Game Theory, Alive

Yuval Peres

with contributions by David B. Wilson

September 27, 2011

Check for updates at `http://dbwilson.com/games`

We are grateful to Alan Hammond, Yun Long, Gábor Pete, and Peter Ralph for scribing early drafts of this book from lectures by the first author. These drafts were edited by Asaf Nachmias, Sara Robinson and Yelena Shvets; Yelena also drew many of the figures. We also thank Ranjit Samra of rojaysoriginalart.com for the lemon figure, and Barry Sinervo for the Lizard picture.

Sourav Chatterjee, Elchanan Mossel, Asaf Nachmias, and Shobhana Stoyanov taught from drafts of the book and provided valuable suggestions. Thanks also to Varsha Dani, Itamar Landau, Mallory Monasterio, 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 and DMS-0104073 is acknowledged.

Contents

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. Here, we outline the content of this course, giving examples.

We will first look at **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 0.1). 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.

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

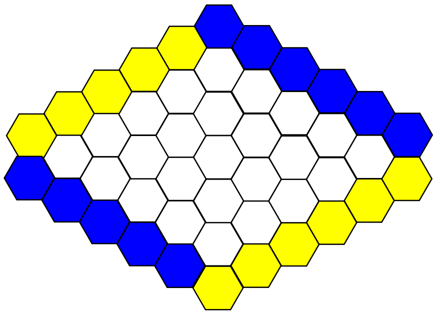

Fig. 0.1. The board for the game of Hex.

to give an explicit description of the optimal strategies of the players. Such **random-turn combinatorial games** are the subject of Chapter 8.

Next, we will turn our attention to games of chance, in which both players move simultaneously. In **two-person zero-sum games**, each player benefits only at the expense of the other. We will show how to find optimal strategies for each player. These strategies will typically turn out to be a randomized choice of the available options.

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.

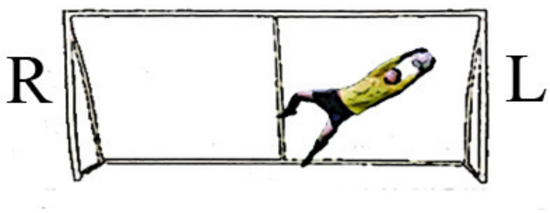

Fig. 0.2. The game of Penalty Kicks.

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.

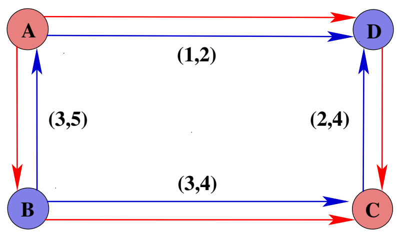

Fig. 0.3. A congestion game. Shown here are the commute times for the four roads connecting four cities. For each road, the first number is the commute time when only one driver uses the road, the second number is the commute time when two drivers use the road.

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.

Another interesting topic is that of 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 4 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, the topic of Chapter 5, is **mechanism design**, the study of how to design a market or scheme that achieves an optimal social outcome when the participating agents act selfishly.

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.

Some of the mathematical results in mechanism design are negative, implying that optimal design is not attainable. For example, a famous theorem by Arrow on **voting schemes** (the topic of Chapter 6) 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.

Another focus of mechanism design is on eliciting truth in auctions. In a standard, sealed-bid auction, there is always a temptation for bidders to bid less than their true value for an item. For example, if an item is worth \$100 to a bidder, then he has no motive to bid more, or even that much, because by exchanging \$100 dollars for an item of equal value, he has not gained anything. The second-price auction is an attempt to overcome this flaw: in this scheme, the lot goes to the highest bidder, but at the price offered by the second-highest bidder. In a second-price auction, as we will show, it is in the interests of bidders to bid their true value for an item, but the mechanism has other shortcomings. The problem of eliciting truth is relevant to the bandwidth auctions held by governments.

In the realm of social choice is the problem of finding **stable matchings**, 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.

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.

1

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 1.0.1. 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 1.1.3 will reduce the process of finding such a strategy to computing a certain recursive function.

We begin with impartial games.

1.1 Impartial games

Before we give formal definitions, let's look at a simple example:

Example 1.1.1 (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:

$$\mathbf{N} = \left\{ x \in \mathbb{N} : \begin{array}{l} \text{the first (``}\underline{\text{n}}\text{ext'') player can ensure a win} \\ \text{if there are } x \text{ chips at the start} \end{array} \right\},$$

$$\mathbf{P} = \left\{ x \in \mathbb{N} : \begin{array}{l} \text{the second (``}\underline{\text{p}}\text{revious'') player can ensure a win} \\ \text{if there are } x \text{ chips at the start} \end{array} \right\}.$$

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 1.1.1. 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 1.1 gives a directed graph corresponding to the Subtraction game with initial position $x = 14$.

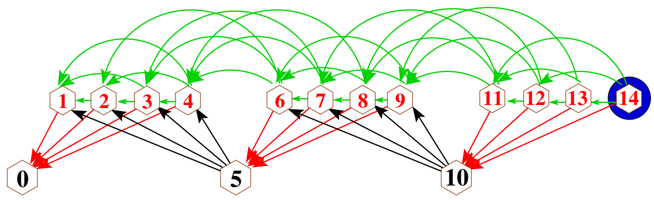

Fig. 1.1. Moves in the Subtraction game. Positions in $\mathbf{N}$ are marked in red and those in $\mathbf{P}$, in black.

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 $\mathbf{P} = \{5n : n \in \mathbb{N}\}$, namely $5\lfloor x/5 \rfloor$.

Let's make a formal definition:

Definition 1.1.2. 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 1.1.3. 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, $\mathbb{N} = \mathbf{N} \cup \mathbf{P}$, 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 § 1.2.

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:

$$\begin{aligned}
\mathbf{N}_0 &= \varnothing \\
\mathbf{P}_0 &= \{\,\text{terminal positions}\,\} \\
\mathbf{N}_{i+1} &= \{\,\text{positions } x \text{ for which there is a move leading to } \mathbf{P}_i\,\} \\
\mathbf{P}_{i+1} &= \{\,\text{positions } y \text{ such that each move leads to } \mathbf{N}_i\,\}
\end{aligned}$$

$$\mathbf{N} = \bigcup_{i \geq 0} \mathbf{N}_i, \qquad \mathbf{P} = \bigcup_{i \geq 0} \mathbf{P}_i.$$

Notice that $\mathbf{P}_0 \subseteq \mathbf{P}_1 \subseteq \mathbf{P}_2 \subseteq \cdots$ and $\mathbf{N}_0 \subseteq \mathbf{N}_1 \subseteq \mathbf{N}_2 \subseteq \cdots$.

In the Subtraction game, we have

$$\begin{aligned} \mathbf{N}_0 &= \varnothing & \mathbf{P}_0 &= \{0\} \\ \mathbf{N}_1 &= \{1,2,3,4\} & \mathbf{P}_1 &= \{0,5\} \\ \mathbf{N}_2 &= \{1,2,3,4,6,7,8,9\} & \mathbf{P}_2 &= \{0,5,10\} \\ &\vdots & &\vdots \\ \mathbf{N} &= \mathbb{N} \setminus 5\mathbb{N} & \mathbf{P} &= 5\mathbb{N} \end{aligned}$$

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 1.1.2 (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.

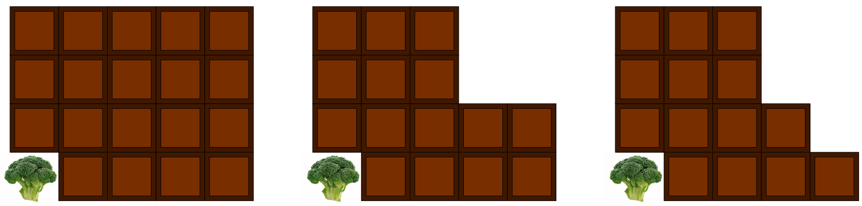

Fig. 1.2. Two moves in a game of Chomp.

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 1.3. However, as the size of the bar increases, the graph becomes very large and a winning strategy difficult to find.

Next we will formally prove that every progressively bounded impartial game has a winning strategy for one of the players.

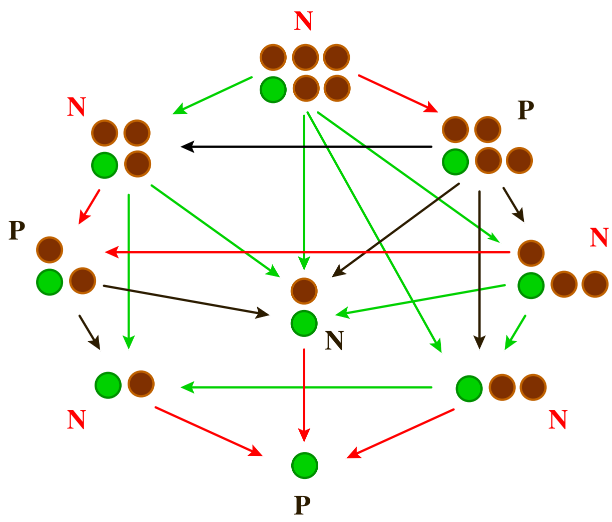

Fig. 1.3. Every move from a **P**-position leads to an **N**-position (bold black links); from every **N**-position there is at least one move to a **P**-position (red links).

Theorem 1.1.1. *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 $\mathbf{N} \cup \mathbf{P}$. $\square$

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 1.1.2. *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 $R \in \mathbf{N}$. $\square$

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 1.3), 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.

1.1.1 Nim and Bouton's solution

Recall the game of Nim from the Introduction.

Example 1.1.3 **(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 **N** or **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 1.1.3 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 $(n_1, n_2, \ldots, n_k)$, 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 1.1.1. *For two Nim positions $X = (x_1, \ldots, x_k)$ and $Y = (y_1, \ldots, y_\ell)$, we denote the position $(x_1, \ldots, x_k, y_1, \ldots, y_\ell)$ 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 **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 **P**-**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 takes chips from one of the piles (assume the pile is in Y without loss of generality). But moving Y from **P**-position always results in a **N**-position, so the resulting game is in a **P**-**N** position with at most n chips, which by the inductive hypothesis is an **N** position. It follows that (X, Y) must be in **P**.

For the final part of the lemma, note that any single pile is in **N**, yet, as we saw above, $(1,1) \in \mathbf{P}$ while $(1,2) \in \mathbf{N}$. □

Going back to our example, $(1,2,3,4,5)$ can be divided into two subgames: $(1,2,3) \in \mathbf{P}$ and $(4,5) \in \mathbf{N}$. By the lemma, we can conclude that $(1,2,3,4,5)$ is in **N**.

The divide-and-sum method (using Lemma 1.1.1) is useful for analyzing Nim positions, but it doesn't immediately determine whether a given position is in **N** or **P**. The following ingenious theorem, proved in 1901 by a Harvard mathematics professor named Charles Bouton, gives a simple and general characterization of **N** and **P** for Nim. Before we state the theorem, we will need a definition.

Definition 1.1.4. 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 $(m_1, m_2, \ldots, m_k)$ 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 1.1.3 (Bouton's Theorem). *A Nim position $x = (x_1, x_2, \ldots, x_k)$ is in* **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	0 1
2	1 0
3	1 1
0	0 0

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 1.1.3. Define Z to be those positions with Nim-sum zero.

Suppose that $x = (x_1, \ldots, x_k) \in Z$, i.e., $x_1 \oplus \cdots \oplus x_k = 0$. 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 $x_1 \oplus \cdots \oplus x_{\ell-1} \oplus x'_\ell \oplus x_{\ell+1} \oplus \cdots \oplus x_k = x'_\ell \oplus x_\ell \neq 0$. Thus any move from a position in Z leads to a position not in Z.

Suppose that $x = (x_1, x_2, \ldots, x_k) \notin Z$. Let $s = x_1 \oplus \cdots \oplus x_k \neq 0$. There are an odd number of values of $i \in \{1, \ldots, k\}$ 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 $(x_1, \ldots, x_{i-1}, x_i \oplus s, x_{i+1}, \ldots, x_k) = 0$, 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 $\mathbf{P}$-position. □

1.1.2 Other impartial games

Example 1.1.4 (**Staircase Nim**)**.** This game is played on a staircase of n steps. On each step j for $j = 1, \ldots, n$ 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.

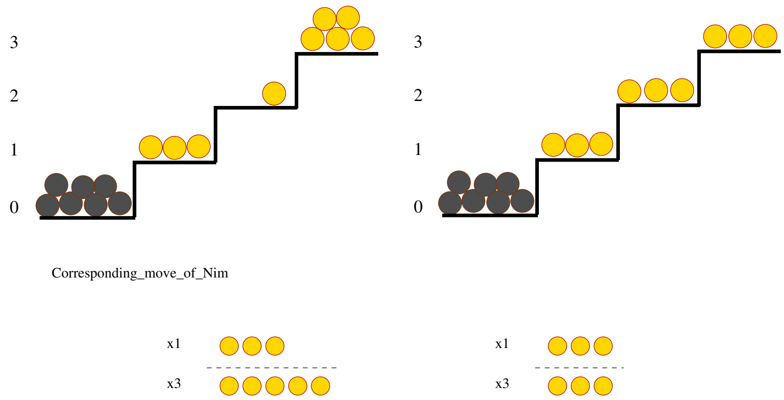

Fig. 1.4. A move in Staircase Nim, in which 2 coins are moved from step 3 to step 2. Considering the odd stairs only, the above move is equivalent to the move in regular Nim from $(3,5)$ to $(3,3)$.

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 1.1.5 (**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.

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

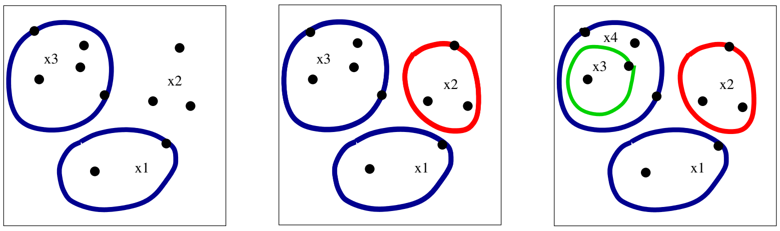

Fig. 1.5. Two moves in a game of Rims.

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.

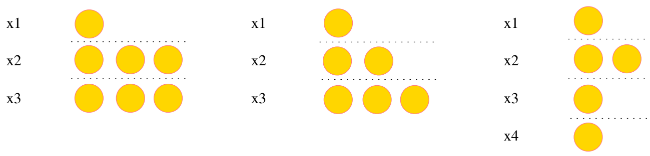

Fig. 1.6. Equivalent sequence of moves in Nim with splittings allowed.

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 1.1.4. *The sets* **N** *and* **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 $\mathbf{P}_{\mathrm{Nim}}$ and $\mathbf{N}_{\mathrm{Nim}}$ for the **P**- and **N**-positions for the game of Nim (these sets are described by Bouton's theorem).

From any position in $\mathbf{N}_{\mathrm{Nim}}$, we may move to $\mathbf{P}_{\mathrm{Nim}}$ by a move in Rims, because each Nim move is legal in Rims.

Next we consider a position $x \in \mathbf{P}_{\mathrm{Nim}}$. 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 $\mathbf{P}_{\mathrm{Nim}}$-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 $\mathbf{P}_{\mathrm{Nim}}$ is to a position in $\mathbf{N}_{\mathrm{Nim}}$.

Thus the strategy of always moving to a position in $\mathbf{P}_{\mathrm{Nim}}$ (if this is possible) will guarantee a win for a player who starts in an $\mathbf{N}_{\mathrm{Nim}}$-position, and if a player starts in a $\mathbf{P}_{\mathrm{Nim}}$-position, this strategy will guarantee a win for the second player. Thus $\mathbf{N}_{\mathrm{Rims}} = \mathbf{N}_{\mathrm{Nim}}$ and $\mathbf{P}_{\mathrm{Rims}} = \mathbf{P}_{\mathrm{Nim}}$. □

The following examples are particularly tricky variants of Nim.

Example 1.1.6 (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 $(n_1, \ldots, n_\ell)$:

$$\begin{aligned} n_1 &= n_1^{(m)} \cdots n_1^{(0)} &= \sum_{j=0}^{m} n_1^{(j)} 2^j, \\ &\vdots \\ n_\ell &= n_\ell^{(m)} \cdots n_\ell^{(0)} &= \sum_{j=0}^{m} n_\ell^{(j)} 2^j, \end{aligned}$$

where each $n_i^{(j)}$ is either 0 or 1.

Theorem 1.1.5 (Moore's Theorem). *For Moore's Nim$_k$,*

$$\mathbf{P} = \Big\{ (n_1, \ldots, n_\ell) : \sum_{i=1}^{\ell} n_i^{(j)} \equiv 0 \bmod (k+1) \textit{ for each } j \Big\}.$$

The notation "$a \equiv b \bmod m$" means that $a - b$ is evenly divisible by m, i.e., that $(a - b)/m$ is an integer.

Proof of Theorem 1.1.5. 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:

pile sizes in binary

$$
\begin{array}{c}
0010010011010011\\
101000\boxed{1}101000011\\
1010010110\boxed{1}01010\\
100\boxed{1}011100100111\\
1010010101000010\\
1000000100010111\\
001\boxed{1}100110100001\\
\hline
5052142633321265
\end{array}
\quad\Rightarrow\quad
\begin{array}{c}
0010010011010011\\
101000\boxed{0}111010111\\
1010010110\boxed{0}10111\\
100\boxed{0}010111010101\\
1010010101000010\\
1000000100010111\\
001\boxed{0}010011000100\\
\hline
5050050555050555
\end{array}
$$

pile sizes in binary

Fig. 1.7. Example move in Moore's Nim_4 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.

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 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 ($0 \le a \le k$), 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 $b \le a$, 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, $a < b \le k$, and $b - a = s \bmod (k + 1) \le s$, 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 1.1.7 (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 $m, n \ge 0$. A legal move is one of the following:

Reduce m to some value between 0 and $m - 1$ without changing m, 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.

To analyze Wythoff Nim (and other games), we define

$$\operatorname{mex}(S) = \min\{n \ge 0 : n \notin S\},$$

for $S \subseteq \{0, 1, \ldots\}$ (the term "mex" stands for "minimal excluded value"). For example, $\operatorname{mex}(\{0, 1, 2, 3, 5, 7, 12\}) = 4$. Consider the following recursive definition of two sequences of natural numbers: For each $k \ge 0$,

$$a_k = \operatorname{mex}(\{a_0, a_1, \ldots, a_{k-1}, b_0, b_1, \ldots, b_{k-1}\}), \text{ and } b_k = a_k + k.$$

Notice that when $k = 0$, we have $a_0 = \operatorname{mex}(\{\}) = 0$ 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, $a_4 = \operatorname{mex}(\{0, 1, 3, 4, 0, 2, 5, 7\}) = 6$ and $b_4 = a_4 + 4 = 10$.)

