURE!!! We show that that $\|\mathbf{v}\| \leq \|\mathbf{u}\|$. We know from (??) in the proof of the separating hyperplane theorem that

$$\mathbf{v}^T(\mathbf{y} - \mathbf{x}) \geq 0$$

and

$$\mathbf{v}^T(\mathbf{x} + \mathbf{u} - \mathbf{y} - \mathbf{v}) \geq 0.$$

Adding these gives $\mathbf{v}^T(\mathbf{u} - \mathbf{v}) \geq 0$, so

$$\|\mathbf{v}\|^2 = \mathbf{v}^T \mathbf{v} \leq \mathbf{v}^T \mathbf{u} \leq \|\mathbf{v}\| \cdot \|\mathbf{u}\|$$

by Cauchy-Schwarz. Thus $\|\mathbf{v}\| \leq \|\mathbf{u}\|$. □

Proof of Brouwer's theorem. Let $K \subset \mathbb{R}^d$ be compact and convex. There is a simplex Δ_0 that contains K . Clearly Δ_0 is homeomorphic to a standard simplex, so it has the f.p.p. by Lemma 3.9.19(i). Then by Lemma 3.9.20, the nearest point map $\Psi : \Delta_0 \rightarrow K$ is a retraction. Thus, Lemma 3.9.19(ii) implies that K has the f.p.p. □

The next corollary follows immediately from Brouwer's Theorem, but is perhaps more intuitively obvious.

Corollary 3.9.21 (No-Retraction Theorem). *Let $\overline{B} = \overline{B(0,1)}$ be the closed ball in \mathbb{R}^d . There is no retraction from \overline{B} to its boundary ∂B .*

Remark. Exercise 3.9.23 below shows that all truly d -dimensional compact, convex sets are homeomorphic to each other. This yields another proof of Brouwer's theorem from the special case of the simplex which avoids retractions.

Exercise 3.9.22. Show that any d -simplex in \mathbb{R}^d contains a ball.

Solution: The d -simplex Δ_0 with vertices the origin and the standard basis $\mathbf{e}_1, \dots, \mathbf{e}_d$ in \mathbb{R}^d , contains the ball $B(\mathbf{y}, \frac{1}{2d})$ where $\mathbf{y} := \frac{1}{2d}(\mathbf{e}_1 + \dots + \mathbf{e}_d)$. Given an arbitrary d -simplex Δ , by translation we may assume its vertices are $0, \mathbf{v}_1, \dots, \mathbf{v}_d$. Let A be the square matrix with columns \mathbf{v}_i for $i \leq d$. Since these columns are linearly independent, A is invertible. Then Δ contains $B(A\mathbf{y}, \varepsilon)$ where $\varepsilon := \min\{\|Ax\| \text{ such that } \|x\| = 1\} > 0$.

Exercise 3.9.23. Let $K \subset \mathbb{R}^d$ be a compact convex set which contains a d -simplex. Show that K is homeomorphic to a closed ball.

Suggested steps:

- (i) Show that K contains a d -simplex and hence contains a ball $B(z, \epsilon)$. By translation, assume without loss of generality that $B(0, \epsilon) \subset K$.
- (ii) Show that $\rho : \mathbb{R}^d \rightarrow \mathbb{R}$ defined by

$$\rho(x) = \inf\{r > 0 : \frac{x}{r} \in K\}$$

is subadditive (i.e., $\rho(x+y) \leq \rho(x) + \rho(y)$) and satisfies

$$\frac{\|x\|}{\text{diam}(K)} \leq \rho(x) \leq \frac{\|x\|}{\epsilon}$$

for all \mathbf{x} . Deduce that ρ is continuous.

(iii) Define

$$h(x) = \frac{\rho(x)}{\|x\|}x$$

for $x \neq 0$ and $h(0) = 0$ and show that $h : K \rightarrow \overline{B(0, 1)}$ is a homeomorphism.

Solution to (ii): Suppose $\frac{x}{r} \in K$ and $\frac{y}{s} \in K$, where $r, s \geq 0$. Then

$$\frac{x+y}{r+s} = \frac{r}{r+s} \cdot \frac{x}{r} + \frac{s}{r+s} \cdot \frac{y}{s} \in K$$

so $\rho(x+y) \leq r+s$. Therefore $\rho(x+y) \leq \rho(x) + \rho(y)$, whence

$$\rho(x+y) - \rho(x) \leq \frac{\|y\|}{\epsilon}.$$

Similarly,

$$\rho(x) - \rho(x+y) \leq \rho(x+y-y) - \rho(x+y) \leq \frac{\|y\|}{\epsilon}.$$

Exercises

3.1 **The game of chicken.** Two drivers are headed for a collision. If both swerve, or Chicken Out, then the payoff to each is 1. If one swerves, and the other displays Iron Will, then the payoffs are -1 and 2 respectively to the players. If both display Iron Will, then a collision occurs, and the payoff is $-a$ to each of them, where $a > 2$. This makes the payoff matrix

		driver II	
		CO	IW
driver I	CO	(1, 1)	(-1, 2)
	IW	(2, -1)	(-a, -a)

Find all the pure and mixed Nash equilibria.

3.2 Modify the game of chicken as follows. There is $p \in (0, 1)$ such that, when a player plays *CO*, the move is changed to *IW* with probability p . Write the matrix for the modified game, and show that, in this case, the effect of increasing the value of a changes from the original version.

3.3 Two smart students form a study group in some Math Class where homeworks are handed in jointly by each study group. In the last homework of the semester, each of the two students can choose to either work (“W”) or defect (“D”). If at least one of them solves the homework that week (chooses “W”), then they will both receive 10 points. But solving the homework incurs an effort worth -7 points for a student doing it alone and an effort worth -2 points for each student if both students work together. Assume that the students do not communicate prior to deciding whether they will work or defect.

Write this situation as a matrix game and determine all Nash equilibria.

3.4 Find all Nash equilibria and determine which of the symmetric equilibria are evolutionarily stable in the following games.

		player II	
		A	B
player I	A	(4, 4)	(2, 5)
	B	(5, 2)	(3, 3)

		player II	
		A	B
player I	A	(4, 4)	(3, 2)
	B	(2, 3)	(5, 5)

3.5 Give an example of a two-player zero-sum game where there are no pure Nash equilibria. Can you give an example where all the entries of the payoff matrix are different?

3.6 **A recursive zero-sum game.** Player I, the Inspector, can inspect a facility on just one occasion, on one of the days $1, \dots, N$. Player II can cheat, or wait, on any given day. The payoff to I is 1 if I inspects while II is cheating. On any given day, the payoff is -1 if II cheats and is not caught. It is also -1 if I inspects but II did not cheat, and there is at least one day left. This leads to the following matrices Γ_n for the game with n days: the matrix Γ_1 is given by

		player II	
		Ch	Wa
player I	In	1	0
	Wa	-1	0

The matrix Γ_n is given by

		player II	
		Ch	Wa
player I	In	1	-1
	Wa	-1	Γ_{n-1}

Final optimal strategies, and the value of Γ_n .

3.7 Consider the following game:

		player II	
		C	D
player I	A	(6, -10)	(0, 10)
	B	(4, 1)	(1, 0)

- Show that this game has a unique mixed Nash equilibrium.
- Show that if player I can commit to playing strategy A with probability slightly more than x^* (the probability she plays A in the mixed Nash equilibrium), then (a) player I can increase her payoff, and (b) player II also benefits, obtaining a greater payoff than he did in the Nash equilibrium.
- Show similarly that if player II can commit to playing strategy C with probability slightly less than y^* (the probability he plays C in the mixed Nash equilibrium), then (a) player II can increase his payoff, and (b) player I also benefits, obtaining a greater payoff than she did in the Nash equilibrium.

3.8 **Two cheetahs and three antelopes:** Two cheetahs each chase one of three antelopes. If they catch the same one, they have to share. The antelopes are Large, Small and Tiny, and their values to the cheetahs are ℓ , s and t . Write the 3×3 matrix for this game. Assume that $t < s < \ell < 2s$, and that

$$\frac{\ell}{2} \left(\frac{2\ell - s}{s + \ell} \right) + s \left(\frac{2s - \ell}{s + \ell} \right) < t.$$

Find the pure equilibria, and the symmetric mixed equilibria.

3.9 Three firms (players I, II, and III) put three items on the market and advertise them either on morning or evening TV. A firm advertises exactly once per day. If more than one firm advertises at the same time, their profits are zero. If exactly one firm advertises in the morning, its profit is \$200K. If exactly one firm advertises in the evening, its profit is \$300K. Firms must make their advertising decisions simultaneously. Find a symmetric mixed Nash equilibrium.

3.10 **CHECK** Consider any two-player game of the following type.

		player II	
		A	B
player I	A	(a, a)	(b, c)
	B	(c, b)	(d, d)

- Compute optimal safety strategies and show that they are not a Nash equilibrium.
 - Compute the mixed Nash equilibrium and show that it results in the same player payoffs as the optimal safety strategies.
- 3.11 Consider the following symmetric game as played by two drivers, both trying to get from Here to There (or, two computers routing messages along cables of different bandwidths). There are two routes from Here to There; one is wider, and therefore faster, but congestion will slow them down if both take the same route. Denote the wide route W and the narrower route N . The payoff matrix is:




Fig. 3.9. The leftmost image shows the payoffs when both drivers drive on the narrower route, the middle image shows the payoffs when both drivers drive on the wider route and the rightmost image shows what happens when the red driver (player I) chooses the wide route and the yellow driver (Player II) chooses the narrow route.

		player II (yellow)	
		W	N
player I (red)	W	$(3, 3)$	$(5, 4)$
	N	$(4, 5)$	$(2, 2)$

Find all Nash equilibria and determine which ones are evolutionarily stable.

- 3.12 Argue that in a symmetric game, if $a_{ii} > b_{i,j}(= a_{j,i})$ for all $j \neq i$, then pure strategy i is an evolutionarily stable strategy.
- 3.13 **The fish-selling game revisited:** A seller sells fish. The fish is fresh with a probability of $2/3$. Whether a given piece of fish is fresh is known to the seller, but the customer knows only the probability. The customer asks, “is this fish fresh?”, and the seller answers, yes

or no. The customer then buys the fish, or leaves the store, without buying it. The payoff to the seller is 6 for selling the fish, and 6 for being truthful. The payoff to the customer is 3 for buying fresh fish, -1 for leaving if the fish is fresh, 0 for leaving if the fish is old, and -8 for buying an old fish.

- 3.14 **The welfare game:** John has no job and might try to get one. Or, he may prefer to take it easy. The government would like to aid John if he is looking for a job, but not if he stays idle. Denoting by T , trying to find work, and by NT , not doing so, and by A , aiding John, and by NA , not doing so, the payoff for each of the parties is given by:


		jobless John	
		try	not try
government	aid	(3,2)	(-1,3)
	no aid	(-1,1)	(0,0)

Find the Nash equilibria.


- 3.15 Show that, in a symmetric game, with $A = B^T$, there is a symmetric Nash equilibrium. One approach is to use the set $D = \{(x, x) : x \in \Delta_n\}$ in place of K in the proof of Nash's theorem.
- 3.16 **The game of Hawks and Doves.** Find the Nash equilibria in the game of Hawks and Doves whose payoffs are given by the matrix:

		player II	
		D	H
player I	D	(1,1)	(0,3)
	H	(3,0)	(-4,-4)

- 3.17 **A sequential congestion game:** Six drivers will travel from A to D , each going via either B or C . The cost in traveling a given road depends on the number of drivers k that have gone before (including the current driver). These costs are displayed in the figure. Each driver moves from A to D in a way that minimizes his or her own cost. Find the total cost. Then consider the variant where a superhighway that leads from A to C is built, whose cost for any driver is 1. Find the total cost in this case also.



- 3.18 **A simultaneous congestion game:** There are two drivers, one who will travel from A to C , the other, from B to D . Each road in the second figure has been marked (x, y) , where x is the cost to any driver who travels the road alone, and y is the cost to each driver who travels the road along with the other. Note that the roads are traveled simultaneously, in the sense that a road is traveled by both drivers if they each use it at some time during their journey. Write the game in matrix form, and find all of the pure Nash equilibria.



- 3.19 Sperner's lemma may be generalized to higher dimensions. In the case of $d = 3$, a simplex with four vertices (think of a pyramid) may be divided up into smaller ones. We insist that on each face of one of the small simplices, there are no edges or vertices of another. Label the four vertices of the big simplex 1, 2, 3, 4. Label those vertices of the small simplices on the boundary of the big one in such a way that each such vertex receives a label of one of the vertices of the big simplex that lies on the same face of the big simplex. Prove that there is a small simplex whose vertices receive distinct labels.
- 3.20 Prove the No-Retraction Theorem directly from Sperner's Lemma and use it to give an alternative proof of Brouwer's Theorem.

Notes

- Discuss to what extent Nash equilibria are a reasonable model for rational behavior.
- Solving polynomial equations Bernd Sturmfels.
- Tragedy of commons and pricing games from AGT chapter 1, example 1.4
- Examples: Investing in communication infrastructure, inspection game from Game Theory chapter, Encyclopedia of Information Systems by Turcocy and von Stengel.
- Regarding ESS definition: In the definition, we only allow the mutant strategies \mathbf{z} to be pure strategies. This definition is sometimes extended to allow any nearby (in some sense) strategy that doesn't differ too much from the population strategy \mathbf{x} , e.g., if the population only uses strategies 1, 3, and 5, then the mutants can introduce no more than one new strategy besides 1, 3, and 5.
- More general definition of what it means for a game to be symmetric.
- Example right before signaling:
Remark. Another situation that would remove the stability of (B, B) is if mutants were allowed to preferentially self-interact.
- potential games: Now, we have the following result due to Monderer and Shapley ([MS96]) and Rosenthal [Ros73]:
- In the absence of a mediator, the players could follow some external signal, like the weather.
- Coloring game from AMS

4

Signaling and asymmetric games

Example 4.0.24 (The car mechanic). A consumer (player I) takes her car to an expert car mechanic (player II) because it is making unusual noises. It is common knowledge that half the time these noises indicate a major repair is required, at a cost of 18 to the consumer, and half the time a minor repair at a cost of 10 suffices. In both cases, the mechanic's profit for doing the job is 8. After examining the car, the mechanic reports that the problem is major or minor. We assume that he always reports the problem truthfully when it is major. When it is minor, he could go either way. Unfortunately, the consumer is not sufficiently knowledgeable to tell whether the mechanic is honest. We assume that she always accepts his advice when he recommends a minor repair, but when he recommends a major repair, she either accepts his advice and pays him accordingly, or rejects his advice and takes the car to a different mechanic at an additional cost of 2.

Thus if the mechanic is honest, and the consumer accepts the advice, then he makes a profit of 8 and she incurs an expected loss of $\frac{1}{2}18 + \frac{1}{2}10 = 14$, whereas if she rejects a major repair, then his expected profit is $\frac{1}{2}8 = 4$ and she incurs an expected loss of $\frac{1}{2}20 + \frac{1}{2}10 = 15$. On the other hand, if the mechanic is dishonest (i.e., always reports that a major repair is necessary), and the consumer accepts, then he makes an expected profit of $\frac{1}{2}(8) + \frac{1}{2}(18 - 10 + 8) = 12$ and she incurs a loss of 18. If he is dishonest and she rejects, then his profit is 0 and she incurs an expected loss of $\frac{1}{2}20 + \frac{1}{2}12 = 16$.

This leads to the following payoff matrix:

		mechanic	
		honest	dishonest
consumer	accept	$(-14, 8)$	$(-18, 12)$
	reject	$(-15, 4)$	$(-16, 0)$

What are good strategies for the consumer and mechanic? We begin by considering safety strategies which, we recall, were optimal for zero-sum games. Since the consumer's payoffs are both lower when the mechanic is dishonest, her safety strategy is to always reject, yielding her a guaranteed payoff of at least -16. Similarly, the mechanic's payoffs are both lower when the consumer rejects, and thus his safety strategy is to always be honest, yielding him a guaranteed payoff of at least 4. However, the safety strategy pair (reject, honest) is *not* a Nash equilibrium. Indeed, knowing the mechanic is being honest, the consumer has an incentive to switch to accepting and would probably do so if the two players were to play the game again. But then, if in the next round they played (accept, honest), knowing that the consumer is accepting, the mechanic would have an incentive to switch to being dishonest in the subsequent round. This cycle would continue if in each round of game-playing, each player were to play a best response to the action of the other in the previous round. Indeed, this argument shows that this game has no pure Nash equilibrium.

There is, however, a mixed Nash equilibrium. Suppose the strategy $(x, 1 - x)$ for the consumer and $(y, 1 - y)$ for the mechanic are a Nash equilibrium. Then each ensures that both possible actions of the opponent yield the same payoff to the opponent and thus $8x + 4(1 - x) = 12x$ and $-14y - 18(1 - y) = -15y - 16(1 - y)$. These equations yield $x = 1/2$ (the consumer rejects with probability 1/2) and $y = 2/3$ (the mechanic is honest with probability 2/3). This equilibrium yields expected payoffs $(-15\frac{1}{3}, 6)$.

We interpret the $(\frac{2}{3}, \frac{1}{3})$ mixed strategy of the mechanic to mean the chance a randomly chosen mechanic will be honest is $q = 2/3$. This could arise from 2/3 of the mechanics being always honest, or from random choices by individual mechanics.

4.1 Signaling and asymmetric information

Example 4.1.1 (Lions and antelopes). In the games we have considered so far, both players are assumed to have access to the same information about the rules of the game. This is not always a valid assumption.

Antelopes have been observed to jump energetically when a lion nearby seems liable to hunt them. Why do they expend energy in this way? One theory was that the antelopes are signaling danger to others at some distance, in a community-spirited gesture. However, the antelopes have been observed doing this all alone. The currently accepted theory is that the signal is intended for the lion, to indicate that the antelope is in good health and is unlikely to be caught in a chase. This is the idea behind **signaling**.




Fig. 4.1. Lone antelope stotting to indicate its good health.

Consider the situation of an antelope catching sight of a lion in the distance. Suppose there are two kinds of antelope, healthy (H) and weak (W); and that a lion has no chance of catching a healthy antelope — but will expend a lot of energy trying — and will be able to catch a weak one. This can be modelled as a combination of two simple games (A^H and A^W), depending on whether the antelope is healthy or weak, in which the antelope has only one strategy (to run if pursued), but the lion has the choice of chasing (C) or ignoring (I).

		antelope	
		run-if-chased	
$A^H =$	lion	chase	(-1, -1)
	ignore		(0, 0)

		antelope	
		run-if-chased	
$A^W =$	lion	chase	(5, -1000)
	ignore		(0, 0)

The lion does not know which game they are playing — and if 20% of the antelopes are weak, then the lion can expect a payoff of $(.8)(-1) + (.2)(5) = .2$ by chasing. However, the antelope does know, and if a healthy antelope can convey that information to the lion by jumping very high, both will be better off — the antelope much more than the lion!

Remark. In this, and many other cases, the act of signaling itself costs something, but less than the expected gain, and there are many examples proposed in biology of such costly signaling.

4.1.1 Examples of signaling (and not)

Example 4.1.2 (A randomized game). For another example, consider the zero-sum two-player game in which the game to be played is randomized

by a fair coin toss. If heads is tossed, the payoff matrix is given by A^H , and if tails is tossed, it is given by A^T .

$$A^H = \begin{array}{c} \text{player I} \\ \begin{array}{cc|cc} & & \text{player II} & \\ & & L & R \\ \hline L & 4 & 1 & \\ R & 3 & 0 & \end{array} \end{array} \qquad A^T = \begin{array}{c} \text{player I} \\ \begin{array}{cc|cc} & & \text{player II} & \\ & & L & R \\ \hline L & 1 & 3 & \\ R & 2 & 5 & \end{array} \end{array}$$

If the players don't know the outcome of the coin flip before playing, they are merely playing the game given by the average matrix, $\frac{1}{2}A^H + \frac{1}{2}A^T$, which has a value of 2.5. If both players know the outcome of the coin flip, then (since A^H has a value of 1 and A^T has a value of 2) the value is 1.5 — player II is able to use the additional information to reduce her losses.

But now suppose that only I is told the result of the coin toss, but I must reveal her move first. If I goes with the simple strategy of picking the best row in whichever game is being played, and II realizes this and counters, then I has a payoff of only 1.5, less than the payoff if she ignores the extra information!

This demonstrates that sometimes the best strategy is to ignore the extra information, and play as if it were unknown. This is illustrated by the following (not entirely verified) story. During World War II, the English had used the Enigma machine to decode the German's communications. They intercepted the information that the Germans planned to bomb Coventry, a smallish city without many military targets. Since Coventry was such a strange target, the English realized that to prepare Coventry for attack would reveal that they had broken the German code, information which they valued more than the higher casualties in Coventry, and chose to not warn Coventry of the impending attack.

Example 4.1.3 (A simultaneous randomized game). Again, the game is chosen by a fair coin toss, the result of which is told to player I, but the players now make simultaneous moves, and a second game, with the same matrix, is played before any payoffs are revealed.

$$A^H = \begin{array}{c} \text{player I} \\ \begin{array}{cc|cc} & & \text{player II} & \\ & & L & R \\ \hline L & -1 & 0 & \\ R & 0 & 0 & \end{array} \end{array} \qquad A^T = \begin{array}{c} \text{player I} \\ \begin{array}{cc|cc} & & \text{player II} & \\ & & L & R \\ \hline L & 0 & 0 & \\ R & 0 & -1 & \end{array} \end{array}$$

Without the extra information, each player will play (L, R) with probabilities $(\frac{1}{2}, \frac{1}{2})$, and the value of the game to I (for the two rounds) is $-\frac{1}{2}$. However, once I knows which game is being played, she can simply choose

the row with all zeros, and lose nothing, regardless of whether II knows the coin toss as well.

Now consider the same story, but with matrices

$$A^H = \begin{array}{c} \text{player I} \\ \begin{array}{c|cc} & \text{player II} & \\ \hline & L & R \\ \hline L & 1 & 0 \\ R & 0 & 0 \end{array} \end{array} \qquad A^T = \begin{array}{c} \text{player I} \\ \begin{array}{c|cc} & \text{player II} & \\ \hline & L & R \\ \hline L & 0 & 0 \\ R & 0 & 1 \end{array} \end{array}$$

Again, without information the value to I is $\frac{1}{2}$. In the second round, I will clearly play the optimal row. The question remains of what I should do in the first round.

Player I has a simple strategy that will get her $\frac{3}{4}$ — this is to ignore the coin flip on the first round (and choose L with probability $\frac{1}{2}$), but then on the second round to choose the row with a 1 in it. In fact, this is the value of the game. If II chooses L with probability $\frac{1}{2}$ on the first round, but on the second round does the following: If I played L on the first round, then choose L or R with probability $\frac{1}{2}$ each; and if I played R on the first round, choose R , then I is restricted to a win of at most $\frac{3}{4}$. This can be shown by checking each of I's four pure strategies (recalling that I will always play the optimal row on the second round).

4.1.2 *The collapsing used car market*

Economist George Akerlof won the Nobel prize for analyzing how a used car market can break down in the presence of asymmetric information. Here is an extremely simplified version of his model. Suppose that there are cars of only two types: good cars (G) and lemons (L), and that both are at first indistinguishable to the buyer, who only discovers what kind of car he bought after a few weeks, when the lemons break down. Suppose that a good car is worth \$9000 to all sellers and \$12000 to all buyers, while a lemon is worth only \$3000 to sellers, and \$6000 to buyers. The fraction p of cars on the market that are lemons is known to all, as are the above values, but only the seller knows whether the car being sold is a lemon. The maximum amount that a rational buyer will pay for a car is $6000p + 12000(1-p) = f(p)$, and a seller who advertises a car at $f(p) - \varepsilon$ will sell it.

However, if $p > \frac{1}{2}$, then $f(p) < \$9000$, and sellers with good cars won't sell them — the market is not good, and they'll keep driving them — and p will increase, $f(p)$ will decrease, and soon only lemons are left on the market. In this case, asymmetric information hurts everyone.



Fig. 4.2. The seller, who knows the type of the car, may misrepresent it to the buyer, who doesn't know the type. (Drawing courtesy of Ranjit Samra.)

4.2 Some further examples

Example 4.2.1 (The fish-selling game).



Fig. 4.3. The seller knows whether the fish is fresh, the customer only knows the probability.

Fish being sold at the market is fresh with probability $\frac{2}{3}$ and old other-

wise, and the customer knows this. The seller knows whether the particular fish on sale now is fresh or old. The customer asks the fish-seller whether the fish is fresh, the seller answers, and then the customer decides to buy the fish, or to leave without buying it. The price asked for the fish is \$12. It is worth \$15 to the customer if fresh, and nothing if it is old. The seller bought the fish for \$6, and if it remains unsold, then he can sell it to another seller for the same \$6 if it is fresh, and he has to throw it out if it is old. On the other hand, if the fish is old, the seller claims it to be fresh, and the customer buys it, then the seller loses \$ R in reputation.

The tree of all possible scenarios, with the net payoffs shown as (seller, customer), is depicted in the figure. This is called the **Kuhn tree** of the game.




Fig. 4.4. The Kuhn tree for the fish-selling game.

The seller clearly should not say “old” if the fish is fresh, hence we should examine two possible pure strategies for him: “FF” means he always says “fresh”; “FO” means he always tells the truth. For the customer, there are four ways to react to what he might hear. Hearing “old” means that the fish is indeed old, so it is clear that he should leave in this case. Thus two rational strategies remain: BL means he buys the fish if he hears “fresh” and leaves if he hears “old”; LL means he just always leaves. Here are the expected payoffs for the two players, with randomness coming from the actual condition of the fish. (Recall that the fish is fresh with probability $2/3$ and old otherwise.)

		customer	
		BL	LL
seller	“FF”	$(6 - R/3, -2)$	$(-2, 0)$
	“FO”	$(2, 2)$	$(-2, 0)$

We see that if losing reputation does not cost too much in dollars, i.e., if $R < 12$, then there is only one pure Nash equilibrium: “FF” against

LL. However, if $R \geq 12$, then the (“FO”, BL) pair also becomes a pure equilibrium, and the payoff for this pair is much higher than the payoff for the other equilibrium.

5

Social choice

As social beings, we frequently find ourselves in situations where a group decision has to be made. Examples range from a simple decision a group of friends makes about picking a movie for the evening, to complex and crucial ones such as electing the president of a nation. Suppose that the individuals in a society are presented with a list of alternatives and have to choose one of them. Can a selection be made so as to truly reflect the preferences of the individuals? What does it mean for a social choice to be fair?

When there are only two options to choose from, **majority rule** can be applied to yield an outcome that more than half of the individuals find satisfactory. When the number of options is three or more, the majority preferences may be inconsistent, i.e. pairwise contests might yield a non-transitive (cyclic) outcome of the obtained by running pairwise contests can be in conflict with each other. This *paradox*, shown in the following figure, was first discovered by the Marquis de Condorcet in the late 18th century.




Fig. 5.1. In one-on-one contests A defeats C , C defeats B , and B defeats A .

5.1 Voting and Ranking Mechanisms

Example 5.1.1 (Plurality Voting). In (extended) plurality voting, each voter submits a rank-ordering of the candidates, and the candidate with the

most first-place votes wins the election (with some tie-breaking rule). It is not required that the winner have a majority of the votes. In the U.S., congressional elections are conducted using plurality voting.

This voting system has many advantages, foremost among them simplicity and transparency. On the other hand, plurality voting has the disturbing property that the candidate that is elected can be the least favorite for a majority of the population! Figure 5.2 gives an example of this with three candidates A , B , and C , and three different types of voters. Under simple




Fig. 5.2. Option A is preferred by 45% of the population, option B by 30% and option C by 25%.

plurality, A wins the election, despite the fact that A is ranked third by 55% of the population. This may motivate voters to misrepresent their preferences: If the 25% of voters who favor C were to move B to the top of their rankings, then B would win the election with a 55% majority, and these voters would be happier with the outcome.




Fig. 5.3. When 25% insincerely switch their votes from C to B , the relative ranking of A and B in the outcome changes.

This example illustrates a phenomenon that at first glance might seem odd: the third type of voters were able to change the outcome from A to B without changing the relative ordering of A and B in the rankings they submitted.

5.1.1 Definitions

We consider settings in which there is a set of candidates \mathcal{A} , a set of voters, and a **rule** that describes how the voters' preferences are used to determine an outcome. We consider two different kinds of rules. A **voting rule** produces a single winner, and a **ranking rule** produces a **social ranking** over

the candidates. Voting rules are obviously used for elections, or more generally, when a group needs to select one of several alternatives. A ranking rule might be used when a university department is ranking faculty candidates based on the preferences of current faculty members.

In both cases, we assume that the ranking of each voter is represented by a preference relation \succ on the set of candidates \mathcal{A} which is **complete** ($\forall A, B, A \succ B$ or $B \succ A$) and **transitive** ($A \succ B$ and $B \succ C$ implies $A \succ C$). Note that this definition does not allow for ties; we discuss rankings with ties in the notes.

We use \succ_i to denote the preference relation of voter i : $A \succ_i B$ if voter i strictly prefers candidate A to candidate B .

Definition 5.1.2. A **voting rule** f maps each **preference profile**, $\pi = (\succ_1, \dots, \succ_n)$ to an element of \mathcal{A} , the winner of the election.

Definition 5.1.3. A **ranking rule** R associates to each **preference profile**, $\pi = (\succ_1, \dots, \succ_n)$, a **social ranking**, another complete and transitive preference relation $\triangleright = R(\pi)$. ($A \triangleright B$ means that A is strictly preferred to B in the social ranking.)

Remark. An obvious way to obtain a voting rule from a ranking rule is to output the top ranked candidate. (For another way, see exercise ??.) Conversely, a voting rule yields an **induced ranking rule** as follows. Apply the voting rule to select the top candidate. Then apply the voting rule to the remaining candidates to select the next candidate and so on. However, not all ranking rules can be obtained this way; see exercise ??.

Two properties that we might desire a ranking rule R to have are:

- **Unanimity:** If for every voter i we have $A \succ_i B$, then $A \triangleright B$. In words, if every voter strictly prefers candidate A to B , then A should be strictly preferred to B in the social ranking.
- **Independence of irrelevant alternatives (IIA):** For any two candidates A and B , the preference between A and B in the social ranking depends only on the voters' preferences between A and B . Formally, if $\pi = \{\succ_i\}$ and $\pi' = \{\succ'_i\}$ are two profiles such that $\{i \mid A \succ_i B\} = \{i \mid A \succ'_i B\}$ and $\{i \mid B \succ_i A\} = \{i \mid B \succ'_i A\}$, then $A \triangleright B$ implies $A \triangleright' B$.

The desirability of unanimity is incontrovertible and indeed it holds for all ranking rules that are used in practice. One motivation for IIA is that, if it fails, then some voter is incentivized to misrepresent his preferences; see the

next definition and lemma. However, almost all ranking rules violate IIA, and we will see why later in the chapter.

Definition 5.1.4. A ranking rule R is **strategically vulnerable** at the profile $\pi = (\succ_1, \dots, \succ_n)$, if there is a voter i and alternatives A and B so that $A \succ_i B$ and $B \triangleright A$ in $R(\pi)$, yet replacing \succ_i by \succ_i^* yields a profile π^* such that $A \triangleright^* B$ in $R(\pi^*)$.

Lemma 5.1.5. *If a ranking rule R violates IIA, then it is strategically vulnerable.*

Proof. Let $\pi = \{\succ_i\}$ and $\pi' = \{\succ'_i\}$ be two profiles such that $\{i \mid A \succ_i B\} = \{i \mid A \succ'_i B\}$ and $\{i \mid B \succ_i A\} = \{i \mid B \succ'_i A\}$, but $A \triangleright B$ in $R(\pi)$ whereas $B \triangleright' A$ in $R(\pi')$. Let $\sigma_i = (\succ'_1, \dots, \succ'_i, \succ_{i+1}, \dots, \succ_n)$, so that $\sigma_0 = \pi$ and $\sigma_n = \pi'$. Let $i \in [1, n]$ be such that $A \triangleright B$ in $R(\sigma_{i-1})$, but $B \triangleright A$ in $R(\sigma_i)$. If $B \succ_i A$, then R is strategically vulnerable at σ_{i-1} since voter i can switch from \succ_i to \succ'_i . Similarly, if $A \succ_i B$, then R is vulnerable at σ_i , since voter i can switch from \succ'_i to \succ_i . \square

For plurality voting, as we saw in the example of Figures 5.2 and 5.3, the induced ranking rule violates IIA.

5.1.2 Instant runoff elections

In **instant runoff elections** (or plurality with elimination), the winner in an election with N candidates is determined by the following procedure. If $N = 2$, then the winner is the candidate with the majority of first-place votes. If $N > 2$, the candidate with the fewest first-place votes is eliminated from consideration, and removed from all the rankings. An instant run-off election is then run on the remaining $N - 1$ candidates. Figure 5.4 shows an example. (Notice that runoff voting and plurality yield different winners in this example.)




Fig. 5.4. In the first round C is eliminated. When votes are redistributed, B gets the majority. The full voter rankings are not revealed in the process.

Instant runoff is used in Australia and Fiji for the House of Representatives, in Ireland to elect the President, and for various municipal elections in Australia, the United States, and New Zealand.

Unfortunately, this method is also vulnerable to strategic manipulation. Consider the scenario depicted in Figure 5.5. If voters in the first group




Fig. 5.5. After *C* is eliminated, *B* gets the majority of votes.

knew the distribution of preferences, they could ensure a victory for *A* by getting some of their constituents to conceal their true preference and move *C* from the bottom to the top of their rankings, as shown in Figure 5.6. In the first round, *B* would be eliminated. Subsequently *A* would win against *C*.




Fig. 5.6. A small group misrepresents their true preferences, ensuring that *B* is eliminated. As a result, *A* wins the election.

This example show that ranking rule induced by instant runoff also violates the IIA criterion since it allows for the relative ranking of *A* and *B* to be switched without changing any of the individual *A-B* preferences.

5.1.3 Borda count

Borda count is a ranking rule in which each voter’s ranking is used to assign points to the candidates. If there are *N* candidates, then *N* points are assigned to each voter’s top-ranked candidate on down to one point for his lowest ranked candidate. The candidates are then ranked in decreasing order of their point totals (with ties broken arbitrarily).

The Borda count is also vulnerable to strategic manipulation. In the example shown in Figure 5.7, *A* has an unambiguous majority of votes and is also the winner.

However, if supporters of *C* were to strategically rank *B* above *A*, they could ensure a victory for *C*. This is also a violation of IIA, since none of the individual *A-C* preferences had been changed.




Fig. 5.7. Alternative A has the overall majority and is the winner under Borda count.




Fig. 5.8. Supporters of C can bury A by moving B up in their rankings.

5.1.4 Dictatorship

A ranking rule is a **dictatorship** if there is a voter v whose preferences are reproduced in the outcome. In other words, for every pair of candidates A and B , $A \succ_v B$ if and only if $A \triangleright B$.

While dictatorship does satisfy unanimity and IIA, most of us would regard this method as unacceptable.

5.2 Arrow's impossibility theorem

In 1951, Kenneth Arrow formulated and proved his famous *Impossibility Theorem*.

Theorem 5.2.1. [Arrow's Impossibility Theorem] *Any ranking rule that satisfies unanimity and independence of irrelevant alternatives is a dictatorship.*

What does the theorem mean? If we want to avoid dictatorship, we must accept the possibility of strategic manipulation in our ranking system; the same applies to voting by Theorem 5.3.2. Thus, strategizing (i.e., game theory) is an inevitable part of ranking and voting.