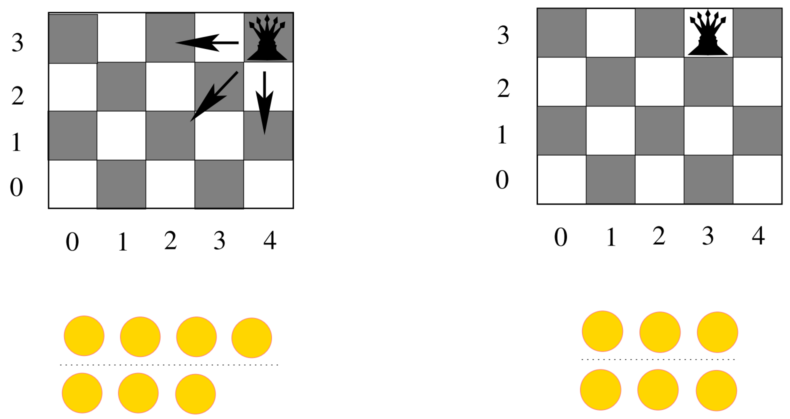

Fig. 1.8. Wythoff Nim can be viewed as the following game played on a chess board. Consider an $m \times n$ section of a chess-board. The players take turns moving a queen, initially positioned in the upper right corner, either left, down, or diagonally toward the lower left. The player that moves the 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.

Theorem 1.1.6. *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 have that both a_j and a_{j-1} are excluded from $\{a_0, \ldots, a_{j-2}, b_0, \ldots, b_{j-2}\}$, 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 $a_j \notin \{b_0, \ldots, b_{j-1}\}$, 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

$$\{1, \ldots, j\} \subset \{a_i\}_{i=1}^{j} \cup \{b_i\}_{i=1}^{j} .$$

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 $a_{j+1} = j + 1$. □

It is easy to check the following theorem:

Theorem 1.1.7. *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 $\mathbf{P} = \hat{P}$. ∎

It turns out that there is there a fast, non-recursive, method to decide if a given position is in $\mathbf{P}$:

Theorem 1.1.8. $a_k = \lfloor k(1 + \sqrt{5})/2 \rfloor$ *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 1.1.8. Consider the following sequences positive integers: Fix any irrational $\theta \in (0, 1)$, and set

$$\alpha_k(\theta) = \lfloor k/\theta \rfloor, \qquad \beta_k(\theta) = \lfloor k/(1 - \theta) \rfloor.$$

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

$$\alpha_k(\theta) = a_k \quad \text{and} \quad \beta_k(\theta) = b_k? \tag{1.1}$$

We will show that there is only one θ for which this is true.

Because $b_k = a_k + k$, (1.1) 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 $k \to \infty$ we find that

$$1/\theta + 1 = 1/(1-\theta). \tag{1.2}$$

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 (1.2) 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

$$\alpha_k = \operatorname{mex}\{\alpha_0, \ldots, \alpha_{k-1}, \beta_0, \ldots, \beta_{k-1}\}.$$

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 $z < \alpha_k \le \alpha_\ell \le \beta_\ell$ for all $\ell \ge k$. Since z is defined as a mex, $z \neq \alpha_i, \beta_i$ for $i \in \{0, \ldots, k-1\}$, 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. □

1.1.3 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 1.1.5. 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 1.1.8. The sum of two Nim games X and Y is the game (X, Y) as defined in Lemma 1.1.1 of the previous section.

It is easy to see that Lemma 1.1.1 generalizes to the sum of any two progressively bounded combinatorial games:

Theorem 1.1.9. *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 1.1.1 for Nim, replace the number of chips with $B(x)$, the maximum number of moves in the game. □

Definition 1.1.6. 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 1.1.9 that in any two progressively bounded impartial combinatorial games, the **P**-positions are equivalent to each other.

In Exercise 1.12 you will prove that this notion of equivalence for games defines an *equivalence relation*. In Exercise 1.13 you will prove that two impartial games are equivalent if and only if there sum is a **P**-position. In Exercise 1.14 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 1.1.9. 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 $(n_1, \ldots, n_k)$ is equivalent to $(n_1 \oplus \cdots \oplus n_k)$ because the Nim-sum of $(n_1, \ldots, n_k, n_1 \oplus \cdots \oplus n_k)$ 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 1.1.7. Let G be a progressively bounded impartial combinatorial game under normal play. Its **Sprague-Grundy function** g is defined recursively as follows:

$$g(x) = \text{mex}(\{g(y) : x \to y \text{ is a legal move}\}).$$

Note that the Sprague-Grundy value of any terminal position is $\text{mex}(\varnothing) = 0$. In general, the Sprague-Grundy function has the following key property:

Lemma 1.1.2. *In a progressively bounded impartial combinatorial game, the Sprague-Grundy value of a position is* 0 *if and only if it is a* $\mathbf{P}$*-position.*

Proof. Proceed as in the proof of Theorem 1.1.3 — 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. 1.15). □

Let's calculate the Sprague-Grundy function for a few examples.

Example 1.1.10 (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 1.1.11 (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 1.1.12. 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 1.1.10 (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 1.1.10. 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 1.1.1. *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 1.1.11 (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 $G = G_1 + G_2$,*

$$g(x_1, x_2) = g_1(x_1) \oplus g_2(x_2), \tag{1.3}$$

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 $g(x_1, x_2) = g(x_1) \oplus g(x_2)$. □

Let's use the Sprague-Grundy and the Sum Theorems to analyze a few games.

Example 1.1.13. (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 1.1.14 (**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 1.1.3 (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 [**?**, 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. □

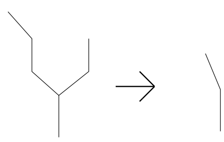

Fig. 1.9. Combining branches in a tree of Green Hackenbush.

As a simple illustration, see Fig. 1.9. 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 [**?**, Part I, Sect. 6] or [**?**].

Next we leave the impartial and discuss a few interesting partisan games.

1.2 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 $E_{\mathrm{I}}, E_{\mathrm{II}}$ be the two edge-sets for players I and II respectively. If (x, y) is a legal move for player $i \in \{\mathrm{I}, \mathrm{II}\}$ 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 1.2.1 (**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.

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 $i \in \{\mathrm{I}, \mathrm{II}\}$ denotes the player

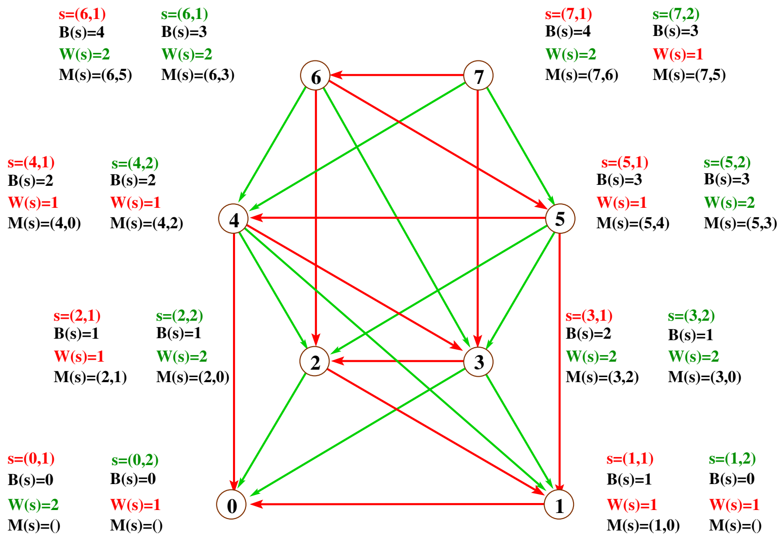

Fig. 1.10. Moves of the partisan Subtraction game. Node 0 is terminal for either player, and node 1 is also terminal with a win for player I.

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 1.2.1. *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 1.2.1. We will recursively define a function W, which specifies the winner for a given state of the game: $W(x,i) = j$ where

$i, j \in \{\mathrm{I}, \mathrm{II}\}$ 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. □

This proof relies essentially on the game being progressively bounded. Next we show that many games have this property.

Lemma 1.2.1. *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 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. □

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.

1.2.1 The game of Hex

Recall the description of Hex from the introduction.

Example 1.2.2 (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.

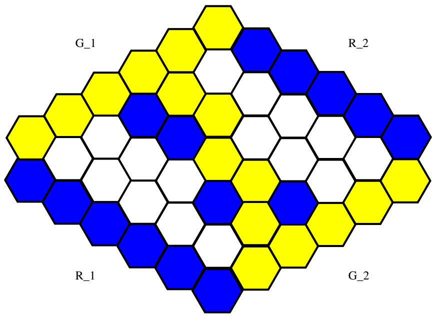

Fig. 1.11. A completed game of Hex with a yellow chain crossing.

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 1.2.2. *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. □

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 [?]. 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 [?], 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.

*1.2.2 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 3.5. 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 1.2.3. *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. 1.12.

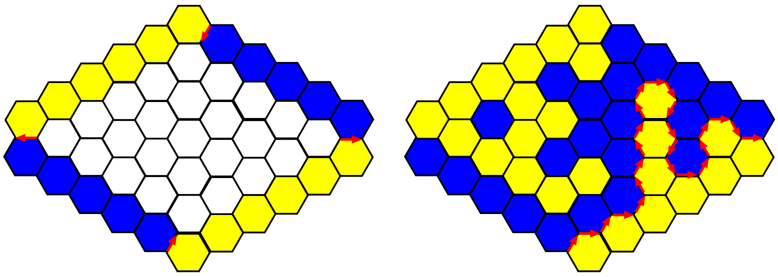

Fig. 1.12. On an empty board the entry and exit arrows are marked. On a completed board, a blue chain lies on the left side of the directed path.

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. 1.13.

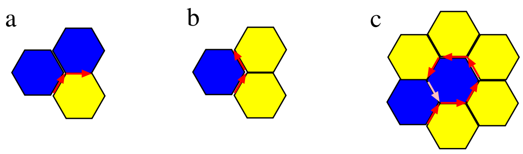

Fig. 1.13. In (a) the third hexagon is blue and the next arrow turns to the right; in (b) — next arrow turns to the left; in (c) we see that in order to close the loop an arrow would have to pass between two hexagons of the same color.

Loops are not possible, as you can see from (c) of Fig. 1.13. A loop circling to the left, for instance, would circle an isolated group of blue hexagons 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. 1.12). If it touches the same blue segment, there is a yellow crossing. □

1.2.3 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 1.2.3 (**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. 1.14 is equivalent to playing Hex in the empty region of the board.

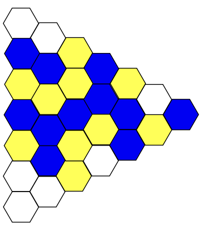

Blue has a winning Y here.

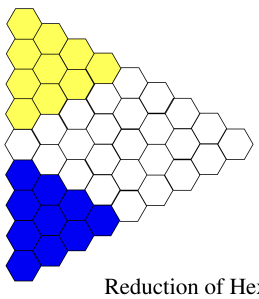

Reduction of Hex to Y.

Fig. 1.14. 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 1.2.4. *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.

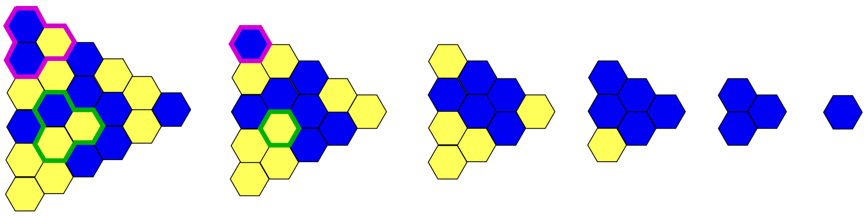

Fig. 1.15. A step-by-step reduction of a colored Y board.

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. 1.3).

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. □

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.

*1.2.4 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 1.2.4 (**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.

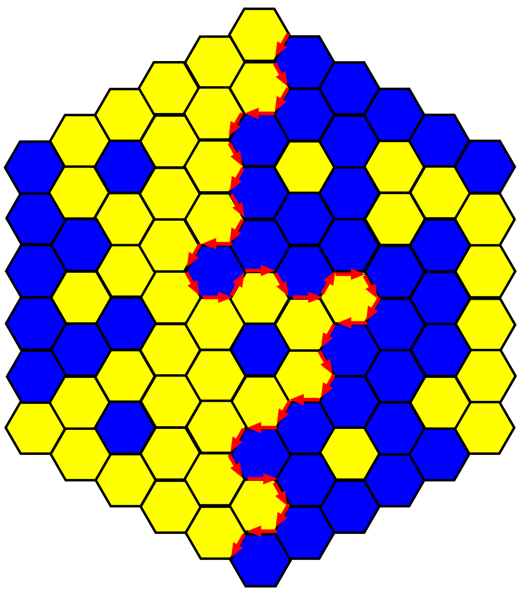

Fig. 1.16. A filled-in Six-sided Hex board can have both blue and yellow crossings. In a game when players take turns to move, one of the crossings will occur first, and that player will be the winner.

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 1.2.5. *For an arbitrarily shaped simply-connected completed Hex* *board with non-empty interior and the boundary partitioned into n blue and and n yellow segments, with $n \geq 2$, there is always at least one crossing between some pair of segments of like color.*

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. 1.16).

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.

1.2.5 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 1.2.5 (**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.

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 1.2.6 (**Bridg-It**)**.** Bridg-It is played on a network of green and black dots (see Fig. 1.18). 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.

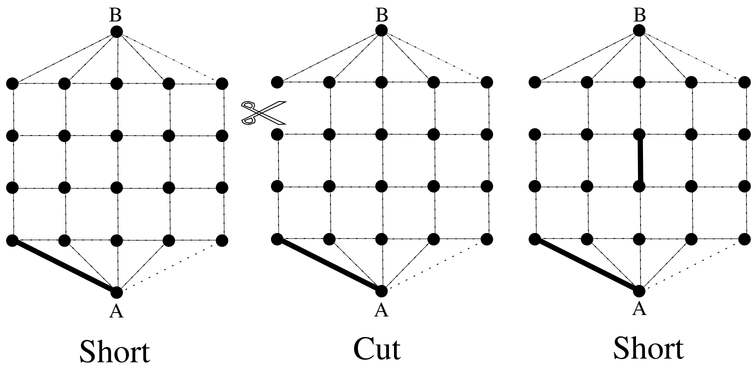

Fig. 1.17. Shannon Switching Game played on a 5×6 grid (the top and bottom rows have been merged to the points A and B). Shown are the first three moves of the game, with Short moving first. Available edges are indicated by dotted lines, and reinforced edges by thick lines. Scissors mark the edge that Short deleted.

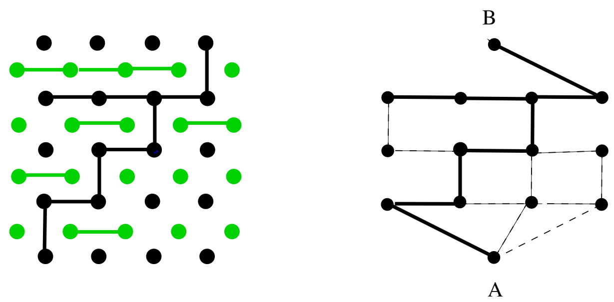

Fig. 1.18. A completed game of Bridg-It and the corresponding Shannon Switching Game. In Bridg-It, the black dots are on the square lattice, and the green dots are on the dual square lattice. Only the black dots appear in the Shannon Switching Game.

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 [**?**] (see also [**?**]), 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 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 ([**?**]).

Before we can describe Short's strategy, we will need a few definitions from graph theory:

Definition 1.2.1. 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. 1.4).

Theorem 1.2.6. *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.

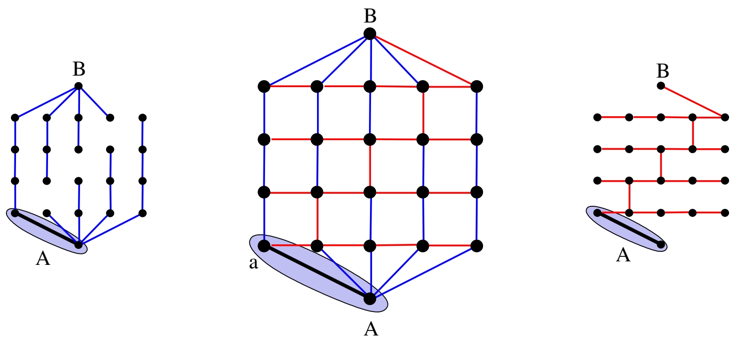

Fig. 1.19. Two spanning trees — the blue one is constructed by first joining top and bottom using the left-most vertical edges, and then adding other vertical edges, omitting exactly one edge in each row along an imaginary diagonal; the red tree contains the remaining edges. The two circled nodes are identified.

Observe that the blue and red subgraphs in the 4×5 grid in Fig. 1.19 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.

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. 1.20), Short can repair the tree as follows: Because

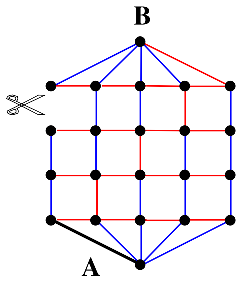

Fig. 1.20. Cut separates the blue tree into two components.

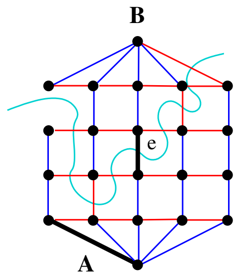

Fig. 1.21. Short reinforces a red edge to reconnect the two components.

the other tree is also a spanning tree, it must have an edge, e, that connects the two components (see Fig. 1.21). 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 reinforced edge each time Cut disconnects them. Thus, Cut will never succeed in disconnecting A from B, and Short will win. $\square$

Example 1.2.7 (**Recursive Majority**)**.** Recursive Majority is played on a complete ternary tree of height h (see Fig. 1.22). 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.

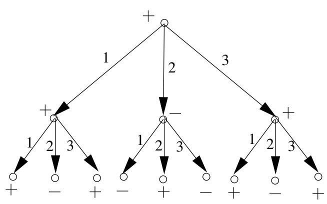

Fig. 1.22. A ternary tree of height 2; the left-most leaf is denoted by 11. Here player I wins the Recursive Majority game.

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 $11\ldots1$, 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 $11\ldots1$ 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 $11\ldots1$, moving one branch to the right, and then retracing the equivalent path back down (see Fig. 1.23). 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: $1^k2w \mapsto 1^k3w$.

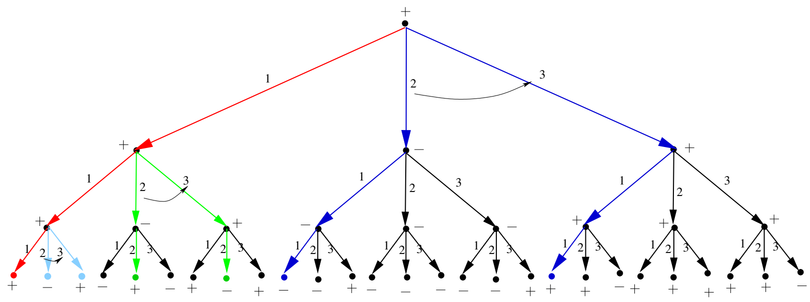

Fig. 1.23. Red marks the left-most leaf and its path. Some sample pair-mates are marked with the same shade of green or blue.

Once the pairs have been identified, for every leaf marked with a "$-$" by player II, player I marks its mate with a "$+$".

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 1.5.

Exercises

1.1 In the game of Chomp, what is the Sprague-Grundy function of the 2×3 rectangular piece of chocolate?

1.2 Recall the game of Y, shown in Fig. 1.14. 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.

1.3 Consider the reduction of a Y board to a smaller one described in section 1.2.1. 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.

1.4 Prove the following statements. Hint: use induction.

(a) Every tree must have a **leaf** — a vertex of degree one.
(b) A tree on n vertices has $n-1$ edges.
(c) A connected graph with n vertices and $n-1$ edges is a *tree*.
(d) A graph with no cycles, n vertices and $n-1$ edges is a *tree*.

1.5 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 1.2.7 is indeed a winning strategy for player I.

1.6 Consider a game of Nim with four piles, of sizes $9, 10, 11, 12$.

(a) Is this position a win for the *next* player or the *previous* player (assuming optimal play)? Describe the winning first move.
(b) 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?

1.7 Consider a game where there are two piles of chips. On a players 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 $(m, n) \in \mathbf{P}$.

1.8 For the game of Moore's Nim, the proof of Lemma 1.1.5 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.

1.9 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.

1.10 Recall that the subtraction game with subtraction set $\{a_1, \ldots, a_m\}$ 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 $i \in \{1, \ldots, m\}$. Find the Sprague-Grundy function for the subtraction game with subtraction set $\{1, 2, 4\}$.

1.11 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 $S_3 = \{1, 2, \ldots, 20\}$. Who has a winning strategy from the starting position $(100, 100, 100)$ in $G_1 + G_2 + G_3$?

1.12 (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.

1.13 Prove that the sum of two progressively bounded impartial combinatorial games is a **P**-position if and only if the games are equivalent.

1.14 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.

1.15 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 1.1.2 and the proof is outlined in the text.

1.16 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?

1.17 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?

1.18 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.

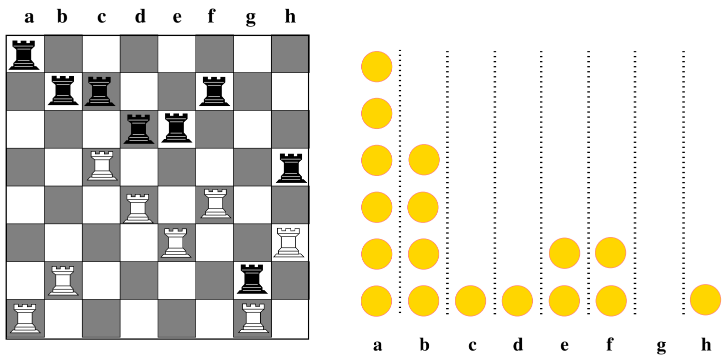

1.19 Two players take turns placing dominos on an $n \times 1$ board of squares, where each domino covers two squares, and dominos cannot overlap. The last player to play wins.

(a) Find the Sprague-Grundy function for $n \le 12$.
(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.

2

Two-person zero-sum games

In the previous chapter, we studied games that are deterministic; nothing is left to chance. In the next two chapters, we will shift our attention to the games in which the players, in essence, move simultaneously, and thus do not have full knowledge of the consequences of their choices. As we will see, chance plays a key role in such games.

In this chapter, we will restrict our attention to **two-person zero-sum games**, in which one player loses what the other gains in every outcome. The central theorem for this class of games says that even if each player's strategy is known to the other, there is an amount that one player can guarantee as his expected gain, and the other, as his maximum expected loss. This amount is known as the *value* of the game.

2.1 Preliminaries

Let's start with a very simple example:

Example 2.1.1 (**Pick-a-hand, a betting game**)**.** There are two players, a chooser (player I), and a hider (player II). The 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, or takes out both and holds them in his right hand. The 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.

We can record all possible outcomes in the form of a **payoff matrix**, whose rows are indexed by player I's possible choices, and whose columns are indexed by player II's choices. Each matrix entry $a_{i,j}$ is the amount that player II loses to player I when I plays i and II plays j. We call this description of a game its **normal** or **strategic** form.

chooser \ hider	$L1$	$R2$
L	1	0
R	0	2

Suppose that hider seeks to minimize his losses by placing one coin in his left hand, ensuring that the most he will lose is that coin. This is a reasonable strategy if he could be certain that chooser has no inkling of what he will choose to do. But suppose chooser learns or reasons out his strategy. Then he loses a coin when his best hope is to lose nothing. Thus, if hider thinks chooser might guess or learn that he will play $L1$, he has an incentive to play $R2$ instead. Clearly, the success of the strategy $L1$ (or $R2$) depends on how much information chooser has. All that hider can guarantee is a maximum loss of one coin.

Similarly, chooser might try to maximize her gain by picking R, hoping to win two coins. If hider guesses or discovers chooser's strategy, however, then he can ensure that she doesn't win anything. Again, without knowing how much hider knows, chooser cannot assure that she will win anything by playing.

Ideally, we would like to find a strategy whose success does not depend on how much information the other player has. The way to achieve this is by introducing some uncertainty into the players' choices. A strategy with uncertainty — that is, a strategy in which a player assigns to each possible move some fixed probability of playing it — is known as a **mixed strategy**. A mixed strategy in which a particular move is played with probability one is known as a **pure strategy**.

Suppose that chooser decides to follow a mixed strategy of choosing R with probability p and L with probability $1-p$. If hider were to play the pure strategy $R2$ (hide two coins in his right hand) his expected loss would be $2p$. If he were to play $L1$ (hide one coin in his left hand), then his expected loss would be $1-p$. Thus, if he somehow learned p, he would play the strategy corresponding to the minimum of $2p$ and $1-p$. Expecting this, chooser would maximize her gains by choosing p so as to maximize $\min\{2p, 1-p\}$.

Note that this maximum occurs at $p = 1/3$, the point at which the two lines cross:

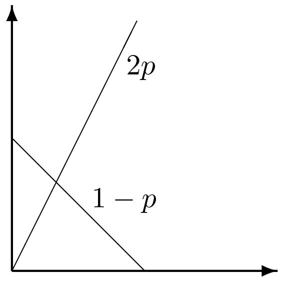

Thus, by following the mixed strategy of choosing R with probability $1/3$ and L with probability $2/3$, chooser assures an expected payoff of $2/3$, regardless of whether hider knows her strategy. How can hider minimize his expected loss?

Hider will play $R2$ with some probability q and $L1$ with probability $1-q$. The payoff for chooser is $2q$ if she picks R, and $1-q$ if she picks L. If she knows q, she will choose the strategy corresponding to the maximum of the two values. If hider, in turn, knows chooser's plan, he will choose $q = 1/3$ to minimize this maximum, guaranteeing that his expected payout is $2/3$ (because $2/3 = 2q = 1-q$).

Thus, chooser can assure an expected gain of $2/3$ and hider can assure an expected loss of no more than $2/3$, regardless of what either knows of the other's strategy. Note that, in contrast to the situation when the players are limited to pure strategies, the assured amounts are equal. Von Neumann's minimax theorem, which we will prove in the next section, says that this is always the case in any two-person, zero-sum game.

Clearly, without some extra incentive, it is not in hider's interest to play *Pick-a-hand* because he can only lose by playing. Thus, we can imagine that chooser pays hider to entice him into joining the game. In this case, $2/3$ is the maximum amount that chooser should pay him in order to gain his participation.

Let's look at another example.

Example 2.1.2 (**Another Betting Game**)**.** A game has the following payoff matrix:

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

Suppose player I plays T with probability p and B with probability $1-p$, and player II plays L with probability q and R with probability $1-q$.

Reasoning from player I's perspective, note that her expected payoff is $2(1-q)$ for playing the pure strategy T, and $4q+1$ for playing the pure strategy B. Thus, if she knows q, she will pick the strategy corresponding to the maximum of $2(1-q)$ and $4q+1$. Player II can choose $q = 1/6$ so as to minimize this maximum, and the expected amount player II will pay player I is $5/3$.

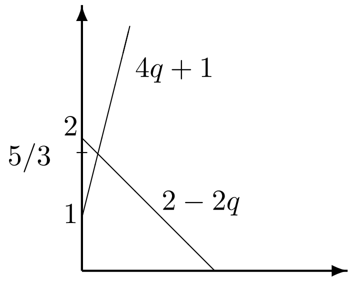

If player II instead chose a higher value of q, say $q = 1/3$, and player I knows this, then player I can play pure strategy B to get an expected payoff of $4q+1 = 7/3 > 5/3$. Similarly, if player II instead chose a smaller value of q, say $q = 1/12$, and player I knows this, then player I can play pure strategy T to get an expected payoff of $2(1-q) = 11/6 > 5/3$.

From player II's perspective, his expected loss is $5(1-p)$ if he plays the pure strategy L and $1+p$ if he plays the pure strategy R, and he will aim to minimize this expected payout. In order to maximize this minimum, player I will choose $p = 2/3$, which again yields an expected gain of $5/3$.

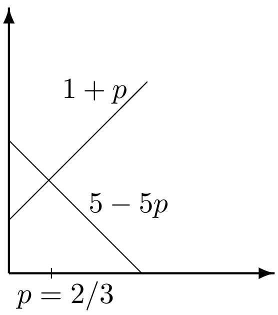

Now, let's set up a formal framework for our theory.

For an arbitrary two-person zero-sum game with $m \times n$ payoff matrix $A = (a_{i,j})_{j=1,\ldots,n}^{i=1,\ldots,m}$, a mixed strategy for player I corresponds to a vector $(x_1, \ldots, x_m)$ where x_i represents the probability of playing pure strategy i. The set of mixed strategies for player I is denoted by

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

(since the probabilities are nonnegative and add up to 1), and the set of mixed strategies for player II by

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

Observe that in this vector notation, pure strategies are represented by the standard basis vectors.

If player I follows a mixed strategy $\mathbf{x}$, and player II follows a mixed strategy $\mathbf{y}$, then with probability $x_i y_j$ player I plays i and player II plays j, resulting in payoff $a_{i,j}$ to player I. Thus the expected payoff to player I is $\sum_{i,j} x_i a_{i,j} y_j = \mathbf{x}^T A \mathbf{y}$.

We refer to $A\mathbf{y}$ as the **payoff vector** for player I corresponding to the mixed strategy $\mathbf{y}$ for player II. The elements of this vector represent the **expected payoffs** to player I corresponding to each of his pure strategies. Similarly, $\mathbf{x}^T A$ is the **payout vector** for player II corresponding to the mixed strategy $\mathbf{x}$ for player I. The elements of this vector represent the **expected payouts** for each of player II's pure strategies.

We say that a vector $\mathbf{w} \in \mathbb{R}^d$ **dominates** another vector $\mathbf{u} \in \mathbb{R}^d$ if $w_i \geq u_i$ for all $i = 1, \ldots, d$. We write $\mathbf{w} \geq \mathbf{u}$.

Next we formally define what it means for a strategy to be optimal for each player:

Definition 2.1.1. A mixed strategy $\tilde{\mathbf{x}} \in \Delta_m$ is **optimal for player I** if

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

Similarly, a mixed strategy $\tilde{\mathbf{y}} \in \Delta_n$ is **optimal for player II** if

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

Notice that in the definiton of an optimal strategy for player I, we give player II the advantage of knowing what strategy player I will play. Similarly, in the definition of an optimal strategy for player II, player I has the advantage of knowing how player II will play. A priori the expected payoffs could be different depending on which player has the advantage of knowing how the other will play. But as we shall see in the next section, these two expected payoffs are the equal at every two-person zero-sum game.

2.2 Von Neumann's minimax theorem

In this section, we will prove that every two-person, zero-sum game has a value. That is, in any two-person zero-sum game, the expected payoff for an optimal strategy for player I equals the expected payout for an optimal strategy of player II. Our proof will rely on a basic theorem from convex geometry.

Definition 2.2.1. 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.

Our proof will make use of the following result about convex sets:

Theorem 2.2.1 (**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 $z^T v$ is the usual dot product $\sum_i z_i v_i$. The theorem says that there is a **hyperplane** (a line in the plane, 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 $\{\mathbf{x} \in \mathbb{R}^d : \mathbf{z}^T \mathbf{x} > c\}$.

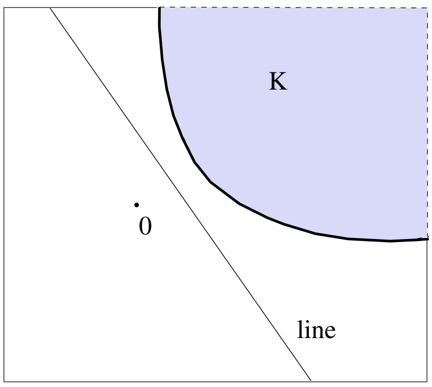

Fig. 2.1. Hyperplane separating the closed convex body K from $\mathbf{0}$.

Recall first that the **(Euclidean) norm of 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 2.2.1. If we pick R so that the ball of radius R centered at $\mathbf{0}$ intersects K, the function $\mathbf{v} \mapsto \|\mathbf{v}\|$, considered as a map from $K \cap \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\| \le R\}$ to $[0, \infty)$, is continuous, with a domain that is nonempty, closed and bounded (see Figure 2.2). Thus the map attains its infimum at some point $\mathbf{z}$ in K. For this $\mathbf{z} \in K$ we have

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

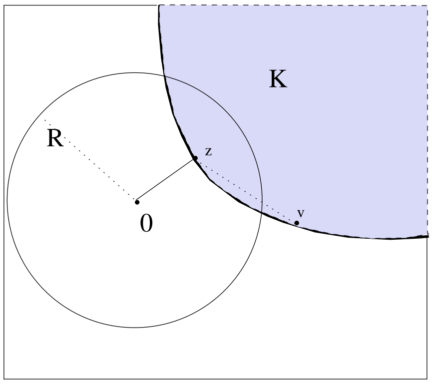

Fig. 2.2. Intersecting K with a ball to get a nonempty closed bounded domain.

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

$$\|\mathbf{z}\|^2 \le \|\varepsilon\mathbf{v} + (1-\varepsilon)\mathbf{z}\|^2.$$

Multiplying this out,

$$\begin{aligned}
\mathbf{z}^T\mathbf{z} &\le \big(\varepsilon\mathbf{v}^T + (1-\varepsilon)\mathbf{z}^T\big)\big(\varepsilon\mathbf{v} + (1-\varepsilon)\mathbf{z}\big) \\
\mathbf{z}^T\mathbf{z} &\le \varepsilon^2\mathbf{v}^T\mathbf{v} + (1-\varepsilon)^2\mathbf{z}^T\mathbf{z} + 2\varepsilon(1-\varepsilon)\mathbf{z}^T\mathbf{v}.
\end{aligned}$$

Rearranging terms we get

$$\varepsilon^2(2\mathbf{z}^T\mathbf{v} - \mathbf{v}^T\mathbf{v} - \mathbf{z}^T\mathbf{z}) \le 2\varepsilon(\mathbf{z}^T\mathbf{v} - \mathbf{z}^T\mathbf{z}).$$

Canceling an ε, and letting ε approach 0, we find

$$0 \le \mathbf{z}^T\mathbf{v} - \mathbf{z}^T\mathbf{z},$$

which means

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

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 $\mathbf{v} \in K$. □

We will also need the following simple lemma:

Lemma 2.2.1. *Let X and Y be closed and bounded sets in $\mathbb{R}^d$. Let $f : X \times Y \to \mathbb{R}$ be continuous. Then*

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

Proof. Let $(\mathbf{x}^*,\mathbf{y}^*) \in X \times Y$. Clearly we have $f(\mathbf{x}^*,\mathbf{y}^*) \le \sup_{\mathbf{x}\in X} f(\mathbf{x},\mathbf{y}^*)$ and $\inf_{\mathbf{y}\in Y} f(\mathbf{x}^*,\mathbf{y}) \le f(\mathbf{x}^*,\mathbf{y}^*)$, which gives us

$$\inf_{\mathbf{y}\in Y} f(\mathbf{x}^*,\mathbf{y}) \le \sup_{\mathbf{x}\in X} f(\mathbf{x},\mathbf{y}^*).$$

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

$$\sup_{\mathbf{x}\in X}\inf_{\mathbf{y}\in Y} f(\mathbf{x},\mathbf{y}) \le \inf_{\mathbf{y}\in Y}\sup_{\mathbf{x}\in X} f(\mathbf{x},\mathbf{y}).$$

Because f is continuous and X and Y are closed and bounded, the minima and maxima are achieved and we have proved the lemma. □

We can now prove:

Theorem 2.2.2 (**Von Neumann's Minimax Theorem**)**.** *Let A be an $m \times n$ payoff matrix, and let $\Delta_m = \{\mathbf{x} \in \mathbb{R}^m : \mathbf{x} \ge \mathbf{0}, \sum_i x_i = 1\}$ and $\Delta_n = \{\mathbf{y} \in \mathbb{R}^n : \mathbf{y} \ge \mathbf{0}, \sum_j y_j = 1\}$. 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.*

By $\mathbf{x} \ge \mathbf{0}$ we mean simply that in each coordinate $\mathbf{x}$ is at least as large as $\mathbf{0}$, i.e., that each coordinate is nonnegative. This condition together with $\sum_i x_i = 1$ ensure that $\mathbf{x}$ is a probability distribution.

Proof. The inequality

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

follows immediately from the lemma 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 game, we have

$$\max_{\mathbf{x}\in\Delta_m} \min_{\mathbf{y}\in\Delta_n} \mathbf{x}^T \hat{A}\mathbf{y} < 0 < \min_{\mathbf{y}\in\Delta_n} \max_{\mathbf{x}\in\Delta_m} \mathbf{x}^T \hat{A}\mathbf{y}. \tag{2.1}$$

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

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

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 fact that $\{\mathbf{v} \in \mathbb{R}^m, \mathbf{v} \geq \mathbf{0}\}$ 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 $\mathbf{x}^T \hat{A}\mathbf{y} \leq 0$, which would contradict the right-hand side of (2.1).

Thus K satisfies the conditions of the separating hyperplane theorem (Theorem 2.2.1), which gives us $\mathbf{z} \in \mathbb{R}^m$ and $c > 0$ such that $c < \mathbf{z}^T\mathbf{w}$ for all $\mathbf{w} \in K$. That is,

$$\mathbf{z}^T(\hat{A}\mathbf{y} + \mathbf{v}) > c > 0 \text{ for all } \mathbf{y} \in \Delta_n \text{ and } \mathbf{v} \geq \mathbf{0}. \tag{2.2}$$

If $z_j < 0$ for some j, we could choose $\mathbf{v} \in \mathbb{R}^m$ so that $\mathbf{z}^T\hat{A}\mathbf{y} + \sum_i z_i v_i$ would be negative for some $\mathbf{y} \in \Delta_n$ (let $v_i = 0$ for $i \neq j$ and $v_j \to \infty$), which would contradict (2.2). Thus $\mathbf{z} \geq \mathbf{0}$.

The same condition (2.2) gives us that not all of the z_i can be zero. This implies that $s = \sum_{i=1}^m z_i$ is strictly positive, so that $\tilde{\mathbf{x}} = \frac{1}{s}(z_1, \ldots, z_m)^T = \mathbf{z}/s \in \Delta_m$, 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 (2.1). □

Note that the above proof merely shows that the value always exists; it doesn't give a way of finding it. Finding the value of a zero-sum game involves solving a linear program, which typically requires a computer for all but the simplest of payoff matrices. In many cases, however, the payoff matrix of a game can be simplified enough to solve it “by hand”. In the next two sections of the chapter, we will look at some techniques for simplifying a payoff matrix.

p.57, the displayed matrix is not aligned, the zero-s do not form a diagonal.