Proof of Arrow's theorem:

Fix a ranking rule R that satisfies unanimity and IIA. The proof we present requires that we consider extremal candidates, those that are either most preferred or least preferred. The proof is written so that it applies verbatim to rankings with ties, as discussed in the notes; therefore, we occasionally refer to "strict" preferences.

Lemma 5.2.2 (Extremal Lemma). *Consider an arbitrary candidate B . For any profile π in which B has an extremal rank for all voters (i.e., B is strictly preferred to all other candidates or all other candidates are strictly preferred to B), B has an extremal rank in the social ranking $R(\pi)$.*

Proof. Suppose not. Then for such a profile π , with $\triangleright = R(\pi)$, there are two candidates A and C such that $A \triangleright B$ and $B \triangleright C$. Consider a new profile $\pi' = (\succ'_1, \dots, \succ'_n)$ obtained from π by having every voter move C just above A in their ranking. None of the AB or BC preferences change since B started out and stays in the same extremal rank. Hence, by IIA, in the outcome $\triangleright' = R(\pi')$, we have $A \triangleright' B$ and $B \triangleright' C$, and hence $A \triangleright' C$. But this violates unanimity, since for all voters i in π' , we have $C \succ'_i A$. \square

Definition 5.2.3. Let B be a candidate. Voter i is said to be B -**pivotal** if there exist profiles π_1 and π_2 such that

- B is extremal for all voters in both profiles;
- The only difference between π_1 and π_2 is that B is strictly lowest ranked by i in π_1 and B is strictly highest ranked by i in π_2 ;
- B is ranked strictly lowest in $R(\pi_1)$ and strictly highest in $R(\pi_2)$.

Such a voter has the “power” to move candidate B from the very bottom of the outcome ranking to the very top.

Lemma 5.2.4. *For every candidate B , there is a B -pivotal voter $v(B)$.*

Proof. Consider an arbitrary profile in which candidate B is ranked strictly lowest by every voter. By unanimity, all other candidates are strictly preferred to B in the social ranking. Now consider a sequence of profiles obtained by letting the voters, one at a time, move B from the bottom to the top of their rankings. By the extremal lemma, for each one of these profiles, B is either at the top or at the bottom of the social ranking. Also, by unanimity, as soon as all the voters put B at the top of their rankings, so must the social ranking. Hence, there is a first voter v whose change in preference precipitates the change in the social ranking of candidate B . This change is illustrated in Figures 5.9 and 5.10, where π_1 is the profile just before v has switched B to the top with $\triangleright_1 = R(\pi_1)$, and π_2 the profile immediately after the switch with $\triangleright_2 = R(\pi_2)$. This voter v is B -pivotal. \square

Lemma 5.2.5. *If voter v is B -pivotal, v is a dictator on $\mathcal{A} \setminus \{B\}$, i.e., for any profile π , if $A \neq B$ and $C \neq B$ satisfy $A \succ_v C$ in π , then $A \triangleright C$ in $R(\pi)$.*




Fig. 5.9.




Fig. 5.10.

Proof. Let π_1 and π_2 be the profiles from the definition of v being B -pivotal. Fix $A \neq B$ and $C \neq B$. Consider the set of profiles \mathcal{P} obtained from π_2 by letting v move A to the top of her ranking, just above B , leaving B in its same extremal position for all other voters, and otherwise, changing the preferences arbitrarily. Let π_3 be any profile in \mathcal{P} , and let $\triangleright_3 = R(\pi_3)$. Then the preferences between A and B in π_3 are the same as in π_1 and thus, by IIA, we have $A \triangleright_3 B$. Also, the preferences between B and C in π_3 are the same as in π_2 and thus, by IIA, we have $B \triangleright_3 C$. Hence, by transitivity, we have $A \triangleright_3 C$. Since the profiles in \mathcal{P} include all possible preferences between A and C for voters other than v , by IIA, for *any* profile π in which $A \succ_v C$, it must be that $A \triangleright C$. □

We can now complete the proof of Theorem 5.2.1. By Lemmas 5.2.4 and 5.2.5, there is a B -pivotal voter $v = v(B)$ that is a dictator on $\mathcal{A} \setminus \{B\}$. Let π_1 and π_2 be the profiles from the definition of v being B -pivotal. We claim that for any other candidate, say C , the C -pivotal voter $v' = v(C)$ is actually the same voter, i.e. $v = v'$.

To see this, consider any $A \neq B$ and $A \neq C$. We know that in \triangleright_1 , we have $A \triangleright B$, and in \triangleright_2 , we have $B \triangleright A$. Moreover, by Lemma 5.2.5, v' dictates the strict preference between A and B in both of these outcomes. But in both profiles, the strict preference between A and B is the same for all voters other than v . Hence $v' = v$, and thus v is a dictator (over all of \mathcal{A}).

5.3 Strategy-proof Voting

We next turn our attention to voting rules. Consider n voters in a society, each with a complete ranking of a set of m alternatives \mathcal{A} , and a voting rule

f mapping each profile $\pi = (\succ_1, \dots, \succ_n)$ of n rankings of \mathcal{A} to an alternative $f(\pi) \in \mathcal{A}$.

What voting rules f have the property that no matter what preferences are submitted by other voters, each voter is incentivized to report their ranking truthfully? Such a voting rule is called **strategy-proof**.

Definition 5.3.1. A voting rule f from profiles to \mathcal{A} is **strategy-proof** if for all profiles π , candidates A and B , and voters i , the following holds: If $A \succ_i B$ and $f(\pi) = B$, then all π' that differ from π only on voter i 's ranking satisfy $f(\pi') \neq A$.

Theorem 5.3.2 (Gibbard-Satterthwaite). *Let f be a strategy-proof voting rule onto \mathcal{A} , where $|\mathcal{A}| \geq 3$. Then f is a **dictatorship**. That is, there is a voter i such that for every profile π , voter i 's highest ranked alternative is equal to $f(\pi)$.*

We prove the theorem as a corollary of Arrow's theorem, by showing that if f is strategy-proof and is not a dictatorship, then it can be extended to a ranking rule that satisfies unanimity, IIA and is not a dictatorship, a contradiction.

The following notation will also be useful.

Definition 5.3.3. For any two profiles $\pi = (\succ_1, \dots, \succ_n)$ and $\pi' = (\succ'_1, \dots, \succ'_n)$, we let $r_i(\pi, \pi')$ denote the profile $(\succ'_1, \dots, \succ'_i, \succ_{i+1}, \dots, \succ_n)$. Thus $r_0(\pi, \pi') = \pi$ and $r_n(\pi, \pi') = \pi'$.

We will repeatedly use the following lemma:

Lemma 5.3.4. *Suppose that f is strategy-proof. Consider two profiles $\pi = (\succ_1, \dots, \succ_n)$ and $\pi' = (\succ'_1, \dots, \succ'_n)$ and two candidates X and Y such that:*

- *all preferences between X and Y in π and π' are the same (i.e., $X \succ_i Y$ iff $X \succ'_i Y$ for all i);*
- *in π' all voters prefer X to all candidates other than possibly Y (i.e., $X \succ'_i Z$ for all $Z \notin \{X, Y\}$);*
- $f(\pi) = X$.

Then $f(\pi') = X$.

Proof. We have $f(r_0) = X$ by assumption. We prove by induction on i , that $f(r_i) = X$, where $r_i = r_i(\pi, \pi')$, or else f is not strategy-proof. To this end, suppose that $f(r_{i-1}) = X$. Observe that r_{i-1} and r_i differ only on voter i 's preferences: in r_{i-1} it is \succ_i and in r_i it is \succ'_i .

There are two cases: If $f(r_i) = Z \notin \{X, Y\}$, then on profile r_i , voter i has an incentive to lie and report \succ_i instead of \succ'_i .

On the other hand, suppose $f(r_i) = Y$. If $X \succ_i Y$, then on profile r_i , voter i has an incentive to lie and report \succ_i instead of \succ'_i . On the other hand, if $Y \succ_i X$, then on profile r_{i-1} , voter i has an incentive to lie and report \succ'_i instead of \succ_i . \square

We also need the following definition.

Definition 5.3.5. Let S be a subset of the alternatives \mathcal{A} , and let π be a ranking of the alternatives \mathcal{A} . Define a new ranking π^S by moving all alternatives in S to the top of the ranking, maintaining the same relative ranking between them, as well as the same relative ranking between all alternatives not in S .

Claim 5.3.6. Let f be strategy-proof and onto \mathcal{A} . Then for any profile π , and any subset S of the alternatives \mathcal{A} , it must be that $f(\pi^S) \in S$.

Proof. Take any $A \in S$. Since f is onto, there is a profile $\tilde{\pi}$ such that $f(\tilde{\pi}) = A$. Consider the sequence of profiles $r_i = r_i(\tilde{\pi}, \pi^S)$, with $0 \leq i \leq n$. We claim that $f(r_{i-1}) \in S$ implies that $f(r_i) \in S$. Otherwise, on profile r_i , voter i has an incentive to lie and report $\tilde{\succ}_i$ instead of \succ_i^S . Thus, since $f(r_0) = f(\tilde{\pi}) \in S$, we conclude that $f(r_n) = f(\pi^S) \in S$ as well. \square

We can now complete the proof of Theorem 5.3.2. Let f be strategy-proof, onto and a non-dictatorship. Define a ranking rule $R(\pi)$ as follows. For each pair of alternatives A and B , let $A \triangleright B$ if $f(\pi^{\{A,B\}}) = A$ and $B \triangleright A$ if $f(\pi^{\{A,B\}}) = B$. (Claim 5.3.6 guarantees that these are the only two possibilities.)

To see that this is a bona fide ranking rule, we observe that these pairwise rankings are transitive. If not, there is a triple of alternatives such that $A \triangleright B$, $B \triangleright C$ and $C \triangleright A$. Let $S = \{A, B, C\}$. We know that $f(\pi^S) \in S$, without loss of generality $f(\pi^S) = A$. Applying Lemma 5.3.4, with $\pi = \pi^S$ and $\pi' = \pi^{\{A,C\}}$, $X = A$ and $Y = C$, we conclude that $f(\pi^{\{A,C\}}) = A$ and $A \triangleright C$, a contradiction.

Finally, we observe that the ranking rule R satisfies unanimity, IIA, and is not a dictatorship.

Unanimity follows from the fact that if in π all voters have $A \succ_i B$, then $(\pi^{\{A,B\}})^A = \pi^{\{A,B\}}$, and thus by Claim 5.3.6, $f(\pi^{\{A,B\}}) = A$.

To see that IIA holds, let π_1 and π_2 be two profiles that agree on all of their AB preferences. Then by Lemma 5.3.4, with $\pi = \pi_1^{\{A,B\}}$ and $\pi' = \pi_2^{\{A,B\}}$, and Claim 5.3.6, we conclude that $f(\pi_1^{\{A,B\}}) = f(\pi_2^{\{A,B\}})$, and hence IIA holds.

Finally, the ranking rule R is not a dictatorship because f is not a dictatorship: For every voter v , there is a profile π for which v 's highest ranked alternative is A , but for which $f(\pi) = B \neq A$. Then, applying Lemma 5.3.4 to the pair of profiles π and $\pi^{\{A,B\}}$, with $X = B$ and $Y = A$, we conclude that $f(\pi^{\{A,B\}}) = B$, and thus $B \triangleright A$ in the outcome of the election. Hence voter v is not a dictator relative to the ranking rule R .

Exercises

- 5.1 Give an example where one of the losing candidates in a runoff election would have a greater support than the winner in a one-on-one contest.
- 5.2 Describe a ranking rule that is not the induced ranking rule of any voting rule.
- 5.3 Another way to go from a ranking rule to a voting rule. Apply this procedure and the one in the text to Vote-counting. What voting rule do you get in the two cases?
- 5.4 For other voting rules and ranking rules, find example violation of IIA or way to manipulate for voting. (e.g. Approval voting):

Notes

The study of voting has a long history....

Voting Rules:

Chevalier de Borda proposed the Borda count in 1770 when he discovered that the plurality method then used by the French Academy of Sciences was vulnerable to strategic manipulation. The Borda count was subsequently used by the Academy for the next two decades.

The method of pairwise contests referred to in the beginning of this chapter was proposed by the Marquis de Condorcet after he demonstrated that the Borda count was also vulnerable to strategic manipulation. He then proceeded to show a vulnerability in his own method — a tie in the presence of a preference cycle [dCMdC90].

Donald G. Saari showed that Borda count is in some sense ??? the least problematic of all single winner mechanisms [Saa90],[Saa06].

We have surveyed only a few of the many voting rules that have been considered. Other voting rules include approval voting, Approval voting is a procedure in which voters can vote for, or approve of, as many candidates as they wish, and candidates are ranked by the number of approval votes they receive.

Arrow's Impossibility Theorem

We have presented here a simplified proof of Arrow's theorem that is due to Geanakoplos [Gea04]. The version in the text assumes that each voter has a complete ranking of all the candidates. However, in many cases voters are indifferent

between certain subsets of candidates. To accommodate this possibility, one can generalize the setting as follows.

Assume that the preferences of each voter are described by a relation \succeq on the set of candidates \mathcal{A} which is **reflexive** ($\forall A, A \succeq A$), **complete** ($\forall A, B, A \succeq B$ or $B \succeq A$ or both) and **transitive** ($A \succeq B$ and $B \succeq C$ implies $A \succeq C$).

As in the chapter, we use \succeq_i to denote the preference relation of voter i : $A \succeq_i B$ if voter i weakly prefers candidate A to candidate B . However, we can now distinguish between strict preferences and indifference. As before, we use the notation $A \succ_i B$ to denote a strict preference, i.e., $A \succeq_i B$ but $B \not\succeq_i A$. (If $A \succeq_i B$ and $B \succeq_i A$, then voter i is indifferent between the two candidates.)

A reflexive, complete and transitive relation \succeq can be described in two other equivalent ways:

- It is a set of equivalence classes (each equivalence class is a set of candidates that the voter is indifferent between), with a total order on the equivalence classes. In other words, it is a ranking that allows for ties.
- It is the ranking induced by a function $g : \mathcal{A} \rightarrow \mathbb{R}$ from the candidates to the reals, such that $A \succeq B$ if and only if $g(A) \geq g(B)$. Obviously, many functions induce the same preference relation.

A **ranking rule** R associates to each **preference profile**, $\pi = (\succeq_1, \dots, \succeq_n)$, another reflexive, complete and transitive preference $\succeq = R(\pi)$.

In this more general setting, the definitions of unanimity and IIA are essentially unchanged. (Formally, IIA states that if $\pi = \{\succeq_i\}$ and $\pi' = \{\succeq'_i\}$ are two profiles such that $\{i \mid A \succeq_i B\} = \{i \mid A \succeq'_i B\}$ and $\{i \mid B \succeq_i A\} = \{i \mid B \succeq'_i A\}$, then $A \succeq B$ implies $A \succeq' B$.)

Arrow's theorem in this setting is virtually identical to the version given in the text: Any ranking rule that satisfies unanimity and IIA is a dictatorship. The only difference is that, in the presence of ties, voters other than the dictator can influence the outcome with respect to candidates that the dictator is indifferent between. Formally, in this more general setting, a dictator is a voter v all of whose *strict* preferences are reproduced in the outcome.

It is straightforward to check that the proof presented in Section 5.2 goes through unchanged.

6

Auctions and Mechanism Design

6.1 Auctions

Auctions are an ancient mechanism for buying and selling goods, and in modern times a huge volume of economic transactions is conducted through auctions: The US government runs auctions to sell treasury bills, spectrum licenses and timber and oil leases, among others. Christie's and Sotheby's run auctions to sell art. In the age of the Internet, we can buy and sell goods and services via auction, using the services of companies like eBay. The advertisement auctions that companies like Google, Yahoo! and Microsoft run in order to sell advertisement slots on their web pages bring in a significant fraction of their revenue.

Why might a seller use an auction as opposed to simply fixing a price? Primarily because sellers often don't know how much buyers value their goods, and don't want to risk setting prices that are either too low, thereby leaving money on the table, or, so high that nobody will want to buy the item. An auction is a technique for dynamically setting prices. Auctions are particularly important these days because of their prevalence in Internet settings where the participants in the auction are computer programs, or individuals with no direct knowledge of or contact with each other. As auctions are games of incomplete information, game theory provides us with the tools to understand their design and analysis.

6.2 Single Item Auctions

We are all familiar with the famous **English or ascending auction** for selling a single item: The auctioneer starts by calling out a low price p . As long as there are at least two people willing to pay the price p , he increases p by a small amount. This continues until there is only one player left willing

to pay the current price, at which point that player “wins” the auction, i.e. receives the item at that price.

When multiple rounds of communication are inconvenient, the English auction is sometimes replaced by other formats. For example, in a **sealed-bid first-price auction**, the participants submit sealed bids to the auctioneer. The auctioneer allocates the item to the highest bidder who pays the amount she bid.

To compare these and other auction formats, we introduce the following model: We assume that each player has a private value v for the item being auctioned off. This means that he would not pay more than v for the item, while if he gets the item at a price $p < v$, his gain is $v - p$. Given the rules of the auction, and any knowledge he has about other players’ bids, he will bid so as to maximize his gain.

In the ascending auction, a player will keep bidding only if the current price is below his value. But how should a player bid in a sealed-bid first price auction? Clearly, bidding one’s value makes no sense, since even upon winning, this would result in a gain of 0! So a bidder will want to bid lower than their true value. But how much lower? Low bidding has the potential to increase a player’s gain, but at the same time increases the risk of losing the auction. In fact, the optimal bid in such an auction depends on how the other players are bidding, which in general, a bidder will not know.

Example 6.2.1. Suppose that two players are competing in a first-price auction, and each of them knows that the other player’s value is a random draw from a $\text{Unif}[0,1]$ distribution.

In this game, each player has a strategy s . A strategy is a mapping $s : [0, 1] \rightarrow [0, 1]$ which describes the bid $b = s(v)$ of a player when his value is v .

We claim that in this setting it is an *equilibrium* for both players to bid half of their value, i.e. $s(v) = v/2$.

To see this, we compute the expected gain of player 1 with value v_1 when his bid is b_1 , assuming the other player’s value v_2 is drawn from $U[0, 1]$ and is using the strategy $s(v_2) = v_2/2$:

$$\begin{aligned} \mathbb{E}_{v_2} [(v_1 - b_1) \mathbb{1}_{b_1 \geq b_2}] &= \mathbb{P}_{v_2} [b_1 \geq s(v_2)](v_1 - b_1) \\ &= \mathbb{P}_{v_2} [b_1 \geq v_2/2](v_1 - b_1). \end{aligned}$$

Since v_2 is $U[0, 1]$, $\mathbb{P}_{v_2} [b_1 \geq v_2/2] = 2b_1$ and therefore

$$\mathbb{E}_{v_2} [\text{gain of player 1 when his value is } v_1] = 2b_1(v_1 - b_1).$$

Choosing b_1 to maximize the above expression, we see that player 1's expected gain is maximized by setting $b_1 = v_1/2 = s(v_1)$.

The symmetric argument for the other player shows that the strategies $s_1(v_1) = v_1/2$ and $s_2(v_2) = v_2/2$ are best responses to each other in expectation over the other player's random value, and thus it is an equilibrium for both players to play this way.

So is strategic bidding a necessary consequence of the convenience of sealed-bid auctions? No. Nobel-prize winner William Vickrey (1960) discovered that one can combine the low communication cost of sealed-bid auctions with the optimality of truthful bidding found in the ascending auction. We can get a hint on how to construct this combination by determining the revenue of the auctioneer in the ascending auction when all players act rationally: The item is sold to the highest bidder when the current price exceeds what other bidders are willing to offer; this threshold price is approximately the value of the item to the second-highest bidder.

Definition 6.2.2. In a **(sealed bid) second price auction**, the highest bidder wins the auction at a price equal to the second highest bid.

Theorem 6.2.3. *The second price auction is truthful. In other words, for each player i , and for any fixed set of bids of all players except for player i , player i 's gain is maximized by bidding their true value v_i .*

Proof. Suppose the maximum of the bids submitted by players other than i is m . If $m > v_i$, bidding truthfully (or bidding any value that is at most m) will result in a utility of 0 for player i . On the other hand, bidding above m would result in a negative utility. Thus, the player cannot gain by lying. On the other hand, if $m \leq v$, then as long as the player wins the auction, his utility will be $v - m \geq 0$. Thus, the only change in utility that can result due to bidding untruthfully occurs if the player bids below m , in which case, his utility will be 0 since he then loses the auction. □

Remark. We emphasize that the theorem statement is not merely saying that truthful bidding is a Nash equilibrium, but rather the much stronger statement that bidding truthfully is a dominant strategy, i.e., it maximizes each player's gain *no matter how other players play*.

6.2.1 Profit in single-item auctions

From the perspective of the bidders in an auction, a second price auction is appealing. They don't need to perform any complex strategic calculations.

The appeal is less clear, however, from the perspective of the auctioneer. Wouldn't the auctioneer make more money running a first price auction?

Example 6.2.4. We return to our early example of two players, each with a value drawn from $U[0,1]$ distribution. From that analysis, we know that if the auctioneer runs a first price auction, then in equilibrium his profit will be

$$E_{v_1, v_2}(\max(v_1/2, v_2/2)) = 1/3.$$

On the other hand, suppose that in the exact same setting, the auctioneer runs a second-price auction. Since the players will bid truthfully, the auctioneer's profit will be the expected value of the second highest bid, which is

$$E_{v_1, v_2}(\min(v_1, v_2)) = 1/3,$$

exactly the same as in the 1st price auction!

Coincidence? No. As we shall see in Section 6.6.5, the very important *revenue equivalence theorem* shows that *any* auction that has same allocation rule in equilibrium yields the same auctioneer revenue! This applies even to weird auctions like the **all-pay auction**, in which the winner is also the bidder who submitted the highest bid, but this time *all* bidders pay the auctioneer what they bid. (While this might seem like an auction nobody would want to participate in, it models a number of interesting scenarios outside the realm of selling a single item. For example, variants on this basic all-pay auction are sometimes used to model political races or patent races among firms, since in these cases there is only one winner, but all the participants put in effort that is ultimately wasted if that participant loses the "auction", and they need to decide how much effort to put in.)

The lesson is that bidders adapt their behavior to the rules of the auction and the information they have about their opponents, and this can lead to outcomes that one would not necessarily expect at first glance.

6.2.2 More profit?

The discussion above implies that auctions that always allocate (in equilibrium) to the player with the highest value make the same amount of profit. Unfortunately, this profit can be very low. A notorious example was the 1990 New Zealand sale of spectrum licenses in which a 2nd price auction was used, the bidder bid \$100,000, but paid only \$6!

Is it possible for the auctioneer to obtain more profit with a different auction format? For example, if there are two players participating, and it is

known that their values are $U[0,1]$, can the auctioneer make a higher expected profit than $1/3$? It turns out he can, by running a **Vickrey auction with a reserve price**. This is a sealed bid auction in which the item is allocated to the highest bidder, but only if her bid is at least the reserve price. Her payment is the maximum of the reserve price and the bid of the next highest bidder. A virtually identical argument to that of Theorem 6.2.3 shows that the Vickrey auction with a reserve price is also truthful.

6.2.3 Exercise: Vickrey with Reserve Price

Determine the expected profit of the auctioneer in a Vickrey auction with a reserve price of r , with two players whose values are drawn independently from a $U[0,1]$ distribution. Show that this expected revenue is $5/12$ when the reserve price is $1/2$, and that this is the optimal choice of reserve price.

Remarkably, this simple auction optimizes the auctioneer's expected revenue over all possible auctions! It is a special case of *Myerson's optimal mechanism*, a broadly applicable technique for maximizing auctioneer revenue when agents values are drawn from known prior distributions. This mechanism is discussed in §Section 6.6.7.

6.3 The general mechanism design problem

Mechanism design is concerned with designing games of *incomplete information* such that, in equilibrium, the mechanism designer's goals are accomplished. In the previous section, we pondered the design of an auction for selling a single item, and asked what auction design best achieves the auctioneer's goals, say profit maximization, *in equilibrium*. There are many other mechanism design problems as well.

Example 6.3.1. The federal government is trying to determine which roads to build to connect a new city C to cities A and B (which already have a road between them). The options are to build a road from A to C or a road from B to C, both roads, or neither. Each road will cost the government 10 million dollars to build. Each city obtains a certain economic/social benefit for each outcome. For example, city A might obtain a 5 million dollar benefit from the creation of a road to city C, but no real benefit from the creation of a road between B and C. City C, on the other hand, currently disconnected from the others, obtains a significant benefit (9 million) from the creation

of each road, but the marginal benefit of adding a second connection is not as great as the benefit creating a first connection. The following table summarizes these values (in millions), and the cost to the government for each option.

	road A-C	road B-C	both	none
City A	5	0	5	0
City B	0	5	5	0
City C	9	9	15	0
Government	-10	-10	-20	0

The government's goal is to choose the option that yields the highest total benefit to society which, for these numbers, is the creation of both roads. However, these numbers are reported to the government by the cities themselves, who may have an incentive to exaggerate their values, so that their preferred option will be selected. Thus, the government would like to employ a mechanism for learning the values and making the decision that provides the correct incentive structure.

Example 6.3.2. Consider a mechanism for the allocation of takeoff and landing slots at a particular airport to competing airlines (the participants). The valuation functions of the airlines are likely to be quite complicated incorporating specifications such as "I want to purchase the right to use a landing slot at 3pm each day and the right to use a takeoff slot at 3:30pm each day for \$X," where X might depend on whether the takeoff and landing slots are the same or close to each other, and might vary depending on the precise timing involved. In such a setting, the mechanism designer (the airport administration or the government) might want to maximize the utilization of these slots, or maximize the total profit taken in, or any of a number of other objectives.

Example 6.3.3. When you perform a query in a search engine, you receive a page of results containing the links the search engine has deemed relevant to the search, together with *sponsored links*, i.e. "advertisements". (See figure.) This enables advertisers to precisely target their ads based on users' search terms. For example, if a travel agent "buys" the search term "Tahiti", then when searching on the word "Tahiti", a user might be shown a link to a web page for that travel agent offering, say, plane tickets to Tahiti. If the user actually clicks on this link, he will be transferred to the aforementioned web page. For each such click, in which the advertiser receives a potential customer, the advertiser pays the search engine.

The process of determining which ads get assigned to which keywords

(search terms) and how much each advertiser pays is resolved via *keyword auctions*. Advertisers choose which keywords they want to bid on and participate in auctions for those keywords. When a user submits a query on that keyword, an instantaneous auction is run to determine which of the advertisers currently bidding on that keyword are allocated advertising slots.

Here, the search engine needs to design the mechanism for allocating these advertisement slots to advertisers and the payment rules. A typical goal for the search engine might be profit maximization, whereas the advertisers seek to maximize their own utility.

The setting for a mechanism design problem is specified by:

- *the number n of participants* (also called players, bidders or agents).
- *the set of possible outcomes \mathcal{A}* : For example, in a single-item auction, the outcome specifies who, if anyone, wins the item.
- *the space of valuation functions for the players*: Player i 's **valuation function** $v_i : \mathcal{A} \rightarrow \mathbb{R}$, maps outcomes to real numbers. The quantity $v_i(a)$ represents the “value” that i assigns to outcome $a \in \mathcal{A}$, measured in a common currency, such as dollars. This valuation function is player i 's *private information*, and is chosen from some specified function space V_i . For the single item auction case, $v_i(a)$ is simply v_i , player i 's value for the item, if a is the outcome in which player i wins the item, and 0 otherwise.
- *the objective of the designer*. In the single-item auction case, this could be, for example, to maximize profit. We will discuss other objectives below.

We consider the following type of **mechanism design problem**: Design a mechanism \mathcal{M} (a game of incomplete information) that takes as input the players valuation functions $\mathbf{v} = (v_1(\cdot), \dots, v_n(\cdot))$, and selects as output an outcome $a = a(\mathbf{v}) \in \mathcal{A}$, and a set of payments $p_i = p_i(\mathbf{v})$ (p_i is the payment by player i), such that, in equilibrium, the mechanism designer's objective is met (or approximately met).

To complete the specification of a mechanism design problem, we need to define the players payoff/utility model and the equilibrium notion of interest.

Before we do that, we introduce some notation that we will use throughout the chapter: Given any vector whose entries represent some quantity for each of the agents, such as, for example, the values of the agents $\mathbf{v} = (v_1, \dots, v_n)$,

we will use \mathbf{v}_{-i} to denote the elements of this vector excluding i , and we will interchangeably refer to the full vector as either \mathbf{v} , (v_1, \dots, v_n) or (v_i, \mathbf{v}_{-i}) .

Definition 6.3.4. The **quasi-linear utility** model linearly combines the players valuation for the outcome with their value for money, and assumes the goal of each agent is to maximize their expected utility. We use $u_i(\mathbf{b}|v_i)$, to represent player i 's expected utility when the mechanism \mathcal{M} is run on the input $\mathbf{b} = (b_i, \mathbf{b}_{-i})$, assuming player i 's actual valuation function is v_i . In a quasi-linear utility model this is

$$u_i(\mathbf{b}|v_i) = v_i(a(\mathbf{b})) - p_i(\mathbf{b}).$$

As for the equilibrium notion, we begin our discussion with the strongest possible notion, namely that it be a dominant strategy for each agent to report their valuation function truthfully.

Definition 6.3.5. We say that a mechanism \mathcal{M} is **truth-eliciting** if, for each player i , each valuation function $v_i(\cdot)$ and each possible report \mathbf{b}_{-i} of the other players, it is a dominant strategy for player i to report their valuation truthfully. Formally, for all \mathbf{b}_{-i} , all i , v_i , and b_i

$$u_i(v_i, \mathbf{b}_{-i}|v_i) \geq u_i(b_i, \mathbf{b}_{-i}|v_i)$$

As we saw, the second price auction for selling a single item is truth-eliciting.

6.4 Social Welfare Maximization

Consider a mechanism design problem where the design goal is to **maximize social welfare**, the total happiness of all the participants including the auctioneer. The happiness of the auctioneer is simply the total payment collected from the participants (minus any cost $C(\cdot)$ he incurs to implement that outcome), and the happiness of a player is their utility for that outcome. Summing these, we obtain that the social welfare of an outcome a is $(\sum_j v_j(a)) - C(a)$. (For most of this section, we assume that $C(a) = 0$. **FIX THIS.**) We use the following mechanism to solve this problem:

Definition 6.4.1. The **Vickrey-Clarke-Groves (VCG) mechanism**, illustrated in Figure 6.1, works as follows: The agents are asked to report their valuation functions. Say they report $\mathbf{b} = (b_1, \dots, b_n)$ (where b_i may or may not equal their true valuation v_i). The outcome $a^* = a(\mathbf{b})$ selected is the one that maximizes social welfare with respect to the reported valuations.

The payment $p_i(\mathbf{b})$ player i makes is the harm his presence causes others (with respect to the reported bids), formally:

$$p_i(\mathbf{b}) = \max_a \sum_{j \neq i} b_j(a) - \sum_{j \neq i} b_j(a^*).$$

The first term is the total reported value the other players would obtain if i was absent, and the term being subtracted is the total reported value the others obtain when i is present.




Fig. 6.1.

Exercise 6.4.2. Check that the Vickrey second price auction is a special case of the VCG mechanism.

Here are two other examples:

Example 6.4.3. Consider the outcome and payments for the VCG mechanism on example 6.3.1, assuming that the cities report truthfully. As the social welfare of each outcome is the sum of the values to each of the participants for that outcome (the final row in the following table), the social welfare maximizing outcome would be to build both roads.

	road A-C	road B-C	both	none
City A	5	0	5	0
City B	0	5	5	0
City C	9	9	15	0
Government	-10	-10	-20	0
Social welfare	4	4	5	0

What about the payments using VCG? For city A, the total value attained by others in A’s absence is 4 (road B-C only would be built), whereas with city A, the total value attained by others is 0, and therefore player A’s payment, the harm his presence causes others is 4. By symmetry, B’s payment is

the same. For city C, the total value attained by other's in C's absence is 0, whereas the total value attained by others in C's presence is -10, and therefore the difference, and C's payment is 10. Notice that the total payment is 18, whereas the government spends 20.

Example 6.4.4. Consider a search engine selling advertising slots on one of its pages. There are three advertising slots with clickthrough rates (probability that an individual viewing the web page will click on the ad) of 0.08, 0.03 and 0.01 respectively, and four advertisers whose values for a click are 10, 8, 2 and 1 respectively. We assume that the expected value for an advertiser to have his ad shown in a particular slot is his value times the clickthrough rate. Suppose the search engine runs a VCG auction in order to decide which advertiser gets which slot. The outcome and payments are shown in the figure.

The beauty of VCG payments is that they ensure that *the players' incentives are precisely aligned with the goal of the mechanism* since when all players report truthfully, that is, $\mathbf{b} = \mathbf{v}$, the outcome a^* selected is the one that maximizes social welfare with respect to the true valuations \mathbf{v} and as a result player i 's utility is:

$$\begin{aligned} u_i(\mathbf{v}) &= v_i(a^*) - p_i(\mathbf{v}) \\ &= v_i(a^*) + \sum_{j \neq i} v_j(a^*) - \max_a \sum_{j \neq i} v_j(a) \\ &= \sum_j v_j(a^*) - \max_a \sum_{j \neq i} v_j(a). \end{aligned}$$

exactly the same quantity the mechanism is trying to maximize (minus a term player i has no control over). Formally:

Theorem 6.4.5. *VCG is a truth-eliciting mechanism for maximizing social welfare.*

Proof. Fix the reports \mathbf{b}_{-i} of all agents except agent i (that may or may not be truthful). Suppose that agent i 's true valuation function is $v_i(\cdot)$. but he lies and reports $b_i(\cdot)$, resulting in outcome

$$a' = \operatorname{argmax}_a \sum_j b_j(a)$$

and payment

$$p_i(\mathbf{b}) = \max_a \sum_{j \neq i} b_j(a) - \sum_{j \neq i} b_j(a') = - \sum_{j \neq i} b_j(a') + C.$$

(Here $C = \max_a \sum_{j \neq i} b_j(a)$ is a constant that agent i 's report has no influence on.) Thus by lying, agent i 's utility is

$$\begin{aligned} u_i(\mathbf{b}|v_i) &= v_i(a') - p_i(b_i, \mathbf{b}_{-i}) \\ &= v_i(a') + \sum_{j \neq i} b_j(a') - C. \end{aligned}$$

Had he reported truthfully, the outcome

$$\tilde{a} = \operatorname{argmax}_a \left(v_i(a) + \sum_{j \neq i} b_j(a) \right)$$

would have been selected. Since

$$v_i(a') + \sum_{j \neq i} b_j(a') - C \leq v_i(\tilde{a}) + \sum_{j \neq i} b_j(\tilde{a}) - C,$$

we obtain

$$u_i(\mathbf{b}|v_i) \leq u_i(v_i, \mathbf{b}_{-i}|v_i).$$

□

The following example illustrates a few of the deficiencies of the VCG mechanism.

Example 6.4.6. Spectrum Auctions: In a spectrum auction, the government is selling licenses for the use of some band of electromagnetic spectrum in a certain geographic area. The participants in the auction are cell phone companies who need such licenses to operate. Company A has recently entered the market and needs two licenses in order to operate efficiently enough to compete with the established companies. Thus, A has no value for a single license, but values a pair of licenses at 1 billion dollars. Companies B and C are already well established and only seek to expand capacity. Thus, each one needs just one license and values that license at 1 billion.

Suppose the government runs a VCG auction to sell 2 licenses. If only companies A and B compete in the auction, the government revenue is 1 billion dollars (either A or B can win). However, if A, B and C all compete, then companies B and C will each receive a license, but pay nothing. Thus, VCG revenue is not necessarily monotonic in participation or bidder values.