2.3 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 2.3.1 (Plus One). Each player chooses a number from $\{1, 2, \ldots, n\}$ and writes it down on a piece of paper; 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	$\cdots$		n
	1	0	-1	2	2	2	2	$\cdots$		2
	2	1	0	-1	2	2	2	$\cdots$		2
	3	-2	1	0	-1	2	2	$\cdots$		2
player I	4	-2	-2	1	0	-1	2	$\cdots$		2
	5	-2	-2	-2	1	0	-1	2		2
	$\vdots$	$\vdots$	$\vdots$					$\ddots$		$\vdots$
	$n-1$	-2	-2	$\cdots$				1	0	-1
	n	-2	-2	$\cdots$					1	0

In general, if each element of row i_1 of a payoff matrix is at least as big as the corresponding element in row i_2, that is, if $a_{i_1,j} \geq a_{i_2,j}$ for each j, then, for the purpose of determining the value of the game, we may erase row i_2. Similarly, there is a notion of domination for player II: If $a_{i,j_1} \leq a_{i,j_2}$ for each i, then we can eliminate column j_2 without affecting the value of the game.

Why is it okay to do this? Assuming that $a_{i,j_1} \leq a_{i,j_2}$ for each i, if player II changes a mixed strategy $\mathbf{y}$ to another $\mathbf{z}$ by letting $z_{j_1} = y_{j_1} + y_{j_2}$, $z_{j_2} = 0$

and $z_\ell = y_\ell$ for all $\ell \neq j_1, j_2$, then

$$\mathbf{x}^T A\mathbf{y} = \sum_{i,\ell} x_i a_{i,\ell} y_\ell \geq \sum_{i,\ell} x_i a_{i,\ell} z_\ell = \mathbf{x}^T A\mathbf{z},$$

because $x_i(a_{i,j_1} y_j + a_{i,j_2} y_{j_2}) \geq x_i a_{i,j_1}(z_j + z_{j_2})$. Therefore, strategy $\mathbf{z}$, in which she didn't use column j_2, is at least as good for player II as $\mathbf{y}$.

In our example, we may eliminate each row and column indexed by four or greater (the reader should verify this) to obtain:

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

To analyze the reduced game, let $\mathbf{x} = (x_1, x_2, x_3)$ correspond to a mixed strategy for player I. The expected payments made by player II for each of her pure strategies 1, 2 and 3 are

$$\big(x_2 - 2x_3, -x_1 + x_3, 2x_1 - x_2\big). \tag{2.3}$$

Player II will try to minimize her expected payment. Player I will choose (x_1, x_2, x_3) so as to maximize the minimum.

For player I's optimal strategy (x_1, x_2, x_3), each component of the payoff vector in (2.3) must be at least the value of the game. For this game, the payoff matrix is antisymmetric, so the value must be 0. Thus $x_2 \geq 2x_3$, $x_3 \geq x_1$, and $2x_1 \geq x_2$. If any one of these inequalities were strict, then combining them we could deduce $x_2 > x_2$, a contradiction, so in fact each of them is an equality. Since the x_i's add up to 1, we find that the optimal strategy for each player is $(1/4, 1/2, 1/4)$.

Remark. It can of course happen in a game that none of the rows dominates another one, but there are two rows, v, w, whose convex combination $pv + (1-p)w$ for some $p \in (0,1)$ does dominate some other rows. In this case the dominated rows can still be eliminated.

2.4 The use of symmetry

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

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. He wins \$1 if he hits the

Example 2.4.1 (Submarine Salvo).

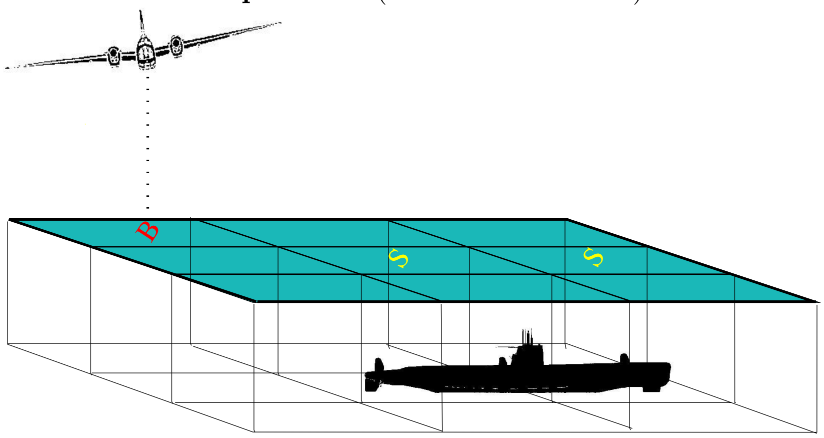

Fig. 2.3.

submarine and \$0 if he 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, but by using symmetry arguments, we can greatly simplify the analysis.

Note that there are three types of essentially equivalent moves that the bomber can make: He can drop a bomb in the center, in the center 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.

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 a little different than in the standard payoff matrix. This is because when the bomber (player I) 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 1/4 probability in the original game. We have a similar situation for 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:

		submarine
		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 bomb drops on one of the four mid-sides with probability 1/4 for each, and the submarine hides in one of the eight possible locations (pairs of adjacent squares) that exclude the center, choosing any given one with a probability of 1/8.

Mathematically, we can think of the symmetry argument as follows. Suppose that we have 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 $a_{i,j}$ satisfy

$$a_{\pi_1(i),\pi_2(j)} = a_{i,j}. \tag{2.4}$$

If this is so, then there are optimal strategies for player I that give equal weight to $\pi_1(i)$ and i for each i. Similarly, there exists a mixed strategy for player II that is optimal and assigns the same weight to the moves $\pi_2(j)$ and j for each j.

2.5 Resistor networks and troll games

In this section we will analyze a zero-sum game played on a road network connecting two cities, A and B. The analysis of this game is related to networks of resistors, where the roads correspond to resistors.

Recall that if two points are connected by a resistor with **resistance** R, and there is a voltage drop of V across the two points, then the current that flows through the resistor is V/R. The **conductance** is the reciprocal of the resistance. When the pair of points are connected by a pair of resistors with resistances R_1 and R_2 arranged in series (see the top of Figure 2.4), the **effective resistance** between the nodes is $R_1 + R_2$, because the current that flows through the resistors is $V/(R_1 + R_2)$. When the resistors are arranged in parallel (see the bottom of Figure 2.4), it is the conductances that add, i.e., the effective conductance between the nodes is $1/R_1 + 1/R_2$, i.e., the effective resistance is

$$\frac{1}{1/R_1 + 1/R_2} = \frac{R_1 R_2}{R_1 + R_2}.$$

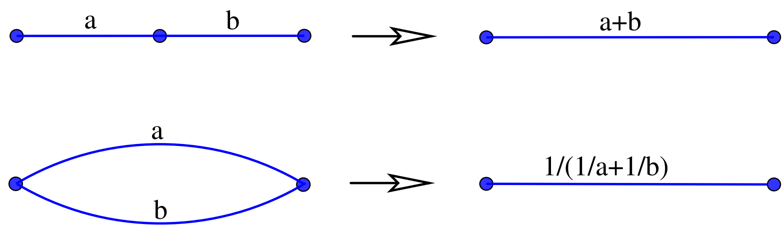

Fig. 2.4. In a network consisting of two resistors with resistances R_1 and R_2 in series (shown on top), the effective resistance is $R_1 + R_2$. When the resistors are in parallel, the effective conductance is $1/R_1 + 1/R_2$, so the effective resistance is $1/(1/R_1 + 1/R_2) = R_1 R_2/(R_1 + R_2)$.

These series and parallel rules for computing the effective resistance can be used in sequence to compute the effective resistance of more complicated networks, as illustrated in Figure 2.5. If the effective resistance between

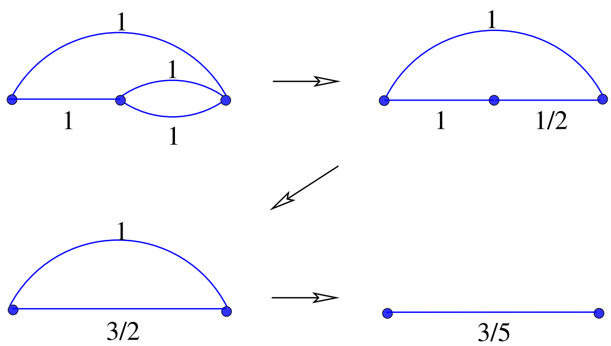

Fig. 2.5. A resistor network, with resistances all equaling to 1, has an effective resistance of 3/5. Here the parallel rule was used first, then the series rule, and then the parallel rule again.

two points can be computed by repeated application of the series rule and parallel rule, then the network is called a **series-parallel network**. Many networks are series-parallel, such as the one shown in Figure 2.6, but some networks are not series-parallel, such as the complete graph on four vertices.

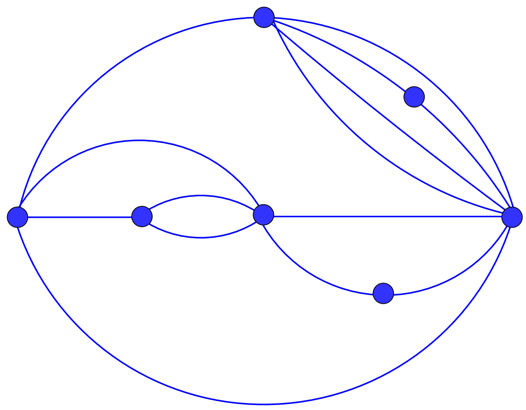

Fig. 2.6. A series-parallel graph, i.e., a graph for which the effective resistance can be computed by repeated application of the series and parallel rules.

For the troll game, we restrict our attention to series-parallel road networks. Given such a network, consider the following game:

Example 2.5.1 (**Troll and Traveler**)**.** A troll and a traveler will each choose a route along which to travel from city A to city B and then they will disclose their routes. Each road has an associated toll. In each case where the troll and the traveler have chosen the same road, the traveler pays the toll to the troll.

This is of course a zero-sum game. As we shall see, there is an elegant and general way to solve this type of a game on series-parallel networks. We may interpret the road network as an electrical circuit, and the tolls as resistances.

We claim that optimal strategies for both players are the same: Under an optimal strategy, a player planning his route, upon reaching a fork in the road, should move along any of the edges emanating from the fork with a probability proportional to the conductance of that edge.

To see why this strategy is optimal we will need some new terminology:

Definition 2.5.1. 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 the 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 as follows:

		player II	
		moves of G_1	moves of G_2
player I	moves of G_1	G_1	0
	moves of G_2	0	G_2

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

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 $v_1/(v_1+v_2)$. (This is also the optimal strategy if v_1 and v_2 are both negative, but if they have opposite signs, say $v_1 < 0 < v_2$, then player I should play in G_2 and II should play in G_1, resulting in a payoff of 0.) Assuming both v_1 and v_2 are positive, the expected payoff of the parallel sum-game is

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

which is the effective resistance of an electrical network with two edges arranged in parallel that have resistances v_1 and v_2. This explains the form of the optimal strategy in troll-traveler games on series-parallel networks.

The troll-and-traveler game could be played on a more general (not necessarily series-parallel) network with two distinguished points A and B. On general networks, we get a similarly elegant solution when we define the game in the following way: If the troll and the traveler traverse an edge in the opposite directions, then the troll pays the cost of the road to the traveler. Then the value of the game turns out to be the effective resistance between A and B.

2.6 Hide-and-seek games

Hide-and-seek games form another class of two-person zero-sum games that we will analyze.

Example 2.6.1 (**Hide-and-seek Game**)**.** The game is played on a matrix whose entries are 0's and 1's. Player II chooses a 1 somewhere in the matrix, and hides there. Player I chooses a row or a column and wins a payoff of 1 if the line that he picks contains the location chosen by player II.

To analyze this game, we will need Hall's marriage theorem, an important result that comes up in many places in graph theory.

Suppose that each member of a group B of boys is acquainted with some subset of a group G of girls. Under what circumstances can we find a pairing of boys to girls so that each boy is matched with a girl with whom he is acquainted?

Clearly, there is no hope of finding such a matching unless for each subset B' of the boys, the collection G' of all girls with whom the boys in B' are acquainted is at least as large as B'. What Hall's theorem says is that this condition is not only necessary but sufficient: As long as the above condition holds, it is always possible to find a matching.

Theorem 2.6.1 (**Hall's marriage theorem**)**.** *Suppose that B is a finite set of boys and G is a finite set of girls. Let $f(b)$ denote the set of girls with whom boy b is acquainted. For a subset $B' \subseteq B$ of the boys, let $f(B')$ denote the set of girls with whom some boy in B' is acquainted, i.e., $f(B') = \cup_{b \in B'} f(b)$. There is a matching between the boys and the girls such that each boy is paired with a girl with whom he is acquainted if and only if for each $B' \subseteq B$ we have $|f(B')| \geq |B'|$.*

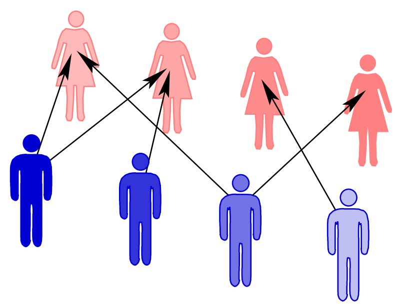

Fig. 2.7. Illustration of Hall's marriage theorem.

Proof. As we stated above, the condition is clearly necessary for there to be a matching. We will prove that the condition is also sufficient by using induction on the number of boys.

The base case when $|B| = 1$ (or even $|B| = 0$) is easy.

Suppose $|f(B')| > |B'|$ for each nonempty $B' \subsetneq B$. Then we can just match an arbitrary boy to any girl he knows. The set of remaining boys and the set of remaining girls still satisfy the condition in the statement of the theorem, so by the inductive hypothesis, we match them up. (Of course this approach does not work for the example in Figure 2.7: there are three sets of boys B' for which $|f(B')| \not> |B'|$, and indeed, if the third boy is paired with the first girl, there is no way to match the remaining boys and girls.)

Otherwise, there is some nonempty set $B' \subsetneq B$ satisfying $|f(B')| = |B'|$. (In the example in Figure 2.7, B' could be the first two boys, or the second boy, or the fourth boy.) Since $|B'| < |B|$, we can use the inductive hypothesis to match up the set of boys B' and the set of girls $f(B')$ with whom they are acquainted. Let A be a set of unmatched boys, i.e., $A \subseteq B \setminus B'$. Then $|f(A \cup B')| = |f(B')| + |f(A) \setminus f(B')|$ and $|f(A \cup B')| \geq |A \cup B'| = |A| + |B'| = |A| + |f(B')|$, so $|f(A) \setminus f(B')| \geq |A|$. Thus each set of unmatched boys is acquainted with at least as many unmatched girls. Since $|B \setminus B'| < |B|$, we can again use the inductive hypothesis to match up the remaining unmatched boys and girls. This completes the induction step. $\square$

Using Hall's theorem, we can prove another useful result. Given a matrix whose entries consist of 0's and 1's, two 1's are said to be **independent** if no row or column contains them both. A **cover** of the matrix is a collection of rows and columns whose union contains each of the 1's.

Lemma 2.6.1 (**König's lemma**). *Given an $n \times m$ matrix whose entries consist of 0's and 1's, the maximal size of a set of independent 1's is equal to the minimal size of a cover.*

Proof. Consider a maximal independent set of 1's (of size k), and a minimal cover consisting of ℓ lines. That $k \leq \ell$ is easy: each 1 in the independent set is covered by a line, and no two are covered by the same line.

For the other direction we make use of Hall's lemma. Suppose that among these ℓ lines, there are r rows and c columns. In applying Hall's lemma, the rows in the cover are the boys, and the columns *not* in the cover are the girls. A boy (row) knows a girl (column) if their intersection contains a 1.

Suppose that j boys (rows in the cover) collectively know s girls (columns not in the cover). We could replace these j rows by these s columns to obtain a new cover. If the cover is minimal, then it must be that $s \geq j$. By Hall's lemma, we can match up the r rows in the cover with r of the columns outside the cover so that each row knows its matched column.

Similarly, we match up the c columns in the cover with c of the rows outside the cover so that each column knows its matched row.

Each of the intersections of the above $\ell = r + c$ pairs of matched rows and columns contains a 1, and these 1's are independent, hence $k \geq \ell$. This completes the proof. □

We now use König's lemma to analyze Hide-and-seek. Recall that in Hide-and-seek, player II chooses a 1 somewhere in the matrix, and hides there, and player I chooses a row or a column and wins a payoff of 1 if the line that he picks contains the location chosen by player II. One strategy for player II is to pick a maximal independent set of 1's, and then hide in a uniformly chosen element of it. Let k be the size of the maximal set of independent 1's. No matter what row or column player I picks, it contains at most one 1 of the independent set, and player II hid there with probability $1/k$, so he is found with probability at most $1/k$. One strategy for player I consists of picking uniformly at random one of the lines of a minimal cover of the matrix. No matter where player II hides, at least one line from the cover will find him, so he is found with probability at least 1 over the size of the minimal cover. Thus König's lemma shows that this is, in fact, a joint optimal strategy, and that the value of the game is k^{-1}, where k is the size of the maximal set of independent 1's.

2.7 General hide-and-seek games

We now analyze a more general version of the game of hide-and-seek.

Example 2.7.1 (Generalized Hide-and-seek). A matrix of values $(b_{i,j})_{n\times n}$ is given. Player II chooses a location (i, j) at which to hide. Player I chooses a row or a column of the matrix. He wins a payment of $b_{i,j}$ if the line he has chosen contains the hiding place of his opponent. We assume that $b_{i,j} > 0$ for all i, j.

First, we propose a strategy for player II, later checking that it is optimal. Player II first chooses a fixed permutation π of the set $\{1, \ldots, n\}$ 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:

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

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

$$d_{i,j} = b_{i,j}^{-1}$$

and

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

because it is this choice that equalizes the expected payments. For the fixed strategy, player I may choose to select row i (for an expected payoff of $p_i b_{i,\pi(i)}$) or column j (for an expected payoff of $p_j b_{\pi^{-1}(j),j}$), so the expected payoff of the game is then

$$\max\left(\max_i p_i b_{i,\pi(i)}, \max_j p_{\pi^{-1}(j)} b_{\pi^{-1}(j),j}\right) = \max\left(\max_i \frac{1}{D_\pi}, \max_j \frac{1}{D_\pi}\right) = \frac{1}{D_\pi}.$$

Thus, if player II is going to use a strategy that consists of picking a permutation π^* and then doing as described, 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 the restricted class of those involving permutations in this way, but over all possible strategies.

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

Lemma 2.7.1. *Consider a minimal covering $(\mathbf{u}^*, \mathbf{w}^*)$ of $D = (d_{i,j})_{n\times n}$ (i.e., one for which $\sum_{i=1}^n \big(u_i + w_i\big)$ is minimal). Then*

$$\sum_{i=1}^{n} \big(u_i^* + w_i^*\big) = \max_\pi D_\pi. \tag{2.5}$$

Proof. Note that a minimal covering exists, because the continuous map

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

defined on the closed and bounded set

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

where $M = \max_{i,j} d_{i,j}$, does indeed attain its infimum.

Note also that we may assume that $\min_i u_i^* > 0$.

That the left-hand-side of (2.5) is at least the right-hand-side is straightforward. Indeed, for any π, we have that $u_i^* + w_{\pi_i}^* \geq d_{i,\pi_i}$. Summing over i, we obtain this inequality.

Showing the other inequality is harder, and requires Hall's marriage lemma, or something similar. We need a definition of "knowing" to use Hall's theorem. We say that row i knows column j if

$$u_i^* + w_j^* = d_{i,j}.$$

Let's check Hall's condition. Suppose that k rows $i_1, \ldots, i_k$ know between them only $\ell < k$ columns $j_1, \ldots, j_\ell$. Define $\tilde{\mathbf{u}}$ from u^* by reducing these rows by a small amount $\varepsilon > 0$. Leave the other rows unchanged. The condition that ε must satisfy is in fact that

$$0 < \varepsilon \leq \min_i u_i^*$$

and also

$$\varepsilon \leq \min \big\{ u_i + w_j - d_{i,j} : (i,j) \text{ such that } u_i + w_j > d_{i,j} \big\}.$$

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

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

We claim that $(\tilde{\mathbf{u}}, \tilde{\mathbf{w}})$ is a covering of the matrix. At places where the equality $d_{i,j} = u_i^* + w_j^*$ holds, we have that $d_{i,j} = \tilde{u}_i + \tilde{w}_j$, by construction. In places where $d_{i,j} < u_i^* + w_j^*$, then

$$\tilde{u}_i + \tilde{w}_j \geq u_i^* - \varepsilon + w_j^* > d_{i,j},$$

the latter inequality is by the assumption on the value of ε.

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 minimal.

We have checked that Hall's condition holds. Hall's theorem provides a matching of columns and rows. This is a permutation π^* such that, for each i, we have that

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

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 gives the other inequality required to prove the lemma. □

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

$$\frac{u_i^* + v_j^*}{D_{\pi^*}} b_{i,j} \geq \frac{d_{i,j} b_{i,j}}{D_{\pi^*}} = D_{\pi^*}^{-1}.$$

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

Example 2.7.2. Consider the Hide-and-seek game with payoff matrix B given by

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

This means that the matrix D is equal to

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

To determine a minimal 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 permutation $\pi = \mathrm{id}$).

An optimal strategy for the hider is to play $p(1,1) = 1/6$ and $p(2,2) = 5/6$. An optimal strategy for the seeker consists of playing $q(\mathrm{row}_1) = 1/6$, $q(\mathrm{row}_2) = 2/3$ and $q(\mathrm{col}_2) = 1/6$. The value of the game is 1/6.

2.8 The bomber and battleship game

Example 2.8.1 (Bomber and Battleship). In this family of games, a battleship is initially located at the origin in $\mathbb{Z}$. At any given time step in $\{0,1,\ldots\}$, 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 $j \in \{0,1,\ldots,n\}$. What is the value of the game?

The answer depends on n. The value of G_n can only increase with larger n, because the bomber has more choices for when to drop the bomb. For each n the value for the bomber is at least $1/3$, since the bomber could pick a uniformly random site in $\{-2, 0, 2\}$ to bomb, and no matter where the battleship goes, there is at least a $1/3$ chance that the bomb will hit it.

The value of G_0 is in fact $1/3$, because the battleship may play the following strategy to ensure that it has a $1/3$ probability of being at any of the sites -2, 0, or 2 at time 2: It moves left or right with equal probability at the first time step, and then turns with probability of $1/3$ or goes on in the same direction with probability $2/3$. No matter what the bomber does, there is only a $1/3$ chance that the battleship is where the bomb was dropped, so the value of G_0 is at most $1/3$ (and hence equal to $1/3$).

The battleship may also manoevre to ensure that the expected payoff for G_1 is also at most $1/3$. What it can do is follow its above strategy for G_0 for its first two moves, and then at time step 2, if it is at location 0 then it continues in the same direction, if it is at location 2 or -2 then it turns with probability $1/2$. If the bomber drops its bomb at time 0, then by our analysis of G_0, the battleship will be where the bomb lands with probability $1/3$. If the bomber drops its bomb at time 1, it sees the battleship's first move, and then drops the bomb. Suppose the battleship moved to 1 on its first move. It moves to 0 and then on to -1 with probability $1/3 \times 1$. It moves to 2 and then on to 3 with probability $2/3 \times 1/2$, or on to 1 with probability $2/3 \times 1/2$. It is at each location with probability no more than $1/3$, so the expected payoff for the bomber is no more than $1/3$ no matter what it does. Similarly, if the battleship's first move was to location -1, the expected payoff for the bomber is no more than $1/3$. Hence the value of G_1 is also $1/3$.

It is impossible for the battleship to pursue this strategy to obtain a value of $1/3$ for the game G_2. Indeed, $v(G_2) > 1/3$.

We now describe a strategy for the game that is due to the mathematician Rufus Isaacs. Isaacs' strategy is not optimal in any given game G_n, but it does have the merit of having the same limiting value, as $n \to \infty$, as optimal play. The strategy is quite simple: on the first move, go in either direction with probability $1/2$, and from then on, turn with probability of $1 - a$, and keep going with probability of a.

We now 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 $a = 2/(1+\sqrt{5})$.

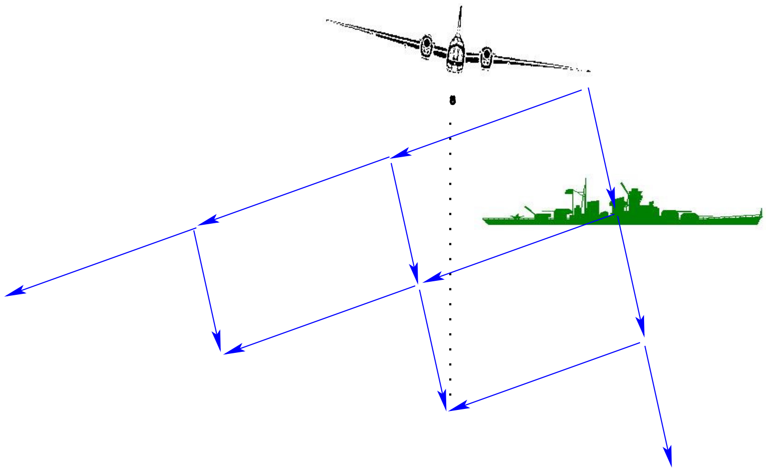

Fig. 2.8. The bomber drops its bomb where it hopes the battleship will be two time units later. The battleship does not see the bomb coming, and randomizes its path to avoid the bomb. (The length of each arrow is 2.)

The payoff for the bomber against this strategy is at most $1-a$. We have proved that the value $v(G_n)$ of the game G_n is at most $1-a$, for each n.

Consider the zero-sum game whose payoff matrix is given by:

$$\begin{array}{cc} & \text{battleship} \\ \text{bomber} & \begin{array}{|ccc|} \hline 1 & 0 & 8 \\ 2 & 3 & -1 \\ \hline \end{array} \end{array}$$

To solve this game, first, we search for saddle points — a value in the matrix that is maximal in its column and minimal in its row. None exist in this case. Nor are there any evident dominations of rows or columns.

Suppose then that player I plays the mixed strategy $(p, 1-p)$. If there is an optimal strategy for player II in which she plays each of her three pure strategies with positive probability, then

$$2 - p = 3 - 3p = 9p - 1.$$

No solution exists, so we consider now mixed strategies for player II in which one pure strategy is never played. If the third column has no weight, then $2 - p = 3 - 3p$ implies that $p = 1/2$. However, the entry 2 in the matrix becomes a saddle point in the 2×2 matrix formed by eliminating the third column, which is not consistent with $p = 1/2$.

Consider instead strategies supported on columns 1 and 3. The equality

$2 - p = 9p - 1$ yields $p = 3/10$, giving payoffs of

$$\left(\frac{17}{10}, \frac{21}{10}, \frac{17}{10}\right)$$

for the three strategies of player II.

If player II plays column 1 with probability q and column 3 otherwise, then player I sees the payoff vector $(8 - 7q, 3q - 1)$. These quantities are equal when $q = 9/10$, so that player I sees the payoff vector $(17/10, 17/10)$. Thus, the value of the game is $17/10$.

Exercises

2.1 Find the value of the following zero-sum game. Find some optimal strategies for each of the players.

$$\begin{array}{cc} & \text{player II} \\ \text{player I} & \begin{array}{|cccc|} \hline 8 & 3 & 4 & 1 \\ 4 & 7 & 1 & 6 \\ 0 & 3 & 8 & 5 \\ \hline \end{array} \end{array}$$

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

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

2.3 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.

2.4 Recall the bomber and battleship game from section 2.8. Set up the payoff matrix and find the value of the game G_2.

2.5 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.

2.6 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.

2.7 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.

2.8 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%.

2.9 **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.)

2.10 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.

2.11 **A recursive zero-sum game.** An inspector can inspect a facility on just one occasion, on one of the days $1, \ldots, n$. 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.

3

General-sum games

We now turn to discussing the theory of **general-sum games**. Such a game is given in **strategic form** 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 there still exists a generalization of the Von Neumann minimax, the so-called Nash equilibrium. These equilibria give the strategies that "rational" players could follow. However, there are often several Nash equilibria, and in choosing one of them, some degree of cooperation between the players may be optimal. 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 evidence held by the police is poor. If one confesses and the other is silent, then the confessor goes free, and the other prisoner is sentenced to ten years. If both confess, they will each spend eight years in prison. If both remain silent, the sentence is one year to each, for some minor crime that the police are able to prove. Writing the negative payoff as the number of years spent in prison, we obtain the following payoff matrix:

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

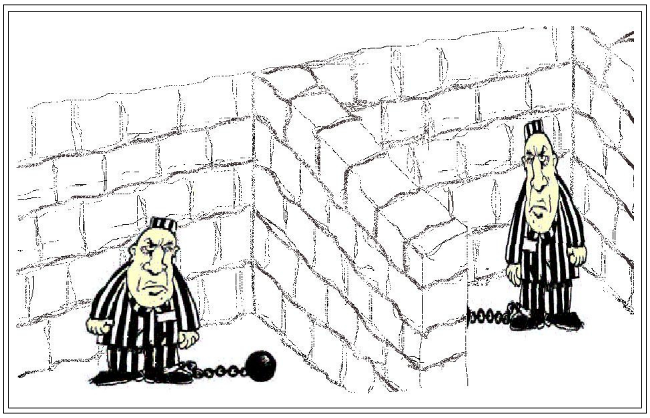

Fig. 3.1. Two prisoners considering whether to confess or remain silent.

The payoff matrices for players I and II are the 2×2 matrices given by the collection of first, or second, entries in each of the vectors in the above matrix.

If the players only play one round, then a domination argument shows that each should confess: the outcome he secures by confessing is preferable to the alternative of remaining silent, whatever the behavior of the other player. However, if they both follow this reasoning, the outcome is much worse for each player than the one achieved by both remaining silent. In a once-only game, the "globally" preferable outcome of each remaining silent could only occur were each player to suppress the desire to achieve the best outcome in selfish terms. In games with repeated play ending at a known time, the same applies, by an argument of backward induction. In games with repeated play ending at a random time, however, the globally preferable solution may arise even with selfish play.

Example 3.1.2 (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)

One might naturally come up with two modifications of Von Neumann's minimax theorem. The first one is that the players do not suppose any rationality about their partner, so they just want to assure a payoff assuming the worst-case scenario. Player I can guarantee a safety value of $\max_{\mathbf{x}\in\Delta_2}\min_{\mathbf{y}\in\Delta_2}\mathbf{x}^T A\mathbf{y}$, where A denotes the matrix of payoffs received by her. This gives the strategy $(1/5, 4/5)$ for her, with an assured expected payoff of $4/5$, regardless of what player II does. The analogous strategy for player II is $(4/5, 1/5)$, with the same assured expected payoff of $4/5$. Note that these values are lower than what each player would get from just agreeing to go where the other prefers.

The second possible adaptation of the minimax approach is that player I announces her probability p of going to the opera, expecting player II to maximize his payoff given this p. Then player I maximizes the result over p. However, in contrast to the case of zero-sum games, the possibility of announcing a strategy and committing to it in a general-sum game might actually raise the payoff for the announcer, and hence it becomes a question how a model can accommodate this possibility. In our game, each player could just announce their favorite choice, and to expect their spouse to behave "rationally" and agree with them. This leads to a disaster, unless one of them manages to make this announcement before the spouse does, and the spouse truly believes that this decision is impossible to change, and takes the effort to act rationally.

In this example, it is quite artificial to suppose that the two players cannot discuss, and that there are no repeated plays. Nevertheless, this example shows clearly that a minimax approach is not suitable anymore.

3.2 Nash equilibria

We now introduce a central notion for the study of general-sum games:

Let A, B be $m \times n$ payoff-matrices, giving the strategic form of a game.

Definition 3.2.1 (**Nash equilibrium**)**.** A pair of vectors $(\mathbf{x}^*, \mathbf{y}^*)$ with $\mathbf{x}^* \in \Delta_m$ and $\mathbf{y}^* \in \Delta_n$ 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 $i, j \in \{1, 2, \ldots, n\}$. A pair $(\mathbf{x}, \mathbf{y})$ of strategies is called **symmetric** if $x_i = y_i$ for all $i = 1, \ldots, n$.

We will see that there always exists a Nash equilibrium; however, there can be many of them. If $\mathbf{x}$ and $\mathbf{y}$ are unit vectors, with a 1 in some coordinate and 0 in all the others, then the equilibrium is called **pure**.

In the above example of the battle of the sexes, there are two pure equilibria: these are BB and OO. There is also a mixed equilibrium, $(4/5, 1/5)$ for player I and $(1/5, 4/5)$ for II, having the value $4/5$, which is very low.

Consider a simple model, where two cheetahs are giving chase to two antelopes. The cheetahs will catch any antelope they choose. If they choose the same one, they must share the spoils. Otherwise, the catch is unshared. There is a large antelope and a small one, that are worth ℓ and s to the cheetahs. Here is the matrix of payoffs:

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

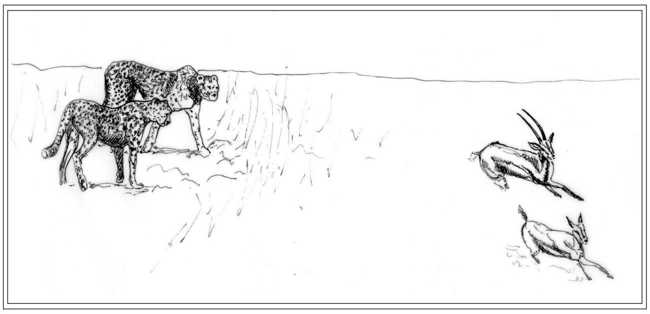

Fig. 3.2. Cheetahs deciding whether to chase the large or the small antelope.

If the larger antelope is worth at least twice as much as the smaller ($\ell \geq 2s$), for player I the first row dominates the second. Similarly for player II, the first column dominates the second. 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 p, then the expected payoff to the second cheetah by chasing the larger antelope is

$$\frac{\ell}{2}p + (1-p)\ell,$$

and the expected payoff arising from chasing the smaller antelope is

$$ps + (1-p)\frac{s}{2}.$$

These expected payoffs are equal when

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

For any other value of p, 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

$$\frac{2\ell - s}{\ell + s},$$

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

Symmetric mixed Nash equilibria are of particular interest. It has been experimentally verified that in some biological situations, systems approach such equilibria, presumably by mechanisms of natural selection. We explain briefly how this might work. First of all, it is natural to consider symmetric strategy pairs, because if the two players are drawn at random from the same large population, then the probabilities with which they follow a particular strategy are the same. Then, among symmetric strategy pairs, Nash equilibria play a special role. Consider the above mixed symmetric Nash equilibrium, in which $p_0 = (2\ell - s)/(\ell + s)$ is the probability of chasing the large antelope. Suppose that a population of cheetahs exhibits an overall probability $p > p_0$ for this behavior (having too many greedy cheetahs, or every single cheetah being slightly too greedy). Now, if a particular cheetah is presented with a competitor chosen randomly from this population, then chasing the small antelope has a higher expected payoff to this particular cheetah than chasing the large one. That is, the more modest a cheetah is, the larger advantage it has over the average cheetah. Similarly, if the cheetah population is too modest on the average, i.e., $p < p_0$, then the more ambitious cheetahs have an advantage over the average. Altogether, the population seems to be forced by evolution to chase antelopes according to the symmetric mixed Nash equilibrium. The related notion of an **evolutionarily stable strategy** is formalized in section 3.7.

Example 3.2.1 (**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 (they both win 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 lose some big value M).

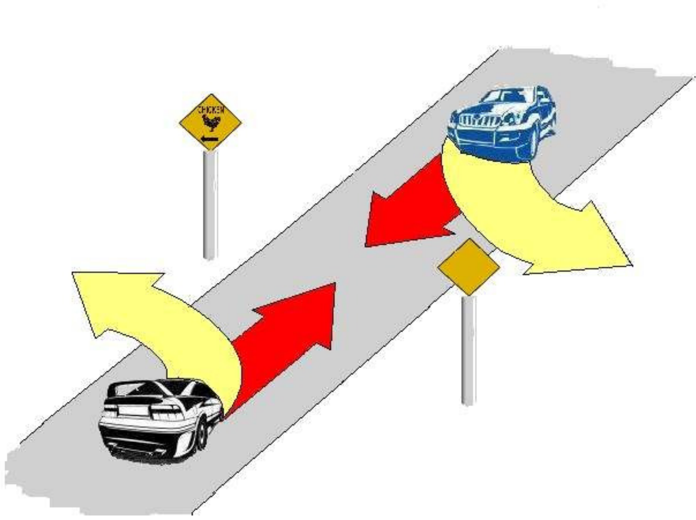

Fig. 3.3. The game of chicken.

We solve the game of chicken. Write C for the strategy of chickening out, D for driving forward. The pure equilibria are (C, D) and (D, C). To determine the mixed equilibria, suppose that player I plays C with probability p and D with probability $1 - p$. This presents player II with expected payoffs of $p \times 1 + (1 - p) \times (-1) = 2p - 1$ if she plays C, and $p \times 2 + (1 - p) \times (-M) = (M + 2)p - M$ if she plays D. We seek an equilibrium where player II has positive probability on each of C and D, and thus one for which

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

That is, $p = 1 - 1/M$. The payoff for player II is $2p - 1$, which equals $1 - 2/M$. Note that, as M increases to infinity, this symmetric mixed equilibrium gets concentrated on (C, C), and the expected payoff increases up to 1.

There is an apparent paradox here. We have a symmetric game with payoff matrices A and B that has a unique symmetric equilibrium with payoff γ. By replacing A and B by smaller matrices $\tilde{A}$ and $\tilde{B}$, we obtain a payoff $\tilde{\gamma} > \gamma$ from a unique symmetric equilibrium. This is impossible in zero-sum games.

However, if the decision of each player gets switched randomly with some small but fixed probability, then letting $M \to \infty$ does not yield total concentration on the strategy pair (C, C).

This is another game in which the possibility of a binding commitment increases the payoff. If one player rips out the steering wheel and throws it out of the car, then he makes it impossible to chicken out. If the other player sees this and believes her eyes, then she has no other choice but to chicken out.

In the battle of sexes and the game of chicken, making a binding commitment pushes the game into a pure Nash equilibrium, and the nature of that equilibrium strongly depends on who managed to commit first. In the game of chicken, the payoff for the one who did not make the commitment is lower than the payoff in the unique mixed Nash equilibrium, while it is higher in the battle of sexes.

Example 3.2.2 (No pure equilibrium). Here is an example where there is no pure Nash equilibrium, only a unique mixed one, and both commitment strategy pairs have the property that the player who did not make the commitment still gets the Nash equilibrium payoff.

player I \ player II	C	D
A	$(6, -10)$	$(0, 10)$
B	$(4, 1)$	$(1, 0)$

In this game, there is no pure Nash equilibrium (one of the players always prefers another strategy, in a cyclic fashion). For mixed strategies, if player I plays (A, B) with probabilities $(p, 1-p)$, and player II plays (C, D) with probabilities $(q, 1-q)$, then the expected payoffs are $1 + 3q - p + 3pq$ for I and $10p + q - 21pq$ for II. We easily get that the unique mixed equilibrium is $p = 1/21$ and $q = 1/3$, with payoffs 2 for I and $10/21$ for II. If player I can make a commitment, then by choosing $p = 1/21 - \varepsilon$ for some small $\varepsilon > 0$, he will make II choose $q = 1$, and the payoffs will be $4 + 2/21 - 2\varepsilon$ for I and $10/21 + 11\varepsilon$ for II. If II can make a commitment, then by choosing $q = 1/3 + \varepsilon$, she will make I choose $p = 1$, and the payoffs will be $2 + 6\varepsilon$ for I and $10/3 - 11\varepsilon$ for II.