A variant on this same setting illustrates another problem with the VCG

mechanism and that is susceptibility to collusion. Suppose that company A's preferences are as above, and companies B and C still only need one license each, but now they only value a license at 25 million. In this case, if companies B and C bid honestly, they lose the auction. However, if they collude and each bid 1 billion, they both win at a price of 0.

6.5 Win/Lose Mechanism Design

We turn now to the objective of profit maximization in settings where the outcomes divide the players into winners and losers. In such settings, the auctioneer is allocating goods or a service to some subset of the agents. Thus, outcomes are represented by binary vectors \mathbf{x} , where $x_i = 1$ if i is a winner and $x_i = 0$ otherwise. Since in these settings, all each agent cares about is whether they win or lose and their value for losing is 0, each agent's private information $v_i(\mathbf{x})$ is defined by a single number $v_i \in [0, \infty)$, such that $v_i(\mathbf{x})$ is v_i if x_i is 1 and 0 otherwise. Thus, in mechanisms for these problems, the agents are asked to simply report their values v_i for allocation. We call their reports **bids**, and observe that agents may or may not bid truthfully. Here are two examples of single-parameter problems:

Example 6.5.1. Concert tickets: A well-known singer is planning a concert in a 10,000 seat arena to raise money for her favorite charity. Her goal is to raise as much money as possible and hence she sells the tickets by auction. In this scenario, the singer is the auctioneer and has a set of 10,000 identical items (tickets) she wishes to sell. Assume that there are 100,000 potential buyers, each interested in exactly one ticket, and each with his or her own private value for the ticket. If the singer wants to maximize social welfare, that is, make sure the tickets end up in the hands of the buyers who value them the most, she can simply run the VCG mechanism from the previous section. It is easy to check that this will result in selling the tickets to the top 10,000 bidders at the price bid by the 10,001st bidder. But if there are a small number (much smaller than 10,000) of buyers that are willing to pay a large sum for a ticket, and vast numbers of buyers willing to pay a tiny amount for a ticket, this could result in a very low profit. This raises the question: what auction format should the singer use if she wishes to make as much money as possible?

Google IPO auction?

Example 6.5.2. Exclusive markets: A massage therapist (the auctioneer in this example) is trying to decide whether to set up shop in Seattle or New

York. To gauge the profit potential in the two different locations, he runs an auction asking a set of individuals from each city to submit a bid indicating how much they will pay for a massage. The constraint in the design of the auction is that winners can only be from one of the markets.

The question we turn to next is how to design a truth-eliciting, profit-maximizing auction in win/lose settings.

6.6 Profit Maximization in Win/Lose Settings

We begin by characterizing truth-eliciting mechanisms in win/lose settings. The following simple theorem is fundamental:

Theorem 6.6.1 (Characterization for Truthfulness). *Let \mathcal{A} be a deterministic mechanism for a win/lose setting. \mathcal{A} is truth-eliciting if and only if, for every agent i and bids \mathbf{b}_{-i}*

- (i) *There is a value t_i such that i wins if her bid is above t_i and loses if her bid is below t_i . (At t_i it can go either way.) The value t_i is called i 's **threshold bid** given \mathbf{b}_{-i} .*
- (ii) *Agent i 's payment is t_i if agent i is a winner and 0 otherwise.*

Proof. We use the following simplified notation in the proof: When i and \mathbf{b}_{-i} are understood, we use $x(b)$ to denote $x_i(b, \mathbf{b}_{-i})$ and $p(b)$ to denote $p_i(b, \mathbf{b}_{-i})$. We also use $u_i(b|v) = vx(b) - p(v)$ to represent agent i 's utility when he bids b and his true value is v . We now proceed to the proof.

It is an easy exercise to check that if conditions (i) and (ii) are satisfied, then the mechanism is truthful. (These conditions are precisely the conditions that made the Vickrey 2nd price auction truthful.)

For the converse, fix \mathbf{b}_{-i} . We observe that \mathcal{A} is truth-eliciting only if for every \mathbf{b}_{-i} , v and w

$$u_i(v|v) = vx(v) - p(v) \geq vx(w) - p(w) = u_i(w|v)$$

and

$$u_i(w|w) = wx(w) - p(w) \geq wx(v) - p(v) = u_i(v|w).$$

Adding these two inequalities, we obtain that for all v and w in $[0, \infty)$

$$(v - w)(x(v) - x(w)) \geq 0.$$

Thus, if $v \geq w$, then $x(v) \geq x(w)$, and $x(z)$ is monotone nondecreasing in z . In other words, condition (i) holds.

As for the payment rule, we assume wlog that the minimum payment p_{\min}

agent i can make is 0; if not, reducing all payments by p_{\min} doesn't change the incentive properties of the mechanism. This implies that no matter what i 's value is, if she loses, she pays 0. Otherwise, she would lie and bid so that her payment is 0. Thus, we have only to argue that agent i 's payment is t_i if she wins. To this end, observe that if there are two winning bids v and v' for which the payments are different, say lower for v , then if her value was v' she would increase her utility by lying and bidding v . Thus, all the winning bids for agent i result in the same payment. Moreover, this payment p has to be at most t_i , since otherwise, a bidder with value v such that $t_i < v < p$ would have an incentive to lie so as to become a loser. On the other hand, if this payment p is strictly below t_i , an agent with value v such that $p < v < t_i$ would have an incentive to lie and become a winner.

□

Remark. This characterization implies that, operationally, a truth-eliciting auction in win/lose settings consists of making an offer at a price of t_i to bidder i , where t_i is a function of all other bids \mathbf{b}_{-i} , but is *independent* of b_i . This offer is then accepted or rejected by i depending on whether b_i is above or below t_i . Thus, truth-eliciting auctions in single-parameter settings are often said to be **bid-independent**.

6.6.1 Profit maximization in digital goods auctions

We now show how to apply Theorem ?? to the design of profit-maximizing digital goods auctions. A *digital goods auction* is an auction to sell digital goods such as mp3's, digital video, pay-per view TV, etc. The unique aspect of digital goods is that the cost of reproducing the items is negligible and therefore the auctioneer effectively has an unlimited supply of the items. This means that there is no constraint on how many of the items can be sold, or to whom.

For digital goods auctions, the VCG mechanism allocates to all of the bidders, and charges them all nothing! Thus, while VCG perfectly maximizes social welfare, it can be disastrous when the goal is to maximize profit.

In this section, we present a truth-eliciting auction that does much better. Specifically, we present an auction that always gets within a factor of four of the profit obtained by the auction that sells the items at a fixed price.

Definition 6.6.2. The **optimal fixed price profit** that can be obtained from bidders with bid vector \mathbf{b} is

$$\text{OFP}(\mathbf{b}) = \max_p \{p \cdot (\text{the number of bids in } \mathbf{b} \text{ at or above } p)\},$$

and the **optimal fixed price** is

$$p^*(\mathbf{b}) = \operatorname{argmax}_p \{p \cdot (\text{the number of bids in } \mathbf{b} \text{ at or above } p)\}.$$

If we knew the true values \mathbf{v} of the agents, a profit of $\text{OFP}(\mathbf{v})$ would be trivial to obtain. We would just offer the price $p^*(\mathbf{v})$ to all the bidders, and sell at that price to all bidders whose values are above p^* . But we can't do this truthfully.

Exercise 6.6.3. Show that no truth-eliciting auction can obtain a profit of $\text{OFP}(\mathbf{v})$ for every bid vector \mathbf{v} .

The following auction is perhaps the first thing one might try as a truth-eliciting alternative:

The Deterministic Optimal Price Auction (DOP): For each bidder i , compute $t_i = p^*(\mathbf{b}_{-i})$, the optimal fixed price for the remaining bidders, and use that as the threshold bid for bidder i .

Unfortunately, this auction does not work well, as the following example shows.

Example 6.6.4. Consider a group of bidders of which 11 bidders have value 100, and 1001 bidders have value 1. Then the best fixed price is 100 – at that price 11 items can be sold for a total profit of 1100.

Unfortunately, if we run the DOT auction on this bid vector, then for each bidder of value 100, the threshold price that will be used is 1, whereas for each bidder of value 1, the threshold price is of value 100, for a total profit of only 11!

In fact, the DOT auction can obtain arbitrarily poor profit compared to the optimal fixed price profit. Moreover, it is possible to prove that *any* deterministic truth-eliciting auction that treats the bidders symmetrically will fail to consistently obtain a constant fraction of the optimal fixed price profit. The key to overcoming this problem is to use randomization. First though, we show how to solve a somewhat easier problem.

6.6.2 Profit Extraction

Suppose that we lower our sights and rather than shooting for the best fixed price profit possible for each input, we set a specific target, say \$1000, and ask if we can design an auction that guarantees us a profit of \$1000, *when the bidders can “afford it”*. Formally:

Definition 6.6.5. A **digital goods profit extractor**, with parameter T , denoted by $\text{pe}_T(\cdot)$, is a *truth-eliciting* auction that, given a set of sealed bids \mathbf{b} and a target profit T , is guaranteed to obtain a profit of T as long as the optimal fixed price profit $\text{OFP}(\mathbf{b})$ is at least T . If the optimal fixed price profit $\text{OFP}(\mathbf{b})$ is less than T , there is no guarantee, and the profit extractor could, in the worst case, obtain no profit.

It turns out that such an auction is easy to design:

Definition 6.6.6 (A Profit Extractor): The digital goods profit extractor $\text{pe}_T(\mathbf{b})$ with target profit T sells to the largest group of k bidders that can equally share the cost T and charges each T/k .

Using Theorem ??, it is straightforward to verify that:

Lemma 6.6.7. *The digital goods profit extractor pe_T is truthful, and guarantees a profit of T on any \mathbf{b} such that $\text{OFP}(\mathbf{b}) \geq T$.*

6.6.3 A profit-making digital goods auction

The following auction is near optimal:

Definition 6.6.8 (RSPE). The *Random Sampling Profit Extraction* auction (RSPE) works as follows:

- Randomly partition the bids \mathbf{b} into two by flipping a fair coin for each bidder and assigning her to \mathbf{b}' or \mathbf{b}'' .
- Compute the optimal fixed price profit for each part: $T' = \text{OFP}(\mathbf{b}')$ and $T'' = \text{OFP}(\mathbf{b}'')$.
- Run the profit extractors: $\text{pe}_{T'}$ on \mathbf{b}'' and $\text{pe}_{T''}$ on \mathbf{b}' .

Our main theorem is the following:

Theorem 6.6.9. *The Random Sampling Profit Extraction (RSPE) auction is truthful, and for all bid vectors \mathbf{v} for which there are at least two values at or above $\mathbf{p}^*(\mathbf{v})$, RSPE obtains at least $1/4$ of the optimal fixed profit $\text{OFP}(\mathbf{v})$.*

Proof. The fact that the RSPE auction is truth-eliciting is straightforward since it is simply randomizing over truth-eliciting auctions, one for each possible partition of the bids. (Note that any target profit used in step 3 of the auction is independent of the bids to which it is applied.) So we have only to lower bound the profit obtained by RSPE on each input \mathbf{v} . The crucial observation is that for any particular partition of the bids, the profit of RSPE is at least $\min(T', T'')$. This follows from the fact that if,




Fig. 6.2. This figure illustrates a possible execution of the RSPE auction when the entire set of bids is $(20, 10, 10, 8, 5, 5, 5, 5, 5, 3, 1)$.

say $T' \leq T''$, then $\text{OFP}(\mathbf{b}'') = T''$ is large enough to ensure the success of $\text{pe}_{T'}(\mathbf{b}'')$, namely the extraction of a profit of T' .

Thus, we just need to analyze $E(\min(T', T''))$.

Assume that $\text{OFP}(\mathbf{b}) = k\mathbf{p}^*$ has with $k \geq 2$ winners at price \mathbf{p}^* . Of the k winners in OFP , let k' be the number of them that are in \mathbf{b}' and k'' the number that are in \mathbf{b}'' . Thus, $T' \geq k'\mathbf{p}^*$ and $T'' \geq k''\mathbf{p}^*$. Therefore

$$\begin{aligned} \frac{E(\text{RSPE}(\mathbf{b}))}{\text{OFP}(\mathbf{b})} &= \frac{E(\min(T', T''))}{k\mathbf{p}^*} \\ &\geq \frac{E(\min(k'\mathbf{p}^*, k''\mathbf{p}^*))}{k\mathbf{p}^*} \\ &= \frac{E(\min(k', k''))}{k} \\ &\geq \frac{k/4}{k} = 1/4. \end{aligned}$$

The last inequality follows from the fact that, for $k \geq 2$,

$$E(\min(k', k'')) = \sum_{0 \leq i \leq k} \min(i, k-i) \binom{k}{i} 2^{-k} = k \left(\frac{1}{2} - \binom{k-1}{\lfloor \frac{k}{2} \rfloor} 2^{-k} \right) \geq \frac{k}{4}.$$

□

6.6.4 With priors: the Bayesian setting

We now return to the Bayesian setting discussed in the introduction. Here, it is assumed that each agent's private value v_i is drawn independently with distribution function F_i and these *prior distributions*, or *priors* for short, are known to both the auctioneer and all the players. Only player i , however, knows her actual draw v_i from distribution F_i . We also assume for now that the distribution functions F_i have support on $[0, \infty)$, and are continuous and differentiable.

In such an auction, each player has a strategy:

Definition 6.6.10. A **strategy** in an auction is a mapping $s : [0, \infty) \rightarrow [0, \infty)$ which specifies agent i 's bid $b = s(v)$ when her value is v . A player chooses her strategy knowing the rules of the auction to be run, and, in this Bayesian setting, the prior distributions from which the other player's values are drawn.

To describe our equilibrium notion, we will use the following notation, in which an expectation or probability taken over subscript “ $-i$ ” means that it is computed with respect to the draws of $v_j \sim F_j$ for each $j \neq i$:

- The probability that agent i receives allocation when all players but i employ strategies \mathbf{s}_{-i} , and agent i bids $s_i(b_i)$ is denoted by

$$\bar{x}_i(b_i) = P_{-i}[x_i(s_i(b_i), \mathbf{s}_{-i}(\mathbf{v}_{-i})) = 1].$$

- The expected payment that agent i will have to make when all players but i employ strategies \mathbf{s}_{-i} and agent i bids $s_i(b_i)$ is denoted by

$$\bar{p}_i(b_i) = E_{-i}[p_i(s_i(b_i), \mathbf{s}_{-i}(\mathbf{v}_{-i}))].$$

- Agent i 's expected utility when all players but i employ strategies \mathbf{s}_{-i} , her value is v_i and she bids $s_i(b_i)$ is denoted by

$$\bar{u}_i(b_i|v_i) = E_{-i}[u_i(s_i(b_i), \mathbf{s}_{-i}(\mathbf{v}_{-i})|v_i)].$$

By definition

$$\bar{u}_i(b_i|v_i) = v_i \bar{x}_i(b_i) - \bar{p}_i(b_i).$$

We use the following equilibrium notion.

Definition 6.6.11. Given a mechanism \mathcal{A} , we say that strategies s_i for $1 \leq i \leq n$ are in **Bayes-Nash equilibrium** if for each agent i and value v_i , strategy s_i is a best response, in expectation over the draws of other players

values, to strategies \mathbf{s}_{-i} for each i . Formally, this means that for all i and values v_i ,

$$\bar{u}_i(b_i|v_i)$$

is maximized when $b_i = v_i$ (which means that agent i bids $s_i(v_i)$).

The key theorem we will need in this setting with priors is the following.

Theorem 6.6.12 (Characterization for Bayes-Nash Equilibrium). *Let \mathcal{A} be a mechanism for a single-parameter allocation problem with priors, where $v_i \sim F_i$ for each i . Then an onto strategy profile $\mathbf{s} = (s_1, \dots, s_n)$ for the agents is in Bayes-Nash equilibrium if and only if*

- (i) *The probability of allocation $\bar{x}_i(b)$ is monotone non-decreasing in b .*
- (ii) *The payment rule is determined by the allocation rule up to an additive constant as follows:*

$$\bar{p}_i(b) = b\bar{x}_i(b) - \int_0^b \bar{x}_i(z)dz + \bar{p}_i(0).$$




Fig. 6.3. This figure shows the monotone non-decreasing curve of $\bar{x}(\cdot)$. The grey area is $\bar{p}(v)$ and the purple area is $\bar{u}(v|v)$.

Remarks:

- (i) To simplify the exposition, we will henceforth assume that $p_i(0) = 0$

for all i . All statements proceed mutatis mutandis if this is not the case.

- (ii) If conditions (i) and (ii) hold then player i 's expected utility in Bayes-Nash equilibrium satisfies

$$\bar{u}(v|v) = \int_0^v \bar{x}_i(z) dz.$$

Figure 6.3 illustrates this fact.




Fig. 6.4. In this figure, the area above the curve $\bar{x}(\cdot)$ up to the line $y = \bar{x}(w)$ is the payment $\bar{p}(w)$ (the teal part together with A , the green part). The picture shows that that $\bar{u}(w|v) = \bar{u}(v|v) - A$. A similar picture shows that $\bar{u}(w|v) \leq \bar{u}(v|v)$ when $w < v$.

Proof. We present a proof by picture that, if conditions (i) and (ii) hold and the strategy profile is onto (for every b , there exists a v such that $s(v) = b$), then the strategy profile s is in Bayes-Nash equilibrium. To see this, observe that since the strategy profile is onto, any deviation by player i with value v from the strategy s_i corresponds to playing strategy s_i with respect to a modified value w . Thus, we need only show that for all v and w ,

$$\bar{u}(v|v) \geq \bar{u}(w|v). \tag{6.1}$$

This is illustrated in Figure 6.4.

For the converse, suppose that inequality (6.1) holds for all v and w .

Then combining $\bar{u}(v|v) \geq \bar{u}(w|v)$ and $\bar{u}(w|w) \geq \bar{u}(v|w)$, and using the fact that $\bar{u}(w|v) = v\bar{x}(w) - \bar{p}(w)$ for all v and w , we obtain (as in the proof of Theorem 6.6.1) that

$$(v - w)(\bar{x}_i(v) - \bar{x}_i(w)) \geq 0, \quad (6.2)$$

meaning that $\bar{x}(z)$ is monotone in z and thus condition (i) holds. It also follows that

$$w(\bar{x}(v) - \bar{x}(w)) \leq \bar{p}(v) - \bar{p}(w) \leq v(\bar{x}(v) - \bar{x}(w)). \quad (6.3)$$

for all v and w . Letting $v = w + dw$ and taking the limit as $dw \rightarrow 0$, we obtain

$$\bar{p}'(v) = vx'(v).$$

Integrating by parts yields condition ((ii)).

□

6.6.5 Revenue Equivalence

One of the most important results in auction theory is a simple corollary of Theorem 6.6.12:

Theorem 6.6.13 (Revenue Equivalence). *Consider any two auctions \mathcal{A} and \mathcal{A}' for a win/lose setting, where players values v_i are independent and identically distributed. Suppose that \mathbf{s} is a Bayes-Nash equilibrium for auction \mathcal{A} and \mathbf{s}' is a BN equilibrium for auction \mathcal{A}' , and that $\mathbf{x}^{\mathcal{A}}(\mathbf{s}(\mathbf{v})) = \mathbf{x}^{\mathcal{A}'}(\mathbf{s}'(\mathbf{v}))$ for all \mathbf{v} . Then \mathcal{A} and \mathcal{A}' have the same expected revenue in equilibrium (assuming that $\bar{p}_i(0) = 0$ for all i in both auctions).*

Proof. By Theorem 6.6.12, the payment to the auctioneer is determined by the allocation rule $\bar{x}_i(\cdot)$. Thus, auctions with the same allocation rule (and with $\bar{p}_i(0) = 0$ for all i), obtain the same expected revenue from each agent, and hence the same expected revenue overall in equilibrium. □

Corollary 6.6.14. *Consider the following three auctions for selling a single item: the first price auction, the second price auction and the all-pay auction, and suppose that each of them is run with agents such that $v_i \sim U[0, 1]$ for all i . Then in Bayes-Nash equilibrium, all three of these auctions have the same expected revenue.*

Proof. By definition $p_i(0) = 0$ for all three auctions. Thus, it will suffice to show that all three auctions have a Bayes-Nash equilibrium strategy in

which the allocation is the same, specifically, the item is allocated to the player with the highest value.

To this end, we show that all three auctions have a BN equilibrium for which the strategy profile is symmetric, ($s_i(v) = s(v)$ for all i), and $s(v)$ is monotone increasing.

For the second price auction, this is immediate from the fact that $s(v) = v$ is a dominant strategy, since every dominant strategy profile is also a BN equilibrium strategy profile.

To find such equilibrium strategies for the first-price auction and all-pay auctions, we observe that if \mathbf{s} is in BNE then $\bar{u}_i(b|v)$ is maximized when $b = v$ for all v . Thus, we have that

$$\frac{\partial \bar{u}(b|v)}{\partial b} \Big|_{(b=v)} = 0. \quad (6.4)$$

Applying this for each of the auctions enables us to derive a differential equation which we can then solve to find the equilibrium strategies.

- For the first-price auction,

$$\bar{u}(b|v) = b^{n-1}(v - s(b)).$$

Applying Equation (6.4), we obtain that

$$(n-1)v^{n-2}(v - s(v)) - v^{n-1}s'(v) = 0,$$

or

$$s'(v) = (n-1) \left(1 - \frac{s(v)}{v} \right).$$

A solution to this differential equation is

$$s(v) = \left(\frac{n-1}{n} \right) v,$$

which can easily be verified to be a BN equilibrium strategy profile for the first-price auction.

- For the all-pay auction

$$\bar{u}(b|v) = b^{n-1}(v - s(b)) + (1 - b^{n-1})(-s(b))$$

Again applying Equation (6.4), we obtain the differential equation

$$s'(v) = (n-1)v^{n-1}$$

which has as a solution

$$s(v) = \frac{n-1}{n} v^n.$$

Again, this can easily be verified to be a BN equilibrium strategy profile for the all-pay auction.

Thus, these two auctions have equilibrium strategy profiles that are symmetric and monotone increasing, from which we conclude that in all three auctions, the item is allocated to the player with the highest value. \square

6.6.6 The revelation principle

We now have almost everything we need in order to derive the profit-maximizing, or *optimal* auction in win/lose settings with priors. The final result we will use is the **revelation principle**, a simple but terribly useful observation that enables us to significantly restrict the class of auctions we consider in our search for the optimal auction.




Fig. 6.5.

Theorem 6.6.15 (The Revelation Principle). *Let \mathcal{M} be a mechanism and \mathbf{s} a strategy profile in Bayes-Nash equilibrium with respect to \mathcal{M} . Then there is another mechanism \mathcal{M}' for which (a) bidding truthfully (i.e. using strategy $s_i(v) = v$ for all i) is a Bayes-Nash equilibrium, and (b) \mathcal{M}' has the same input/output relationship as \mathcal{M} .*

Proof. The mechanism \mathcal{M}' operates as follows: On each input \mathbf{b} , \mathcal{M}' computes $\mathbf{s}(\mathbf{b}) = (s_1(b_1), s_2(b_2), \dots, s_n(b_n))$, and then runs \mathcal{M} on $\mathbf{s}(\mathbf{b})$ to compute the output and payments. (See Figure 6.5.) It is straightforward to check that if \mathbf{s} is in Bayes-Nash equilibrium for \mathcal{M} , then bidding truthfully is a Bayes-Nash equilibrium for \mathcal{M}' . \square

Definition 6.6.16. If bidding truthfully is a Bayes-Nash equilibrium for mechanism \mathcal{M} , then \mathcal{M} is said to be **Bayes-Nash incentive compatible (BIC)**.

6.6.7 Optimal Auctions

Example 6.6.17. Consider an auctioneer with a single item to sell to an agent whose private value is publicly known to be drawn from distribution F . What price should the auctioneer offer the agent to maximize her profit? If she offers a price of v , the agent will accept the offer if her value is at least that price, i.e. with probability $1 - F(v)$. Thus, the auctioneer should choose the price to maximize $R(v) = v(1 - F(v))$, her expected revenue at price v . If $R(v)$ is concave, then $R(v)$ is maximized when $R'(v) = 0$, i.e., at the value at which the agents “marginal revenue” is 0. Note that if $R(v)$ is concave, then $R'(v)$ is monotone nonincreasing in v . ♣

We now derive the famous *Myerson optimal mechanism* for maximizing auctioneer profit in win/lose settings where players’ priors are drawn from known distributions.

By the revelation principle (Theorem 6.6.15), we need consider optimizing only over BIC mechanisms. Moreover, by Theorem 6.6.12, we need consider only how to determine the allocation rule, since the payment rule is determined by the allocation rule (and we will fix $p_i(0) = 0$ for all i).

Consider a mechanism \mathcal{M} with truthtelling ($s_i(v) = v$ for all i) a Bayes-Nash equilibrium, and suppose that its allocation rule is $\mathbf{x}(\mathbf{b})$. Then the goal of the auctioneer is to choose $\mathbf{x}(\mathbf{b})$ to maximize

$$\mathbb{E}_i \left[\sum_i \bar{p}(v_i) \right].$$

It will be convenient to write everything in terms of the probability of sale q at a given price. If a player whose value is drawn from distribution F is offered a price of v , he will accept this price (i.e. have a value higher than this price) with probability.

$$q(v) = 1 - F(v).$$

(We will assume that F is invertible and refer also to the value at a given probability of sale q which is $v(q) = F^{-1}(1 - q)$.) Using this notation, the payment rule (condition (ii) from Theorem 6.6.12) can be rewritten as fol-

lows:

$$\bar{p}_i(q) = - \int_q^1 \bar{x}'_i(\ell)v(\ell)d\ell + \bar{p}_i(1).$$

Notice that since $\bar{x}_i(v)$ is monotone non-decreasing in v , $\bar{x}_i(q)$ is monotone non-increasing in q . Moreover, drawing v from the distribution F is equivalent to drawing q from $U[0, 1]$.

Fix \mathbf{v}_{-i} . Then the expected revenue from agent i is

$$\mathbb{E}_q [p(v(q), \mathbf{v}_{-i})] = - \int_0^1 \int_q^1 x'(\ell)v(\ell)d\ell dq.$$

Reversing the order of integration, we get

$$\begin{aligned} &= - \int_0^1 x'(\ell)v(\ell) \int_0^\ell dq d\ell \\ &= - \int_0^1 x'(\ell)\ell v(\ell)d\ell \\ &= - \int_0^1 x'(\ell)R(\ell)d\ell. \end{aligned}$$

where $R(q) = qv(q)$ is the profit from offering a price of v to an agent whose value is drawn from F . Finally, integrating by parts, we obtain

$$= - \int_0^1 x(q)R'(q)dq.$$

In the absence

Lemma 6.6.18. *Suppose that \mathcal{A} has allocation rule $x(\mathbf{v})$ and is Bayes-Nash incentive-compatible auction when $v_i \sim F_i$ for all i and $q_i = 1 - F_i(v_i)$. Let $R_i(q_i) = q_i v_i(q_i)$, the profit the auctioneer would make if he offered agent i a price of $v_i(q_i)$. Then the expected profit of the auctioneer is*

$$E_{\mathbf{v}}(\text{profit of } \mathcal{A}) = -E_{\mathbf{q}} \left(\sum_i x_i(\mathbf{q}(\mathbf{v}))R'_i(q_i) \right) \quad (6.5)$$

Notice that we now have a formula for the expected profit of *any* BIC auction in terms of its allocation rule! We can also now see the path to designing an auction that maximizes expected profit in equilibrium. If we want to maximize equation (6.5), we should choose our allocation rule $\mathbf{x}(\mathbf{v})$ so that pointwise, on each input \mathbf{v} , it chooses the feasible outcome that maximizes the right hand side of (6.5), in other words, *choose the allocation that maximizes marginal revenue*. That is the best we can do!

Unfortunately, such an allocation rule is not always monotone, and hence does not always yield a BIC mechanism. There is, however, a very large class of settings for which it is a monotone allocation rule: whenever $R'_i(q_i)$ is monotone non-increasing in q_i for all i .

Definition 6.6.19. A probability distribution with distribution function F is *regular* if $R'_i(q_i)$ is monotone nonincreasing in q_i .

For the purposes of summarizing the resulting mechanism, it will be convenient to switch back from “probability space”.

Definition 6.6.20. For agent i whose value v is drawn from distribution F_i , define the **virtual value** of agent i to be $\phi_i(v) = v - \frac{1-F_i(v)}{f_i(v)}$.

Exercise 6.6.21. Show that $\phi_i(v) = R'(q(v))$, and that $\mathbb{R}'(q)$ is monotone nonincreasing in q if and only if $\phi_i(v)$ is monotone nondecreasing in v .

Thus a probability distribution is regular if the corresponding virtual values are monotone nondecreasing in v . Many common distributions, such as Gaussian, exponential and even many heavy-tailed distributions, are regular.

We can now summarize Myerson’s optimal mechanism (for regular distributions):

Definition 6.6.22 (Myerson’s mechanism for regular distributions (Mye)). The Myerson mechanism for regular distributions is defined by the following steps:

- (i) Solicit a bid vector \mathbf{b} from the agents.
- (ii) Compute the virtual value b'_i of each agent, where $b'_i = \phi_i(b_i)$, and choose the feasible win/loss vector \mathbf{x} for which $\sum_i b'_i x_i$ is maximized.
- (iii) Charge each winning agent her threshold bid.

Theorem 6.6.23. *If each distribution F_i is regular, Mye is truth-eliciting and maximizes the auctioneer’s expected profit over values drawn from \mathbf{F} . In other words, for any BIC mechanism \mathcal{M} :*

$$\mathbb{E}_{\mathbf{v}} [\text{Mye}(\mathbf{v})] \geq \mathbb{E}_{\mathbf{v}} [\mathcal{M}(\mathbf{v})]$$

Remark. Note that not only is Myerson’s mechanism BIC – it is actually dominant strategy truthful! Indeed it meets the conditions of Theorem 6.6.1.

7

Stable matching

7.1 Introduction

Stable matching was introduced by Gale and Shapley in 1962. The problem is described as follows.

Suppose we have n men and n women. Every man has a preference order over the n women, while every woman also has a preference order over the n men. A **matching** is a one-to-one mapping between the men and women, and it is **perfect** if all men and women are matched. A matching M is **unstable** if there exists a man and a woman who are not matched to each other in M , but prefer each other to their partners in M . Otherwise, the matching is called **stable**.




Fig. 7.1.

Consider the following example with three men x , y and z , and three women a , b and c . Their preference lists are:

$$x : a > b > c, \quad y : b > c > a, \quad z : a > c > b.$$

$$a : y > z > x, \quad b : y > z > x, \quad c : x > y > z.$$

Then, $x \longleftrightarrow a$, $y \longleftrightarrow b$, $z \longleftrightarrow c$ is an unstable matching, since z and a prefer each other to their partners.

Our questions are, whether there always exist stable matchings and how can we find one.

7.2 Algorithms for finding stable matchings

The following algorithm which is called the **men-proposing algorithm** was introduced by Gale and Shapley.

- (i) Initially each woman is not tentatively matched.
- (ii) Each man proposes to his most preferred woman.
- (iii) Each woman evaluates her proposers, including the man she is tentatively matched to, if there is one, and rejects all but the most preferred one. She becomes tentatively matched to this latter man.
- (iv) Each rejected man proposes to his next preferred woman.
- (v) Repeat step (ii) and (iii) until each woman has a tentative match. At that point the tentative matches become final.




Fig. 7.2. Arrows indicate proposals, cross indicates rejection.




Fig. 7.3. Stable matching is achieved in the second stage.

Similarly, we could define a women-proposing algorithm.

Theorem 7.2.1. *The men-proposing algorithm yields a stable matching.*

Proof. First, observe that the algorithm terminates, because a man proposes to each woman at most once, and he can only make a total of n proposals. In the worst case, a man is rejected every round, and thus the number of rounds is upper bounded by n^2 .

Next, we argue that when the algorithm terminates, a perfect matching of

men to women has been found. We claim that from the first time a woman is proposed to, she remains tentatively matched for the rest of the execution of the algorithm (and permanently matched at the end). If the algorithm terminates without finding a perfect matching, then some man has been rejected by all women. But if some man ends up unmatched, then some woman, to whom he proposed at some point, is unmatched as well. This is a contradiction to the previous claim. Hence, the algorithm terminates with a perfect matching, which we call M .

Finally, we prove that M is a stable matching. To this end, consider any man Bob and woman Alice, not matched to each other, such that Bob prefers Alice to his match in M . This means that he proposed to Alice before he proposed to his final match, and she, at some point, rejected him. But whenever a woman rejects a man during the execution of the algorithm, she rejects him for someone she prefers over him. Moreover, as discussed above, her tentative matches just get better and better over time, from her perspective. Thus, it can't be that Alice prefers Bob to her final match in M . \square

7.3 Properties of stable matchings

We say a woman a is **attainable** for a man x if there exists a stable matching M with $M(x) = a$.

Theorem 7.3.1. *Let M be the stable matching produced by Gale-Shapley men-proposing algorithm. Then,*

- (a) *For every man i , $M(i)$ is the most preferred attainable woman for i .*
- (b) *For every woman j , $M^{-1}(j)$ is the least preferred attainable man for j .*

Proof. We prove (a) by contradiction. Suppose that M does not match each man with his most preferred attainable woman. Consider the first time during the execution of the algorithm that a man m is rejected by his most preferred attainable woman w , and suppose that w rejects m at that moment for m' who she prefers to m . Since this is the first time a man is rejected by his most preferred attainable woman, we know that m' likes w at least as much as his most preferred attainable woman.

Also, since w is m 's most preferred attainable woman, there is another stable matching M' in which they are matched. In M' , m' is matched to someone other than w . But now we have derived a contradiction: m' likes w at least as much as his most preferred attainable woman and hence more than his match in M' and w prefers m' to m . Thus M' is unstable.

We also prove part (b) by contradiction. Suppose that in M , woman w ends up matched to a man m she prefers over her least preferred attainable man m' . Then there is another stable matching M' in which m' and w are matched, and m is matched to a different woman. Then in M' , w prefers m to her match m' . Also, by part (a), in M , m is matched to his most preferred attainable woman. Thus, m prefers w to the woman he is matched with in M' , which is a contradiction to the stability of M' . \square

Corollary 7.3.2. *If Alice is assigned to the same man in both the man-proposing and the woman-proposing version of algorithms, then this is the only attainable man for her.*

7.4 A special preference order case

Suppose we seek stable matchings for n men and n women with preference order determined by a matrix $A = (a_{i,j})_{n \times n}$ where all entries in each row are distinct, and all entries in each column are distinct. If in the i^{th} row of the matrix, we have

$$a_{i,j_1} < a_{i,j_2} < \cdots < a_{i,j_n},$$

then the preference order of man i is: $j_1 > j_2 > \cdots > j_n$. Similarly, if in the j^{th} column, we have

$$a_{i_1,j} < a_{i_2,j} < \cdots < a_{i_n,j}$$

then the preference order of woman j is: $i_1 > i_2 > \cdots > i_n$.

Lemma 7.4.1. *In this case, there exists a unique stable matching.*

Proof. By Theorem 7.3.1, we get that the men-proposing algorithm produces a stable matching which maximizes $\sum_i a_{i,M(i)}$ among all the stable matchings M . Moreover, this stable matching reaches the unique maximum of $\sum_i a_{i,M(i)}$. Similarly, the women-proposing algorithm produces a stable matching which maximizes $\sum_j a_{M^{-1}(j),j}$ among all stable matchings M . Thus the stable matchings produced by the two algorithms are exactly the same. By Corollary 7.3.2, there exists a unique stable matching. \square

Exercises

- 7.1 There are 3 men, called a, b, c and 3 women, called x, y, z , with the following preference lists (most preferred on left):

$$\begin{array}{ll} \text{for } a : & x > y > z & \text{for } x : & c > b > a \\ \text{for } b : & y > x > z & \text{for } y : & a > b > c \\ \text{for } c : & y > x > z & \text{for } z : & c > a > b \end{array}$$

Find the stable matchings that will be produced by the men-proposing and by the women-proposing Gale-Shapley algorithm.