An amusing real-life example of binding commitments comes from 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 Correlated equilibria

Recall the "battle of the sexes":

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

Here, there are two pure Nash equilibria: both go to the opera or both watch baseball. What would be a good way to decide between them?

One way to do this would be to pick a joint action based on a flip of a single coin. For example, if a coin lands head then both go to the opera, otherwise both watch baseball. This is different from mixed strategies where each player independently randomized over individual strategies. In contrast, here a single coin-flip determines the strategies for both.

This idea was introduced in 1974 by Aumann ([?]) and is now called a **correlated equilibrium**. It generalizes Nash equilibrium and can be, surprisingly, easier to find in large games.

Definition 3.3.1 (**Correlated Equilbrium**)**.** A joint distribution on strategies for all players is called a **correlated equilibrium** if no player gains by deviating unilaterally from it. More formally, in a two-player general-sum game with $m \times n$ payoff matrices A and B, a correlated equilibrium is given by an $m \times n$ matrix $\mathbf{z}$. This matrix represents a joint density and has the following properties:

$$z_{i,j} \geq 0, \quad \text{for all } 1 \leq i \leq m, 1 \leq j \leq n$$

and

$$\sum_{i=1}^{m} \sum_{j=1}^{n} z_{i,j} = 1.$$

We say that no player benefits from unilaterally deviating provided:

$$(\mathbf{z})_i A\mathbf{z} \geq \mathbf{x}^T A\mathbf{z}$$

for all $i \in \{1, \ldots, m\}$ and all $\mathbf{x} \in \Delta_m$; while

$$\mathbf{z}B(\mathbf{z})^j \geq \mathbf{z}B\mathbf{y}$$

for all $j \in \{1, \ldots, n\}$ and all $\mathbf{y} \in \Delta_n$.

Observe that Nash equilibrium provides a correlated equilibrium where the joint distribution is the product of the two independent individual distributions. In the example of the battle of the sexes, where Nash equilibrium is of the form $(4/5, 1/5)$ for player I and $(1/5, 4/5)$ for player II, when players follow a Nash equilibrium they are, in effect, flipping a biased coin with probability of heads 4/5 and tails 1/5 twice — if head-tail, both go to the opera; tail-head, both watch baseball, etc. The joint density matrix looks like:

		husband: opera	husband: baseball
wife	opera	4/25	16/25
wife	baseball	1/25	4/25

Let's now go back to the Game of Chicken.

		player II: C	player II: D
player I	C	$(1, 1)$	$(-1, 2)$
player I	D	$(2, -1)$	$(-100, -100)$

There is no dominant strategy here and the pure equilibria are (C, D) and (D, C) with the payoffs of $(-1, 2)$ and $(2, -1)$ respectively. There is a symmetric mixed Nash equilibrium which puts probability $p = 1 - \frac{1}{100}$ on C and $1 - p = \frac{1}{100}$ on D, giving the expected payoff of $\frac{98}{100}$.

If one of the players could commit to D, say by ripping out the steering wheel, then the other would do better to swerve and the payoffs are: 2 to the one that committed first and 1 to the other one.

Another option would be to enter a binding agreement. They could, for instance, use a correlated equilibrium and flip a coin between (C, D) and (D, C). Then the expected payoff is 1.5. This is the average between the payoff to the one that commits first and the other player. It is higher than the expected payoff to a mixed strategy.

Finally, they could select a mediator and let her suggest a strategy to each. Suppose that a mediator chooses $(C, D), (D, C), (C, C)$ with probability $\frac{1}{3}$ each. Next the mediator 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 following the mediator's suggestion is a correlated equilibrium. Notice that the strategies are dependent, so this is not a Nash equilibrium.

Suppose mediator tells player I to play D, in that case he knows that player II was told to swerve and player I does best by complying to collect the payoff of 2. He has no incentive to deviate.

On the other hand, if the mediator tells him to play C, he is uncertain about what player II is told, so (C, C) and (C, D) are equally likely. We have expected payoff to following the suggestion of $\frac{1}{2} - \frac{1}{2} = 0$, while the expected payoff from switching is $2 \times \frac{1}{2} - 100 \times \frac{1}{2} = -49$, so the player is better off following the suggestion.

Overall the expected payoff to player I when both follow the suggestion is $-1 \times \frac{1}{3} + 2 \times \frac{1}{3} + 1 \times \frac{1}{3} = \frac{2}{3}$. This is better than they could do by following an uncorrelated Nash equilibrium.

Surprisingly, finding a correlated equilibrium in large scale problems is actually easier than finding a Nash equilibrium. The problem reduces to linear programming.

In the absence of a mediator, the players could follow some external signal, like the weather.

3.4 General-sum games with more than two players

It does not make sense to talk about zero-sum games when there are more than two players. The notion of a Nash equilibrium for general-sum games, however, can be used in this context. We now describe formally the set-up of a game with $k \geq 2$ players. Each player i has a set S_i of pure strategies. We are given functions $F_j : S_1 \times S_2 \times \cdots \times S_k \to \mathbb{R}$, for $j \in \{1, \ldots, k\}$. If, for each $i \in \{1, \ldots, k\}$, player i uses strategy $\ell_i \in S_i$, then player j has a payoff of $F_j(\ell_1, \ldots, \ell_k)$.

Example 3.4.1 (**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 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)

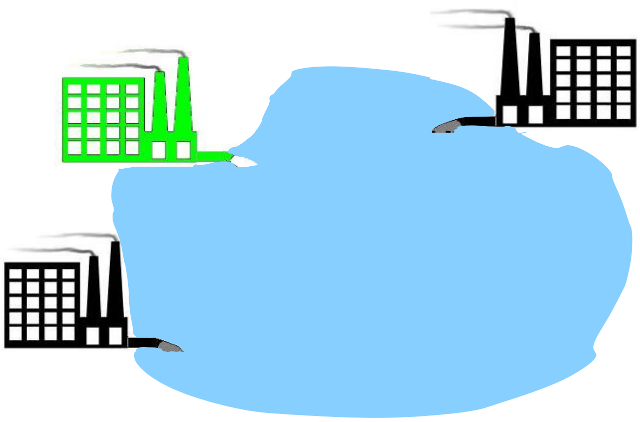

Fig. 3.4.

To discuss the game, we firstly introduce the notion of Nash equilibrium in the context of games with several players:

Definition 3.4.1. A **pure Nash equilibrium** in a k-person game is a set of pure strategies for each of the players,

$$(\ell_1^*, \ldots, \ell_k^*) \in S_1 \times \cdots \times S_k$$

such that, for each $j \in \{1, \ldots, k\}$ and $\ell_j \in S_j$,

$$F_j(\ell_1^*, \ldots, \ell_{j-1}^*, \ell_j, \ell_{j+1}^*, \ldots, \ell_k^*) \leq F_j(\ell_1^*, \ldots, \ell_{j-1}^*, \ell_j^*, \ell_{j+1}^*, \ldots, \ell_k^*).$$

More generally, a **mixed Nash equilibrium** is a collection of k probability vectors $\tilde{\mathbf{x}}^i$, each of length $|S_i|$, such that

$$F_j(\tilde{\mathbf{x}}^1, \ldots, \tilde{\mathbf{x}}^{j-1}, x, \tilde{\mathbf{x}}^{j+1}, \ldots, \tilde{\mathbf{x}}^k) \leq F_j(\tilde{\mathbf{x}}^1, \ldots, \tilde{\mathbf{x}}^{j-1}, \tilde{\mathbf{x}}^j, \tilde{\mathbf{x}}^{j+1}, \ldots, \tilde{\mathbf{x}}^k),$$

for each $j \in \{1, \ldots, k\}$ and each probability vector $\mathbf{x}$ of length $|S_j|$. Here

$$F_j(\mathbf{x}^1, \mathbf{x}^2, \ldots, \mathbf{x}^k) := \sum_{\ell_1 \in S_1, \ldots, \ell_k \in S_k} \mathbf{x}^1(\ell_1) \ldots \mathbf{x}^k(\ell_k) F_j(\ell_1, \ldots, \ell_k).$$

Definition 3.4.2. A game is **symmetric** if, for every $i_0, j_0 \in \{1, \ldots, k\}$, there is a permutation π of the set $\{1, \ldots, k\}$ such that $\pi(i_0) = j_0$ and

$$F_{\pi(i)}(\ell_{\pi(1)}, \ldots, \ell_{\pi(k)}) = F_i(\ell_1, \ldots, \ell_k).$$

For this definition to make sense, we are in fact requiring that the strategy sets of the players coincide.

We will prove the following result:

Theorem 3.4.1 (Nash's theorem). *Every game has a Nash equilibrium.*

Note that the equilibrium may be mixed.

Corollary 3.4.1. *In a symmetric game, there is a symmetric Nash equilibrium.*

Returning to the ecology game, note that the pure equilibria consist of all three firms polluting, or one of the three firms polluting, and the remaining two purifying. We now seek mixed equilibria. Let p_1, p_2, p_3 be the probability that firm I, II, III purifies, respectively. If firm III purifies, then its expected cost is $p_1 p_2 + p_1(1 - p_2) + p_2(1 - p_1) + 4(1 - p_1)(1 - p_2)$. If it pollutes, then the cost is $3p_1(1 - p_2) + 3p_2(1 - p_1) + 3(1 - p_1)(1 - p_2)$. If we want an equilibrium with $0 < p_3 < 1$, then these two expected values must coincide, which gives $1 = 3(p_1 + p_2 - 2p_1 p_2)$. Similarly, assuming $0 < p_2 < 1$ we get $1 = 3(p_1 + p_3 - 2p_1 p_3)$, and assuming $0 < p_1 < 1$ we get $1 = 3(p_2 + p_3 - 2p_2 p_3)$. Subtracting the second equation from the first one we get $0 = 3(p_2 - p_3)(1 - 2p_1)$. If $p_2 = p_3$, then the third equation becomes quadratic in p_2, with two solutions, $p_2 = p_3 = (3 \pm \sqrt{3})/6$, both in $(0, 1)$. Substituting these solutions into the first equation, both yield $p_1 = p_2 = p_3$, so there are 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.

3.5 The proof of Nash's theorem

Recall Nash's theorem:

Theorem 3.5.1. *For any general-sum game with $k \geq 2$ players, there exists at least one Nash equilibrium.*

To prove this theorem, we will use:

Theorem 3.5.2 (Brouwer's fixed-point theorem). *If $K \subseteq \mathbb{R}^d$ is closed, convex and bounded, and $T : K \to K$ is continuous, then there exists $\mathbf{x} \in K$ such that $T(\mathbf{x}) = \mathbf{x}$.*

Remark. We will prove this fixed-point theorem in section 3.6.3, but observe now that the proof is easy in case the dimension $d = 1$, and K is a closed interval $[a, b]$. Defining $f(\mathbf{x}) = T(\mathbf{x}) - \mathbf{x}$, note that $[a, b] \ni T(a) \geq a$ implies that $f(a) \geq 0$, while $[a, b] \ni T(b) \leq b$ implies that $f(b) \leq 0$. The intermediate value theorem assures the existence of $\mathbf{x} \in [a, b]$ for which $f(\mathbf{x}) = 0$, so $T(\mathbf{x}) = \mathbf{x}$. Note also that each of the hypotheses on K in the theorem is required, as the following examples show:

(i) $K = \mathbb{R}$ (closed, convex, not bounded) with $T(\mathbf{x}) = \mathbf{x} + 1$
(ii) $K = (0, 1)$ (bounded, convex, not closed) with $T(\mathbf{x}) = \mathbf{x}/2$
(iii) $K = \{\mathbf{z} \in \mathbb{C} : |\mathbf{z}| \in [1, 2]\}$ (bounded, closed, not convex) with $T(\mathbf{z}) = -\mathbf{z}$.

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. Put $K = \Delta_m \times \Delta_n$ and we will define a map $T : K \to K$ from a pair of strategies for the two players to another such pair. Note firstly that K is convex, closed and bounded. Define, for $\mathbf{x} \in \Delta_m$ and $\mathbf{y} \in \Delta_n$,

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

where $A_{(i)}$ denotes the i^{th} row of the matrix A. That is, c_i equals the gain for player I obtained by switching from strategy $\mathbf{x}$ to pure strategy i, if this gain is positive: otherwise, it is zero. Similarly, we define

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

where $B^{(j)}$ denotes the j^{th} column of B. The quantities d_j have the same interpretation for player II as the c_i do for player I. We now define the map T; it is given by $T(\mathbf{x}, \mathbf{y}) = \left(\hat{\mathbf{x}}, \hat{\mathbf{y}}\right)$, where

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

for $i \in \{1, \ldots, m\}$, and

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

for $j \in \{1, \dots, n\}$. The map $T : K \to K$ since

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

and $\hat{x}_i \geq 0$ for all $i \in \{1, \dots, m\}$, and similarly for $\hat{\mathbf{y}}$. Note that T is continuous, because c_i and d_j are. Applying Brouwer's theorem, we find that there exists $(\mathbf{x}, \mathbf{y}) \in K$ for which $(\mathbf{x}, \mathbf{y}) = (\hat{\mathbf{x}}, \hat{\mathbf{y}})$. We now claim that, for this choice of $\mathbf{x}$ and $\mathbf{y}$, each $c_i = 0$ for $i \in \{1, \dots, m\}$, and $d_j = 0$ for $j \in \{1, \dots, n\}$. To see this, suppose, for example, that $c_1 > 0$. There must exist $\ell \in \{1, \dots, m\}$ for which $\mathbf{x}_\ell > 0$ and $\mathbf{x}^T A\mathbf{y} \geq A_{(\ell)}\mathbf{y}$. (Otherwise

$$\mathbf{x}^T A\mathbf{y} = \sum_{i=1}^{m} x_i A_{(i)}\mathbf{y} = \sum_{\{\ell : x_\ell > 0\}} x_\ell A_{(\ell)}\mathbf{y} > (\sum_{\{\ell : x_\ell > 0\}} x_{ell})\mathbf{x}^T A\mathbf{y} = \mathbf{x}^T A\mathbf{y},$$

which is a contradiction.) For this ℓ, we have that $c_\ell = 0$, by definition. This implies that

$$\hat{x}_\ell = \frac{x_\ell}{1 + \sum_{k=1}^{m} c_k} < x_\ell,$$

because $c_1 > 0$. That is, the assumption that $c_1 > 0$ has given us a contradiction.

We may repeat this argument for each $i \in \{1, \dots, m\}$, thereby proving that each $c_i = 0$. Similarly, each $d_j = 0$. We deduce that $\mathbf{x}^T A\mathbf{y} \geq A_{(i)}\mathbf{y}$ for all $i \in \{1, \dots, m\}$. This implies that

$$\mathbf{x}^T A\mathbf{y} \geq \mathbf{x}'^T A\mathbf{y}$$

for all $\mathbf{x}' \in \Delta_m$. Similarly,

$$\mathbf{x}^T B\mathbf{y} \geq \mathbf{x}^T B\mathbf{y}'$$

for all $\mathbf{y}' \in \Delta_n$. Thus, $(\mathbf{x}, \mathbf{y})$ is a Nash equilibrium.

For $k > 2$ players, we still can consider the functions

$$c_i^{(j)}(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(k)}) \text{ for } i, j = 1, \dots, k,$$

where $\mathbf{x}^{(j)} \in \Delta_{n(j)}$ is a mixed strategy for player j, and $c_i^{(j)}$ is the gain for player j obtained by switching from strategy $\mathbf{x}^{(j)}$ to pure strategy i, if this gain is positive. The simple notation for $c_i^{(j)}$ is lost, but the proof carries over. $\square$

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

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

The image of D under T is again in D, because, in a symmetric game, $c_i^{(1)}(\mathbf{x}, \ldots, \mathbf{x}) = \cdots = c_i^{(k)}(\mathbf{x}, \ldots, \mathbf{x})$ for all $i = 1, \ldots, k$ 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.

3.6 Fixed-point theorems*

We now discuss various fixed-point theorems, beginning with a few easier ones.

3.6.1 *Easier fixed-point theorems*

Theorem 3.6.1 (**Banach's fixed-point theorem**). *Let K be a complete metric space. Suppose that $T : K \to K$ 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 in K.*

Remark. Recall that a metric space is **complete** if each Cauchy sequence therein converges to a point in the space. Consider, for example, any metric space that is a subset of $\mathbb{R}^n$ together with the metric d which is the Euclidean distance:

$$d(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\| = \sqrt{(x_1 - y_1)^2 + \cdots + (x_n - y_n)^2}.$$

See [**?**] for a discussion of general metric spaces.

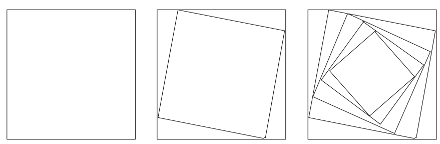

Fig. 3.5. Under the transformation T a square is mapped to a smaller square, rotated with respect to the original. When iterated repeatedly, the map produces a sequence of nested squares. If we were to continue this process indefinitely, a single point (fixed by T) would emerge.

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,

$$d(\mathbf{x}_n, \mathbf{x}_k) \leq d(\mathbf{x}_n, \mathbf{x}_{n+1}) + \cdots + d(\mathbf{x}_{k-1}, \mathbf{x}_k) \leq a\big(\lambda^n + \cdots + \lambda^{k-1}\big) \leq \frac{a\lambda^n}{1-\lambda}.$$

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

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

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

Example 3.6.1 (A map that decreases distances but has no fixed points). Consider the map $T : \mathbb{R} \to \mathbb{R}$ 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

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

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 [**?**].

Theorem 3.6.2 (Compact fixed-point theorem). *If X is a compact metric space and $T : X \to X$ satisfies $d(T(\mathbf{x}), T(\mathbf{y})) < d(\mathbf{x}, \mathbf{y})$ for all $\mathbf{x} \neq \mathbf{y} \in X$, then T has a fixed point.*

Proof. Let $f : X \to \mathbb{R}$ 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 X is compact, there exists $\mathbf{x}_0 \in X$ such that

$$f(\mathbf{x}_0) = \min_{\mathbf{x} \in X} f(\mathbf{x}). \tag{3.1}$$

If $T\mathbf{x}_0 \neq \mathbf{x}_0$, then $f(T(\mathbf{x}_0)) = d(T\mathbf{x}_0, T^2\mathbf{x}_0) < d(\mathbf{x}_0, T\mathbf{x}_0) = f(\mathbf{x}_0)$, and we have a contradiction to the minimizing property (3.1) of $\mathbf{x}_0$. This implies that $T\mathbf{x}_0 = \mathbf{x}_0$. $\square$

3.6.2 Sperner's lemma

We now state and prove a tool to be used in the proof of Brouwer's fixed-point theorem.