- 7.2 Consider an instance of the stable matching problem, and suppose that M and M' are two distinct stable matchings. Show that the men who prefer their match in M to their match in M' are matched in M to women that prefer their match in M' to their match in M .
- 7.3 Give an instance of the stable matching problem in which, by lying about her preferences during the execution of the Gale-Shapley algorithm, a woman can end up with a man that she prefers over the man she would have ended up with had she told the truth.
- 7.4 Consider using stable matching in the National Resident Matching Program, for the problem of assigning medical residents to hospitals. In this setting, there are n hospitals and m students that can be assigned as medical residents. Each hospital has a certain number of positions for residents, say o_i for hospital i . Suppose also that $m > \sum_i p_i$, i.e., there is an oversupply of students. Each hospital has a ranking of all the students, and each student has a ranking of all the hospitals.

Construct an assignment of students to hospitals such that each student is assigned to at most one hospital, no hospital is assigned more students than it has slots, and the assignment is stable in the sense that: (a) there is no student s and hospital h that are not matched, and for which hospital h prefers s to some other student s' assigned to h , and s prefers h to the hospital she was assigned (or she simply wasn't assigned).

- 7.5 Consider the following integer programming[†] formulation of the stable matching problem. To describe the program, we use the following notation. Let m be a particular man and w a particular women. Then $j >_m w$ represents the set of all women j that m prefers over w , and $i >_w m$ represents the set of all men i that w prefers over m .

[†] In Section 2.7 we introduced linear programming. Integer programming is linear programming in which the variables are required to take integer values.

In the following program the variable x_{ij} will be selected to be 1 if man i and woman j are matched in the matching selected:

$$\begin{aligned}
 & \text{maximize} && \sum_{i,j} x_{ij} \\
 \text{subject to} &&& \sum_j x_{m,j} \leq 1 \text{ for all men } m && \text{(E7.1)} \\
 &&& \sum_i x_{i,w} \leq 1 \text{ for all women } w \\
 &&& \sum_{j>_m w} x_{m,j} + \sum_{i>_w m} x_{i,w} + x_{m,w} \geq 1 \text{ for all pairs } (m, w) \\
 &&& x_{m,w} \in \{0, 1\} \text{ for all pairs } (m, w)
 \end{aligned}$$

- Prove that this integer program is a correct formulation of the stable matching problem.
- Consider the relaxation of the integer program that allows *fractional* stable matchings. It is identical to the above program, except that instead of each $x_{m,w}$ being either 0 or 1, $x_{m,w}$ is allowed to take any real value in $[0, 1]$. Show that the following program is the dual program to the relaxation of E7.1.

$$\begin{aligned}
 & \text{minimize} && \sum_i \alpha_i + \sum_j \beta_j - \sum_{i,j} \gamma_{ij} \\
 \text{subject to} &&& \alpha_m + \beta_w - \sum_{j<_m w} \gamma_{m,j} - \sum_{i<_w m} \gamma_{i,w} - \gamma_{m,w} \geq 1 \\
 &&& \text{for all pairs } (m, w) \\
 &&& \alpha_i, \beta_j, \gamma_{i,j} \geq 0 \text{ for all } i \text{ and } j.
 \end{aligned}$$

- Use complementary slackness (Theorem ??) to show that every feasible fractional solution to the relaxation of E7.1 is optimal and that setting

$$\begin{aligned}
 \alpha_m &= \sum_j x_{m,j} \text{ for all } m, \\
 \beta_w &= \sum_i x_{i,w} \text{ for all } w
 \end{aligned}$$

and

$$\gamma_{ij} = x_{ij} \text{ for all } i, j$$

is optimal for the dual program.

8

Coalitions and Shapley value

In this chapter, we consider **cooperative game theory**, in which players form coalitions to work toward a common goal. In these settings, there is a set $n > 2$ of players that can achieve a common goal yielding an overall payoff of v , if they all cooperate with each other. However, subsets of these players, so-called *coalitions*, have the option of going off on their own and collaborating only with each other, rather than working as part of a grand coalition. Questions addressed by this theory include: How should rewards be shared among the players so as to discourage subgroups from defecting? What power or influence does a player have in the game?

8.1 The Shapley value and the glove market

We review the example discussed in the Chapter (1). Suppose that three people are selling their wares in a market. Two of them are selling a single, left-handed glove, while the third is selling a right-handed one. A wealthy tourist arrives at the market in dire need of a pair of gloves, willing to pay \$100 for a pair of gloves. She refuses to deal with the glove-bearers individually, and thus, these sellers have to come to some agreement as to how to make a sale of a left- and right-handed glove to her and how to then split the \$100 amongst themselves. Clearly, the third player has an advantage, because his commodity is in scarcer supply. This means that he should be able to obtain a higher fraction of the payment than either of the other players. However, if he holds out for too high a fraction of the earnings, the other players may agree between them to refuse to deal with him at all, blocking any sale, thereby risking his earnings. Finding a solution for such a game involves a mathematical concept known as the **Shapley value**.

The question then is, in their negotiations prior to the purchase, how

much can each player realistically demand out of the total payment made by the customer?




Fig. 8.1.

To resolve this question, we introduce a **characteristic function** v , defined on subsets of the player set. This characteristic function captures for each subset S of the players, whether or not they are able between them to effect their aim. In our example of the glove market, $v(S)$, where S is a subset of the three players, is 1 if, just amongst themselves, the players in S have both a left glove and a right glove. Thus, in our example

$$v(123) = v(12) = v(13) = 1,$$

and the value is 0 on every other subset of $\{1, 2, 3\}$. (We abuse notation in this chapter and write $v(12)$ instead of $v(\{1, 2\})$, etc.)

More generally, a **cooperative game** is defined by a set N of n players and a characteristic function v on subsets of the n players, where $v(S)$ is the value or payoff that subset S of players can achieve on their own regardless of what the remaining players do. The characteristic function satisfies the following properties:

- $v(\emptyset) = 0$.
- The characteristic function is monotone nondecreasing. That is, if $S \subseteq T$, then $v(S) \leq v(T)$. This is because players in T always have the option of achieving at least what subset S can achieve on their own.
- The characteristic function is **superadditive**, that is: $v(S \cup T) \geq v(S) + v(T)$ if S and T are disjoint. This is because subsets S

and T always have the option of simply cooperating each amongst themselves, and ignoring the other group.

The outcome of the game is a set of “shares”, one per player, where $\psi_i(v)$ is the share player i gets when the characteristic function is v . We think of $\psi_i(v)$ as reflecting player i 's power in the game.

How should these shares be determined? A first natural property is **efficiency**, i.e.

$$\sum_i \psi_i(v) = v(N). \quad (8.1)$$

But beyond this, what properties might we desire that the shares have? Shapley analyzed this question by considering the following axioms:

- (i) **Symmetry**: if $v(S \cup \{i\}) = v(S \cup \{j\})$ for all S with $i, j \notin S$, then $\psi_i(v) = \psi_j(v)$.
- (ii) **Dummy**: A player that doesn't add value gets nothing: if $v(S \cup \{i\}) = v(S)$ for all S , then $\psi_i(v) = 0$.
- (iii) **Additivity**: $\psi_i(v + u) = \psi_i(v) + \psi_i(u)$.
- (iv) **Efficiency**: $\sum_{i=1}^n \psi_i(v) = v(\{1, \dots, n\})$.

What is fascinating is that it turns out that there is a *unique* choice of ψ , given these axioms. This unique choice for each ψ_i is called the **Shapley value** of player i in the game defined by characteristic function v .

Before we prove this theorem in general, let's see why it's true in an example.

Example 8.1.1. The S -veto game: Consider a coalitional game with n players, in which a fixed subset S of the players hold all the power. We will denote the characteristic function here by w_S , defined as: $w_S(T)$ is 1 if T contains S and 0 otherwise. We will show that, under Shapley's axioms, we have

$$\psi_i(w_S) = \frac{1}{|S|} \quad \text{if } i \in S,$$

and 0 otherwise. To see this observe first that, using the dummy axiom,

$$\psi_i(w_S) = 0 \quad \text{if } i \notin S.$$

Then, for $i, j \in S$, the “symmetry” axiom gives $\psi_i(w_S) = \psi_j(w_S)$. Finally, the “efficiency” axiom implies that

$$\psi_i(w_S) = \frac{1}{|S|} \quad \text{if } i \in S.$$

Note that to derive this, we did not use the additivity axiom. However,

using the additivity axiom, we can also derive that $\psi_i(cw_S) = c\psi_i(w_S)$ for any $c \in [0, \infty)$.

The glove market game, again: We can now use our understanding of the S -veto game to solve for the unique shares in the glove game under the above axioms. The observation is that the glove market game has the same payoffs as $w_{12} + w_{13}$, except for the case of the set $\{1, 2, 3\}$. In fact, we have that

$$w_{12} + w_{13} = v + w_{123},$$

where, as you recall, v is the characteristic function of the glove market game. Thus, the additivity axiom gives

$$\psi_i(w_{12}) + \psi_i(w_{13}) = \psi_i(v) + \psi_i(w_{123}).$$

We conclude from this that for player 1, $1/2 + 1/2 = \psi_1(v) + 1/3$, whereas for player 3, $0 + 1/2 = \psi_3(v) + 1/3$. Hence $\psi_1(v) = 2/3$ and $\psi_2(v) = \psi_3(v) = 1/6$. Thus, under Shapley's axioms, player 1 obtains a two-thirds share of the payoff, while players 2 and 3 equally share one-third between them.

Example 8.1.2. Four Stockholders: Four people own stock in ACME. Player i holds i units of stock, for each $i \in \{1, 2, 3, 4\}$. Six shares are needed to pass a resolution at the board meeting. Here $v(S)$ is 1 if subset S of players have enough shares of stock between them to pass a resolution. Thus,

$$1 = v(1234) = v(24) = v(34),$$

while $v = 1$ on any 3-tuple, and $v = 0$ in each other case. What power share does each of the players have under our axioms?

We will assume that the characteristic function v may be written in the form

$$v = \sum_{S \neq \emptyset} c_S w_S.$$

Later (in the proof of Theorem 8.2.1), we will see that there always exists such a way of writing v . For now, however, we assume this, and compute the coefficients c_S . Note first that

$$0 = v(1) = c_1$$

Similarly, $0 = c_2 = c_3 = c_4$. Also,

$$0 = v(12) = c_1 + c_2 + c_{12},$$

implying that $c_{12} = 0$. Similarly, $c_{13} = c_{14} = c_{23} = 0$. Next,

$$1 = v(24) = c_2 + c_4 + c_{24} = 0 + 0 + c_{24},$$

implying that $c_{24} = 1$. Similarly, $c_{34} = 1$. Proceeding, we have

$$1 = v(123) = c_{123},$$

and

$$1 = v(124) = c_{24} + c_{124} = 1 + c_{124},$$

implying that $c_{124} = 0$. Similarly, $c_{134} = 0$, and

$$1 = v(234) = c_{24} + c_{34} + c_{234} = 1 + 1 + c_{234},$$

implying that $c_{234} = -1$. Finally,

$$\begin{aligned} 1 = v(1234) &= c_{24} + c_{34} + c_{123} + c_{124} + c_{134} + c_{234} + c_{1234} \\ &= 1 + 1 + 1 + 0 + 0 - 1 + c_{1234}, \end{aligned}$$

implying that $c_{1234} = -1$. Thus,

$$v = w_{24} + w_{34} + w_{123} - w_{234} - w_{1234},$$

whence

$$\psi_1(v) = 1/3 - 1/4 = 1/12,$$

and

$$\psi_2(v) = 1/2 + 1/3 - 1/3 - 1/4 = 1/4,$$

while $\psi_3(v) = 1/4$, by symmetry with player 2. Finally, $\psi_4(v) = 5/12$. It is interesting to note that the person with 2 shares and the person with 3 shares have equal power.

8.2 The Shapley value

Consider a fixed ordering of the players, defined by a permutation π of $[1..n]$. Imagine the players arriving one by one according to this permutation π , and define $\phi_i(v, \pi)$ to be the marginal contribution of player i at the time of his arrival assuming players arrive in this order. Thus, if $\pi(k) = i$, we have

$$\phi_i(v, \pi) = v(\pi(1), \dots, \pi(k)) - v(\pi(1), \dots, \pi(k-1)). \quad (8.2)$$

Notice that if we were to set $\psi_i(v) = \phi_i(v, \pi)$ for any fixed π , the “dummy”, “efficiency” and “additivity” axioms would be satisfied.

To satisfy the “symmetry” axiom as well, we will instead imagine that the

players arrive in a random order and define $\psi_i(v)$ to be the expected value of $\phi_i(v, \pi)$ when π is chosen uniformly at random.

Remark. If we apply the approach just described to the four stockholders example, then there exists a moment when, with the arrival of the next stockholder, the coalition already present in the board-room becomes effective. The Shapley value of a given player turns out to be precisely the probability of that player being the one to make the existing coalition effective when the stockholders arrive in a random order.

Theorem 8.2.1. *Shapley's four axioms uniquely determine the functions ψ_i which follow the random arrival formula:*

$$\psi_i(v) = \frac{1}{n!} \sum_{k=1}^n \sum_{\pi \in S_n: \pi(k)=i} \left(v(\pi(1), \dots, \pi(k)) - v(\pi(1), \dots, \pi(k-1)) \right)$$

Remark. Note that this formula indeed specifies the probability just mentioned.

Proof. First, we prove that the shares $\psi_i(v)$'s are uniquely determined by v and the four axioms. To prove this, we show that any characteristic function v can be uniquely represented as a linear combination of S-veto characteristic functions w_S for different subsets S of $[1..n]$. Recall that $\psi_i(w_S) = 1/|S|$ if $i \in S$, and $\psi_i(w_S) = 0$ otherwise.

We claim that, given v , there are coefficients $\{c_S\}_{S \subseteq [n], S \neq \emptyset}$ such that for all $T \subseteq [n]$

$$v(T) = \sum_{\emptyset \neq S \subseteq [n]} c_S w_S(T) = \sum_{\emptyset \neq S \subseteq T} c_S w_S(T). \quad (8.3)$$

To see that this system of $2^n - 1$ equations in $2^n - 1$ unknowns has a solution, we construct the coefficients c_S inductively, in order of increasing cardinality. To begin, choose c_i as follows:

$$v(i) = \sum_{\emptyset \neq S \subseteq \{i\}} c_S w_S(i) = c_i w_i(i) = c_i. \quad (8.4)$$

Now suppose that we have defined c_S for all S with $|S| < \ell$, and have shown that (8.3) holds for any T with $|T| < \ell$. To determine $c_{\tilde{S}}$ for some \tilde{S} with $|\tilde{S}| = \ell$, we observe that

$$v(\tilde{S}) = \sum_{\emptyset \neq S \subseteq \tilde{S}} c_S w_S(\tilde{S}) = \sum_{S \subseteq \tilde{S}, |S| < \ell} c_S + c_{\tilde{S}}, \quad (8.5)$$

and thus (8.3) is satisfied for $c_{\tilde{S}}$ if we choose

$$c_{\tilde{S}} = v(\tilde{S}) - \sum_{S \subseteq \tilde{S}: |S| < \ell} c_S.$$

Next, we apply the additivity axiom and conclude that $\psi_i(v)$ is uniquely determined:

$$\psi_i(v) = \psi_i\left(\sum_{\emptyset \neq S \subseteq [n]} c_S w_S\right) = \sum_{\emptyset \neq S \subseteq [n]} \psi_i(c_S w_S) = \sum_{S \subseteq [n], i \in S} \frac{c_S}{|S|}.$$

We complete the proof by showing that the specific values given in the statement of the theorem satisfy all of the axioms. Recall the definition of $\phi_i(v, \pi)$ from (8.2). By averaging over all permutations π , and then defining

$$\psi_i(v) = \frac{1}{n!} \sum_{\{\pi: \pi(k)=i\}} \phi_i(v, \pi),$$

we claim that all four axioms are satisfied. Since averaging preserves the “dummy”, “efficiency” and “additivity” axioms, we only need to prove the intuitive fact that by averaging over all permutations, we obtain symmetry.

To this end, suppose that i and j are such that

$$v(S \cup \{i\}) = v(S \cup \{j\})$$

for all $S \subseteq [n]$ with $S \cap \{i, j\} = \emptyset$. For every permutation π , define π^* to be the same as π except that the positions of i and j are switched. Then

$$\phi_i(v, \pi) = \phi_j(v, \pi^*).$$

Using the fact that the map $\pi \mapsto \pi^*$ is a one-to-one map from S_n to itself for which $\pi^{**} = \pi$, we obtain

$$\begin{aligned} \psi_i(v) &= \frac{1}{n!} \sum_{\pi \in S_n} \phi_i(v, \pi) = \frac{1}{n!} \sum_{\pi \in S_n} \phi_j(v, \pi^*) \\ &= \frac{1}{n!} \sum_{\pi^* \in S_n} \phi_j(v, \pi^*) = \psi_j(v). \end{aligned}$$

Therefore, $\psi_i(v)$ is indeed the unique Shapley value. □

8.3 Two more examples

A fish without intrinsic value. A seller has a fish having no intrinsic value to him, i.e., he values it at \$0. A buyer values the fish at \$10. Writing S and B for the seller and buyer, we have that $v(S) = 0$ and $v(B) = 0$, since

separately neither can obtain any positive payoff. However, if the seller sells the fish for $\$x$, then the seller obtains a reward of x and the buyer a reward of $10 - x$ (for $0 < x \leq 10$). Thus, $v(S, B) = (10 - x) + x = 10$. In this game we have the following Shapley values: $\psi_S(v) = \psi_B(v) = 5$.

Notice however that the buyer's value is private to her, and if this is how the buyer and seller split her value for the fish, then she will have an incentive to underreport her desire for the fish to the party that arbitrates the transaction.

Many right gloves. Consider the following variant of the glove game. There are $n = r + 2$ players. Players 1 and 2 have left gloves. The remaining players each have a right glove. Thus, the characteristic function $v(S)$ is the maximum number of proper and disjoint pairs of gloves owned by players in S . Note that $\psi_1(v) = \psi_2(v)$, and $\psi_r(v) = \psi_3(v)$, for each $r \geq 3$. By the efficiency property ((8.1)), we have

$$2\psi_1(v) + r\psi_3(v) = 2$$

provided that $r \geq 2$. To determine the Shapley value of the third player, we consider all permutations π with the property that the third player adds value to the group of players that precede him in π . These are the following orders:

$$13, 23, \{1, 2\}3, \{1, 2, j\}3,$$

where j is any value in $\{4, \dots, n\}$, and the curly brackets mean that each permutation of the elements in curly brackets is included. The number of permutations corresponding to each of these possibilities is: $r!$, $r!$, $2(r - 1)!$, and $6(r - 1) \cdot (r - 2)!$. Thus,

$$\psi_3(v) = \frac{2r! + 8(r - 1)!}{(r + 2)!} = \frac{2r + 8}{(r + 2)(r + 1)r}.$$

Exercises

8.1 **The glove market revisited.** A proper pair of gloves consists of a left glove and a right glove. There are n players. Player 1 has two left gloves, while each of the other $n - 1$ players has one right glove. The payoff $v(S)$ for a coalition S is the number of proper pairs that can be formed from the gloves owned by the members of S .

(a) For $n = 3$, determine $v(S)$ for each of the 7 nonempty sets $S \subset \{1, 2, 3\}$. Then find the Shapley value $\varphi_i(v)$ for each of the players $i = 1, 2, 3$.

(b) For a general n , find the Shapley value $\varphi_i(v)$ for each of the n players $i = 1, 2, \dots, n$.

9

Interactive Protocols

So far we have studied how different players should play a given game. The goal of mechanism design is to construct a mechanism (a game) through which the participants interact with one another (“play the game”), so that when the participants act in their own self interest (“play strategically”), the resulting “game play” has desirable properties. For example, an auctioneer will wish to set up the rules of an auction so that the players will play against one another and drive up the price. Another example is cake cutting, where the participants wish to divvy up a cake so that everyone feels like he or she received a fair share of the best parts of the cake. Zero-knowledge proofs are another example: here one of the participants (Alice) has a secret, and wishes to prove that to another participant (Bob) that she knows the secret, but without giving the secret away. If Alice follows the protocol, she is assured that her secret is safe, and if Bob follows the protocol, he is assured that Alice knows the secret.

9.1 Keeping the meteorologist honest

The employer of a weatherman is determined that he should provide a good prediction of the weather for the following day. The weatherman’s instruments are good, and he can, with sufficient effort, tune them to obtain the correct value for the probability of rain on the next day. There are many days, and on the i^{th} day the true probability of rain is called p_i . On the evening of the $(i - 1)^{\text{th}}$ day, the weatherman submits his estimate \hat{p}_i for the probability of rain on the following day, the i^{th} one. Which scheme should we adopt to reward or penalize the weatherman for his predictions, so that he is motivated to correctly determine p_i (that is, to declare $\hat{p}_i = p_i$)? The employer does not know what p_i is because he has no access to technical

equipment, but he does know the \hat{p}_i values that the weatherman provides, and he knows whether or not it is raining on each day.

One suggestion is to pay the weatherman on the i^{th} day the amount \hat{p}_i (or some dollar multiple of that amount) if it rains, and $1 - \hat{p}_i$ if it shines. If $\hat{p}_i = p_i = 0.6$, then the payoff is

$$\begin{aligned} \hat{p}_i \Pr(\text{rainy}) + (1 - \hat{p}_i) \Pr(\text{sunny}) &= \hat{p}_i p_i + (1 - \hat{p}_i)(1 - p_i) \\ &= 0.6 \times 0.6 + 0.4 \times 0.4 = 0.52. \end{aligned}$$

But in this case, even if the weatherman does correctly compute that $p_i = 0.6$, he is tempted to report the \hat{p}_i value of 1 because, by the same formula, in this case, his earnings are 0.6.

Another idea is to pay the weatherman a fixed salary over a term, say, one year. At the end of the term, penalize the weatherman according to how accurate his predictions have been on the average. More concretely, suppose for the sake of simplicity that the weatherman is only able to report \hat{p}_i values on a scale of $\frac{1}{10}$, so that he has eleven choices, namely $\{k/10 : k \in \{0, \dots, 10\}\}$. When a year has gone by, the days of that year may be divided into eleven types according to the \hat{p}_i -value that the weatherman declared. Suppose there are n_k days that the predicted value \hat{p}_i is $\frac{k}{10}$, while according to the actual weather, r_k days out of these n_k days rained. Then, we give the penalty as

$$\sum_{k=0}^{10} \left(\frac{r_k}{n_k} - \frac{k}{10} \right)^2.$$

A scheme like this seems quite reasonable, but in fact, it can be quite disastrous. If the weather doesn't fluctuate too much from year to year and the weatherman knows that on average it rained on $\frac{3}{10}$ of the days last year, he will be able to ignore his instruments completely and still do reasonably well. To see this, suppose the weatherman simply sets $\hat{p} = \frac{3}{10}$; then $n_3 = 365$ and $n_{k \neq 3} = 0$. In this case his penalty will be

$$\left(\frac{r_3}{365} - \frac{3}{10} \right)^2,$$

where r_3 is simply the overall number of rainy days in a year, which is expected to be quite close to $365 \times \frac{3}{10}$. By the Law of Large Numbers, as the number of observations increases, the penalty is likely to be close to zero.

It turns out that even if the weatherman doesn't know the average rainfall, he can still do quite well as the following theorem indicates.

Theorem 9.1.1. *Suppose the weatherman is restricted to report \hat{p}_i values on a scale of $\frac{1}{10}$. Even if he knows nothing about the weather, he can devise a strategy so that over a period of n days his penalty is, on average, within $\frac{1}{20}$, in each slot.*

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{10} \left| r_k - \frac{k}{10} n_k \right| \leq \frac{1}{20}.$$

One proof of this can be found in ([FV99]), and an explicit strategy has been constructed in (need ref Dean Foster). Since then, the result has been recast as a consequence of minimax theorem (see [HMC00]), by considering the situation as a zero-sum game between the weatherman and a certain adversary. In this case the adversary is obtained from the combined effects of the employer and the weather.

There are two players, the weatherman W and the adversary A . Each day, A can play a mixed strategy randomizing between Rain and Shine. The problem is to devise an optimal response for W , which consists of a prediction for each day. Such a prediction can also be viewed as a mixed strategy, randomizing between Rain and Shine. At the end of the term, the weatherman W pays the adversary A a penalty as described above.

In this case, there is no need for instruments: the minimax theorem guarantees that there is an optimal response strategy. We can go even further and give a specific prescription: On each day, compute a probability of rain, conditional on what the weather had been up to now.

A solution to the problem

The above examples cast the situation in a somewhat pessimistic light — so far we have shown that the scheme encourages the weatherman to ignore his instruments. Is it possible to give him an incentive to tune them up? In fact, it is possible to design a scheme whereby we decide day-by-day how to reward the weatherman only on the basis of his declaration from the previous evening, without encountering the kind of problem that the last scheme had [Win69].

Suppose that we pay $f(\hat{p}_i)$ to the weatherman if it rains, and $f(1 - \hat{p}_i)$ if it shines on day i . If $p_i = p$ and $\hat{p}_i = x$, then the expected payment made on day i is equal to

$$g_p(x) := pf(x) + (1 - p)f(1 - x).$$

Our aim is to reward the weatherman if his \hat{p}_i equals p_i , in other words, to

ensure that the expected payout is maximized when $x = p$. This means that the function $g_p : [0, 1] \rightarrow \mathbb{R}$ should satisfy $g_p(p) > g_p(x)$ for all $x \in [0, 1] \setminus \{p\}$.

One good choice is to let $f(x) = \log x$. In this case, the derivative of $g_p(x)$ is:

$$g'_p(x) = pf'(x) + (1-p)f'(1-x) = \frac{p}{x} - \frac{1-p}{1-x}.$$

The derivative is positive if $x < p$, and negative if $x > p$, and the function $g_p(x)$ achieves its maximum at $x = p$, as we wished.

9.2 Secret sharing

In the introduction, we talked about the problem of sharing a secret between two people. Suppose we do not trust either of them entirely, but want the secret to be known to each of them, provided that they co-operate. More generally, we can ask the same question about n people.

Think of this in a computing context: Suppose that the secret is a password that is represented as an integer S that lies between 0 and some large value, for example, $0 \leq S < M = 10^{15}$.

We might take the password and split it in n chunks, giving one chunk to each of the players. However, this would force the length of the password to be high, if none of the chunks are to be guessed by repeated tries. Moreover, as more players put together their chunks, the size of the unknown chunk goes down, making it more likely to be guessed by repeated trials.

A more ambitious goal is to split the secret S among n people in such a way that all of them together can reconstruct S , but no coalition of size $\ell < n$ has any information about S . We need to clarify what we mean when we say that a coalition has no information about S :

Definition 9.2.1. Let $A = \{i_1, \dots, i_\ell\} \subset \{1, \dots, n\}$ be any subset of size $\ell < n$. We say that a coalition of ℓ people holding a random vector $(X_{i_1}, \dots, X_{i_\ell})$ has **no information** about a secret S provided $(X_{i_1}, \dots, X_{i_\ell})$ is a random vector on $\{0, \dots, M-1\}^\ell$, whose distribution is independent of S , that is

$$\Pr(X_{i_1} = x_1, \dots, X_{i_\ell} = x_\ell | S = s) = \Pr(X_{i_1} = x_1, \dots, X_{i_\ell} = x_\ell).$$

The simplest way to ensure that the distribution of $(X_{i_1}, \dots, X_{i_\ell})$ does not depend upon S is to make its distribution uniformly random. Recall that a random variable X has a **uniform** distribution on a space of size N , denoted by Ω , provided each of the N possible outcomes is equally likely:

$$\Pr(X = x) = \frac{1}{N} \quad \forall x \in \Omega.$$

In the case of an ℓ -dimensional vector with elements in $\{0, \dots, M-1\}$, we have $\Omega = \{0, \dots, M-1\}^\ell$, of size M^ℓ .

9.2.1 A simple secret sharing method

The following scheme allows the secret holder to split a secret $S \in \{0, \dots, M-1\}$ among n individuals in such a way that any coalition of size $\ell < n$ has no information about S : The secret holder, produces a random $(n-1)$ -dimensional vector $(X_1, X_2, \dots, X_{n-1})$, whose distribution is uniform on $\{0, \dots, M-1\}^{n-1}$. She gives the number X_i to the i^{th} person for $1 \leq i \leq n-1$, and the number

$$X_n = \left(S - \sum_{i=1}^{n-1} X_i \right) \pmod{M} \quad (9.1)$$

to the last person. Notice that with this definition, X_n is also a uniformly random variable on $\{0, \dots, M-1\}$, you will prove this in Ex. 9.2.

It is enough to show that any coalition of size $n-1$ has no useful information. For $\{i_1, \dots, i_{n-1}\} = \{1, \dots, n-1\}$, the coalition of the first $n-1$ people, this is clear from the definition. What about those that include the last one? To proceed further we'll need an elementary lemma, whose proof is left as an Ex. 9.1:

Lemma 9.2.2. *Let Ω be a finite set of size N . Let T be a one-to-one and onto function from Ω to itself. If a random variable X has a uniform distribution over Ω , then so does $Y = T(X)$.*

Consider a coalition that omits the j^{th} person: $A = \{1, \dots, j-1, j+1, \dots, n\}$. Let $T_j((X_1, \dots, X_{n-1})) = (X_1, \dots, X_{j-1}, X_{j+1}, \dots, X_n)$, where X_n is defined by Eq. (9.1). This map is one-to-one and onto for each j since we can explicitly define its inverse:

$$T_j^{-1}((Z_1, \dots, Z_{j-1}, Z_{j+1}, \dots, Z_n)^T) = (Z_1, \dots, Z_{j-1}, Z_j, Z_{j+1}, \dots, Z_{n-1})^T,$$

where $Z_j = S - \sum_{1 \leq i \neq j \leq n-1} Z_i$.

Thus, by Lemma 9.2.2, a coalition (that does not include all players) that puts together all its available information, still has only a uniformly random vector. Since they could generate a uniformly random vector themselves without knowing anything about S , the coalition has the same chance of guessing the secret S as if it had no information at all.

All together, however, the players can add the values they had been given, reduce the answer mod M , and obtain the secret S .

9.2.2 Polynomial method

We now consider a generalization of the previous method, due to Adi Shamir [Sha79]. This generalization also provides a method for splitting a secret between n individuals, but now guarantees that any coalition of at least m individuals can recover it, while a group of a smaller size cannot. This could be useful if a certain action required a quorum of m individuals, less than the total number of people in the group.

Let S be the secret to be shared. Let p be a prime number such that $0 \leq S < p$ and $n < p$. We define a polynomial of order $m - 1$:

$$F(z) = \sum_{i=0}^{m-1} A_i z^i \pmod{p},$$

where A_0 is the secret S and (A_1, \dots, A_{m-1}) is a uniform random vector on $\{0, \dots, p-1\}^{m-1}$.

Let z_1, \dots, z_n be distinct numbers in $\{1, \dots, p-1\}$. To split the secret we give the j^{th} person the number $F(z_j)$ (together with z_j, p , and m). We claim that

Theorem 9.2.3. *A coalition of size m or bigger can reconstruct the secret S , but a coalition of size $\ell < m$ has no useful information:*

$$\Pr(F(z_1) = x_1, \dots, F(z_\ell) = x_\ell \mid S) = \frac{1}{p^\ell}, \quad x_i \in \{0, \dots, p-1\}.$$

Intuitively, the reason this works is that a polynomial of degree $m - 1$ is uniquely determined by its value at m points, and thus, any coalition of size m can determine A_0 . On the other hand, if we know $(z, F(z))$ for only $m - 1$ values of z , there are still p possibilities for what $A_0 = S$ can be.

Proof. Clearly it's enough to consider the case $\ell = m - 1$. We will show that for any fixed distinct non-zero integers $z_1, \dots, z_m \in \{0, \dots, p-1\}$,

$$T((A_0, \dots, A_{m-1})) = (F(z_1), \dots, F(z_m))$$

is an invertible linear map on $\{0, \dots, p-1\}^m$, and hence m people together can recover all the coefficients of F , including $A_0 = S$.

Let's construct these maps explicitly:

$$T \begin{pmatrix} A_0 \\ \vdots \\ A_{m-1} \end{pmatrix} = \begin{pmatrix} \sum_{i=0}^{m-1} A_i z_1^i \pmod{p} \\ \vdots \\ \sum_{i=0}^{m-1} A_i z_m^i \pmod{p} \end{pmatrix}.$$

We see that T is a linear transformation on $\{0, \dots, p-1\}^m$ that is equivalent to multiplying on the left with the following $m \times m$ matrix M , known as the **Vandermonde matrix**:

$$M = \begin{pmatrix} 1 & z_1 & \dots & z_1^{m-1} \\ 1 & z_2 & \dots & z_2^{m-1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & z_{m-1} & \dots & z_{m-1}^{m-1} \\ 1 & z_m & \dots & z_m^{m-1} \end{pmatrix}.$$

You will prove in Ex. 9.3 that

$$\det(M) = \prod_{1 \leq i < j \leq m} (z_j - z_i).$$

Recall that the numbers $\{0, \dots, p-1\}$ (recall that p is a prime) together with addition and multiplication (mod p) form a finite field. (Recall that a field is a set \mathcal{S} with operations called $+$ and \times which are associative and commutative, for which multiplication distributes over addition, which contains an additive identity called 0 and a multiplicative identity called 1, for which each element has an additive inverse, and each non-zero element contains a multiplicative inverse. Because multiplicative inverses of non-zero elements are defined, there are no zero divisors, i.e., a pair of elements whose product is zero.)

Since the z_i 's are all distinct and p is a prime number, the Vandermonde determinant $\det M$ is non-zero modulo p , so the transformation is invertible.

This shows that any coalition of m people can recover the secret S . Almost the same argument shows that any coalition of $m-1$ people have no information about S . Let the $m-1$ people be z_1, \dots, z_{m-1} , and let $z_m = 0$. We have shown that the map

$$T((A_0, \dots, A_{m-1})) = (F(z_1), \dots, F(z_{m-1}), A_0 = F(z_m))$$

is invertible. Thus, for any fixed value of A_0 , the map

$$T((A_1, \dots, A_{m-1})) = (F(z_1), \dots, F(z_{m-1}))$$

is invertible. Since A_1, \dots, A_{m-1} are uniformly random and independent of $A_0 = S$, it follows that $(F(z_1), \dots, F(z_{m-1}))$ is uniformly random and independent of S .

The proof is complete, however, it is quite instructive to construct the inverse map T^{-1} explicitly. We use the method of **Lagrange interpolation**

to reconstruct the polynomial:

$$F(z) = \sum_{j=1}^m F(z_j) \prod_{\substack{1 \leq i \leq m \\ i \neq j}} \frac{z - z_i}{z_j - z_i} \pmod{p}.$$

Once we expand the right-hand side and bring it to the standard form, (A_0, \dots, A_{m-1}) will appear as the coefficients of the corresponding powers of the indeterminate z . Evaluating at $z = 0$ gives back the secret. \square

9.3 Private computation

An applied physics professor at Harvard posed the following problem to his fellow faculty during tea hour: Suppose that all the faculty members would like to know the average salary in their department. How can they compute it without revealing the individual salaries? Since there is no disinterested third party who could be trusted by all the faculty members, they hit upon the following scheme:

All the faculty members gather around a table. A designated first person picks a very large integer M (which he keeps private), adds his salary to that number, and passes the result to his neighbor on the right. She, in turn, adds her salary and passes the result to her right. The intention is that the total should eventually return to the designated first person, who would then subtract M , compute and reveal the average. However, before the physicists could finish the computation, a Nobel laureate, who was flanked by two junior faculty, refused to participate when he realized that the two could collude to find out his salary.

Luckily, the physicists shared their tea-room with computer scientists who, after some thought, proposed the following ingenious scheme that is closely related to the secret sharing method described in section 9.2.1: A very large integer M is picked and announced to the entire faculty, consisting of n individuals. An individual with salary s_i generates $n - 1$ random numbers $X_{i,1}, \dots, X_{i,n-1}$, uniformly distributed in the set $\{0, 1, 2, \dots, M - 1\}$, and produces $X_{i,n}$, such that $X_{i,1} + \dots + X_{i,n} = s_i \pmod{M}$. He then forwards $X_{i,j}$ to the j^{th} faculty member. In this manner each person receives n uniform random numbers mod M , adds them up and reports the result. These are tallied mod M and divided by n .

Here a coalition of $n - 1$ faculty can deduce the last professor's salary, if for no other reason than that they know their own salaries and also the average salary. This holds for any scheme that the faculty adopt. Similarly, for any scheme for computing the average salary, a coalition of $n - j$ faculty

could deduce the sum of the salaries of the remaining j faculty. You will show in Ex. 9.5 that the above scheme leaks no additional information about the salaries.

9.4 Cake cutting

Recall from the introduction the problem of cutting a cake with several different toppings. The game has two or more players, each with a particular preference regarding which parts of the cake they would most like to have. We assume that all parts of the cake are divisible.

If there are just two players, there is a well-known method for dividing the cake: One splits it into two halves, and the other chooses which he would like. Each obtains at least one-half of the cake, as measured according to his own preferences. But what if there are three or more players? This can still be done, but requires some new notions.

Let's denote the cake by Ω . Then \mathcal{F} denotes the **algebra** of measurable subsets of Ω . Roughly speaking, these are all the subsets into which the cake can be subdivided by repeated cutting.

Definition 9.4.1 (Algebra of sets). More formally, we say that a collection \mathcal{F} of subsets of Ω forms an algebra if:

- (i) $\emptyset \in \mathcal{F}$;
- (ii) if $A \in \mathcal{F}$ then $A^c \in \mathcal{F}$;
- (iii) if $A, B \in \mathcal{F}$ then $A \cup B \in \mathcal{F}$.

The sets in \mathcal{F} are called **measurable**.

We will need a tool to measure the “desirability” of any possible piece of the cake for any given individual.

Definition 9.4.2. A non-negative real-valued set function μ defined on \mathcal{F} is called a **finite measure** if:

- (i) $\mu(\emptyset) = 0$ and $\mu(\Omega) = M < \infty$;
- (ii) if $A, B \in \mathcal{F}$ and $A \cap B = \emptyset$ then $\mu(A \cup B) = \mu(A) + \mu(B)$.

The triple $(\Omega, \mathcal{F}, \mu)$ is called a **finite measure space**.

In addition we will require that the measure space should have the **intermediate value** property: For every measurable set $A \in \mathcal{F}$ and any real number $\beta \in (0, \mu(A))$, there is a measurable set $B \in \mathcal{F}$ such that $B \subset A$ and $\mu(B) = \beta$. This ensures that there are no indivisible elements in the cake such as hard nuts that cannot be cut into two.

Now let μ_j be the measure on the cake which reflects the preferences of

the j^{th} person. Notice that each person gives a personal value to the whole cake. For each person, however, the value of the “empty slice” is 0, and the value of any slice is bigger than or equal to that of any of its parts.

Our task is to divide the cake into K slices $\{A_1^*, \dots, A_K^*\}$, such that for each individual i ,

$$\mu_i(A_i^*) \geq \frac{\mu_i(\Omega)}{K}.$$

In this case, we say that the division is fair. Notice that this notion addresses fairness from the point of view of each individual: She is assured a slice that is at least $\frac{1}{K}$ of her particular valuation of the cake.

The following algorithm provides such a subdivision: The first person is asked to mark a slice A_1 such that $\mu_1(A_1) = \frac{\mu_1(\Omega)}{K}$, and this slice becomes the “current proposal”. Each person j in turn looks at the current proposed slice of cake A , and if $\mu_j(A) > \mu_j(\Omega)/K$, person j proposes a smaller slice of cake $A_j \subset A$ such that $\mu_j(A_j) = \mu_j(\Omega)/K$, which then becomes the current proposal, and otherwise person j passes on the slice. After each person has had a chance to propose a smaller slice, the proposed slice of cake is cut and goes to the person k who proposed it (call the slice A_k^*). This person is happy because $\mu_k(A_k^*) = \mu_k(\Omega)/K$. Let $\tilde{\Omega} = \Omega \setminus A_k^*$ be the rest of the cake. Notice that for each of the remaining $K - 1$ individuals $\mu_j(A_k^*) \leq \mu_j(\Omega)/K$, and hence for the remainder of the cake

$$\mu_j(\tilde{\Omega}) \geq \mu_j(\Omega) \left(1 - \frac{1}{K}\right) = \mu_j(\Omega) \frac{K-1}{K}.$$

We can repeat the process on $\tilde{\Omega}$ with the remaining $K - 1$ individuals. By induction, each person m obtains a slice A_m^* with

$$\mu_m(A_m^*) \geq \mu_m(\tilde{\Omega}) \frac{1}{K-1} \geq \frac{\mu_m(\Omega)}{K}.$$

This is true if each person j carries out the instructions faithfully. After all, since we do not know his measure μ_j , we cannot judge whether he had marked off a fair slice at every stage of the game. However, since everyone’s measure has the intermediate property, a person who chooses to comply can ensure that she gets her fair share.

9.5 Zero-knowledge proofs

Determining whether or not a graph is 3-colorable, i.e., whether or not it is possible to color the vertices red, green, and blue, so that

each edge in the graph connects vertices with different colors, is a classic NP-hard problem. Solving 3-colorability for general graphs is at least as hard

as factoring integers, solving the traveling salesman problem, or solving any of a number of other hard problems. We describe a simple zero-knowledge proof of 3-colorability, which means that any of these other problems also has a zero-knowledge proof.

Suppose that Alice knows a 3-coloring of a graph G , and wishes to prove to Bob that the graph is 3-colorable, but does not wish to reveal the 3-coloring. What she can do is randomly permute the 3 colors red, green, and blue, and then write down the new color of each vertex in a sealed envelope, and place the envelopes on a table. Bob then picks a random edge (u, v) of the graph, and Alice then gives the envelopes for u and v to Bob, who opens them and checks that the colors are different. If the graph G has E edges, this protocol is then repeated tE times, where t might be 20.

There are three things to check: (1) completeness: if Alice knows a 3-coloring, she can convince Bob, (2) soundness: if there is no 3-coloring, then Bob catches her with high probability, and (3) zero-knowledge: Bob learns nothing about the 3-coloring other than that it exists.

Completeness here is trivial: if Alice knows a 3-coloring, and follows the protocol, then when Bob opens the two envelopes, he will always see different colors.

Soundness is straightforward too: If there is no 3-coloring, then there is always at least one edge of the graph whose endpoints have the same color. With probability $1/E$ Bob will pick that edge, and discover that Alice was cheating. Since this protocol is repeated tE times, the probability that Alice is not caught is at most $(1 - 1/E)^{tE} < e^{-t}$. For $t = 20$, this probability is about 2×10^{-9} .

Zero-knowledge: Suppose Alice knows a 3-coloring and follows the protocol, can Bob learn anything about the 3-coloring about it? Because Alice randomly permuted the labels of the colors, for any edge that Bob selects, each of the 6 possible 2-colorings of that edge are equally likely. At the end of the protocol, Bob sees tE random 2-colorings of edges. But Bob was perfectly able to randomly 2-color these edges on his own without Alice's help. Therefore, this communication from Alice did not reveal anything about her 3-coloring.

In a computer implementation, rather than use envelopes, Alice would use some cryptography to conceal the colors of the vertices but commit to their values. With a cryptographic implementation, the zero-knowledge property is not perfect zero-knowledge, but relies on Bob not being able to break the cryptosystem.

9.6 Remote coin tossing

Suppose, while speaking on the phone, two people would like to make a decision that depends on an outcome of a coin toss. How can they imitate such a set-up?

The standard way to do this before search-engines was for one of them to pick an arbitrary phone number from the phone-book, announce it to the other person and then ask him to decide whether this number is on an even- or odd-numbered page. Once the other person announces the guess, the first supplies the name of the person, whose phone number was used. In this way, the parity of the page number can be checked and the correctness of the phone number verified.

With the advent of fast search engines this has become impractical, since, from a phone number, the name (and hence the page number) can easily be looked up. A modification of this scheme that is somewhat more search-engine resistant is for one person to give a sequence of say 20 digits that occur in the 4th position on twenty consecutive phone numbers from the same page, and then to ask whether this page is even or odd.

If the two people have computers and email, another method can be used. One person could randomly pick two large prime numbers, multiply them, and mail the result to the other person. The other person guesses whether or not the two primes have the same parity of their middle digit, at which point the first person mails the primes. If the guess was right, the coin was heads, otherwise it is tails.

Exercises

- 9.1 Let Ω be a finite set of size N . Let T be a one-to-one and onto function from Ω to itself. Show that if a random variable X has a uniform distribution over Ω , then so does $Y = T(X)$.
- 9.2 Given a random $(n-1)$ -dimensional vector $(X_1, X_2, \dots, X_{n-1})$, with a uniform distribution on $\{0, \dots, M-1\}^{n-1}$. Show that
- Each X_i is a uniform random variable on $\{0, \dots, M-1\}$.
 - X_i 's are independent random variables.
 - Let $S \in \{0, \dots, M-1\}$ be given then

$$X_n = \left(S - \sum_{i=1}^{n-1} X_i \right) \bmod M$$

is also a uniform random variable on $\{0, \dots, M - 1\}$.

- 9.3 Prove that the Vandermonde matrix has the following determinant:

$$\det \begin{bmatrix} 1 & z_1 & \dots & z_1^{m-1} \\ 1 & z_2 & \dots & z_2^{m-1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & z_{m-1} & \dots & z_{m-1}^{m-1} \\ 1 & z_m & \dots & z_m^{m-1} \end{bmatrix} = \prod_{1 \leq i < j \leq m} (z_j - z_i).$$

Hint: the determinant is a multivariate polynomial. Show that the determinant is 0 when $z_i = z_j$ for $i \neq j$, show that the polynomial on the right divides the determinant, show that they have the same degree, and show that the constant factor is correct.

- 9.4 Evaluate the following determinant, known as a **Cauchy determinant**:

$$\det \begin{bmatrix} \frac{1}{x_1 - y_1} & \frac{1}{x_1 - y_2} & \dots & \frac{1}{x_1 - y_m} \\ \frac{1}{x_2 - y_1} & \frac{1}{x_2 - y_2} & \dots & \frac{1}{x_2 - y_m} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{1}{x_m - y_1} & \frac{1}{x_m - y_2} & \dots & \frac{1}{x_m - y_m} \end{bmatrix}.$$

Hint: find the zeros and poles and the constant factor. It is helpful to consider the limit $x_i \rightarrow y_j$.

- 9.5 Show that for the scheme for computing average salary described in section 9.3, a coalition $n - j$ faculty learn nothing about the salaries of the remaining j faculty beyond the sum of their salaries (which is what they could deduce knowing the average salary of everybody).

10

Combinatorial games

In this chapter, we will look at **combinatorial games**, a class of games that includes some popular two-player board games such as Nim and Hex, discussed in the introduction. In a combinatorial game, there are two players, a set of positions, and a set of legal moves between positions. Some of the positions are terminal. The players take turns moving from position to position. The goal for each is to reach the terminal position that is winning for that player. Combinatorial games generally fall into two categories:

Those for which the winning positions and the available moves are the same for both players are called **impartial**. The player who first reaches one of the terminal positions wins the game. We will see that all such games are related to Nim.

All other games are called **partisan**. In such games the available moves, as well as the winning positions, may differ for the two players. In addition, some partisan games may terminate in a **tie**, a position in which neither player wins decisively.

Some combinatorial games, both partisan and impartial, can also be **drawn** or go on forever.

For a given combinatorial game, our goal will be to find out whether one of the players can always force a win, and if so, to determine the winning strategy — the moves this player should make under every contingency. Since this is extremely difficult in most cases, we will restrict our attention to relatively simple games.

In particular, we will concentrate on the combinatorial games that terminate in a finite number of steps. Hex is one example of such a game, since each position has finitely many uncolored hexagons. Nim is another example, since there are finitely many chips. This class of games is important enough to merit a definition:

Definition 10.0.1. A combinatorial game with a position set X is said to be **progressively bounded** if, starting from any position $x \in X$, the game must terminate after a finite number $B(x)$ of moves.

Here $B(x)$ is an upper bound on the number of steps it takes to play a game to completion. It may be that an actual game takes fewer steps.

Note that, in principle, Chess, Checkers and Go need not terminate in a finite number of steps since positions may recur cyclically; however, in each of these games there are special rules that make them effectively progressively bounded games.

We will show that in a progressively bounded combinatorial game that cannot terminate in a tie, one of the players has a winning strategy. For many games, we will be able to identify that player, but not necessarily the strategy. Moreover, for all progressively bounded impartial combinatorial games, the Sprague-Grundy theory developed in section 10.1.3 will reduce the process of finding such a strategy to computing a certain recursive function.

We begin with impartial games.

10.1 Impartial games

Before we give formal definitions, let's look at a simple example:

Example 10.1.1 (A Subtraction game). Starting with a pile of $x \in \mathbb{N}$ chips, two players alternate taking one to four chips. The player who removes the last chip wins.

Observe that starting from any $x \in \mathbb{N}$, this game is progressively bounded with $B(x) = x$.

If the game starts with 4 or fewer chips, the first player has a winning move: he just removes them all. If there are five chips to start with, however, the second player will be left with between one and four chips, regardless of what the first player does.

What about 6 chips? This is again a winning position for the first player because if he removes one chip, the second player is left in the losing position of 5 chips. The same is true for 7, 8, or 9 chips. With 10 chips, however, the second player again can guarantee that he will win.

Let's make the following definition:

$$\mathbf{N} = \left\{ x \in \mathbb{N} : \begin{array}{l} \text{the first ("next") player can ensure a win} \\ \text{if there are } x \text{ chips at the start} \end{array} \right\},$$

$$\mathbf{P} = \left\{ x \in \mathbb{N} : \begin{array}{l} \text{the second ("previous") player can ensure a win} \\ \text{if there are } x \text{ chips at the start} \end{array} \right\}.$$

So far, we have seen that $\{1, 2, 3, 4, 6, 7, 8, 9\} \subseteq \mathbf{N}$, and $\{0, 5\} \subseteq \mathbf{P}$. Continuing with our line of reasoning, we find that $\mathbf{P} = \{x \in \mathbb{N} : x \text{ is divisible by five}\}$ and $\mathbf{N} = \mathbb{N} \setminus \mathbf{P}$.

The approach that we used to analyze the Subtraction game can be extended to other impartial games. To do this we will need to develop a formal framework.

Definition 10.1.2. An **impartial combinatorial game** has two players, and a set of possible positions. To make a **move** is to take the game from one position to another. More formally, a move is an ordered pair of positions. A terminal position is one from which there are no legal moves. For every non-terminal position, there is a set of legal moves, the same for both players. Under **normal play**, the player who moves to a terminal position wins.

We can think of the game positions as **nodes** and the moves as directed **links**. Such a collection of nodes (vertices) and links (edges) between them is called a **graph**. If the moves are reversible, the edges can be taken as undirected. At the start of the game, a token is placed at the node corresponding to the initial position. Subsequently, players take turns placing the token on one of the neighboring nodes until one of them reaches a terminal node and is declared the winner.

With this definition, it is clear that the Subtraction game is an impartial game under normal play. The only terminal position is $x = 0$. Figure 10.1 gives a directed graph corresponding to the Subtraction game with initial position $x = 14$.




Fig. 10.1. Moves in the Subtraction game. Positions in \mathbf{N} are marked in red and those in \mathbf{P} , in black.

We saw that starting from a position $x \in \mathbf{N}$, the next player to move can force a win by moving to one of the elements in $\mathbf{P} = \{5n : n \in \mathbf{N}\}$, namely $5\lfloor x/5 \rfloor$.

Let's make a formal definition:

Definition 10.1.3. A (**memoryless**) **strategy** for a player is a function that assigns a legal move to each non-terminal position. A **winning strategy** from a position x is a strategy that, starting from x , is guaranteed to result in a win for that player in a finite number of steps.

We say that the strategy is memoryless because it does not depend on the history of the game, i.e., the previous moves that led to the current game position. For games which are not progressively bounded, where the game might never end, the players may need to consider more general strategies that depend on the history in order to force the game to end. But for games that are progressively bounded, this is not an issue, since as we will see, one of the players will have a winning memoryless strategy.

We can extend the notions of \mathbf{N} and \mathbf{P} to any impartial game.

Definition 10.1.4. For any impartial combinatorial game, we define \mathbf{N} (for “next”) to be the set of positions such that the first player to move can guarantee a win. The set of positions for which every move leads to an \mathbf{N} -position is denoted by \mathbf{P} (for “previous”), since the player who can force a \mathbf{P} -position can guarantee a win.

In the Subtraction game, $\mathbf{N} = \mathbf{N} \cup \mathbf{P}$, and we were easily able to specify a winning strategy. This holds more generally: If the set of positions in an impartial combinatorial game equals $\mathbf{N} \cup \mathbf{P}$, then from any initial position one of the players must have a winning strategy. If the starting position is in \mathbf{N} , then the first player has such a strategy, otherwise, the second player does.

In principle, for any progressively bounded impartial game it is possible, working recursively from the terminal positions, to label every position as either belonging to \mathbf{N} or to \mathbf{P} . Hence, starting from any position, a winning strategy for one of the players can be determined. This, however, may be algorithmically hard when the graph is large. In fact, a similar statement also holds for progressively bounded partisan games. We will see this in §Section 10.2.

We get a recursive characterization of \mathbf{N} and \mathbf{P} under normal play by letting \mathbf{N}_i and \mathbf{P}_i be the positions from which the first and second players respectively can win within $i \geq 0$ moves:

$$\begin{aligned} \mathbf{N}_0 &= \emptyset \\ \mathbf{P}_0 &= \{ \text{terminal positions} \} \\ \mathbf{N}_{i+1} &= \{ \text{positions } x \text{ for which there is a move leading to } \mathbf{P}_i \} \\ \mathbf{P}_{i+1} &= \{ \text{positions } y \text{ such that each move leads to } \mathbf{N}_i \} \end{aligned}$$

$$\mathbf{N} = \bigcup_{i \geq 0} \mathbf{N}_i, \quad \mathbf{P} = \bigcup_{i \geq 0} \mathbf{P}_i.$$

Notice that $\mathbf{P}_0 \subseteq \mathbf{P}_1 \subseteq \mathbf{P}_2 \subseteq \dots$ and $\mathbf{N}_0 \subseteq \mathbf{N}_1 \subseteq \mathbf{N}_2 \subseteq \dots$.

In the Subtraction game, we have

$$\begin{array}{ll} \mathbf{N}_0 = \emptyset & \mathbf{P}_0 = \{0\} \\ \mathbf{N}_1 = \{1, 2, 3, 4\} & \mathbf{P}_1 = \{0, 5\} \\ \mathbf{N}_2 = \{1, 2, 3, 4, 6, 7, 8, 9\} & \mathbf{P}_2 = \{0, 5, 10\} \\ \vdots & \vdots \\ \mathbf{N} = \mathbb{N} \setminus 5\mathbb{N} & \mathbf{P} = 5\mathbb{N} \end{array}$$

Let's consider another impartial game that has some interesting properties. The game of Chomp was invented in the 1970's by David Gale, now a professor emeritus of mathematics at the University of California, Berkeley.

Example 10.1.5 (Chomp). In Chomp, two players take turns biting off a chunk of a rectangular bar of chocolate that is divided into squares. The bottom left corner of the bar has been removed and replaced with a broccoli floret. Each player, in his turn, chooses an uneaten chocolate square and removes it along with all the squares that lie above and to the right of it. The person who bites off the last piece of chocolate wins and the loser has to eat the broccoli.



Fig. 10.2. Two moves in a game of Chomp.

In Chomp, the terminal position is when all the chocolate is gone.

The graph for a small (2×3) bar can easily be constructed and \mathbf{N} and \mathbf{P} (and therefore a winning strategy) identified, see Figure 10.3. However, as the size of the bar increases, the graph becomes very large and a winning strategy difficult to find.




Fig. 10.3. Every move from a \mathbf{P} -position leads to an \mathbf{N} -position (bold black links); from every \mathbf{N} -position there is at least one move to a \mathbf{P} -position (red links).

Next we will formally prove that every progressively bounded impartial game has a winning strategy for one of the players.

Theorem 10.1.6. *In a progressively bounded impartial combinatorial game under normal play, all positions x lie in $\mathbf{N} \cup \mathbf{P}$.*

Proof. We proceed by induction on $B(x)$, where $B(x)$ is the maximum number of moves that a game from x might last (not just an upper bound).

Certainly, for all x such that $B(x) = 0$, we have that $x \in \mathbf{P}_0 \subseteq \mathbf{P}$. Assume the theorem is true for those positions x for which $B(x) \leq n$, and consider any position z satisfying $B(z) = n + 1$. Any move from z will take us to a position in $\mathbf{N} \cup \mathbf{P}$ by the inductive hypothesis.

There are two cases:

Case 1: Each move from z leads to a position in \mathbf{N} . Then $z \in \mathbf{P}_{n+1}$ by definition, and thus $z \in \mathbf{P}$.

Case 2: If it is not the case that every move from z leads to a position in \mathbf{N} , it must be that there is a move from z to some \mathbf{P}_n -position. In this case, by definition, $z \in \mathbf{N}_{n+1} \subseteq \mathbf{N}$.

Hence, all positions lie in $\mathbf{N} \cup \mathbf{P}$. □

Now, we have the tools to analyze Chomp. Recall that a legal move (for either player) in Chomp consists of identifying a square of chocolate and removing that square as well as all the squares above and to the right of it. There is only one terminal position where all the chocolate is gone and only broccoli remains.

Chomp is progressively bounded because we start with a finite number of squares and remove at least one in each turn. Thus, the above theorem implies that one of the players must have a winning strategy.

We will show that it's the first player that does. In fact, we will show something stronger: that starting from any position in which the remaining chocolate is rectangular, the next player to move can guarantee a win. The idea behind the proof is that of *strategy-stealing*. This is a general technique that we will use frequently throughout the chapter.

Theorem 10.1.7. *Starting from a position in which the remaining chocolate bar is rectangular of size greater than 1×1 , the next player to move has a winning strategy.*

Proof. Given a rectangular bar of chocolate R of size greater than 1×1 , let R^- be the result of chomping off the upper-right corner of R .

If $R^- \in \mathbf{P}$, then $R \in \mathbf{N}$, and a winning move is to chomp off the upper-right corner.

If $R^- \in \mathbf{N}$, then there is a move from R^- to some position S in \mathbf{P} . But if we can chomp R^- to get S , then chomping R in the same way will also give S , since the upper-right corner will be removed by any such chomp. Since there is a move from R to the position S in \mathbf{P} , it follows that $R \in \mathbf{N}$. □

Note that the proof does *not* show that chomping the upper-right hand corner is a winning move. In the 2×3 case, chomping the upper-right corner happens to be a winning move (since this leads to a move in \mathbf{P} , see Figure 10.3), but for the 3×3 case, chomping the upper-right corner is *not* a winning move. The strategy-stealing argument merely shows that a winning strategy for the first player must exist; it does not help us identify the strategy. In fact, it is an open research problem to describe a general winning strategy for Chomp.

Next we analyze the game of Nim, a particularly important progressively bounded impartial game.

10.1.1 Nim and Bouton's solution

Recall the game of Nim from the Introduction.

Example 10.1.8 (Nim). In Nim, there are several piles, each containing finitely many chips. A legal move is to remove any number of chips from a single pile. Two players alternate turns with the aim of removing the last chip. Thus, the terminal position is the one where there are no chips left.

Because Nim is progressively bounded, all the positions are in \mathbf{N} or \mathbf{P} , and one of the players has a winning strategy. We will be able to describe the winning strategy explicitly. We will see in section 10.1.3 that any progressively bounded impartial game is equivalent to a single Nim pile of a certain size. Hence, if the size of such a Nim pile can be determined, a winning strategy for the game can also be constructed explicitly.

As usual, we will analyze the game by working backwards from the terminal positions. We denote a position in the game by (n_1, n_2, \dots, n_k) , meaning that there are k piles of chips, and that the first has n_1 chips in it, the second has n_2 , and so on.

Certainly $(0, 1)$ and $(1, 0)$ are in \mathbf{N} . On the other hand, $(1, 1) \in \mathbf{P}$ because either of the two available moves leads to $(0, 1)$ or $(1, 0)$. We see that $(1, 2), (2, 1) \in \mathbf{N}$ because the next player can create the position $(1, 1) \in \mathbf{P}$. More generally, $(n, n) \in \mathbf{P}$ for $n \in \mathbb{N}$ and $(n, m) \in \mathbf{N}$ if $n, m \in \mathbb{N}$ are not equal.

Moving to three piles, we see that $(1, 2, 3) \in \mathbf{P}$, because whichever move the first player makes, the second can force two piles of equal size. It follows that $(1, 2, 3, 4) \in \mathbf{N}$ because the next player to move can remove the fourth pile.

To analyze $(1, 2, 3, 4, 5)$, we will need the following lemma:

Lemma 10.1.9. *For two Nim positions $X = (x_1, \dots, x_k)$ and $Y = (y_1, \dots, y_\ell)$, we denote the position $(x_1, \dots, x_k, y_1, \dots, y_\ell)$ by (X, Y) .*

- (i) *If X and Y are in \mathbf{P} , then $(X, Y) \in \mathbf{P}$.*
- (ii) *If $X \in \mathbf{P}$ and $Y \in \mathbf{N}$ (or vice versa), then $(X, Y) \in \mathbf{N}$.*
- (iii) *If $X, Y \in \mathbf{N}$, however, then (X, Y) can be either in \mathbf{P} or in \mathbf{N} .*

Proof. If (X, Y) has 0 chips, then X, Y , and (X, Y) are all \mathbf{P} -positions, so the lemma is true in this case.

Next, we suppose by induction that whenever (X, Y) has n or fewer chips,

$$X \in \mathbf{P} \text{ and } Y \in \mathbf{P} \text{ implies } (X, Y) \in \mathbf{P}$$

and

$$X \in \mathbf{P} \text{ and } Y \in \mathbf{N} \text{ implies } (X, Y) \in \mathbf{N}.$$

Suppose (X, Y) has at most $n + 1$ chips.

If $X \in \mathbf{P}$ and $Y \in \mathbf{N}$, then the next player to move can reduce Y to a position in \mathbf{P} , creating a \mathbf{P} - \mathbf{P} configuration with at most n chips, so by the inductive hypothesis it must be in \mathbf{P} . It follows that (X, Y) is in \mathbf{N} .

If $X \in \mathbf{P}$ and $Y \in \mathbf{P}$, then the next player to move must take chips from one of the piles (assume the pile is in Y without loss of generality). But moving Y from \mathbf{P} -position always results in a \mathbf{N} -position, so the resulting game is in a \mathbf{P} - \mathbf{N} position with at most n chips, which by the inductive hypothesis is an \mathbf{N} position. It follows that (X, Y) must be in \mathbf{P} .

For the final part of the lemma, note that any single pile is in \mathbf{N} , yet, as we saw above, $(1, 1) \in \mathbf{P}$ while $(1, 2) \in \mathbf{N}$. \square

Going back to our example, $(1, 2, 3, 4, 5)$ can be divided into two sub-games: $(1, 2, 3) \in \mathbf{P}$ and $(4, 5) \in \mathbf{N}$. By the lemma, we can conclude that $(1, 2, 3, 4, 5)$ is in \mathbf{N} .

The divide-and-sum method (using Lemma 10.1.9) is useful for analyzing Nim positions, but it doesn't immediately determine whether a given position is in \mathbf{N} or \mathbf{P} . The following ingenious theorem, proved in 1901 by a Harvard mathematics professor named Charles Bouton, gives a simple and general characterization of \mathbf{N} and \mathbf{P} for Nim. Before we state the theorem, we will need a definition.

Definition 10.1.10. The **Nim-sum** of $m, n \in \mathbb{N}$ is the following operation: Write m and n in binary form, and sum the digits in each column modulo 2. The resulting number, which is expressed in binary, is the Nim-sum of m and n . We denote the Nim-sum of m and n by $m \oplus n$.

Equivalently, the Nim-sum of a collection of values (m_1, m_2, \dots, m_k) is the sum of all the powers of 2 that occurred an odd number of times when each of the numbers m_i is written as a sum of powers of 2.

If $m_1 = 3$, $m_2 = 9$, $m_3 = 13$, in powers of 2 we have:

$$\begin{aligned} m_1 &= 0 \times 2^3 + 0 \times 2^2 + 1 \times 2^1 + 1 \times 2^0 \\ m_2 &= 1 \times 2^3 + 0 \times 2^2 + 0 \times 2^1 + 1 \times 2^0 \\ m_3 &= 1 \times 2^3 + 1 \times 2^2 + 0 \times 2^1 + 1 \times 2^0. \end{aligned}$$

The powers of 2 that appear an odd number of times are $2^0 = 1$, $2^1 = 2$, and $2^2 = 4$, so $m_1 \oplus m_2 \oplus m_3 = 1 + 2 + 4 = 7$.

We can compute the Nim-sum efficiently by using binary notation:

decimal	binary
3	0011
9	1001
13	1101
7	0111

Theorem 10.1.11 (Bouton's Theorem). *A Nim position $x = (x_1, x_2, \dots, x_k)$ is in \mathbf{P} if and only if the Nim-sum of its components is 0.*

To illustrate the theorem, consider the starting position (1, 2, 3):

decimal	binary
1	01
2	10
3	11
0	00

Summing the two columns of the binary expansions modulo two, we obtain 00. The theorem affirms that $(1, 2, 3) \in \mathbf{P}$. Now, we prove Bouton's theorem.

Proof of Theorem 10.1.11. Define Z to be those positions with Nim-sum zero.

Suppose that $x = (x_1, \dots, x_k) \in Z$, i.e., $x_1 \oplus \dots \oplus x_k = 0$. Maybe there are no chips left, but if there are some left, suppose that we remove some chips from a pile ℓ , leaving $x'_\ell < x_\ell$ chips. The Nim-sum of the resulting piles is $x_1 \oplus \dots \oplus x_{\ell-1} \oplus x'_\ell \oplus x_{\ell+1} \oplus \dots \oplus x_k = x'_\ell \oplus x_\ell \neq 0$. Thus any move from a position in Z leads to a position not in Z .

Suppose that $x = (x_1, x_2, \dots, x_k) \notin Z$. Let $s = x_1 \oplus \dots \oplus x_k \neq 0$. There are an odd number of values of $i \in \{1, \dots, k\}$ for which the binary expression for x_i has a 1 in the position of the left-most 1 in the expression for s . Choose one such i . Note that $x_i \oplus s < x_i$, because $x_i \oplus s$ has no 1 in this left-most position, and so is less than any number whose binary expression does. Consider the move in which a player removes $x_i - x_i \oplus s$ chips from the i^{th} pile. This changes x_i to $x_i \oplus s$. The Nim-sum of the resulting position $(x_1, \dots, x_{i-1}, x_i \oplus s, x_{i+1}, \dots, x_k) = 0$, so this new position lies in Z . Thus, for any position $x \notin Z$, there exists a move from x leading to a position in Z .


For any Nim-position that is not in Z , the first player can adopt the strategy of always moving to a position in Z . The second player, if he has any moves, will necessarily always move to a position not in Z , always leaving the first player with a move to make. Thus any position that is not in Z is an \mathbf{N} -position. Similarly, if the game starts in a position in Z , the

second player can guarantee a win by always moving to a position in Z when it is his turn. Thus any position in Z is a **P**-position. \square

10.1.2 Other impartial games

Example 10.1.12 (Staircase Nim). This game is played on a staircase of n steps. On each step j for $j = 1, \dots, n$ is a stack of coins of size $x_j \geq 0$.

Each player, in his turn, moves one or more coins from a stack on a step j and places them on the stack on step $j - 1$. Coins reaching the ground (step 0) are removed from play. The game ends when all coins are on the ground, and the last player to move wins.



Corresponding move in Nim on odd-numbered steps.



Fig. 10.4. A move in Staircase Nim, in which 2 coins are moved from step 3 to step 2. Considering the odd stairs only, the above move is equivalent to the move in regular Nim from $(3, 5)$ to $(3, 3)$.

As it turns out, the **P**-positions in Staircase Nim are the positions such that the stacks of coins on the odd-numbered steps correspond to a **P**-position in Nim.

We can view moving y coins from an odd-numbered step to an even-numbered one as corresponding to the legal move of removing y chips in Nim. What happens when we move coins from an even numbered step to an odd numbered one?

If a player moves z coins from an even numbered step to an odd numbered one, his opponent may then move the coins to the next even-numbered step; that is, she may repeat her opponent's move at one step lower. This move restores the Nim-sum on the odd-numbered steps to its previous value, and ensures that such a move plays no role in the outcome of the game.

Now, we will look at another game, called Rims, which, as we will see, is also just Nim in disguise.

Example 10.1.13 (Rims). A starting position consists of a finite number of dots in the plane and a finite number of continuous loops that do not intersect. Each loop may pass through any number of dots, and must pass through at least one.

Each player, in his turn, draws a new loop that does not intersect any other loop. The goal is to draw the last such loop.




Fig. 10.5. Two moves in a game of Rims.

For a given position of Rims, we can divide the dots that have no loop through them into equivalence classes as follows: Each class consists of a set of dots that can be reached from a particular dot via a continuous path that does not cross any loops.

To see the connection to Nim, think of each class of dots as a pile of chips. A loop, because it passes through at least one dot, in effect, removes at least one chip from a pile, and splits the remaining chips into two new piles. This last part is not consistent with the rules of Nim unless the player draws the loop so as to leave the remaining chips in a single pile.




Fig. 10.6. Equivalent sequence of moves in Nim with splittings allowed.

Thus, Rims is equivalent to a variant of Nim where players have the option of splitting a pile into two piles after removing chips from it. As the following theorem shows, the fact that players have the option of splitting piles has no impact on the analysis of the game.

Theorem 10.1.14. *The sets \mathbf{N} and \mathbf{P} coincide for Nim and Rims.*

Proof. Thinking of a position in Rims as a collection of piles of chips, rather than as dots and loops, we write \mathbf{P}_{Nim} and \mathbf{N}_{Nim} for the \mathbf{P} - and \mathbf{N} -positions for the game of Nim (these sets are described by Bouton's theorem).

From any position in \mathbf{N}_{Nim} , we may move to \mathbf{P}_{Nim} by a move in Rims, because each Nim move is legal in Rims.

Next we consider a position $x \in \mathbf{P}_{\text{Nim}}$. Maybe there are no moves from x , but if there are, any move reduces one of the piles, and possibly splits it into two piles. Say the ℓ^{th} pile goes from x_ℓ to $x'_\ell < x_\ell$, and possibly splits into u, v where $u + v < x_\ell$.