Lemma 3.6.1 (Sperner)**.** *In $d = 1$: Suppose that the unit interval is subdivided $0 = t_0 < t_1 < \cdots < t_n = 1$, with each t_i being marked zero or one. If t_0 is marked zero and t_n is marked one, then the number of adjacent pairs (t_j, t_{j+1}) with different markings is odd.*

In $d = 2$: Subdivide a triangle into smaller triangles in such a way that a vertex of any of the small triangles may not lie in the interior of an edge of another. Assume that the division consists of at least one step. Label the vertices of the small triangles 0, 1 or 2: the three vertices of the big triangle must be labelled 0, 1, and 2; vertices of the small triangles that lie on an edge of the big triangle must receive the label of one of the endpoints of that edge. Then the number of small triangles with three differently labelled vertices is odd; in particular, it is non-zero.

Remark. Sperner's lemma holds in any dimension. In the general case d, we replace the triangle by a d-simplex, use $d + 1$ labels, with analogous restrictions on the labels used.

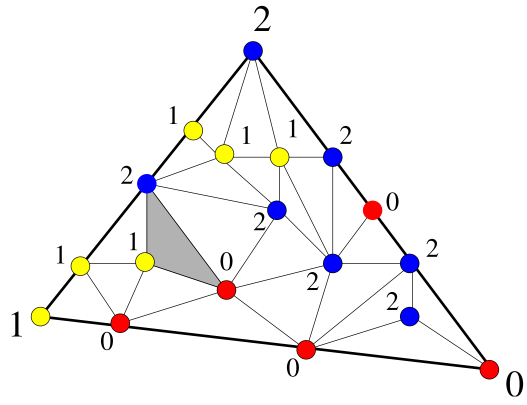

Fig. 3.6. Sperner's lemma when $d = 2$.

Proof. For $d = 1$, this is obvious (and can be proven by induction on n). For $d = 2$, we will count in two ways the set Q of pairs consisting of a small triangle and an edge on that triangle. Let A_{12} denote the number of 12-type edges of small triangles that lie in the boundary of the big triangle. Let B_{12} be the number of such edges in the interior. Let N_{abc} denote the number of small triangles where the three labels are a, b and c. Note that

$$N_{012} + 2N_{112} + 2N_{122} = A_{12} + 2B_{12},$$

because each side of this equation is equal to the number of pairs of triangle and edge, where the edge is of type (12). From the case $d = 1$ of the lemma, we know that A_{12} is odd, and hence N_{012} is odd, too. (In general, we may induct on the dimension, and use the inductive hypothesis to find that this quantity is odd.) $\square$

Corollary 3.6.1 (No-Retraction Theorem)**.** *Let $K \subseteq \mathbb{R}^d$ be compact and convex, and with non-empty interior. There is no continuous map $F : K \to \partial K$ whose restriction to ∂K is the identity.*

Case $d = 2$. First, we show that it suffices to take $K = \Delta$, where Δ is an equilateral triangle. Otherwise, because K has a non-empty interior, we may locate $\mathbf{x} \in K$ such that there exists a small triangle centered at $\mathbf{x}$ and contained in K. We call this triangle Δ for convenience. Construct a map $H : K \to \Delta$ as follows: For each $y \in \partial K$, define $H(y)$ to be equal to the element of $\partial\Delta$ that the line segment from $\mathbf{x}$ through y intersects. Setting $H(\mathbf{x}) = \mathbf{x}$, define $H(z)$ for other $z \in K$ by a linear interpolation of the values $H(\mathbf{x})$ and $H(q)$, where q is the element of ∂K lying on the line segment from

$\mathbf{x}$ through z. Note that ∂K is not empty since K is not empty and does not equal $\mathbb{R}^d$ since bounded.

Note that, if $F : K \to \partial K$ is a retraction from K to ∂K, then $H \circ F \circ H^{-1} : \Delta \to \partial\Delta$ is a retraction of Δ. This is the reduction we claimed.

Now suppose that $F_\Delta : \Delta \to \partial\Delta$ is a retraction of the equilateral triangle with side length 1. Since $F = F_\Delta$ is continuous on the compact Δ, it is uniformly continuous, in particular there exists $\delta > 0$ such that for all $\mathbf{x}, \mathbf{y} \in \Delta$ satisfying $\|\mathbf{x} - \mathbf{y}\| < \delta$ we have $\|F(\mathbf{x}) - F(\mathbf{y})\| < \frac{\sqrt{3}}{4}$. We can assume that $\delta < 1$.

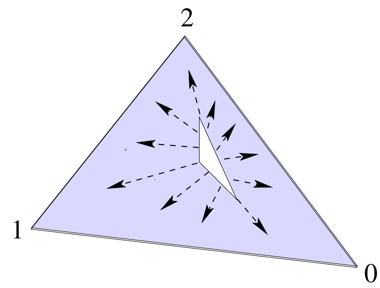

Fig. 3.7. Candidate for a retraction.

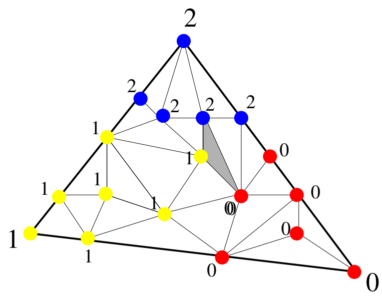

Fig. 3.8. A triangle with multicolored vertices indicates a discontinuity.

Label the three vertices of Δ by $0, 1, 2$. Triangulate Δ into triangles of side length less than δ. In this subdivision, label any vertex $\mathbf{x}$ according to the label of the vertex of Δ nearest to $F(\mathbf{x})$, with an arbitrary choice being made to break ties.

By Sperner's lemma, there exists a small triangle whose vertices are labelled $0, 1, 2$. The condition that $\|F(\mathbf{x}) - F(\mathbf{y})\| < \frac{\sqrt{3}}{4}$ implies that any pair of these vertices must be mapped under F to interior points of one of the sides of Δ, with a different side of Δ for each pair. This is impossible, implying that no retraction of Δ exists. $\square$

Remark. We should note, that the Brouwer's fixed-point theorem fails if the convexity assumption is completely omitted. This is also true for the above corollary. However, the main property of K that we used was not convexity; it is enough if there is a homeomorphism (a one-to-one continuous map with continuous inverse) between K and Δ.

3.6.3 Brouwer's fixed-point theorem

First proof of Brouwer's fixed-point theorem. Recall that we are given a continuous map $T : K \to K$, with K a closed, convex and bounded set. If K is contained in an affine hyperplane of $\mathbb{R}^d$ then, by the induction assumption, T must have a fixed point. Hence, by Lemma 3.6.2 below, we can assume that the interior of K is not empty. Suppose that T has no fixed points. Then we can define a continuous map $F : K \to \partial K$ as follows. For each $\mathbf{x} \in K$, we draw a ray from $T(\mathbf{x})$ through $\mathbf{x}$ until it meets ∂K. We set $F(\mathbf{x})$ equal to this point of intersection. If $T(\mathbf{x}) \in \partial K$, we set $F(\mathbf{x})$ equal that intersection point of the ray with ∂K which is not equal to $T(\mathbf{x})$. In the case of the domain $K = \{(x_1, x_2) \in \mathbb{R}^2 : x_1^2 + x_2^2 \leq 1\}$, for instance, the map F could have been written explicitly in terms of T:

$$F(\mathbf{x}) = \frac{T(\mathbf{x}) - \mathbf{x}}{\|T(\mathbf{x}) - \mathbf{x}\|}.$$

With some checking, it follows that $F : K \to \partial K$ is continuous. Thus, F is a retraction of K – but this contradicts the No-Retraction Theorem 3.6.1, so T must have a fixed point. $\square$

Lemma 3.6.2. *Let $K \subset \mathbb{R}^d$ be compact and convex. Then either K has an interior point or K is contained in an affine hyperplane of $\mathbb{R}^d$.*

Proof. Without loss of generality, $\mathbf{0} \in K$. If K contains d linearly independent vectors $v_1, \ldots, v_d \in \mathbb{R}^d$ then the convex set K contains the simplex $\mathrm{conv}\{\mathbf{0}, v_1, \ldots, v_d\}$ which equals $A\mathrm{conv}\{\mathbf{0}, e_1, \ldots, e_d\}$ for some matrix A ($A = (v_1, \ldots, v_d)$). Here $\{e_1, \ldots, e_d\}$ denotes the standard basis of $\mathbb{R}^d$. Note that

$$\mathrm{conv}\{\mathbf{0}, e_1, \ldots, e_d\} = \{(x_1, \ldots, x_d) : x_i \geq 0, \sum_{i=1}^{d} x_i \leq 1\},$$

of which $(\frac{1}{d+1}, \ldots, \frac{1}{d+1})$ is an interior point. Otherwise, there is a maximal independent set $v_1, \ldots, v_\ell$, with $\ell < d$, in K such that $K \subset \{v_1, \ldots, v_d\}$. $\square$

3.6.4 Brouwer's fixed-point theorem 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.

This lattice is generated by two vectors $u, v \in \mathbb{R}^2$ as shown in the left

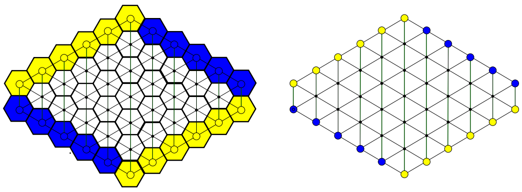

Fig. 3.9. Hexagonal lattice and its dual triangular lattice.

of Figure 3.10. 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 $\|x - y\| = 1$.

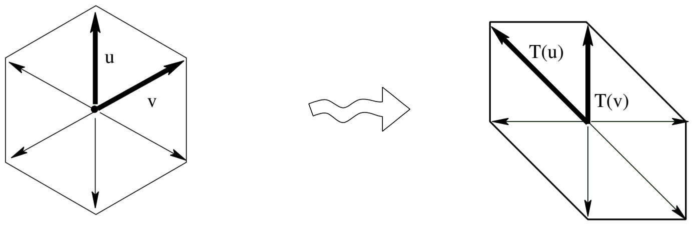

Fig. 3.10. Action of G on the generators of the lattice.

We can obtain a more convenient representation of this lattice by applying a linear transformation G defined by:

$$G(u) = \left(-\frac{\sqrt{2}}{2}, \frac{\sqrt{2}}{2}\right); \quad G(v) = (0, 1).$$

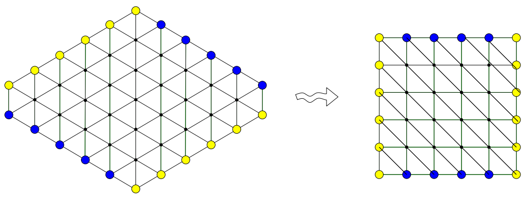

Fig. 3.11. Under G an equilateral triangular lattice is transformed to an equivalent lattice.

The game of Hex can be thought of as a game on the corresponding graph (see Fig. 3.11). 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 1.2.1, 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 1.2.3) 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 $T : [0,1]^2 \longrightarrow [0,1]^2$. 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. 3.12. 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,

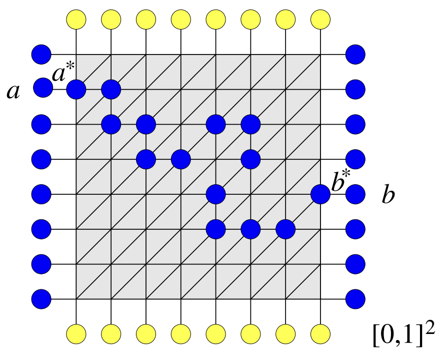

Fig. 3.12.

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. $\square$

3.7 Evolutionary game theory

We begin by introducing a new variant of our old game of Chicken:

3.7.1 Hawks and Doves

This game is a simple model for two behaviors — one bellicose, the other pacifistic — in the population of a single species (not the interactions between a predator and its prey).

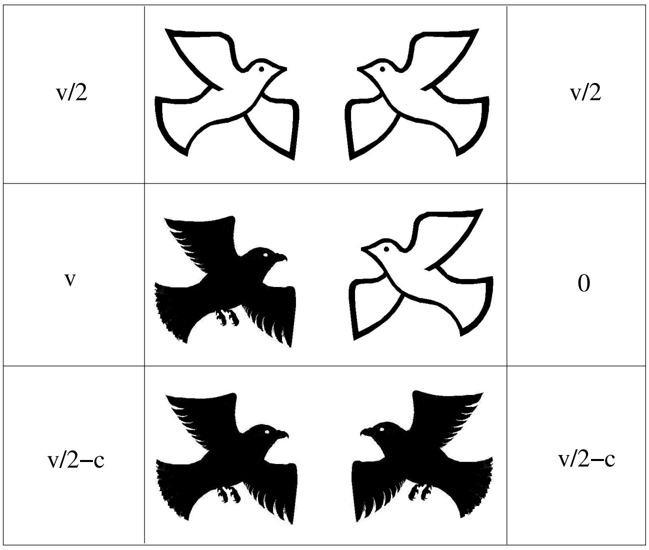

Fig. 3.13. Two players play this game, for a prize of value $v > 0$. They confront each other, and each chooses (simultaneously) to fight or to flee; these two strategies are called the "hawk" and the "dove" strategies, respectively. If they both choose to fight (two hawks), then each pays a cost c to 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 in Figure 3.13 has the payoff matrix

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

Now imagine a large population, each of whose members are hardwired genetically either as hawks or as doves, and assume that those who do better at this game have more offspring. It will turn out that the Nash equilibrium is also an equilibrium for the population, in the sense that a population composition of hawks and doves in the proportions specified by the Nash equilibrium (it is a symmetric game, so these are the same for both players) is locally stable — small changes in composition will return it to the equilibrium.

Next, we investigate the Nash equilibria. There are two cases, depending on the relative values of c and v.

If $c < \frac{v}{2}$, then simply by comparing rows, it is clear that player I always prefers to play H (hawk), no matter what player II does. By comparing columns, the same is true for player II. This implies that (H, H) is a pure Nash equilibrium. Are there mixed equilibria? Suppose I plays the mixed strategy $\{H : p, D : (1-p)\}$. Then II's payoff if playing H is $p(v/2-c) + (1-p)v$, and if playing D is $(1-p)v/2$. Since $c < \frac{v}{2}$, the payoff for H is always greater, and by symmetry, there are no mixed equilibria.

Note that in this case, Hawks and Doves is a version of Prisoner's Dilemma. If both players were to play D, they'd do better than at the Nash equilibrium — but without binding commitments, they can't get there. Suppose that instead of playing one game of Prisoner's Dilemma, they are to play many. If they are to play a fixed, known, number of games, the situation does not change. (*proof:* The last game is equivalent to playing one game only, so for this game both players play H. Since both know what will happen on the last game, the second-to-last game is also equivalent to playing one game only, so both play H here as well... and so forth, by "backwards induction".) However, if the number of games is random, the situation can change. In this case, the equilibrium strategy can be "tit-for-tat" — in which I play D as long as you do, but if you play H, I counter by playing H on the next game (only). All this, and more, is covered in a book by Axelrod, *Evolution of Cooperation*, see [**?**].

The case $c > \frac{v}{2}$ is more interesting. This is the case that is equivalent to Chicken. 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 p. To be a Nash equilibrium, we need the payoffs for player II to play H and D to be equal:

$$\text{(L)} \qquad p(\frac{v}{2} - c) + (1-p)v = (1-p)\frac{v}{2} \quad \text{(R)}. \tag{3.2}$$

For this to be true, we need $p = \frac{v}{2c}$, 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 p 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 p and D with probability $(1-p)$ — and so also will go the proportion of Hawks and Doves in the next generation.

If $p < \frac{v}{2c}$, then in equation (3.2), (L)>(R) — the expected payoff for a Hawk is greater than that for a Dove, and so in the next generation, p will increase.

On the other hand, if $p > \frac{v}{2c}$, then (L)<(R), so in the next generation, p will decrease. This case might seem strange — in a population of hawks, how could a few doves possibly do well? Recall that we are examining local stability, so the proportion of doves must be significant (a single dove in a population of hawks is not allowed); and imagine that the hawks are always getting injured fighting each other.

Some more work needs to be done — in particular, specifying the population dynamics more completely — to show that the mixed Nash equilibrium is a population equilibrium, but this certainly suggests it.

Example 3.7.1 (Sex Ratios). A standard example of this in nature is the case of sex ratios. In mostly monogamous species, a ratio close to $1:1$ males to females seems like a good idea, but what about sea lions, in which a single male gathers a large harem of females, while the majority of males never reproduce? Game theory provides an explanation for this. In a stable population, the expected number of offspring that live to adulthood per adult individual per lifetime is 2. The number of offspring a female sea lion produces in her life probably doesn't vary too much from 2. However, there is a large probability a male sea lion won't produce any offspring, balanced by a small probability that he gets a harem and produces a prodigious number. If the percentage of males in a (stable) population decreases, then since the number of harems is fixed, the expected number of offspring per male increases, and payoff (in terms of second-generation offspring) of producing a male increases.

3.7.2 Evolutionarily stable strategies

Consider a symmetric, two-player game with n pure strategies each, and payoff matrices ($A_{i,j} = B_{j,i}$), where $A_{i,j}$ is the payoff of player I when playing strategy i if player II plays strategy j, and $B_{i,j}$ is the payoff of player II when playing strategy i if player I plays strategy j.

Definition 3.7.1 (). 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 $\mathbf{z}^t A\mathbf{z} < \mathbf{x}^t A\mathbf{z}$.

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.

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

$$\begin{aligned} &\varepsilon\mathbf{x}^t A\mathbf{z} + (1-\varepsilon)\mathbf{x}^t A\mathbf{x} \quad (\text{for } \mathbf{x}\text{'s}) \\ &\varepsilon\mathbf{z}^t A\mathbf{z} + (1-\varepsilon)\mathbf{z}^t A\mathbf{x} \quad (\text{for } \mathbf{z}\text{'s}). \end{aligned}$$

The two criterions for $\mathbf{x}$ to be ESS imply that, for small enough ε, the average payoff for $\mathbf{x}$ will be strictly greater than that for $\mathbf{z}$, so the invaders will disappear.

Note also that criterion (i) in the definition of an ESS looks unlikely to occur in practice, but recall that if a mixed Nash equilibrium is found by averaging, then any mutant not introducing a new strategy will have $\mathbf{z}^t A\mathbf{x} = \mathbf{x}^t A\mathbf{x}$.

Example 3.7.2 (Hawks and Doves). We will check that the mixed Nash equilibrium in Hawks and Doves is an ESS when $c > \frac{v}{2}$. Let $\mathbf{x} = \frac{v}{2c}H + (1 - \frac{v}{2c})D$.

- if $\mathbf{z} = (1,0)$ ("H") then $\mathbf{z}^t A\mathbf{z} = \frac{v}{2} - c$, which is strictly less than $\mathbf{x}^t A\mathbf{z} = p(\frac{v}{2} - c) + (1-p)0$.
- if $\mathbf{z} = (0,1)$ ("D") then $\mathbf{z}^t A\mathbf{z} = \frac{v}{2} < \mathbf{x}^t A\mathbf{z} = pv + (1-p)\frac{v}{2}$.

The mixed Nash equilibrium for Hawks and Doves (when it exists) is an ESS.

Example 3.7.3 (Rock-Paper-Scissors)**.** The unique Nash equilibrium in Rock-Paper-Scissors, $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, is **not** evolutionarily stable. Under appropriate notions of population dynamics, this leads to cycling: 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 [**?**].