Because our starting position x was a \mathbf{P}_{Nim} -position, its Nim-sum was

$$x_1 \oplus \cdots \oplus x_\ell \oplus \cdots \oplus x_k = 0.$$

The Nim-sum of the new position is either

$$x_1 \oplus \cdots \oplus x'_\ell \oplus \cdots \oplus x_k = x_\ell \oplus x'_\ell \neq 0,$$

(if the pile was not split), or else

$$x_1 \oplus \cdots \oplus (u \oplus v) \oplus \cdots \oplus x_k = x_\ell \oplus u \oplus v.$$

Notice that the Nim-sum $u \oplus v$ of u and v is at most the ordinary sum $u + v$: This is because the Nim-sum involves omitting certain powers of 2 from the expression for $u + v$. Hence, we have

$$u \oplus v \leq u + v < x_\ell.$$

Thus, whether or not the pile is split, the Nim-sum of the resulting position is nonzero, so any Rims move from a position in \mathbf{P}_{Nim} is to a position in \mathbf{N}_{Nim} .

Thus the strategy of always moving to a position in \mathbf{P}_{Nim} (if this is possible) will guarantee a win for a player who starts in an \mathbf{N}_{Nim} -position, and if a player starts in a \mathbf{P}_{Nim} -position, this strategy will guarantee a win for the second player. Thus $\mathbf{N}_{\text{Rims}} = \mathbf{N}_{\text{Nim}}$ and $\mathbf{P}_{\text{Rims}} = \mathbf{P}_{\text{Nim}}$. \square

The following examples are particularly tricky variants of Nim.

Example 10.1.15 (Moore's Nim_k). This game is like Nim, except that each player, in his turn, is allowed to remove any number of chips from at most k of the piles.

Write the binary expansions of the pile sizes (n_1, \dots, n_ℓ) :

$$\begin{aligned} n_1 &= n_1^{(m)} \cdots n_1^{(0)} = \sum_{j=0}^m n_1^{(j)} 2^j, \\ &\vdots \\ n_\ell &= n_\ell^{(m)} \cdots n_\ell^{(0)} = \sum_{j=0}^m n_\ell^{(j)} 2^j, \end{aligned}$$

where each $n_i^{(j)}$ is either 0 or 1.

Theorem 10.1.16 (Moore's Theorem). *For Moore's Nim_k ,*

$$\mathbf{P} = \left\{ (n_1, \dots, n_\ell) : \sum_{i=1}^{\ell} n_i^{(j)} \equiv 0 \pmod{k+1} \text{ for each } j \right\}.$$

The notation " $a \equiv b \pmod{m}$ " means that $a - b$ is evenly divisible by m , i.e., that $(a - b)/m$ is an integer.

Proof of Theorem 10.1.16. Let Z denote the right-hand-side of the above expression. We will show that every move from a position in Z leads to a position not in Z , and that for every position not in Z , there is a move to a position in Z . As with ordinary Nim, it will follow that a winning strategy is to always move to position in Z if possible, and consequently $\mathbf{P} = Z$.

Take any move from a position in Z , and consider the left-most column for which this move changes the binary expansion of at least one of the pile numbers. Any change in this column must be from one to zero. The existing sum of the ones and zeros (mod $(k + 1)$) is zero, and we are adjusting at most k piles. Because ones are turning into zeros in this column, we are decreasing the sum in that column and by at least 1 and at most k , so the resulting sum in this column cannot be congruent to 0 modulo $k + 1$. We have verified that no move starting from Z takes us back to Z .

We must also check that for each position x not in Z , we can find a move to some y that is in Z . The way we find this move is a little bit tricky, and we illustrate it in the following example:

We write the pile sizes of x in binary, and make changes to the bits so that the sum of the bits in each column congruent to 0 modulo $k + 1$. For these changes to correspond to a valid move in Moore's Nim_k , we are constrained to change the bits in at most k rows, and for any row that we change, the left-most bit that is changed must be a change from a 1 to a 0.

To make these changes, we scan the bits columns from the most significant




Fig. 10.7. Example move in Moore’s Nim₄ from a position not in Z to a position in Z . When a row becomes activated, the bit is boxed, and active rows are shaded. The bits in only 4 rows are changed, and the resulting column sums are all divisible by 5.

to the least significant. When we scan, we can “activate” a row if it contains a 1 in the given column which we change to a 0, and once a row is activated, we may change the remaining bits in the row in any fashion.

At a given column, let a be the number of rows that have already been activated ($0 \leq a \leq k$), and let s be the sum of the bits in the rows that have not been activated. Let $b = (s + a) \bmod (k + 1)$. If $b \leq a$, then we can set the bits in b of the active rows to 0 and $a - b$ of the active rows to 1. The new column sum is then $s + a - b$, which is evenly divisible by $k + 1$. Otherwise, $a < b \leq k$, and $b - a = s \bmod (k + 1) \leq s$, so we may activate $b - a$ inactive rows that have a 1 in that column, and set the bits in all the active rows in that column to 0. The column sum is then $s - (b - a)$, which is again evenly divisible by $k + 1$, and the number of active rows remains at most k . Continuing in this fashion results in a position in Z , by reducing at most k of the piles. □

Example 10.1.17 (Wythoff Nim). A position in this game consists of two piles of sizes m and n . The legal moves are those of Nim, with one addition: players may remove equal numbers of chips from both piles in a single move. This extra move prevents the positions $\{(n, n) : n \in \mathbb{N}\}$ from being **P**-positions.

This game has a very interesting structure. We can say that a position consists of a pair (m, n) of natural numbers, such that $m, n \geq 0$. A legal move is one of the following:

Reduce m to some value between 0 and $m - 1$ without changing n , reducing n to some value between 0 and $n - 1$ without changing m , or reducing each of m and n by the same amount. The one who reaches $(0, 0)$ is the winner.

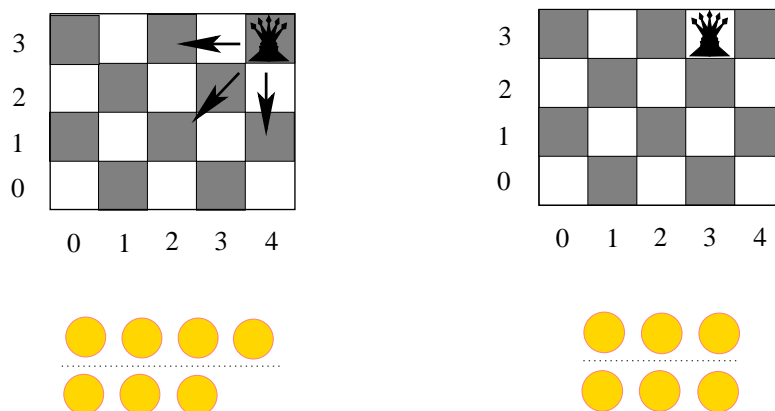


Fig. 10.8. Wythoff Nim can be viewed as the following game played on a chess board. Consider an $m \times n$ section of a chess-board. The players take turns moving a queen, initially positioned in the upper right corner, either left, down, or diagonally toward the lower left. The player that moves the queen into the bottom left corner wins. If the position of the queen at every turn is denoted by (x, y) , with $1 \leq x \leq m, 1 \leq y \leq n$, we see that the game corresponds to Wythoff Nim.

To analyze Wythoff Nim (and other games), we define

$$\text{mex}(S) = \min\{n \geq 0 : n \notin S\},$$

for $S \subseteq \{0, 1, \dots\}$ (the term “mex” stands for “minimal excluded value”). For example, $\text{mex}(\{0, 1, 2, 3, 5, 7, 12\}) = 4$. Consider the following recursive definition of two sequences of natural numbers: For each $k \geq 0$,

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

Notice that when $k = 0$, we have $a_0 = \text{mex}(\{\}) = 0$ and $b_0 = a_0 + 0 = 0$. The first few values of these two sequences are

k	0	1	2	3	4	5	6	7	8	9	...
a_k	0	1	3	4	6	8	9	11	12	14	...
b_k	0	2	5	7	10	13	15	18	20	23	...

(For example, $a_4 = \text{mex}(\{0, 1, 3, 4, 0, 2, 5, 7\}) = 6$ and $b_4 = a_4 + 4 = 10$.)

Theorem 10.1.18. *Each natural number greater than zero is equal to precisely one of the a_i 's or b_i 's. That is, $\{a_i\}_{i=1}^\infty$ and $\{b_i\}_{i=1}^\infty$ form a partition of \mathbb{N}^* .*

Proof. First we will show, by induction on j , that $\{a_i\}_{i=1}^j$ and $\{b_i\}_{i=1}^j$ are disjoint strictly increasing subsets of \mathbb{N}^* . This is vacuously true when

$j = 0$, since then both sets are empty. Now suppose that $\{a_i\}_{i=1}^{j-1}$ is strictly increasing and disjoint from $\{b_i\}_{i=1}^{j-1}$, which, in turn, is strictly increasing. By the definition of the a_i 's, we have that both a_j and a_{j-1} are excluded from $\{a_0, \dots, a_{j-2}, b_0, \dots, b_{j-2}\}$, but a_{j-1} is the smallest such excluded value, so $a_{j-1} \leq a_j$. By the definition of a_j , we also have $a_j \neq a_{j-1}$ and $a_j \notin \{b_0, \dots, b_{j-1}\}$, so in fact $\{a_i\}_{i=1}^j$ and $\{b_i\}_{i=1}^{j-1}$ are disjoint strictly increasing sequences. Moreover, for each $i < j$ we have $b_j = a_j + j > a_i + j > a_i + i = b_i > a_i$, so $\{a_i\}_{i=1}^j$ and $\{b_i\}_{i=1}^j$ are strictly increasing and disjoint from each other, as well.

To see that every integer is covered, we show by induction that

$$\{1, \dots, j\} \subset \{a_i\}_{i=1}^j \cup \{b_i\}_{i=1}^j.$$

This is clearly true when $j = 0$. If it is true for j , then either $j + 1 \in \{a_i\}_{i=1}^j \cup \{b_i\}_{i=1}^j$ or it is excluded, in which case $a_{j+1} = j + 1$. \square

It is easy to check the following theorem:

Theorem 10.1.19. *The set of \mathbf{P} -positions for Wythoff Nim is exactly $\hat{P} := \{(a_k, b_k) : k = 0, 1, 2, \dots\} \cup \{(b_k, a_k) : k = 0, 1, 2, \dots\}$.*

Proof. First we check that any move from a position $(a_k, b_k) \in \hat{P}$ is to a position not in \hat{P} . If we reduce both piles, then the gap between them remains k , and the only position in \hat{P} with gap k is (a_k, b_k) . If we reduce the first pile, the number b_k only occurs with a_k in \hat{P} , so we are taken to a position not in \hat{P} , and similarly, reducing the second pile also leads to a position not in \hat{P} .

Let (m, n) be a position not in \hat{P} , say $m \leq n$, and let $k = n - m$. If $(m, n) > (a_k, b_k)$, we can reduce both piles of chips to take the configuration to (a_k, b_k) , which is in \hat{P} . If $(m, n) < (a_k, b_k)$, then either $m = a_j$ or $m = b_j$ for some $j < k$. If $m = a_j$, then we can remove $k - j$ chips from the second pile to take the configuration to $(a_j, b_j) \in \hat{P}$. If $m = b_j$, then $n \geq m = b_j > a_j$, so we can remove chips from the second pile to take the state to $(b_j, a_j) \in \hat{P}$.

Thus $\mathbf{P} = \hat{P}$. \square

It turns out that there is there a fast, non-recursive, method to decide if a given position is in \mathbf{P} :

Theorem 10.1.20. $a_k = \lfloor k(1 + \sqrt{5})/2 \rfloor$ and $b_k = \lfloor k(3 + \sqrt{5})/2 \rfloor$.

$\lfloor x \rfloor$ denotes the “floor of x ,” i.e., the greatest integer that is $\leq x$. Similarly, $\lceil x \rceil$ denotes the “ceiling of x ,” the smallest integer that is $\geq x$.

Proof of Theorem 10.1.20. Consider the following sequences positive integers: Fix any irrational $\theta \in (0, 1)$, and set

$$\alpha_k(\theta) = \lfloor k/\theta \rfloor, \quad \beta_k(\theta) = \lfloor k/(1 - \theta) \rfloor.$$

We claim that $\{\alpha_k(\theta)\}_{k=1}^\infty$ and $\{\beta_k(\theta)\}_{k=1}^\infty$ form a partition of \mathbb{N}^* . Clearly, $\alpha_k(\theta) < \alpha_{k+1}(\theta)$ and $\beta_k(\theta) < \beta_{k+1}(\theta)$ for any k . Observe that $\alpha_k(\theta) = N$ if and only if

$$k \in I_N := [N\theta, N\theta + \theta),$$

and $\beta_\ell(\theta) = N$ if and only if

$$-\ell + N \in J_N := (N\theta + \theta - 1, N\theta].$$

These events cannot both happen with $\theta \in (0, 1)$ unless $N = 0$, $k = 0$, and $\ell = 0$. Thus, $\{\alpha_k(\theta)\}_{k=1}^\infty$ and $\{\beta_k(\theta)\}_{k=1}^\infty$ are disjoint. On the other hand, so long as $N \neq -1$, at least one of these events must occur for some k or ℓ , since $J_N \cup I_N = ((N + 1)\theta - 1, (N + 1)\theta)$ contains an integer when $N \neq -1$ and θ is irrational. This implies that each positive integer N is contained in either $\{\alpha_k(\theta)\}_{k=1}^\infty$ or $\{\beta_k(\theta)\}_{k=1}^\infty$.

Does there exist a $\theta \in (0, 1)$ for which

$$\alpha_k(\theta) = a_k \quad \text{and} \quad \beta_k(\theta) = b_k? \tag{10.1}$$

We will show that there is only one θ for which this is true.

Because $b_k = a_k + k$, (10.1) implies that $\lfloor k/\theta \rfloor + k = \lfloor k/(1 - \theta) \rfloor$. Dividing by k we get

$$\frac{1}{k} \lfloor k/\theta \rfloor + 1 = \frac{1}{k} \lfloor k/(1 - \theta) \rfloor,$$

and taking a limit as $k \rightarrow \infty$ we find that

$$1/\theta + 1 = 1/(1 - \theta). \tag{10.2}$$

Thus, $\theta^2 + \theta - 1 = 0$. The only solution in $(0, 1)$ is $\theta = (\sqrt{5} - 1)/2 = 2/(1 + \sqrt{5})$.

We now fix $\theta = 2/(1 + \sqrt{5})$ and let $\alpha_k = \alpha_k(\theta)$, $\beta_k = \beta_k(\theta)$. Note that (10.2) holds for this particular θ , so that

$$\lfloor k/(1 - \theta) \rfloor = \lfloor k/\theta \rfloor + k.$$

This means that $\beta_k = \alpha_k + k$. We need to verify that

$$\alpha_k = \text{mex} \{ \alpha_0, \dots, \alpha_{k-1}, \beta_0, \dots, \beta_{k-1} \}.$$

We checked earlier that α_k is not one of these values. Why is it equal to

their mex? Suppose, toward a contradiction, that z is the mex, and $\alpha_k \neq z$. Then $z < \alpha_k \leq \alpha_\ell \leq \beta_\ell$ for all $\ell \geq k$. Since z is defined as a mex, $z \neq \alpha_i, \beta_i$ for $i \in \{0, \dots, k-1\}$, so z is missed and hence $\{\alpha_k\}_{k=1}^\infty$ and $\{\beta_k\}_{k=1}^\infty$ would not be a partition of \mathbb{N}^* , a contradiction. \square

10.1.3 Impartial games and the Sprague-Grundy theorem

In this section, we will develop a general framework for analyzing all progressively bounded impartial combinatorial games. As in the case of Nim, we will look at sums of games and develop a tool that enables us to analyze any impartial combinatorial game under normal play as if it were a Nim pile of a certain size.

Definition 10.1.21. The sum of two combinatorial games, G_1 and G_2 , is a game G in which each player, in his turn, chooses one of G_1 or G_2 in which to play. The terminal positions in G are (t_1, t_2) , where t_i is a terminal position in G_i for $i \in \{1, 2\}$. We write $G = G_1 + G_2$.

Example 10.1.22. The sum of two Nim games X and Y is the game (X, Y) as defined in Lemma 10.1.9 of the previous section.

It is easy to see that Lemma 10.1.9 generalizes to the sum of any two progressively bounded combinatorial games:

Theorem 10.1.23. Suppose G_1 and G_2 are progressively bounded impartial combinatorial games.

- (i) If $x_1 \in \mathbf{P}_{G_1}$ and $x_2 \in \mathbf{P}_{G_2}$, then $(x_1, x_2) \in \mathbf{P}_{G_1+G_2}$.
- (ii) If $x_1 \in \mathbf{P}_{G_1}$ and $x_2 \in \mathbf{N}_{G_2}$, then $(x_1, x_2) \in \mathbf{N}_{G_1+G_2}$.
- (iii) If $x_1 \in \mathbf{N}_{G_1}$ and $x_2 \in \mathbf{N}_{G_2}$, then (x_1, x_2) could be in either $\mathbf{N}_{G_1+G_2}$ or $\mathbf{P}_{G_1+G_2}$.

Proof. In the proof for Lemma 10.1.9 for Nim, replace the number of chips with $B(x)$, the maximum number of moves in the game. \square

Definition 10.1.24. Consider two arbitrary progressively bounded combinatorial games G_1 and G_2 with positions x_1 and x_2 . If for any third such game G_3 and position x_3 , the outcome of (x_1, x_3) in $G_1 + G_3$ (i.e., whether it's an **N**- or **P**-position) is the same as the outcome of (x_2, x_3) in $G_2 + G_3$, then we say that (G_1, x_1) and (G_2, x_2) are **equivalent**.

It follows from Theorem 10.1.23 that in any two progressively bounded impartial combinatorial games, the **P**-positions are equivalent to each other.

In Exercise 10.12 you will prove that this notion of equivalence for games

defines an *equivalence relation*. In Exercise 10.13 you will prove that two impartial games are equivalent if and only if their sum is a **P**-position. In Exercise 10.14 you will show that if G_1 and G_2 are equivalent, and G_3 is a third game, then $G_1 + G_3$ and $G_2 + G_3$ are equivalent.

Example 10.1.25. The Nim game with starting position $(1, 3, 6)$ is equivalent to the Nim game with starting position (4) , because the Nim-sum of the sum game $(1, 3, 4, 6)$ is zero. More generally, the position (n_1, \dots, n_k) is equivalent to $(n_1 \oplus \dots \oplus n_k)$ because the Nim-sum of $(n_1, \dots, n_k, n_1 \oplus \dots \oplus n_k)$ is zero.

If we can show that an arbitrary impartial game (G, x) is equivalent to a single Nim pile (n) , we can immediately determine whether (G, x) is in **P** or in **N**, since the only single Nim pile in **P** is (0) .

We need a tool that will enable us to determine the size n of a Nim pile equivalent to an arbitrary position (G, x) .

Definition 10.1.26. Let G be a progressively bounded impartial combinatorial game under normal play. Its **Sprague-Grundy function** g is defined recursively as follows:

$$g(x) = \text{mex}(\{g(y) : x \rightarrow y \text{ is a legal move}\}).$$

Note that the Sprague-Grundy value of any terminal position is $\text{mex}(\emptyset) = 0$. In general, the Sprague-Grundy function has the following key property:

Lemma 10.1.27. *In a progressively bounded impartial combinatorial game, the Sprague-Grundy value of a position is 0 if and only if it is a **P**-position.*

Proof. Proceed as in the proof of Theorem 10.1.11 — define \hat{P} to be those positions x with $g(x) = 0$, and \hat{N} to be all other positions. We claim that

$$\hat{P} = \mathbf{P} \quad \text{and} \quad \hat{N} = \mathbf{N}.$$

To show this, we need to show first that $t \in \hat{P}$ for every terminal position t . Second, that for all $x \in \hat{N}$, there exists a move from x leading to \hat{P} . Finally, we need to show that for every $y \in \hat{P}$, all moves from y lead to \hat{N} .

All these are a direct consequence of the definition of mex . The details of the proof are left as an exercise (Ex. 10.15). \square

Let's calculate the Sprague-Grundy function for a few examples.

Example 10.1.28 (The m -Subtraction game). In the m -subtraction game with subtraction set $\{a_1, \dots, a_m\}$, a position consists of a pile of chips, and a legal move is to remove from the pile a_i chips, for some $i \in \{1, \dots, m\}$. The player who removes the last chip wins.

Consider a 3-subtraction game with subtraction set $\{1, 2, 3\}$. The following table summarizes a few values of its Sprague-Grundy function:

x	0	1	2	3	4	5	6
$g(x)$	0	1	2	3	0	1	2

In general, $g(x) = x \bmod 4$.

Example 10.1.29 (The Proportional Subtraction game). A position consists of a pile of chips. A legal move from a position with n chips is to remove any positive number of chips that is at most $\lceil n/2 \rceil$.

Here, the first few values of the Sprague-Grundy function are:

x	0	1	2	3	4	5	6
$g(x)$	0	1	0	2	1	3	0

Example 10.1.30. Note that the Sprague-Grundy value of any Nim pile (n) is just n .

Now we are ready to state the Sprague-Grundy theorem, which allows us relate impartial games to Nim:

Theorem 10.1.31 (Sprague-Grundy Theorem). *Let G be a progressively bounded impartial combinatorial game under normal play with starting position x . Then G is equivalent to a single Nim pile of size $g(x) \geq 0$, where $g(x)$ is the Sprague-Grundy function evaluated at the starting position x .*

Proof. We let $G_1 = G$, and G_2 be the Nim pile of size $g(x)$. Let G_3 be any other combinatorial game under normal play. One player or the other, say player A, has a winning strategy for $G_2 + G_3$. We claim that player A also has a winning strategy for $G_1 + G_3$.

For each move of $G_2 + G_3$ there is an associated move in $G_1 + G_3$: If one of the players moves in G_3 when playing $G_2 + G_3$, this corresponds to the same move in G_3 when playing $G_1 + G_3$. If one of the players plays in G_2 when playing $G_2 + G_3$, say by moving from a Nim pile with y chips to a Nim pile with $z < y$ chips, then the corresponding move in $G_1 + G_3$ would be to move in G_1 from a position with Sprague-Grundy value y to a position with Sprague-Grundy value z (such a move exists by the definition of the Sprague-Grundy function). There may be extra moves in $G_1 + G_3$ that do not correspond to any move $G_2 + G_3$, namely, it may be possible to play in G_1 from a position with Sprague-Grundy value y to a position with Sprague-Grundy value $z > y$.

When playing in $G_1 + G_3$, player A can pretend that the game is really

$G_2 + G_3$. If player A's winning strategy is some move in $G_2 + G_3$, then A can play the corresponding move in $G_1 + G_3$, and pretends that this move was made in $G_2 + G_3$. If A's opponent makes a move in $G_1 + G_3$ that corresponds to a move in $G_2 + G_3$, then A pretends that this move was made in $G_2 + G_3$. But player A's opponent could also make a move in $G_1 + G_3$ that does not correspond to any move of $G_2 + G_3$, by moving in G_1 and increasing the Sprague-Grundy value of the position in G_1 from y to $z > y$. In this case, by the definition of the Sprague-Grundy value, player A can simply play in G_1 and move to a position with Sprague-Grundy value y . These two turns correspond to no move, or a pause, in the game $G_2 + G_3$. Because $G_1 + G_3$ is progressively bounded, $G_2 + G_3$ will not remain paused forever. Since player A has a winning strategy for the game $G_2 + G_3$, player A will win this game that A is pretending to play, and this will correspond to a win in the game $G_1 + G_3$. Thus whichever player has a winning strategy in $G_2 + G_3$ also has a winning strategy in $G_1 + G_3$, so G_1 and G_2 are equivalent games. \square

We can use this theorem to find the **P**- and **N**-positions of a particular impartial, progressively bounded game under normal play, provided we can evaluate its Sprague-Grundy function.

For example, recall the 3-subtraction game we considered in Example 10.1.28. We determined that the Sprague-Grundy function of the game is $g(x) = x \bmod 4$. Hence, by the Sprague-Grundy theorem, 3-subtraction game with starting position x is equivalent to a single Nim pile with $x \bmod 4$ chips. Recall that $(0) \in \mathbf{P}_{\text{Nim}}$ while $(1), (2), (3) \in \mathbf{N}_{\text{Nim}}$. Hence, the **P**-positions for the Subtraction game are the natural numbers that are divisible by four.

Corollary 10.1.32. *Let G_1 and G_2 be two progressively bounded impartial combinatorial games under normal play. These games are equivalent if and only if the Sprague-Grundy values of their starting positions are the same.*

Proof. Let x_1 and x_2 denote the starting positions of G_1 and G_2 . We saw already that G_1 is equivalent to the Nim pile $(g(x_1))$, and G_2 is equivalent to $(g(x_2))$. Since equivalence is transitive, if the Sprague-Grundy values $g(x_1)$ and $g(x_2)$ are the same, G_1 and G_2 must be equivalent. Now suppose $g(x_1) \neq g(x_2)$. We have that $G_1 + (g(x_1))$ is equivalent to $(g(x_1)) + (g(x_1))$ which is a **P**-position, while $G_2 + (g(x_1))$ is equivalent to $(g(x_2)) + (g(x_1))$, which is an **N**-position, so G_1 and G_2 are not equivalent. \square

The following theorem gives a way of finding the Sprague-Grundy function of the sum game $G_1 + G_2$, given the Sprague-Grundy functions of the component games G_1 and G_2 .

Theorem 10.1.33 (Sum Theorem). *Let G_1 and G_2 be a pair of impartial combinatorial games and x_1 and x_2 positions within those respective games. For the sum game $G = G_1 + G_2$,*

$$g(x_1, x_2) = g_1(x_1) \oplus g_2(x_2), \quad (10.3)$$

where g , g_1 , and g_2 respectively denote the Sprague-Grundy functions for the games G , G_1 , and G_2 , and \oplus is the Nim-sum.

Proof. It is straightforward to see that $G_1 + G_1$ is a **P**-position, since the second player can always just make the same moves that the first player makes but in the other copy of the game. Thus $G_1 + G_2 + G_1 + G_2$ is a **P**-position. Since G_1 is equivalent to $(g(x_1))$, G_2 is equivalent to $(g(x_2))$, and $G_1 + G_2$ is equivalent to $(g(x_1, x_2))$, we have that $(g(x_1), g(x_2), g(x_1, x_2))$ is a **P**-position. From our analysis of Nim, we know that this happens only when the three Nim piles have Nim-sum zero, and hence $g(x_1, x_2) = g(x_1) \oplus g(x_2)$. \square

Let's use the Sprague-Grundy and the Sum Theorems to analyze a few games.

Example 10.1.34. (4 or 5) There are two piles of chips. Each player, in his turn, removes either one to four chips from the first pile or one to five chips from the second pile.

Our goal is to figure out the **P**-positions for this game. Note that the game is of the form $G_1 + G_2$ where G_1 is a 4-subtraction game and G_2 is a 5-subtraction game. By analogy with the 3-subtraction game, $g_1(x) = x \bmod 5$ and $g_2(y) = y \bmod 6$. By the Sum Theorem, we have that $g(x, y) = (x \bmod 5) \oplus (y \bmod 6)$. We see that $g(x, y) = 0$ if and only if $x \bmod 5 = y \bmod 6$.

The following example bears no obvious resemblance to Nim, yet we can use the Sprague-Grundy function to analyze it.

Example 10.1.35 (Green Hackenbush). Green Hackenbush is played on a finite graph with one distinguished vertex r , called the root, which may be thought of as the base on which the rest of the structure is standing. (Recall that a graph is a collection of vertices and edges that connect unordered pairs of vertices.) In his turn, a player may remove an edge from the graph. This causes not only that edge to disappear, but all of the structure that relies on it — the edges for which every path to the root travels through the removed edge.

The goal for each player is to remove the last edge from the graph.

We talk of “Green” Hackenbush because there is a partisan variant of the game in which edges are colored red, blue, or green, and one player can remove red or green edges, while the other player can remove blue or green edges.

Note that if the original graph consists of a finite number of paths, each of which ends at the root, then Green Hackenbush is equivalent to the game of Nim, in which the number of piles is equal to the number of paths, and the number of chips in a pile is equal to the length of the corresponding path.

To handle the case in which the graph is a tree, we will need the following lemma:

Lemma 10.1.36 (Colon Principle). *The Sprague-Grundy function of Green Hackenbush on a tree is unaffected by the following operation: For any two branches of the tree meeting at a vertex, replace these two branches by a path emanating from the vertex whose length is the Nim-sum of the Sprague-Grundy functions of the two branches.*

Proof. We will only sketch the proof. For the details, see Ferguson [Fer08, I-42].

If the two branches consist simply of paths, or “stalks,” emanating from a given vertex, then the result follows from the fact that the two branches form a two-pile game of Nim, using the direct sum theorem for the Sprague-Grundy functions of two games. More generally, we may perform the replacement operation on any two branches meeting at a vertex by iterating replacing pairs of stalks meeting inside a given branch until each of the two branches itself has become a stalk. \square




Fig. 10.9. Combining branches in a tree of Green Hackenbush.

As a simple illustration, see Fig. 10.9. The two branches in this case are stalks of lengths 2 and 3. The Sprague-Grundy values of these stalks are 2 and 3, and their Nim-sum is 1.

For a more in-depth discussion of Hackenbush and references, see Ferguson [Fer08, Part I, Sect. 6] or [BCG82a].

Next we leave the impartial and discuss a few interesting partisan games.

10.2 Partisan games

A combinatorial game that is not impartial is called **partisan**. In a partisan games the legal moves for some positions may be different for each player. Also, in some partisan games, the terminal positions may be divided into those that have a win for player I and those that have a win for player II.

Hex is an important partisan game that we described in the introduction. In Hex, one player (Blue) can only place blue tiles on the board and the other player (Yellow) can only place yellow tiles, and the resulting board configurations are different, so the legal moves for the two players are different. One could modify Hex to allow both players to place tiles of either color (though neither player will want to place a tile of the other color), so that both players will have the same set of legal moves. This modified Hex is still partisan because the winning configurations for the two players are different: positions with a blue crossing are winning for Blue and those with a yellow crossing are winning for Yellow.

Typically in a partisan game not all positions may be reachable by every player from a given starting position. We can illustrate this with the game of Hex. If the game is started on an empty board, the player that moves first can never face a position where the number of blue and yellow hexagons on the board is different.

In some partisan games there may be additional terminal positions which mean that neither of the players wins. These can be labelled “ties” or “draws” (as in Chess, when there is a stalemate).

While an impartial combinatorial game can be represented as a graph with a single edge-set, a partisan game is most often given by a single set of nodes and two sets of edges that represent legal moves available to either player. Let X denote the set of positions and E_I, E_{II} be the two edge-sets for players I and II respectively. If (x, y) is a legal move for player $i \in \{I, II\}$ then $(x, y) \in E_i$ and we say that y is a **successor** of x . We write $S_i(x) = \{y : (x, y) \in E_i\}$. The edges are directed if the moves are irreversible.

A partisan game follows the **normal play condition** if the first player who cannot move loses. The **misère play condition** is the opposite, i.e., the first player who cannot move wins. In games such as Hex, some terminal nodes are winning for one player or the other, regardless of whose turn it is when the game arrived in that position. Such games are equivalent to normal play games on a closely related graph (you will show this in an exercise).

A strategy is defined in the same way as for impartial games; however, a complete specification of the state of the game will now, in addition to the

position, require an identification of which player is to move next (which edge-set is to be used).

We start with a simple example:

Example 10.2.1 (A partisan Subtraction game). Starting with a pile of $x \in \mathbb{N}$ chips, two players, I and II, alternate taking a certain number of chips. Player I can remove 1 or 4 chips. Player II can remove 2 or 3 chips. The last player who removes chips wins the game.

This is a progressively bounded partisan game where both the terminal nodes and the moves are different for the two players.




Fig. 10.10. Moves of the partisan Subtraction game. Node 0 is terminal for either player, and node 1 is also terminal with a win for player I.

From this example we see that the number of steps it takes to complete the game from a given position now depends on the **state of the game**, $s = (x, i)$, where x denotes the position and $i \in \{I, II\}$ denotes the player that moves next. We let $B(x, i)$ denote the maximum number of moves to complete the game from state (x, i) .

We next prove an important theorem that extends our previous result to include partisan games.

Theorem 10.2.2. *In any progressively bounded combinatorial game with no ties allowed, one of the players has a winning strategy which depends only upon the current state of the game.*

At first the statement that the winning strategy only depends upon the

current state of the game might seem odd, since what else could it depend on? A strategy tells a player which moves to make when playing the game, and *a priori* a strategy could depend upon the history of the game rather than just the game state at a given time. In games which are not progressively bounded, if the game play never terminates, typically one player is assigned a payoff of $-\infty$ and the other player gets $+\infty$. There are examples of such games (which we don't describe here), where the optimal strategy of one of the players must take into account the history of the game to ensure that the other player is not simply trying to prolong the game. But such issues do not exist with progressively bounded games.

Proof of Theorem 10.2.2. We will recursively define a function W , which specifies the winner for a given state of the game: $W(x, i) = j$ where $i, j \in \{I, II\}$ and $x \in X$. For convenience we let $o(i)$ denote the opponent of player i .

When $B(x, i) = 0$, we set $W(x, i)$ to be the player who wins from terminal position x .

Suppose by induction, that whenever $B(y, i) < k$, the $W(y, i)$ has been defined. Let x be a position with $B(x, i) = k$ for one of the players. Then for every $y \in S_i(x)$ we must have $B(y, o(i)) < k$ and hence $W(y, o(i))$ is defined. There are two cases:

Case 1: For some successor state $y \in S_i(x)$, we have $W(y, o(i)) = i$. Then we define $W(x, i) = i$, since player i can move to state y from which he can win. Any such state y will be a winning move.

Case 2: For all successor states $y \in S_i(x)$, we have $W(y, o(i)) = o(i)$. Then we define $W(x, i) = o(i)$, since no matter what state y player i moves to, player $o(i)$ can win.

In this way we inductively define the function W which tells which player has a winning strategy from a given game state. \square

This proof relies essentially on the game being progressively bounded. Next we show that many games have this property.

Lemma 10.2.3. *In a game with a finite position set, if the players cannot move to repeat a previous game state, then the game is progressively bounded.*

Proof. If there are n positions x in the game, there are $2n$ possible game states (x, i) , where i is one of the players. When the players play from position (x, i) , the game can last at most $2n$ steps, since otherwise a state would be repeated. \square

The games of Chess and Go both have special rules to ensure that the

game is progressively bounded. In Chess, whenever the board position (together with whose turn it is) is repeated for a third time, the game is declared a draw. (Thus the real game state effectively has built into it all previous board positions.) In Go, it is not legal to repeat a board position (together with whose turn it is), and this has a big effect on how the game is played.

Next we go on to analyze some interesting partisan games.

10.2.1 The game of Hex

Recall the description of Hex from the introduction.

Example 10.2.4 (Hex). Hex is played on a rhombus-shaped board tiled with hexagons. Each player is assigned a color, either blue or yellow, and two opposing sides of the board. The players take turns coloring in empty hexagons. The goal for each player is to link his two sides of the board with a chain of hexagons in his color. Thus, the terminal positions of Hex are the full or partial colorings of the board that have a chain crossing.




Fig. 10.11. A completed game of Hex with a yellow chain crossing.

Note that Hex is a partisan game where both the terminal positions and the legal moves are different for the two players. We will prove that any fully-colored, standard Hex board contains either a blue crossing or a yellow crossing but not both. This topological fact guarantees that in the game of Hex ties are not possible.

Clearly, Hex is progressively bounded. Since ties are not possible, one of the players must have a winning strategy. We will now prove, again using a strategy-stealing argument, that the first player can always win.

Theorem 10.2.5. *On a standard, symmetric Hex board of arbitrary size, the first player has a winning strategy.*

Proof. We know that one of the players has a winning strategy. Suppose that the second player is the one. Because moves by the players are symmetric, it is possible for the first player to adopt the second player's winning strategy as follows: The first player, on his first move, just colors in an arbitrarily chosen hexagon. Subsequently, for each move by the other player, the first player responds with the appropriate move dictated by second player's winning strategy. If the strategy requires that first player move in the spot that he chose in his first turn and there are empty hexagons left, he just picks another arbitrary spot and moves there instead.

Having an extra hexagon on the board can never hurt the first player — it can only help him. In this way, the first player, too, is guaranteed to win, implying that both players have winning strategies, a contradiction. \square

In 1981, Stefan Reisch, a professor of mathematics at the Universität Bielefeld in Germany, proved that determining which player has a winning move in a general Hex position is PSPACE-complete for arbitrary size Hex boards [Rei81]. This means that it is unlikely that it's possible to write an efficient computer program for solving Hex on boards of arbitrary size. For small boards, however, an Internet-based community of Hex enthusiasts has made substantial progress (much of it unpublished). Jing Yang [Yan], a member of this community, has announced the solution of Hex (and provided associated computer programs) for boards of size up to 9×9 . Usually, Hex is played on an 11×11 board, for which a winning strategy for player I is not yet known.

We will now prove that any colored standard Hex board contains a monochromatic crossing (and all such crossings have the same color), which means that the game always ends in a win for one of the players. This is a purely topological fact that is independent of the strategies used by the players.

In the following two sections, we will provide two different proofs of this result. The first one is actually quite general and can be applied to non-standard boards. The section is optional, hence the *. The second proof has the advantage that it also shows that there can be no more than one crossing, a statement that seems obvious but is quite difficult to prove.

10.2.2 Topology and Hex: a path of arrows*

The claim that any coloring of the board contains a monochromatic crossing is actually the discrete analog of the 2-dimensional Brouwer fixed-point

theorem, which we will prove in section 3.8. In this section, we provide a direct proof.

In the following discussion, pre-colored hexagons are referred to as **boundary**. Uncolored hexagons are called **interior**. Without loss of generality, we may assume that the edges of the board are made up of pre-colored hexagons (see figure). Thus, the interior hexagons are surrounded by hexagons on all sides.

Theorem 10.2.6. *For a completed standard Hex board with non-empty interior and with the boundary divided into two disjoint yellow and two disjoint blue segments, there is always at least one crossing between a pair of segments of like color.*

Proof. Along every edge separating a blue hexagon and a yellow one, insert an arrow so that the blue hexagon is to the arrow's left and the yellow one to its right. There will be four paths of such arrows, two directed toward the interior of the board (call these entry arrows) and two directed away from the interior (call these exit arrows), see Fig. 10.12.




Fig. 10.12. On an empty board the entry and exit arrows are marked. On a completed board, a blue chain lies on the left side of the directed path.

Now, suppose the board has been arbitrarily filled with blue and yellow hexagons. Starting with one of the entry arrows, we will show that it is possible to construct a continuous path by adding arrows tail-to-head always keeping a blue hexagon on the left and a yellow on the right.

In the interior of the board, when two hexagons share an edge with an arrow, there is always a third hexagon which meets them at the vertex toward which the arrow is pointing. If that third hexagon is blue, the next arrow will turn to the right. If the third hexagon is yellow, the arrow will turn to the left. See (a,b) of Fig. 10.13.

Loops are not possible, as you can see from (c) of Fig. 10.13. A loop circling to the left, for instance, would circle an isolated group of blue hexagons




Fig. 10.13. In (a) the third hexagon is blue and the next arrow turns to the right; in (b) — next arrow turns to the left; in (c) we see that in order to close the loop an arrow would have to pass between two hexagons of the same color.

surrounded by yellow ones. Because we started our path at the boundary, where yellow and blue meet, our path will never contain a loop. Because there are finitely many available edges on the board and our path has no loops, it eventually must exit the board using via of the exit arrows.

All the hexagons on the left of such a path are blue, while those on the right are yellow. If the exit arrow touches the same yellow segment of the boundary as the entry arrow, there is a blue crossing (see Fig. 10.12). If it touches the same blue segment, there is a yellow crossing. \square

10.2.3 *Hex and Y*

That there cannot be more than one crossing in the game of Hex seems obvious until you actually try to prove it carefully. To do this directly, we would need a discrete analog of the Jordan curve theorem, which says that a continuous closed curve in the plane divides the plane into two connected components. The discrete version of the theorem is slightly easier than the continuous one, but it is still quite challenging to prove.

Thus, rather than attacking this claim directly, we will resort to a trick: We will instead prove a similar result for a related, more general game — the game of Y, also known as Tripod. Y was introduced in the 1950s by the famous information theorist, Claude Shannon.

Our proof for Y will give us a second proof of the result of the last section, that each completed Hex board contains a monochromatic crossing. Unlike that proof, it will also show that there cannot be more than one crossing in a complete board.

Example 10.2.7 (Game of Y). Y is played on a triangular board tiled with hexagons. As in Hex, the two players take turns coloring in hexagons, each using his assigned color. The goal for both players is to establish a Y, a monochromatic connected region that meets all three sides of the triangle. Thus, the terminal positions are the ones that contain a monochromatic Y.

We can see that Hex is actually a special case of Y: Playing Y, starting from the position shown in Fig. 10.14 is equivalent to playing Hex in the empty region of the board.




Fig. 10.14. Hex is a special case of Y.

We will first show below that a filled-in Y board always contains a single Y. Because Hex is equivalent to Y with certain hexagons pre-colored, the existence and uniqueness of the chain crossing is inherited by Hex from Y.

Once we have established this, we can apply the strategy-stealing argument we gave for Hex to show that the first player to move has a winning strategy.

Theorem 10.2.8. *Any blue/yellow coloring of the triangular board contains either contains a blue Y or a yellow Y, but not both.*

Proof. We can reduce a colored board with sides of size n to one with sides of size $n - 1$ as follows: Think of the board as an arrow pointing right. Except for the left-most column of cells, each cell is the tip of a small arrow-shaped cluster of three adjacent cells pointing the same way as the board. Starting from the right, recolor each cell the majority color of the arrow that it tips, removing the left-most column of cells altogether.

Continuing in this way, we can reduce the board to a single, colored cell.




Fig. 10.15. A step-by-step reduction of a colored Y board.

We claim that the color of this last cell is the color of a winning Y on the original board. Indeed, notice that any chain of connected blue hexagons on a board of size n reduces to a connected blue chain of hexagons on the board of size $n - 1$. Moreover, if the chain touched a side of the original board, it also touches the corresponding side of the smaller board.

The converse statement is harder to see: if there is a chain of blue hexagons connecting two sides of the smaller board, then there was a corresponding blue chain connecting the corresponding sides of the larger board. The proof is left as an exercise (Ex. 10.3).

Thus, there is a Y on a reduced board if and only if there was a Y on the original board. Because the single, colored cell of the board of size one forms a winning Y on that board, there must have been a Y of the same color on the original board. \square

Because any colored Y board contains one and only one winning Y, it follows that any colored Hex board contains one and only one crossing.

10.2.4 *More general boards**

The statement that any colored Hex board contains exactly one crossing is stronger than the statement that every sequence of moves in a Hex game always leads to a terminal position. To see why it's stronger, consider the following variant of Hex, called Six-sided Hex.

Example 10.2.9 (Six-sided Hex). Six-sided Hex is just like ordinary Hex, except that the board is hexagonal, rather than square. Each player is assigned 3 non-adjacent sides and the goal for each player is to create a crossing in his color between any pair of his assigned sides.

Thus, the terminal positions are those that contain one and only one monochromatic crossing between two like-colored sides.

Note that in Six-sided Hex, there can be crossings of both colors in a completed board, but the game ends before a situation with these two crossings can be realized.

The following general theorem shows that, as in standard Hex, there is always at least one crossing.

Theorem 10.2.10. *For an arbitrarily shaped simply-connected completed Hex board with non-empty interior and the boundary partitioned into n blue and n yellow segments, with $n \geq 2$, there is always at least one crossing between some pair of segments of like color.*




Fig. 10.16. A filled-in Six-sided Hex board can have both blue and yellow crossings. In a game when players take turns to move, one of the crossings will occur first, and that player will be the winner.

The proof is very similar to that for standard Hex; however, with a larger number of colored segments it is possible that the path uses an exit arrow that lies on the boundary between a different pair of segments. In this case there is both a blue and a yellow crossing (see Fig. 10.16).

Remark. We have restricted our attention to simply-connected boards (those without holes) only for the sake of simplicity. With the right notion of entry and exit points the theorem can be extended to practically any finite board with non-empty interior, including those with holes.

10.2.5 Other partisan games played on graphs

We now discuss several other partisan games which are played on graphs. For each of our examples, we can describe an explicit winning strategy for the first player.

Example 10.2.11 (The Shannon Switching Game). The Shannon Switching Game, a partisan game similar to Hex, is played by two players, Cut and Short, on a connected graph with two distinguished nodes, A and B . Short, in his turn, reinforces an edge of the graph, making it immune to being cut. Cut, in her turn, deletes an edge that has not been reinforced. Cut wins if

she manages to disconnect A from B . Short wins if he manages to link A to B with a reinforced path.

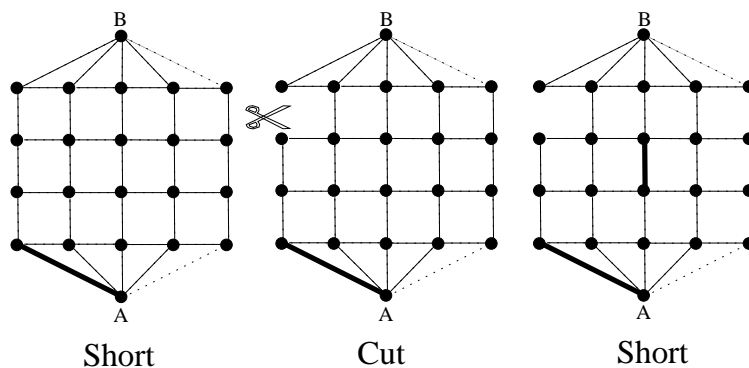


Fig. 10.17. Shannon Switching Game played on a 5×6 grid (the top and bottom rows have been merged to the points A and B). Shown are the first three moves of the game, with Short moving first. Available edges are indicated by dotted lines, and reinforced edges by thick lines. Scissors mark the edge that Short deleted.

There is a solution to the general Shannon Switching Game, but we will not describe it here. Instead, we will focus our attention on a restricted, simpler case: When the Shannon Switching Game is played on a graph that is an $L \times (L + 1)$ grid with the vertices of the top side merged into a single vertex, A , and the vertices on the bottom side merged into another node, B , then it is equivalent to another game, known as Bridg-It (it is also referred to as Gale, after its inventor, David Gale).

Example 10.2.12 (Bridg-It). Bridg-It is played on a network of green and black dots (see Fig. 10.18). Black, in his turn, chooses two adjacent black dots and connects them with a line. Green tries to block Black's progress by connecting an adjacent pair of green dots. Connecting lines, once drawn, may not be crossed.

Black's goal is to make a path from top to bottom, while Green's goal is to block him by building a left-to-right path.

In 1956, Oliver Gross, a mathematician at the RAND Corporation, proved that the player who moves first in Bridg-It has a winning strategy. Several years later, Alfred B. Lehman [Leh64] (see also [Man96]), a professor of computer science at the University of Toronto, devised a solution to the general Shannon Switching Game.

Applying Lehman's method to the restricted Shannon Switching Game




Fig. 10.18. A completed game of Bridg-It and the corresponding Shannon Switching Game. In Bridg-It, the black dots are on the square lattice, and the green dots are on the dual square lattice. Only the black dots appear in the Shannon Switching Game.

that is equivalent to Bridg-It, we will show that Short, if he moves first, has a winning strategy. Our discussion will elaborate on the presentation found in ([BCG82b]).

Before we can describe Short's strategy, we will need a few definitions from graph theory:

Definition 10.2.13. A **tree** is a connected undirected graph without cycles.

- (i) Every tree must have a *leaf*, a vertex of degree one.
- (ii) A tree on n vertices has $n - 1$ edges.
- (iii) A connected graph with n vertices and $n - 1$ edges is a tree.
- (iv) A graph with no cycles, n vertices, and $n - 1$ edges is a tree.

The proofs of these properties of trees are left as an exercise (Ex. 10.4).

Theorem 10.2.14. *In a game of Bridg-It on an $L \times (L + 1)$ board, Short has a winning strategy if he moves first.*

Proof. Short begins by reinforcing an edge of the graph G , connecting A to an adjacent dot, a . We identify A and a by "fusing" them into a single new A . On the resulting graph, there are two edge-disjoint trees such that each tree *spans* (contains all the nodes of) G .

Observe that the blue and red subgraphs in the 4×5 grid in Fig. 10.19 are such a pair of spanning trees: The blue subgraph spans every node, is connected, and has no cycles, so it is a spanning tree by definition. The red subgraph is connected, touches every node, and has the right number of edges, so it is also a spanning tree by property (iii). The same construction could be repeated on an arbitrary $L \times (L + 1)$ grid.




Fig. 10.19. Two spanning trees — the blue one is constructed by first joining top and bottom using the left-most vertical edges, and then adding other vertical edges, omitting exactly one edge in each row along an imaginary diagonal; the red tree contains the remaining edges. The two circled nodes are identified.




Fig. 10.20. Cut separates the blue tree into two components.




Fig. 10.21. Short reinforces a red edge to reconnect the two components.

Using these two spanning trees, which necessarily connect A to B , we can define a strategy for Short.

The first move by Cut disconnects one of the spanning trees into two components (see Fig. 10.20), Short can repair the tree as follows: Because the other tree is also a spanning tree, it must have an edge, e , that connects the two components (see Fig. 10.21). Short reinforces e .

If we think of a reinforced edge e as being both red and blue, then the resulting red and blue subgraphs will still be spanning trees for G . To see this, note that both subgraphs will be connected, and they will still have n edges and $n - 1$ vertices. Thus, by property (iii) they will be trees that span every vertex of G .

Continuing in this way, Short can repair the spanning trees with a rein-

forced edge each time Cut disconnects them. Thus, Cut will never succeed in disconnecting A from B , and Short will win. \square

Example 10.2.15 (Recursive Majority). Recursive Majority is played on a complete ternary tree of height h (see Fig. 10.22). The players take turns marking the leaves, player I with a “+” and player II with a “-.” A parent node acquires the majority sign of its children. Because each interior (non-leaf) has an odd number of children, its sign is determined unambiguously. The player whose mark is assigned to the root wins.

This game always ends in a win for one of the players, so one of them has a winning strategy.




Fig. 10.22. A ternary tree of height 2; the left-most leaf is denoted by 11. Here player I wins the Recursive Majority game.

To describe our analysis, we will need to give each node of the tree a name: Label each of the three branches emanating from a single node in the following way: 1 denotes the left-most edge, 2 denotes the middle edge and 3, the right-most edge. Using these labels, we can identify each node below the root with the “zip-code” of the path from the root that leads to it. For instance, the left-most edge is denoted by $11 \dots 1$, a word of length h consisting entirely of ones.

A strategy-stealing argument implies that the first player to move has the advantage. We can describe his winning strategy explicitly: On his first move, player I marks the leaf $11 \dots 1$ with a plus. For the remaining even number of leaves, he uses the following algorithm to pair them: The partner of the left-most unpaired leaf is found by moving up through the tree to the first common ancestor of the unpaired leaf with the leaf $11 \dots 1$, moving one branch to the right, and then retracing the equivalent path back down (see Fig. 10.23). Formally, letting 1^k be shorthand for a string of ones of fixed length $k \geq 0$ and letting w stand for an arbitrary fixed word of length $h - k - 1$, player I pairs the leaves by the following map: $1^k 2w \mapsto 1^k 3w$.

Once the pairs have been identified, for every leaf marked with a “-” by player II, player I marks its mate with a “+”.




Fig. 10.23. Red marks the left-most leaf and its path. Some sample pair-mates are marked with the same shade of green or blue.

We can show by induction on h that player I is guaranteed to be the winner in the left subtree of depth $h - 1$.

As for the other two subtrees of the same depth, whenever player II wins in one, player I wins the other because each leaf in one of those subtrees is paired with the corresponding leaf in the other. Hence, player I is guaranteed to win two of the three subtrees, thus determining the sign of the root. A rigorous proof of this statement is left to Exercise 10.5.

10.3 Brouwer’s fixed-point theorem via Hex

In this section, we present a proof of Theorem 3.8.2 via Hex. Thinking of a Hex board as a hexagonal lattice, we can construct what is known as a **dual lattice** in the following way: The nodes of the dual are the centers of the hexagons and the edges link every two neighboring nodes (those are a unit distance apart).

Coloring the hexagons is now equivalent to coloring the nodes.




Fig. 10.24. Hexagonal lattice and its dual triangular lattice.

This lattice is generated by two vectors $u, v \in \mathbb{R}^2$ as shown in the left of Figure 10.25. The set of nodes can be described as $\{au + bv : a, b \in \mathbb{Z}\}$. Let’s put $u = (0, 1)$ and $v = (\frac{\sqrt{3}}{2}, \frac{1}{2})$. Two nodes x and y are neighbors if $\|x - y\| = 1$.




Fig. 10.25. Action of G on the generators of the lattice.

We can obtain a more convenient representation of this lattice by applying a linear transformation G defined by:

$$G(u) = \left(-\frac{\sqrt{2}}{2}, \frac{\sqrt{2}}{2} \right); \quad G(v) = (0, 1).$$




Fig. 10.26. Under G an equilateral triangular lattice is transformed to an equivalent lattice.

The game of Hex can be thought of as a game on the corresponding graph (see Fig. 10.26). There, a Hex move corresponds to coloring of one of the nodes. A player wins if she manages to create a connected subgraph consisting of nodes in her assigned color, which also includes at least one node from each of the two sets of her boundary nodes.

The fact that any colored graph contains one and only one such subgraph is inherited from the corresponding theorem for the original Hex board.

Proof of Brouwer's theorem using Hex. As we remarked in section 10.2.1, the fact that there is a winner in any play of Hex is the discrete analogue of the two-dimensional Brouwer fixed-point theorem. We now use this fact about Hex (proved as Theorem 10.2.6) to prove Brouwer's theorem, at least in dimension two. This is due to David Gale.

By an argument similar to the one in the proof of the No-Retraction Theorem, we may restrict our attention to a unit square. Consider a continuous map $T : [0, 1]^2 \rightarrow [0, 1]^2$. Component-wise we write: $T(\mathbf{x}) = (T_1(\mathbf{x}), T_2(\mathbf{x}))$. Suppose it has no fixed points. Then define a function

$f(\mathbf{x}) = T(\mathbf{x}) - \mathbf{x}$. The function f is never zero and continuous on a compact set, hence $\|f\|$ has a positive minimum $\varepsilon > 0$. In addition, as a continuous map on a compact set, T is uniformly continuous, hence $\exists \delta > 0$ such that $\|\mathbf{x} - \mathbf{y}\| < \delta$ implies $\|T(\mathbf{x}) - T(\mathbf{y})\| < \varepsilon$. Take such a δ with a further requirement $\delta < (\sqrt{2} - 1)\varepsilon$. (In particular, $\delta < \frac{\varepsilon}{\sqrt{2}}$.)

Consider a Hex board drawn in $[0, 1]^2$ such that the distance between neighboring vertices is at most δ , as shown in Fig. 10.27. Color a vertex \mathbf{v} on the board blue if $|f_1(\mathbf{v})|$ is at least $\varepsilon/\sqrt{2}$. If a vertex \mathbf{v} is not blue, then $\|f(\mathbf{v})\| \geq \varepsilon$ implies that $|f_2(\mathbf{v})|$ is at least $\varepsilon/\sqrt{2}$; in this case, color \mathbf{v} yellow. We know from Hex that in this coloring, there is a winning path, say, in blue,




Fig. 10.27.

between certain boundary vertices \mathbf{a} and \mathbf{b} . For the vertex \mathbf{a}^* , neighboring \mathbf{a} on this blue path, we have $0 < a_1^* \leq \delta$. Also, the range of T is in $[0, 1]^2$. Hence, since $|T_1(\mathbf{a}^*) - a_1^*| \geq \varepsilon/\sqrt{2}$ (as \mathbf{a}^* is blue), and by the requirement on δ , we necessarily have $T_1(\mathbf{a}^*) - a_1^* \geq \varepsilon/\sqrt{2}$. Similarly, for the vertex \mathbf{b}^* , neighboring \mathbf{b} , we have $T_1(\mathbf{b}^*) - b_1^* \leq -\varepsilon/\sqrt{2}$. Examining the vertices on this blue path one-by-one from \mathbf{a}^* to \mathbf{b}^* , we must find neighboring vertices \mathbf{u} and \mathbf{v} such that $T_1(\mathbf{u}) - u_1 \geq \varepsilon/\sqrt{2}$ and $T_1(\mathbf{v}) - v_1 \leq -\varepsilon/\sqrt{2}$. Therefore,

$$T_1(\mathbf{u}) - T_1(\mathbf{v}) \geq 2\frac{\varepsilon}{\sqrt{2}} - (v_1 - u_1) \geq \sqrt{2}\varepsilon - \delta > \varepsilon.$$

However, $\|\mathbf{u} - \mathbf{v}\| \leq \delta$ should also imply $\|T(\mathbf{u}) - T(\mathbf{v})\| < \varepsilon$, a contradiction. \square

Exercises

- 10.1 In the game of Chomp, what is the Sprague-Grundy function of the 2×3 rectangular piece of chocolate?

- 10.2 Recall the game of Y , shown in Fig. 10.14. Blue puts down blue hexagons, and Yellow puts down yellow hexagons. This exercise is to prove that the first player has a winning strategy by using the idea of strategy stealing that was used to solve the game of Chomp. The first step is to show that from any position, one of the players has a winning strategy. In the second step, assume that the second player has a winning strategy, and derive a contradiction.
- 10.3 Consider the reduction of a Y board to a smaller one described in section 10.2.1. Show that if there is a Y of blue hexagons connecting the three sides of the smaller board, then there was a corresponding blue Y connecting the sides of the larger board.
- 10.4 Prove the following statements. Hint: use induction.
- Every tree must have a **leaf** — a vertex of degree one.
 - A tree on n vertices has $n - 1$ edges.
 - A connected graph with n vertices and $n - 1$ edges is a *tree*.
 - A graph with no cycles, n vertices and $n - 1$ edges is a *tree*.
- 10.5 For the game of Recursive majority on a ternary tree of depth h , use induction on the depth to prove that the strategy described in Example 10.2.15 is indeed a winning strategy for player I.
- 10.6 Consider a game of Nim with four piles, of sizes 9, 10, 11, 12.
- Is this position a win for the *next* player or the *previous* player (assuming optimal play)? Describe the winning first move.
 - Consider the same initial position, but suppose that each player is allowed to remove at most 9 chips in a single move (other rules of Nim remain in force). Is this an **N**- or **P**-position?
- 10.7 Consider a game where there are two piles of chips. On a player's turn, he may remove between 1 and 4 chips from the first pile, or else remove between 1 and 5 chips from the second pile. The person, who takes the last chip wins. Determine for which $m, n \in \mathbb{N}$ it is the case that $(m, n) \in \mathbf{P}$.
- 10.8 For the game of Moore's Nim, the proof of Lemma 10.1.16 gave a procedure which, for **N**-position x , finds a y which is **P**-position and for which it is legal to move to y . Give an example of a legal

move from an **N**-position to a **P**-position which is *not* of the form described by the procedure.

- 10.9 In the game of **Nimble**, a finite number of coins are placed on a row of slots of finite length. Several coins can occupy a given slot. In any given turn, a player may move one of the coins to the left, by any number of places. The game ends when all the coins are at the left-most slot. Determine which of the starting positions are **P**-positions.
- 10.10 Recall that the subtraction game with subtraction set $\{a_1, \dots, a_m\}$ is that game in which a position consists of a pile of chips, and in which a legal move is to remove a_i chips from the pile, for some $i \in \{1, \dots, m\}$. Find the Sprague-Grundy function for the subtraction game with subtraction set $\{1, 2, 4\}$.
- 10.11 Let G_1 be the subtraction game with subtraction set $S_1 = \{1, 3, 4\}$, G_2 be the subtraction game with $S_2 = \{2, 4, 6\}$, and G_3 be the subtraction game with $S_3 = \{1, 2, \dots, 20\}$. Who has a winning strategy from the starting position $(100, 100, 100)$ in $G_1 + G_2 + G_3$?
- 10.12 (a) Find a direct proof that **equivalence** for games is a transitive relation.
 (b) Show that it is reflexive and symmetric and conclude that it is indeed an equivalence relation.
- 10.13 Prove that the sum of two progressively bounded impartial combinatorial games is a **P**-position if and only if the games are equivalent.
- 10.14 Show that if G_1 and G_2 are equivalent, and G_3 is a third game, then $G_1 + G_3$ and $G_2 + G_3$ are equivalent.
- 10.15 By using the properties of mex, show that a position x is in **P** if and only if $g(x) = 0$. This is the content of Lemma 10.1.27 and the proof is outlined in the text.
- 10.16 Consider the game which is played with piles of chips like Nim, but with the additional move allowed of breaking one pile of size $k > 0$ into two nonempty piles of sizes $i > 0$ and $k - i > 0$. Show that the Sprague-Grundy function g for this game, when evaluated at positions with a single pile, satisfies $g(3) = 4$. Find $g(1000)$, that is,


g evaluated at a position with a single pile of size 1000.

Given a position consisting of piles of sizes 13, 24, and 17, how would you play?

- 10.17 Yet another relative of Nim is played with the additional rule that the number of chips taken in one move can only be 1, 3 or 4. Show that the Sprague-Grundy function g for this game, when evaluated at positions with a single pile, is periodic: $g(n + p) = g(n)$ for some fixed p and all n . Find $g(75)$, that is, g evaluated at a position with a single pile of size 75.

Given a position consisting of piles of sizes 13, 24, and 17, how would you play?

- 10.18 Consider the game of up-and-down rooks played on a standard chessboard. Player I has a set of white rooks initially located at level 1, while player II has a set of black rooks at level 8. The players take turns moving their rooks up and down until one of the players has no more moves, at which point the other player wins. This game is not progressively bounded. Yet an optimal strategy exists and can be obtained by relating this game to a Nim with 8 piles.



- 10.19 Two players take turns placing dominos on an $n \times 1$ board of squares, where each domino covers two squares, and dominos cannot overlap. The last player to play wins.

- (a) Find the Sprague-Grundy function for $n \leq 12$.
- (b) Where would you place the first domino when $n = 11$?
- (c) Show that for n even and positive, the first player can guarantee a win.

11

Random-turn and auctioned-turn games

In Chapter 10 we considered combinatorial games, in which the right to move alternates between players; and in Chapters 2 and 3 we considered matrix-based games, in which both players (usually) declare their moves simultaneously, and possible randomness decides what happens next. In this chapter, we consider some games which are combinatorial in nature, but the right to make the next move depends on randomness or some other procedure between the players. In a random-turn game the right to make a move is determined by a coin-toss; in a Richman game, each player offers money to the other player for the right to make the next move, and the player who offers more gets to move. (At the end of the Richman game, the money has no value.) This chapter is based on the work in [LLP⁺99] and [PSSW07].

11.1 Random-turn games defined

Suppose we are given a finite directed graph — a set of vertices V and a collection of arrows leading between pairs of vertices — on which a distinguished subset ∂V of the vertices are called the **boundary** or the **terminal vertices**, and each terminal vertex v has an associated payoff $f(v)$. Vertices in $V \setminus \partial V$ are called the **internal vertices**. We assume that from every node there is a path to some terminal vertex.

Play a two-player, zero-sum game as follows. Begin with a token on some vertex. At each turn, players flip a fair coin, and the winner gets to move the token along some directed edge. The game ends when a terminal vertex v is reached; at this point II pays I the associated payoff $f(v)$.

Let $u(x)$ denote the value of the game begun at vertex x . (Note that since there are infinitely many strategies if the graph has cycles, it should be proved that this exists.) Suppose that from x there are edges to x_1, \dots, x_k .

Claim:

$$u(x) = \frac{1}{2} \left(\max_i \{u(x_i)\} + \min_j \{u(x_j)\} \right). \quad (11.1)$$

More precisely, if S_I denotes strategies available to player I, and S_{II} those available to player II, τ is the time the game ends, and X_τ is the terminal state reached, write

$$u_I(x) = \begin{cases} \sup_{S_I} \{ \inf_{S_{II}} \{ \mathbb{E} f(X_\tau) \} \}, & \text{if } \tau < \infty \\ -\infty, & \text{if } \tau = \infty. \end{cases}$$

Likewise, let

$$u_{II}(x) = \begin{cases} \inf_{S_{II}} \{ \sup_{S_I} \{ \mathbb{E} f(X_\tau) \} \}, & \text{if } \tau < \infty \\ +\infty, & \text{if } \tau = \infty. \end{cases}$$

Then both u_I and u_{II} satisfy (11.1).

We call functions satisfying (11.1) “infinity-harmonic”. In the original paper by Lazarus, Loeb, Propp, and Ullman, [LLP⁺99] they were called “Richman functions”.

11.2 Random-turn selection games

Now we describe a general class of games that includes the famous game of Hex. *Random-turn Hex* is the same as ordinary Hex, except that instead of alternating turns, players toss a coin before each turn to decide who gets to place the next stone. Although ordinary Hex is famously difficult to analyze, the optimal strategy for random-turn Hex turns out to be very simple.

Let S be an n -element set, which will sometimes be called the *board*, and let f be a function from the 2^n subsets of S to \mathbb{R} . A *selection game* is played as follows: the first player selects an element of S , the second player selects one of the remaining $n - 1$ elements, the first player selects one of the remaining $n - 2$, and so forth, until all elements have been chosen. Let S_1 and S_2 signify the sets chosen by the first and second players respectively. Then player I receives a payoff of $f(S_1)$ and player II a payoff of $-f(S_1)$. (Selection games are zero-sum.) The following are examples of selection games:

11.2.1 Hex

Here S is the set of hexagons on a rhombus-shaped $L \times L$ hexagonal grid, and $f(S_1)$ is 1 if S_1 contains a left-right crossing, -1 otherwise. In this case, once

S_1 contains a left-right crossing or S_2 contains an up-down crossing (which precludes the possibility of S_1 having a left-right crossing), the outcome is determined and there is no need to continue the game.




Fig. 11.1. A game between a human player and a program by David Wilson on a 15×15 board.

We will also sometimes consider Hex played on other types of boards. In the general setting, some hexagons are given to the first or second players before the game has begun. One of the reasons for considering such games is that after a number of moves are played in ordinary Hex, the remaining game has this form.

11.2.2 *Bridg-It*

Bridg-It is another example of a selection game. The random-turn version is just like regular Bridg-It, but the right to move is determined by a coin-toss. Player I attempts to make a vertical crossing by connecting the blue dots and player II — a horizontal crossing by bridging the red ones.



Fig. 11.2. The game of random-turn Bridg-It and the corresponding Shannon's edge-switching game; circled numbers give the order of turns.

In the corresponding Shannon's edge-switching game, S is a set of edges connecting the nodes on an $(L + 1) \times L$ grid with top nodes merged into one (similarly for the bottom nodes). In this case, $f(S_1)$ is 1 if S_1 contains a top-to-bottom crossing and -1 otherwise.

11.2.3 Surround

The famous game of “Go” is not a selection game (for one, a player can remove an opponent’s pieces), but the game of “Surround,” in which, as in Go, surrounding area is important, is a selection game. In this game S is the set of n hexagons in a hexagonal grid (of any shape). At the end of the game, each hexagon is recolored to be the color of the outermost cluster surrounding it (if there is such a cluster). The payoff $f(S_1)$ is the number of hexagons recolored black minus the number of hexagons recolored white. (Another natural payoff function is $f^*(S_1) = \text{sign}(f(S_1))$.)




Fig. 11.3. A completed game of Surround before recoloring surrounded territory (on left), and after recoloring (on right). 10 black spaces were recolored white, and 12 white spaces were recolored black, so $f(S_1) = 2$.

11.2.4 Full-board Tic-Tac-Toe

Here S is the set of spaces in a 3×3 grid, and $f(S_1)$ is the number of horizontal, vertical, or diagonal lines in S_1 minus the number of horizontal, vertical, or diagonal lines in $S \setminus S_1$. This is different from ordinary tic-tac-toe in that the game does not end after the first line is completed.

11.2.5 Recursive majority

Suppose we are given a complete ternary tree of depth h . S is the set of leaves. Players will take turns marking the leaves, player I with a + and player II with a -. A parent node acquires the same sign as the majority of its children. The player whose mark is assigned to the root wins. In the random-turn version the sequence of moves is determined by a coin-toss.

Let $S_1(h)$ be a subset of the leaves of the complete ternary tree of depth h (the nodes that have been marked by I). Inductively, let $S_1(j)$ be the set of nodes at level j such that the majority of their children at level $j + 1$ are in




Fig. 11.4. Random-turn tic-tac-toe played out until no new rows can be constructed. $f(S_1) = 1$.




Fig. 11.5. Here player II wins; the circled numbers give the order of the moves.

$S_1(j + 1)$. The payoff function $f(S_1)$ for the recursive three-fold majority is -1 if $S_1(0) = \emptyset$ and $+1$ if $S_1(0) = \{\text{root}\}$.

11.2.6 Team captains

Two team captains are choosing baseball teams from a finite set S of n players for the purpose of playing a single game against each other. The payoff $f(S_1)$ for the first captain is the probability that the players in S_1 (together with the first captain) would beat the players in S_2 (together with the second captain). The payoff function may be very complicated (depending on which players know which positions, which players have played together before, which players get along well with which captain, etc.). Because we have not specified the payoff function, this game is as general as the class of selection games.

Every selection game has a random-turn variant in which at each turn a fair coin is tossed to decide who moves next.

Consider the following questions:

- (i) What can one say about the probability distribution of S_1 after a typical game of optimally played random-turn Surround?
- (ii) More generally, in a generic random-turn selection game, how does the probability distribution of the final state depend on the payoff function f ?
- (iii) Less precise: Are the teams chosen by random-turn Team captains “good teams” in any objective sense?

The answers are surprisingly simple.

11.3 Optimal strategy for random-turn selection games

A (pure) *strategy* for a given player in a random-turn selection game is a function M which maps each pair of disjoint subsets (T_1, T_2) of S to an element of S . Thus, $M(T_1, T_2)$ indicates the element that the player will pick if given a turn at a time in the game when player I has thus far picked the elements of T_1 and player II — the elements of T_2 . Let’s denote by $T_3 = S \setminus (T_1 \cup T_2)$ the set of available moves.

Denote by $E(T_1, T_2)$ the expected payoff for player I at this stage in the game, assuming that both players play optimally with the goal of maximizing expected payoff. As is true for all finite perfect-information, two-player games, E is well defined, and one can compute E and the set of possible optimal strategies inductively as follows. First, if $T_1 \cup T_2 = S$, then $E(T_1, T_2) = f(T_1)$. Next, suppose that we have computed $E(T_1, T_2)$ whenever $|T_3| \leq k$. Then if $|T_3| = k + 1$, and player I has the chance to move, player I will play optimally if and only if she chooses an s from T_3 for which $E(T_1 \cup \{s\}, T_2)$ is maximal. (If she chose any other s , her expected payoff would be reduced.) Similarly, player II plays optimally if and only if she minimizes $E(T_1, T_2 \cup \{t\})$ at each stage. Hence

$$E(T_1, T_2) = \frac{1}{2} \left(\max_{s \in T_3} E(T_1 \cup \{s\}, T_2) + \min_{t \in T_3} E(T_1, T_2 \cup \{t\}) \right).$$

We will see that the maximizing and the minimizing moves are actually the same.

The foregoing analysis also demonstrates a well-known fundamental fact about finite, turn-based, perfect-information games: both players have optimal pure strategies (i.e., strategies that do not require flipping coins), and knowing the other player’s strategy does not give a player any advantage when both players play optimally. (This contrasts with the situation in which the players play “simultaneously” as they do in Rock-Paper-Scissors.) We should remark that for games such as Hex the terminal position need

not be of the form $T_1 \cup T_2 = S$. If for some (T_1, T_2) for any \tilde{T} such that $\tilde{T} \supset T_1$ and $\tilde{T} \cap T_2 = \emptyset$ we have that $f(\tilde{T}) = C$, then $E(T_1, T_2) = C$.

Theorem 11.3.1. *The value of a random-turn selection game is the expectation of $f(T)$ when a set T is selected randomly and uniformly among all subsets of S . Moreover, any optimal strategy for one of the players is also an optimal strategy for the other player.*

Proof. If player II plays any optimal strategy, player I can achieve the expected payoff $\mathbb{E}[f(T)]$ by playing exactly the same strategy (since, when both players play the same strategy, each element will belong to S_1 with probability $1/2$, independently). Thus, the value of the game is at least $\mathbb{E}[f(T)]$. However, a symmetric argument applied with the roles of the players interchanged implies that the value is no more than $\mathbb{E}[f(T)]$.

Suppose that M is an optimal strategy for the first player. We have seen that when both players use M , the expected payoff is $\mathbb{E}[f(T)] = E(\emptyset, \emptyset)$. Since M is optimal for player I, it follows that when both players use M player II always plays optimally (otherwise, player I would gain an advantage, since she is playing optimally). This means that $M(\emptyset, \emptyset)$ is an optimal first move for player II, and therefore every optimal first move for player I is an optimal first move for player II. Now note that the game started at any position is equivalent to a selection game. We conclude that every optimal move for one of the players is an optimal move for the other, which completes the proof. \square

If f is identically zero, then all strategies are optimal. However, if f is *generic* (meaning that all of the values $f(S_1)$ for different subsets S_1 of S are linearly independent over \mathbb{Q}), then the preceding argument shows that the optimal choice of s is always unique and that it is the same for both players. We thus have the following result:

Theorem 11.3.2. *If f is generic, then there is a unique optimal strategy and it is the same strategy for both players. Moreover, when both players play optimally, the final S_1 is equally likely to be any one of the 2^n subsets of S .*

Theorem 11.3.1 and Theorem 11.3.2 are in some ways quite surprising. In the baseball team selection, for example, one has to think very hard in order to play the game optimally, knowing that at each stage there is exactly one correct choice and that the adversary can capitalize on any miscalculation. Yet, despite all of that mental effort by the team captains, the final teams look no different than they would look if at each step both captains chose players uniformly at random.

Also, for illustration, suppose that there are only two players who know how to pitch and that a team without a pitcher always loses. In the alternating turn game, a captain can always wait to select a pitcher until just after the other captain selects a pitcher. In the random-turn game, the captains must try to select the pitchers in the opening moves, and there is an even chance the pitchers will end up on the same team.

Theorem 11.3.1 and Theorem 11.3.2 generalize to random-turn selection games in which the player to get the next turn is chosen using a biased coin. If player I gets each turn with probability p , independently, then the value of the game is $\mathbb{E}[f(T)]$, where T is a random subset of S for which each element of S is in T with probability p , independently. For the corresponding statement of the proposition to hold, the notion of “generic” needs to be modified. For example, it suffices to assume that the values of f are linearly independent over $\mathbb{Q}[p]$. The proofs are essentially the same.

11.4 Win-or-lose selection games

We say that a game is a *win-or-lose* game if $f(T)$ takes on precisely two values, which we may as well assume to be -1 and 1 . If $S_1 \subset S$ and $s \in S$, we say that s is *pivotal* for S_1 if $f(S_1 \cup \{s\}) \neq f(S_1 \setminus \{s\})$. A selection game is *monotone* if f is monotone; that is, $f(S_1) \geq f(S_2)$ whenever $S_1 \supset S_2$. Hex is an example of a monotone, win-or-lose game. For such games, the optimal moves have the following simple description.

Lemma 11.4.1. *In a monotone, win-or-lose, random-turn selection game, a first move s is optimal if and only if s is an element of S that is most likely to be pivotal for a random-uniform subset T of S . When the position is (S_1, S_2) , the move s in $S \setminus (S_1 \cup S_2)$ is optimal if and only if s is an element of $S \setminus (S_1 \cup S_2)$ that is most likely to be pivotal for $S_1 \cup T$, where T is a random-uniform subset of $S \setminus (S_1 \cup S_2)$.*

The proof of the lemma is straightforward at this point and is left to the reader.

For win-or-lose games, such as Hex, the players may stop making moves after the winner has been determined, and it is interesting to calculate how long a random-turn, win-or-lose, selection game will last when both players play optimally. Suppose that the game is a monotone game and that, when there is more than one optimal move, the players break ties in the same way. Then we may take the point of view that the playing of the game is a (possibly randomized) decision procedure for evaluating the payoff function f when the items are randomly allocated. Let \vec{x} denote the allocation of the

items, where $x_i = \pm 1$ according to whether the i^{th} item goes to the first or second player. We may think of the x_i as input variables, and the playing of the game is one way to compute $f(\vec{x})$. The number of turns played is the number of variables of \vec{x} examined before $f(\vec{x})$ is computed. We may use some inequalities from the theory of Boolean functions to bound the average length of play.

Let $I_i(f)$ denote the influence of the i^{th} bit on f (i.e., the probability that flipping x_i will change the value of $f(\vec{x})$). The following inequality is from O'Donnell and Servedio [OS04]:

$$\begin{aligned} \sum_i I_i(f) &= \mathbb{E} \left[\sum_i f(\vec{x}) x_i \right] = \mathbb{E} \left[f(\vec{x}) \sum_i x_i 1_{x_i \text{ examined}} \right] \\ &\leq (\text{by Cauchy-Schwarz}) \sqrt{\mathbb{E}[f(\vec{x})^2] \mathbb{E} \left[\left(\sum_{i: x_i \text{ examined}} x_i \right)^2 \right]} \\ &= \sqrt{\mathbb{E} \left[\left(\sum_{i: x_i \text{ examined}} x_i \right)^2 \right]} = \sqrt{\mathbb{E}[\# \text{ bits examined}]}. \quad (11.2) \end{aligned}$$

The last equality is justified by noting that $\mathbb{E}[x_i x_j 1_{x_i \text{ and } x_j \text{ both examined}}] = 0$ when $i \neq j$, which holds since conditioned on x_i being examined before x_j , conditioned on the value of x_i , and conditioned on x_j being examined, the expected value of x_j is zero. By (11.2) we have

$$\mathbb{E}[\# \text{ turns}] \geq \left[\sum_i I_i(f) \right]^2.$$

We shall shortly apply this bound to the game of random-turn Recursive Majority. An application to Hex can be found in the notes for this chapter.

11.4.1 Length of play for random-turn Recursive Majority

In order to compute the probability that flipping the sign of a given leaf changes the overall result, we can compute the probability that flipping the sign of a child will flip the sign of its parent along the entire path that connects the given leaf to the root. Then, by independence, the probability at the leaf will be the product of the probabilities at each ancestral node on the path.

For any given node, the probability that flipping its sign will change the

sign of the parent is just the probability that the signs of the other two siblings are distinct.




Fig. 11.6.

When none of the leaves are filled this probability is $p = 1/2$. This holds all along the path to the root, so the probability that flipping the sign of leaf i will flip the sign of the root is just $I_i(f) = (\frac{1}{2})^h$. By symmetry this is the same for every leaf.

We now use (11.2) to produce the bound:

$$\mathbb{E}[\# \text{ turns}] \geq \left[\sum_i I_i(f) \right]^2 = \left(\frac{3}{2} \right)^{2h} .$$

11.5 Richman games

Richman games were suggested by the mathematician David Richman, and analyzed by Lazarus, Loeb, Propp, and Ullman in 1995 [LLPU96]. Begin with a finite, directed, acyclic graph, with two distinguished terminal vertices, labeled b and r . Player Blue tries to reach b , and player Red tries to reach r . Call the payoff function R , and let $R(b) = 0$, $R(r) = 1$. Play as in the random-turn game setup above, except instead of a coin flip, players bid for the right to make the next move. The player who bids the larger amount pays that amount to the other, and moves the token along a directed edge of her choice. In the case of a tie, they flip a coin to see who gets to buy the next move. In these games there is also a natural infinity-harmonic (Richman) function, the optimal bids for each player.

Let $R^+(v) = \max_{v \rightsquigarrow w} R(w)$ and $R^-(v) = \min_{v \rightsquigarrow w} R(w)$, where the maxima and minima are over vertices w for which there exists a directed path leading from v to w . Extend R to the interior vertices by

$$R(v) = \frac{1}{2}(R^+(v) + R^-(v)).$$

Note that R is a Richman function.




Fig. 11.7.

Theorem 11.5.1. *Suppose Blue has \$x, Red has \$y, and the current position is v. If*

$$\frac{x}{x+y} > R(v) \tag{11.3}$$

holds before Blue bids, and Blue bids $[R(v) - R(u)](x+y)$, where $v \rightsquigarrow u$ and $R(u) = R^-(v)$, then the inequality (11.3) holds after the next player moves, provided Blue moves to u if he wins the bid.

Proof. There are two cases to analyze.

Case I: Blue wins the bid. After this move, Blue has $\$x' = x - [R(v) - R(u)](x+y)$ dollars. We need to show that $\frac{x'}{x+y} > R(u)$.

$$\frac{x'}{x+y} > R(u) = \frac{x}{x+y} - [R(v) - R(u)] > R(v) - [R(v) - R(u)] = R(u).$$

Case II: Red wins the bid. Now Blue has $\$x' \geq x + [R(v) - R(u)](x+y)$ dollars. Note that if $R(w) = R^+(v)$, then $[R(v) - R(u)] = [R(w) - R(v)]$.

$$\frac{x'}{x+y} \geq \frac{x}{x+y} + [R(w) - R(v)] \geq R(w),$$

and by definition of w, if z is Red's choice, $R(w) \geq R(z)$. □

Corollary 11.5.2. *If (11.3) holds at the beginning of the game, Blue has a winning strategy.*

Proof. When Blue loses, $R(v) = 1$, but $\frac{x}{x+y} \leq 1$. □

Corollary 11.5.3. *If*

$$\frac{x}{x+y} < R(v)$$

holds at the beginning of the game, Red has a winning strategy.

Proof. Recolor the vertices, and replace R with $1 - R$. □

Remark. The above strategy is, in effect, to assume the opponent has the critical amount of money, and apply the first strategy. There are, in fact, many winning strategies if (11.3) holds.

Exercises

11.1 Generalize the proofs of Theorem 11.3.1 and Theorem 11.3.2 further so as to include the following two games:

a) Restaurant selection

Two parties (with opposite food preferences) want to select a dinner location. They begin with a map containing 2^n distinct points in \mathbb{R}^2 , indicating restaurant locations. At each step, the player who wins a coin toss may draw a straight line that divides the set of remaining restaurants exactly in half and eliminate all the restaurants on one side of that line. Play continues until one restaurant z remains, at which time player I receives payoff $f(z)$ and player II receives $-f(z)$.

b) Balanced team captains

Suppose that the captains wish to have the final teams equal in size (i.e., there are $2n$ players and we want a guarantee that each team will have exactly n players in the end). Then instead of tossing coins, the captains may shuffle a deck of $2n$ cards (say, with n red cards and n black cards). At each step, a card is turned over and the captain whose color is shown on the card gets to choose the next player.

11.2 Recursive Majority on b-ary trees Let $b = 2r + 1$, $r \in \mathbb{N}$. Consider the game of recursive majority on a b-ary tree of depth h . For each leaf, determine the probability that flipping the sign of that leaf would change the overall result.

11.3 Even if y is unknown, but (11.3) holds, Blue still has a winning

strategy, which is to bid

$$\left(1 - \frac{R(u)}{R(v)}\right) x.$$

Prove this.

11.6 Additional notes on random-turn Hex

11.6.1 Odds of winning on large boards under biased play.

In the game of Hex, the propositions discussed earlier imply that the probability that player I wins is given by the probability that there is a left-right crossing in independent Bernoulli percolation on the sites (i.e., when the sites are independently and randomly colored black or white). One perhaps surprising consequence of the connection to Bernoulli percolation is that, if player I has a slight edge in the coin toss and wins the coin toss with probability $1/2 + \varepsilon$, then for any $r > 0$ and any $\varepsilon > 0$ and any $\delta > 0$, there is a strategy for player I that wins with probability at least $1 - \delta$ on the $L \times rL$ board, provided that L is sufficiently large.

We do not know if the correct move in random-turn Hex can be found in polynomial time. On the other hand, for any fixed ε a computer can sample $O(L^4 \varepsilon^{-2} \log(L^4/\varepsilon))$ percolation configurations (filling in the empty hexagons at random) to estimate which empty site is most likely to be pivotal given the current board configuration. Except with probability $O(\varepsilon/L^2)$, the computer will pick a site that is within $O(\varepsilon/L^2)$ of being optimal. This simple randomized strategy provably beats an optimal opponent $(50 - \varepsilon)\%$ of time.




Fig. 11.8. Random-turn Hex on boards of size 11×11 and 63×63 under (near) optimal play.

Typical games under optimal play.

What can we say about how long an average game of random-turn Hex will last, assuming that both players play optimally? (Here we assume that the game is stopped once a winner is determined.) If the side length of the board is L , we wish to know how the expected length of a game grows with L (see Figure 11.8 for games on a large board). Computer simulations on a variety of board sizes suggest that the exponent is about 1.5–1.6. As far as rigorous bounds go, a trivial upper bound is $O(L^2)$. Since the game does not end until a player has found a crossing, the length of the shortest crossing in percolation is a lower bound, and empirically this distance grows as $L^{1.1306 \pm 0.0003}$ [Gra99], where the exponent is known to be strictly larger than 1. We give a stronger lower bound:

Theorem 11.6.1. *Random-turn Hex under optimal play on an order L board, when the two players break ties in the same manner, takes at least $L^{3/2+o(1)}$ time on average.*

Proof. To use the O’Donnell-Servedio bound (11.2), we need to know the influence that the sites have on whether or not there is a percolation crossing (a path of black hexagons connecting the two opposite black sides). The influence $I_i(f)$ is the probability that flipping site i changes whether there is a black crossing or a white crossing. The “4-arm exponent” for percolation is $5/4$ [SW01] (as predicted earlier in [Con89]), so $I_i(f) = L^{-5/4+o(1)}$ for sites i “away from the boundary,” say in the middle ninth of the region. Thus $\sum_i I_i(f) \geq L^{3/4+o(1)}$, so $\mathbb{E}[\# \text{ turns}] \geq L^{3/2+o(1)}$. \square

An optimally played game of random-turn Hex on a small board may occasionally have a move that is disconnected from the other played hexagons, as the game in Figure 11.9 shows. But this is very much the exception rather than the rule. For moderate- to large-sized boards, it appears that in almost every optimally played game, the set of played hexagons remains a connected set throughout the game (which is in sharp contrast to the usual game of Hex). We do not have an explanation for this phenomenon, nor is it clear to us if it persists as the board size increases beyond the reach of simulations.

11.7 Random-turn Bridg-It

Next we consider the random-turn version of Bridg-It or the Shannon Switching Game. Just as random-turn Hex is connected to site percolation on the triangular lattice, where the vertices of the lattice (or equivalently faces of




Fig. 11.9. A rare occurrence — a game of random-turn Hex under (near) optimal play with a disconnected play.

the hexagonal lattice) are independently colored black or white with probability $1/2$, random-turn Bridg-It is connected to bond percolation on the square lattice, where the edges of the square lattice are independently colored black or white with probability $1/2$. We don't know the optimal strategy for random-turn Bridg-It, but as with random-turn Hex, one can make a randomized algorithm that plays near optimally. Less is known about bond percolation than site percolation, but it is believed that the crossing probabilities for these two processes are asymptotically the same on “nice” domains [LPSA94], so that the probability that Cut wins in random-turn Bridg-It is well approximated by the probability that a player wins in random-turn Hex on a similarly shaped board.

Bibliography

- [AH81] Robert Axelrod and William D. Hamilton. The evolution of cooperation. *Science*, 211(4489):1390–1396, 1981.
- [Ansa] Hex information. <http://www.cs.unimaas.nl/icga/games/hex/>.
- [Ansb] V. V. Anshelevich. The game of Hex: An automatic theorem proving approach to game programming. <http://home.earthlink.net/~vanshel/VAnshelevich-01.pdf>.
- [Arr51] Kenneth J. Arrow. *Social Choice and Individual Values*. Cowles Commission Monograph No. 12. John Wiley & Sons Inc., New York, N. Y., 1951.
- [Arr02] A. Arratia. On the descriptive complexity of a simplified game of Hex. *Log. J. IGPL*, 10:105–122, 2002.
- [Aum87] Robert J. Aumann. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55(1):1–18, jan 1987. <http://www.jstor.org/pss/1911154>. 71
- [Axe85] Robert Axelrod. *The Evolution of Cooperation*. Basic Books, 387 Park Avenue So., New York, NY 10016, 1985.
- [Axe97] Robert Axelrod. The evolution of strategies in the iterated prisoner’s dilemma. In *The dynamics of norms*, Cambridge Stud. Probab. Induc. Decis. Theory, pages 1–16. Cambridge Univ. Press, Cambridge, 1997.
- [BBC69] A. Beck, M. Bleicher, and J. Crow. *Excursions into Mathematics*. Worth, 1969.
- [BCG82a] E. R. Berlekamp, J. H. Conway, and R. K. Guy. *Winning Ways for Your Mathematical Plays*, volume 1. Academic Press, 1982. 191
- [BCG82b] E. R. Berlekamp, J. H. Conway, and R. K. Guy. *Winning Ways for Your Mathematical Plays*, volume 2. Academic Press, 1982. 203
- [Bou02] Charles L. Bouton. Nim, a game with a complete mathematical theory. *Ann. of Math. (2)*, 3(1-4):35–39, 1901/02.
- [Bro00] C. Browne. *Hex Strategy: Making the Right Connections*. A. K. Peters, 2000.
- [BSW05] Itai Benjamini, Oded Schramm, and David B. Wilson. Balanced Boolean functions that can be evaluated so that every input bit is unlikely to be read. In *Proc. 37th Symposium on the Theory of Computing*, 2005.
- [Car92] John L. Cardy. Critical percolation in finite geometries. *J. Phys. A*, 25(4):L201–L206, 1992.
- [Con89] Antonio Coniglio. Fractal structure of Ising and Potts clusters: exact results. *Phys. Rev. Lett.*, 62(26):3054–3057, 1989. 225

- [Con01] J. H. Conway. *On Numbers and Games*. A K Peters Ltd., Natick, MA, second edition, 2001.
- [dCMdC90] Marie Jean Antoine Nicolas de Caritat (Marquis de Condorcet). Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix (eng: Essay on the application of analysis to the probability of majority decisions), 1785. In *The French Revolution Research Collection*. Pergamon Press, Headington Hill Hall, Oxford OX3 0BW UK, 1990. <http://gallica.bnf.fr/scripts/ConsultationTout.exe?E=0&O=N041718>. 112
- [Eve57] H. Everett. Recursive games. In *Contributions to the theory of games, vol. 3*, Annals of Mathematics Studies, no. 39, pages 47–78. Princeton University Press, Princeton, N. J., 1957.
- [Fer08] Thomas Ferguson. *Game Theory*. 2008. http://www.math.ucla.edu/~tom/Game_Theory/Contents.html. 191
- [FV93] Dean P. Foster and Rakesh V. Vohra. A randomization rule for selecting forecasts. *Operations Research*, 41(4):704–709, Jul 1993. <http://www.jstor.org/pss/171965>.
- [FV97] Dean P. Foster and Rakesh V. Vohra. Calibrated learning and correlated equilibrium. *Games Econom. Behav.*, 21(1-2):40–55, 1997.
- [FV99] Dean P. Foster and Rakesh V. Vohra. Calibration, expected utility and local optimality. Discussion Papers 1254, Northwestern University, Center for Mathematical Studies in Economics and Management Science, March 1999. <http://ideas.repec.org/p/nwu/cmsems/1254.html>. 157
- [Gal79] D. Gale. The game of Hex and the Brouwer fixed-point theorem. *Amer. Math. Monthly*, 86:818–827, 1979.
- [Gar59] M. Gardner. The game of Hex. In *Hexaflexagons and Other Mathematical Diversions: The First Scientific American Book of Puzzles and Games*, pages 73–83. Simon and Schuster, 1959.
- [GBOW88] S. Goldwasser, M. Ben-Or, and A. Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computing. In *Proc. of the 20th STOC*, pages 1–10, 1988.
- [Gea04] J. Geanakoplos. *Three Brief Proofs of Arrow's Impossibility Theorem*. Cowles Commission Monograph No. 1123R4. Cowles Foundation for Research in Economics, Yale University, Box 208281, New Haven, Connecticut 06520-8281., 1996 (updated 2004). <http://cowles.econ.yale.edu/>. 112
- [Gin00] H. Gintis. *Game Theory Evolving; A problem-centered introduction to modeling strategic interaction*. Princeton University Press, Princeton, New Jersey, 2000.
- [Gra99] Peter Grassberger. Pair connectedness and shortest-path scaling in critical percolation. *J. Phys. A*, 32(35):6233–6238, 1999. 225
- [GS62] D. Gale and L. S. Shapley. College Admissions and the Stability of Marriage. *Amer. Math. Monthly*, 69(1):9–15, 1962.
- [HMC00] Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, sep 2000. <http://www.jstor.org/pss/2999445>. 157
- [Kar59] S. Karlin. *Mathematical methods and theory in games, programming and economics*, volume 2. Addison-Wesley, 1959.
- [KN02] H. W. Kuhn and S. Nasar, editors. *The Essential John Nash*. Princeton University Press, 2002.
- [Las31] E. Lasker. *Brettspiele der Völker, Rätsel und mathematische Spiele*. Berlin, 1931.

- [Leh64] Alfred Lehman. A solution of the Shannon switching game. *J. Soc. Indust. Appl. Math.*, 12:687–725, 1964. 202
- [LLP⁺99] Andrew J. Lazarus, Daniel E. Loeb, James G. Propp, Walter R. Stromquist, and Daniel H. Ullman. Combinatorial games under auction play. *Games Econom. Behav.*, 27(2):229–264, 1999. 212, 213
- [LLPU96] Andrew J. Lazarus, Daniel E. Loeb, James G. Propp, and Daniel Ullman. Richman games. In Richard J. Nowakowski, editor, *Games of No Chance*, volume 29 of *MSRI Publications*, pages 439–449. Cambridge Univ. Press, Cambridge, 1996. 221
- [LPSA94] Robert Langlands, Philippe Pouliot, and Yvan Saint-Aubin. Conformal invariance in two-dimensional percolation. *Bull. Amer. Math. Soc. (N.S.)*, 30(1):1–61, 1994. 226
- [Man96] Richard Mansfield. Strategies for the Shannon switching game. *Amer. Math. Monthly*, 103(3):250–252, 1996. 202
- [Mar98] Donald A. Martin. The determinacy of Blackwell games. *J. Symbolic Logic*, 63(4):1565–1581, 1998.
- [Moo10] E. H. Moore. A generalization of the game called nim. *Ann. of Math. (Ser. 2)*, 11:93–94, 1909–1910.
- [MS96] Dov Monderer and Lloyd S. Shapley. Potential games. *Games Econom. Behav.*, 14(1):124–143, 1996. 93
- [NS03] Abraham Neyman and Sylvain Sorin, editors. *Stochastic games and applications*, volume 570 of *NATO Science Series C: Mathematical and Physical Sciences*, Dordrecht, 2003. Kluwer Academic Publishers.
- [OS04] Ryan O’Donnell and Rocco Servedio. On decision trees, influences, and learning monotone decision trees. Technical Report CU-CS-023-04, Columbia University, Dept. of Computer Science, 2004. <http://www1.cs.columbia.edu/~library/2004.html>. 220
- [OSSS05] Ryan O’Donnell, Mike Saks, Oded Schramm, and Rocco Servedio. Every decision tree has an influential variable. In *Proc. of the 46th Annual Symposium on Foundations of Computer Science (FOCS)*, 2005. <http://arxiv.org/abs/cs/0508071>.
- [PSSW07] Yuval Peres, Oded Schramm, Scott Sheffield, and David B. Wilson. Random-turn Hex and other selection games. *Amer. Math. Monthly*, 114(5):373–387, 2007. <http://arxiv.org/math/0508580>. 212
- [PSSW09] Yuval Peres, Oded Schramm, Scott Sheffield, and David B. Wilson. Tug-of-war and the infinity Laplacian. *J. Amer. Math. Soc.*, 22(1):167–210, 2009. <http://arxiv.org/abs/math.AP/0605002>.
- [Rei81] S. Reisch. Hex ist PSPACE-vollständig. *Acta Inform.*, 15:167–191, 1981. 196
- [Ros73] R.W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973. 93
- [Rud64] Walter Rudin. *Principles of Mathematical Analysis*. MH, NY, third edition, 1964. 77, 79
- [Saa90] D. G. Saari. The Borda dictionary. *Social Choice and Welfare*, 7(4):279–317, 1990. 112
- [Saa06] Donald G. Saari. Which is better: the Condorcet or Borda winner? *Social Choice and Welfare*, 26(1):107–129, 2006. 112
- [Sha53] C. E. Shannon. Computers and automata. *Proc. Inst. Radio Eng.*, 41:1234–1241, 1953.
- [Sha79] Adi Shamir. How to share a secret. *Comm. ACM*, 22(11):612–613, 1979.

160

- [SL96] B. Sinervo and C. M. Lively. The rock-paper-scissors game and the evolution of alternative male strategies. *Nature*, 380:240–243, March 1996. http://adsabs.harvard.edu/cgi-bin/nph-bib_query?bibcode=1996Natur.380..240S&db_key=GEN. 66
- [Smi01] S. Smirnov. Critical percolation in the plane. I. Conformal invariance and Cardy’s formula. II. Continuum scaling limit, 2001. <http://www.math.kth.se/~stas/papers/percol.ps>.
- [Spr36] R. Sprague. Über mathematische kampfspiele. *Tôhoku Math. J.*, 41:438–444, 1935–36.
- [Spr37] R. Sprague. Über zwei abarten von Nim. *Tôhoku Math. J.*, 43:351–359, 1937.
- [SS05] Oded Schramm and Jeffrey E. Steif. Quantitative noise sensitivity and exceptional times for percolation, 2005. <http://arxiv.org/abs/math/0504586>.
- [Ste00] I. Stewart. Hex marks the spot. *Sci. Amer.*, 283:100–103, Sep. 2000.
- [SW57] Maurice Sion and Philip Wolfe. On a game without a value. In *Contributions to the theory of games, vol. 3*, Annals of Mathematics Studies, no. 39, pages 299–306. Princeton University Press, Princeton, N. J., 1957.
- [SW01] Stanislav Smirnov and Wendelin Werner. Critical exponents for two-dimensional percolation. *Math. Res. Lett.*, 8(5-6):729–744, 2001. <http://arxiv.org/abs/math/0109120>. 225
- [SW05] Oded Schramm and David B. Wilson. SLE coordinate changes. *New York J. Math.*, 11:659–669, 2005.
- [vNM53] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behaviour*. Princeton University Press, Princeton, NJ., 3rd edition, 1953.
- [vR00] J. van Rijswijk. Computer Hex: Are bees better than fruitflies? Master’s thesis, University of Alberta, 2000.
- [vR02] J. van Rijswijk. Search and evaluation in Hex, 2002. <http://www.cse.iitb.ac.in/~nikhilh/NASH/y-hex.pdf>.
- [Wil86] J. D. Williams. *The Compleat Strategyst: Being a primer on the theory of games of strategy*. Dover Publications Inc., New York, second edition, 1986.
- [Win69] Robert L. Winkler. Scoring rules and the evaluation of probability assessors. *Journal of the American Statistical Association*, 64:1073–1078, 1969. <http://www.jstor.org/pss/2283486>. 157
- [WM68] Robert L. Winkler and Allan H. Murphy. “good” probability assessors. *J. Applied Meteorology*, 7(5):751–758, 1968.
- [Wyt07] W. A. Wythoff. A modification of the game of Nim. *Nieuw Arch. Wisk.*, 7:199–202, 1907.
- [Yan] Jing Yang. Hex solutions. <http://www.ee.umanitoba.ca/~jingyang/>. 196
- [YLP03] J. Yang, S. Liao, and M. Pawlak. New winning and losing positions for Hex. In J. Schaeffer, M. Muller, and Y. Bjornsson, editors, *Computers and Games: Third International Conference, CG 2002, Edmonton, Canada, July 25-27, 2002, Revised Papers*, pages 230–248. Springer-Verlag, 2003.
- [Zif99] Robert M. Ziff. Exact critical exponent for the shortest-path scaling function in percolation. *J. Phys. A*, 32(43):L457–L459, 1999.

Index

- S-veto game, 148
- alternative theorem, 35
- antisymmetric game, 13
- Arrow's Impossibility Theorem, 107
- asymmetric information, 3
- auction, 4
 - concert tickets, 125
 - digital goods, 127
 - exclusive markets, 125
 - strategy, 131
- auctioned-turn games, 212
- auctions
 - digital goods
 - optimal fixed price profit, 127
- Banach's fixed-point theorem, 77
- battle of the sexes, 71
- Bayes-Nash equilibrium, 131
- Bayes-Nash incentive compatible, 137
- betting game, 6
- binding commitment, 56
- bipartite graph, 18
- bomber and battleship game, 25
- Bridg-It, 202, 203
 - random-turn, 214, 225
- Characterization for BN equilibrium, 132
- cheetahs and antelopes, 53
- chicken, game of, 54, 62, 87
- Chomp, 172
- chromatic number, 70
- Colon Principle, 191
- combinatorial game, 4, 168
 - impartial, 186–188
 - partisan, 201
 - progressively bounded, 189
 - progressively bounded, 186
- congestion game, 2, 91, 92
- cooperative game theory, 3
- cooperative game, 147
- cooperative game theory, 146
- correlated equilibrium, 72
- correlated strategy pair, 72
- dictatorship, 107, 110
- digital goods auction
 - profit extraction, 129
- domination, 12
- dual LP, 34
- ecology game, 56, 59
- edge, 170
- effective resistance, 46
- efficiency, 148
- equivalence relation, 187
- evolutionarily stable strategy, 54, 65
- evolutionary game theory, 62
- extreme point, 37
- Farkas' Lemma, 35
- feasible, 33
- feasible LP, 34
- feasible set, 32
- fish-selling game, 90, 99, 100
- fixed-point theorem
 - Brouwer, 2, 75, 76, 80, 86
 - 2-dimensional, 196
 - Brouwer for simplex, 84
 - compact, 79
- fixed-point theorems
 - Banach, 77
- Four Stockholders, 149
- Gale, *see* Bridg-It
- general-sum game, 2, 49, 75
- Gibbard-Satterthwaite Theorem, 110
- glove market game, 149
- graph, 170
- graph coloring game, 69
- Hall's condition, 18
- Hawks and Doves, 62, 63, 65, 91
- Hex, 5, 195–200, 206–208
 - random-turn, 213, 224, 225
- hide-and-seek, 20, 24

- generalized, 22
- homeomorphism, 84
- impartial combinatorial game, 170
- independence of irrelevant alternatives (IIA), 104
- Inspection game, 51
- Investing in communication infrastructure, 50
- Jensen's Inequality, 42
- Jordan curve theorem, 198
- König's lemma, 21
- keyword auctions, 120
- Kuhn tree, 100
- Lagrange interpolation, 161
- line-cover, 21
- linear programming duality, 33
- linear program (LP), 31
- link, 170
- Lions and antelopes, 95
- matching, 18
- mechanism design, 3
- mechanism design problem, 120
- metric space
 - compact, 79
 - complete, 77
- mex, 183
- minimum covering, 23
- Moore's theorem, 181
- Nash equilibrium, 68
 - mixed, 54
- Nash equilibrium, 2, 11, 49, 52, 56–58, 60, 63, 65, 68, 75
 - mixed, 56, 58, 63, 65, 66
 - pure, 53
- Nash's theorem, 58, 75
- Nim, 4, 175, 176
 - Moore's Nim, 180
 - Staircase Nim, 178
 - Wythoff Nim, 182, 184
- Nim-sum, 176–178, 180, 187, 190, 191
- no-retraction theorem, 86, 207
- node, 170
- objective function, 32
- partisan, 192
- partisan subtraction game, 193
- payoff matrix, 8, 52
- penalty kick game, 1
- play
 - normal, 170, 173, 186–188
- potential function, 68, 69
- potential game, 68
- preference profile, 104
- pricing game, 61
- primal LP, 34
- priors, 131
- prisoner's dilemma, 49, 63
- profit extractor, 129
- progressively bounded game, 169, 173
- proper coloring, 70
- protein problem, 32
- pure strategy, 68
- quasi-linear utility, 121
- randomized game, 96
- ranking rule, 103, 104
 - strategically vulnerable, 105
- recursive majority, 215, 220
- retraction, 84
- revelation principle, 136
- Rock-Paper-Scissors, 65
- saddle point, 10
- second-price auction, 4
- secret sharing, 158
- separating hyperplane, 26
- sex ratio, 64
- Shannon switching game, 201–203
- Shapley value, 3, 146, 150, 153
- signaling, 3, 95, 96
- simplex, 80
 - barycentric subdivision, 80
 - face, 80
 - proper labeling, 81
 - subdivision of a simplex, 80
- single-item auction
 - all-pay, 117
 - English or ascending, 114
 - first price, 115
 - second price, 116
- single-parameter auction
 - bid-independence, 127
- social choice, 3
- social ranking, 103, 104
- social welfare, 4
- spectrum auction, 124
- Sperner Labeling, 81
- Sperner's lemma, 81
- sponsored search, 119
- Sprague-Grundy function, 187, 189–191
- Sprague-Grundy theorem, 186, 188, 189
- strategy
 - memoryless, 171
 - mixed, 8, 52
 - pure, 8, 52
- strategy stealing, 174, 195, 199, 205, 209
- subtraction game, 169, 187–190
- sum of games, 186
- sum theorem, 190, 191
- sum-game
 - parallel, 17
 - series, 17
- symmetric game, 52, 58, 63, 65, 76
- symmetry, 14
- terminal position, 170

- The car mechanic, 94
- the duality theorem, 33
- Tragedy of the commons, 60
- truthful, truthfulness, 121
- Truthfulness
 - characterization
 - single-parameter domains, 126
- two-person zero-sum games, 6

- unanimity, 104
- utility function, 56

- valuation function, 120
- value of zero-sum game, 9, 29
- Vandermonde matrix, 161, 167
- VCG, 121
- vertex, 37
- Vickrey auction
 - with reserve price, 118
- virtual value, 139
- voting
 - strategy-proof, 110
- voting preferences, 104
 - transitivity, 104
- voting rule, 103, 104
- voting rules
 - Borda count, 106
 - dictatorship, 107
 - instant runoff elections, 105
 - plurality voting, 102

- winning strategy, 4

- Y, 198–200, 209

- zero-sum game, 1