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 3.7.4 (**Congestion Game**)**.** 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

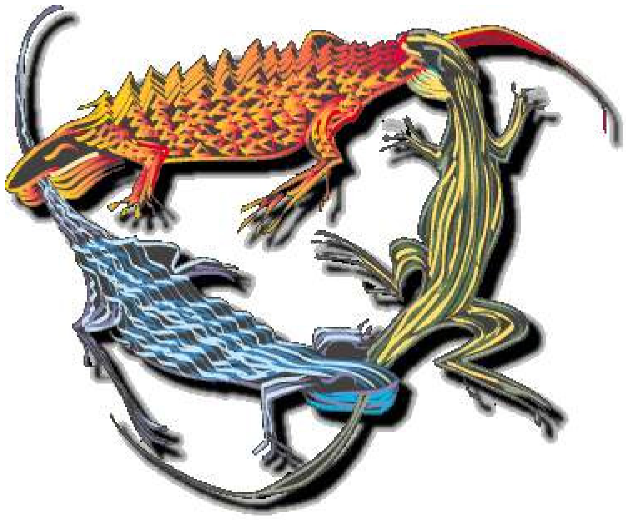

Fig. 3.14. The three types of male of the lizard *Uta stansburiana*. Picture courtesy of Barry Sinervo; see `http://bio.research.ucsc.edu/~barrylab`.

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:

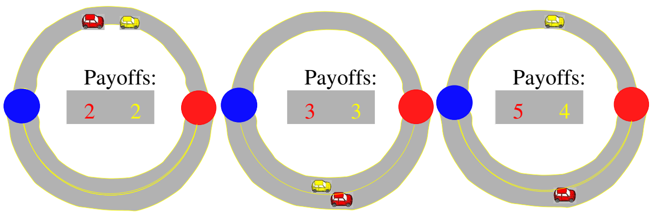

Fig. 3.15.

player II

player I	W	N
W	$(3,3)$	$(5,4)$
N	$(4,5)$	$(2,2)$

There are two pure Nash equilibria: (W,N) and (N,W).

If player I chooses W with probability p, II's payoff for choosing W is $3p+5(1-p)$, and for choosing N is $4p+2(1-p)$. Equating these, we get that the symmetric Nash equilibrium is when both players take the wide route with probability $p = \frac{3}{4}$.

Is this a stable equilibrium? Let $\mathbf{x} = (.75, .25)$ be our equilibrium strategy. We already checked that $\mathbf{x}^t A\mathbf{x} = \mathbf{z}^t A\mathbf{x}$ for all pure strategies $\mathbf{z}$, we need only check that $\mathbf{x}^t A\mathbf{z} > \mathbf{z}^t A\mathbf{z}$. For $\mathbf{z} = (1, 0)$, $\mathbf{x}^t A\mathbf{z} = 3.25 > \mathbf{z}^t A\mathbf{z} = 3$, and for $\mathbf{z} = (0, 1)$, $\mathbf{x}^t A\mathbf{z} = 4.25 < \mathbf{z}^t A\mathbf{z} = 2$, implying that $\mathbf{x}$ is evolutionarily stable.

Remark. For the above game to make sense in a population setting, one could suppose that only two randomly chosen drivers may travel at once — although one might also imagine that a driver's payoff on a day when a proportion x of the population are taking the wide route is proportional to her expected payoff when facing a single opponent who chooses W with probability x.

The symmetric Nash equilibrium may represent a stable partition of the population — in this case implying that if driver preferences are such that on an average day, one-quarter of the population of drivers prefer the narrow route, then any significant shift in driver preferences will leave those who changed going slower than they had before.

To be true in general, the statement above should read "any *small but* significant shift in driver preferences will leave those who changed going slower than they had before". The fact that it is a Nash equilibrium means that the choice of route to a *single* driver does not matter (if the population is large). However, since the strategy is evolutionarily stable, if enough drivers change their preferences so that they begin to interact with each other, they will go slower than those who did not change, on average. In this case, there is only one evolutionarily stable strategy, and this is true no matter the size of the perturbation. In general, there may be more than one, and a large enough change in strategy may move the population to a different ESS.

This is another game where binding commitments will change the outcome — and in this case, both players will come out better off!

Example 3.7.5 (A symmetric game). If in the above game, the payoff matrix was instead

player I \ player II	W	N
W	$(4, 4)$	$(5, 3)$
N	$(3, 5)$	$(2, 2)$

then the only Nash equilibrium is (W, W), which is also evolutionarily stable.

This is an example of the following **general fact:** In a symmetric game, if $a_{ii} > a_{i,j}$ for all $j \neq i$, then pure strategy i is an evolutionarily stable strategy. This is clear, since if I plays i, then II's best response is also the pure strategy i.

Example 3.7.6 (**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.

Remark. In this game, if a large enough population of mutant As invades a population of Bs, then the "stable" population will in fact shift to being entirely composed of As. Another situation that would remove the stability of (B, B) is if mutants were allowed to preferentially self-interact.

3.8 Signaling and asymmetric information

Example 3.8.1 (**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**.

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 to catch 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)

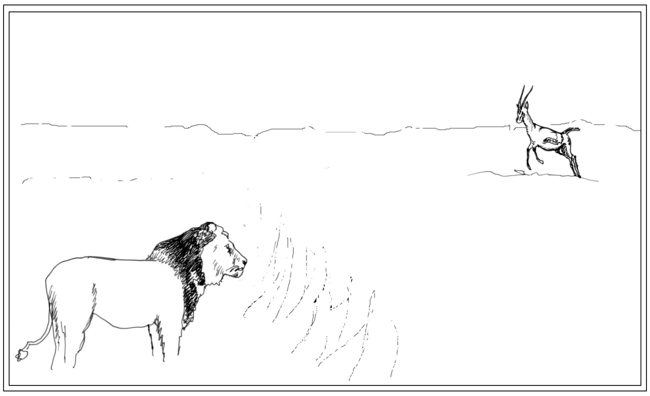

Fig. 3.16. Lone antelope stotting to indicate its good health.

or ignoring (I).

$A^H =$

		antelope
		run-if-chased
lion	chase	$(-1, -1)$
	ignore	$(0, 0)$

$A^W =$

		antelope
		run-if-chased
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.

3.8.1 Examples of signaling (and not)

Example 3.8.2 (**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.

$A^H =$

		player II	
		L	R
player I	L	4	1
	R	3	0

$A^T =$

		player II	
		L	R
player I	L	1	3
	R	2	5

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 payoff of 2.5. If both players know the outcome of the coin flip, then (since A^H has a payoff of 1 and A^T has a payoff of 2) the payoff is 1.5 — player II has been 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, but 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 3.8.3 (**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.

$$A^H = \text{player I}\ \begin{array}{c|cc} & \multicolumn{2}{c}{\text{player II}} \\ & L & R \\ \hline L & -1 & 0 \\ R & 0 & 0 \end{array} \qquad\qquad A^T = \text{player I}\ \begin{array}{c|cc} & \multicolumn{2}{c}{\text{player II}} \\ & L & R \\ \hline L & 0 & 0 \\ R & 0 & -1 \end{array}$$

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

$$A^H = \text{player I}\ \begin{array}{c|cc} & \multicolumn{2}{c}{\text{player II}} \\ & L & R \\ \hline L & 1 & 0 \\ R & 0 & 0 \end{array} \qquad\qquad A^T = \text{player I}\ \begin{array}{c|cc} & \multicolumn{2}{c}{\text{player II}} \\ & L & R \\ \hline L & 0 & 0 \\ R & 0 & 1 \end{array}$$

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).

3.8.2 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. This is an extremely simplified version. 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.

3.9 Some further examples

Fish being sold at the market is fresh with probability 2/3 and old otherwise, 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

Fig. 3.17. 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.)

Example 3.9.1 (The fish-selling game).

Fig. 3.18. The seller knows whether the fish is fresh, the customer only knows the probability.

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.

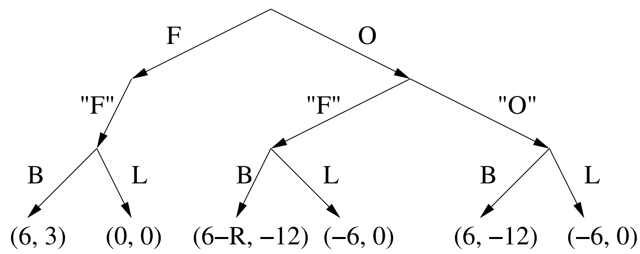

Fig. 3.19. The Kuhn tree for the fish-selling game.

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.

3.10 Potential games

We now discuss a collection of games called **potential games**, which are k-players general-sum games that have a special feature. Let $F_i(s_1, s_2, \ldots, s_k)$ denote the payoff to player i if the players adopt the pure strategies $s_1, s_2, \ldots, s_k$, respectively. In a potential game, there is a function $\psi : S_1 \times \cdots \times S_k \to \mathbb{R}$, defined on the product of the players' strategy spaces, such that

$$F_i\big(s_1, \ldots, s_{i-1}, \tilde{s}_i, s_{i+1}, \ldots, s_k\big) - F_i\big(s_1, \ldots, s_k\big) \\ = \psi\big(s_1, \ldots, s_{i-1}, \tilde{s}_i, s_{i+1}, \ldots, s_k\big) - \psi\big(s_1, \ldots, s_k\big), \quad (3.3)$$

for each i. We assume that each S_i is finite. We call the function ψ the **potential function** associated with the game.

Example 3.10.1 (**A simultaneous congestion game**). In this sort of game, the cost of using each road depends on the number of users of the road. In the game depicted in the figure, for road 1 connecting A to B, it is $C(1, i)$ if there are i users, with $i \in \{1, 2\}$. Note that the cost paid by a given driver depends only on the number of users, not on which user she is.

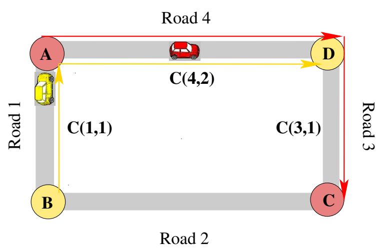

Fig. 3.20. Red car is travelling from A to C via D; yellow — from B to D via A.

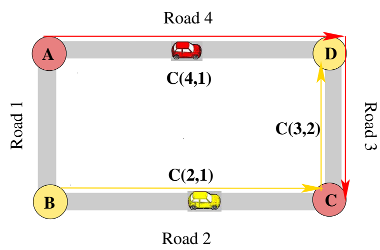

Fig. 3.21. Red car is travelling from A to C via D; yellow — from B to D via C.

More generally, for k drivers and R roads, we may define $\mathbb{R}$-valued map C on the product space of the road-index set and the set $\{1, \ldots, k\}$, so that $C(j, u_j)$ is equal to the cost incurred by any driver using road j in the case that the total number of drivers using this road is equal to u_j. Note that the strategy vector $\mathbf{s} = (s_1, s_2, \ldots, s_k)$ determines the usage of each road. That is, it determines $u_i(s)$ for each $i \in \{1, \ldots R\}$, where

$$u_i(s) = \Big|\Big\{ j \in \{1, \ldots, k\} : \text{ player } j \text{ uses road } i \text{ under strategy } s_j \Big\}\Big|.$$

In the case of the game depicted in the figure, we suppose that two drivers, I (red) and II (yellow), have to travel from A to C, or from B to D, respectively.

In general, we set

$$\psi\big(s_1,\ldots,s_k\big) = -\sum_{r=1}^{R}\sum_{\ell=1}^{u_r(\mathbf{s})} C(r,\ell).$$

We claim that ψ is a potential function for such a game. We show why this is so in the specific example. Suppose that driver II, using roads 1 and 4, makes a decision to use roads 2 and 3 instead. What will be the effect on her cost? The answer is a change of

$$\Big(C\big(2,u_2(\mathbf{s})+1\big)+C\big(3,u_3(\mathbf{s})+1\big)\Big) - \Big(C\big(1,u_1(\mathbf{s})\big)+C\big(4,u_4(\mathbf{s})\big)\Big).$$

How did the potential function change as a result of her decision? We find that, in fact,

$$\psi(\mathbf{s}) - \psi(\tilde{\mathbf{s}}) = C\big(2,u_2(\mathbf{s})+1\big)+C\big(3,u_3(\mathbf{s})+1\big) - C\big(1,u_1(\mathbf{s})\big) - C\big(4,u_4(\mathbf{s})\big)$$

where $\tilde{\mathbf{s}}$ denotes the new joint strategy (after her decision), and $\mathbf{s}$ denotes the previous one. Noting that payoff is the negation of cost, we find that the change in payoff is equal to the change in the value of ψ. To show that ψ is indeed a potential function, it would be necessary to reprise this argument in the case of a general change in strategy by one of the players.

Now, we have the following result due to Monderer and Shapley ([**?**]) and Rosenthal [**?**]:

Theorem 3.10.1. *Every potential game has a Nash equilibrium in pure strategies.*

Proof. By the finiteness of the set $S_1 \times \cdots \times S_k$, there exists some $\mathbf{s}$ that maximizes $\psi(\mathbf{s})$. Note that for this $\mathbf{s}$ the expression in (3.3) is at most zero, for any $i \in \{1,\ldots,k\}$ and any choice of $\tilde{\mathbf{s}}_i$. This implies that $\mathbf{s}$ is a Nash equilibrium. □

It is interesting to note that the very natural idea of looking for a Nash equilibrium by minimizing $\sum_{r=1}^{R} u_r C(r,u_r)$ does not work.

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 $1, \ldots, N$. Player II can cheat, or wait, on any given day. The payoff to I if 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 I		player II: Ch	Wa
	In	1	0
	Wa	-1	0

The matrix Γ_n is given by

player I		player II: Ch	Wa
	In	1	-1
	Wa	-1	Γ_{n-1}

Final optimal strategies, and the value of Γ_n.

3.7 **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

$$\frac{\ell}{2}\left(\frac{2l - s}{s + \ell}\right) + s\left(\frac{2s - \ell}{s + \ell}\right) < t.$$

Find the pure equilibria, and the symmetric mixed equilibria.

3.8 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.

3.9 **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 is the fish is old, and -8 for buying an old fish.

3.10 **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.

3.11 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.

3.12 **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)	$(-4,-4)$

3.13 **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.

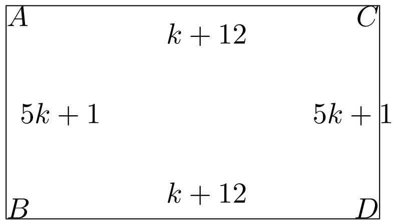

3.14 **A simultaneous congestion game:** There are two drivers, one who will travel from A to C, the other, from B to D. Each road in the second figure has been marked (x, y), where x is the cost to any driver who travels the road alone, and y is the cost to each driver who travels the road along with 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.

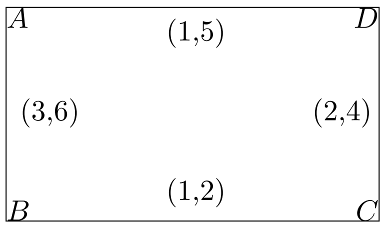

3.15 Sperner's lemma may be generalized to higher dimensions. In the case of $d = 3$, a simplex with four vertices (think of a pyramid) may be divided up into smaller ones. We insist that on each face of one of the small simplices, there are no edges or vertices of another. Label the four vertices of 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.

4
Coalitions and Shapley value

The topic we now turn to is that of games involving coalitions. Suppose we have a group of $k > 2$ players. Each seeks a part of a given prize, but may achieve that prize only by joining forces with some of the other players. The players have varying influence — but how much power does each have? This is a pretty general summary. We describe the theory in the context of an example.

4.1 The Shapley value and the glove market

We discuss an example, mentioned in the introduction. A customer enters a shop seeking to buy a pair of gloves. In the store are the three players. Player I has a left glove and players II and III each have a right glove. The customer will make a payment of \$100 for a pair of gloves. In their negotiations prior to the purchase, how much can each player realistically demand of the payment made by the customer?

To resolve this question, we introduce a **characteristic function** v, defined on subsets of the player set. By an abuse of notation, we will write v_{12} in place of $v(\{1,2\})$, and so on. The function v will take the values 0 or 1, and will take the value 1 precisely when the subset of players in question are able between them to effect their aim. In this case, this means that the subset includes one player with a left glove, and one with a right one — so that, between them, they may offer the customer a pair of gloves. Thus, the values are

$$v_{123} = v_{12} = v_{13} = 1,$$

and the value is 0 on every other subset of $\{1,2,3\}$. Note that v is a $\{0,1\}$-valued monotone function: if $S \subseteq T$, then $v_S \leq v_T$. Such a function is always **superadditive**: $v(S \cup T) \geq v(S) + v(T)$ if S and T are disjoint.

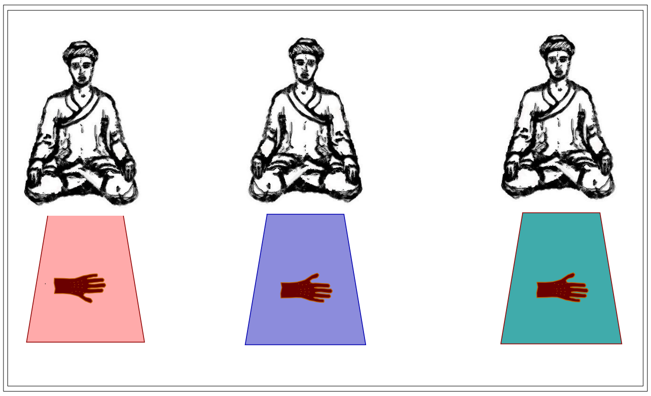

Fig. 4.1.

In general, a characteristic function is just a superadditive function with $v(\varnothing) = 0$. Shapley was searching for a value function ψ_i, $i \in \{1, \ldots, k\}$, such that $\psi_i(v)$ would be the **arbitration value** (now called **Shapley value**) for player i in a game whose characteristic function is v. Shapley analyzed this problem by introducing the following axioms:

(i) **Symmetry**: if $v\big(S \cup \{i\}\big) = v\big(S \cup \{j\}\big)$ for all S with $i, j \notin S$, then $\psi_i(v) = \psi_j(v)$.
(ii) **No power / no value**: if $v\big(S \cup \{i\}\big) = v(S)$ for all S, then $\psi_i(v) = 0$.
(iii) **Additivity**: $\psi_i(v + u) = \psi_i(v) + \psi_i(u)$.
(iv) **Efficiency**: $\sum_{i=1}^{k} \psi_i(v) = v\big(\{1, \ldots, k\}\big)$.

The second one is also called the "dummy" axiom. The third axiom is the most problematic: it assumes that for any of the players, there is no effect of earlier games on later ones.

Theorem 4.1.1 (Shapley)**.** *There exists a unique solution for ψ.*

A simpler example first: For a fixed subset $S \subseteq \{1, \ldots, n\}$, consider the **$S$-veto game**, in which the effective coalitions are those that contain each member of S. This game has characteristic function w_S, given by $w_S(T) = 1$ if and only if $S \subseteq T$. It is easy to find the unique function that is a Shapley value. Firstly, the "dummy" axiom gives that

$$\psi_i\big(w_S\big) = 0 \qquad \text{if } i \notin S.$$

Then, for $i, j \in S$, the "symmetry" axiom gives $\psi_i(w_S) = \psi_j(w_S)$. This and the "efficiency" axiom imply

$$\psi_i(w_S) = \frac{1}{|S|} \qquad \text{if } i \in S,$$

and we have determined the Shapley value (without using the additivity axiom). Moreover, we have that $\psi_i(c\,w_S) = c\,\psi_i(w_S)$ for any $c \in [0, \infty)$.

Now, note 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

$$w_{12} + w_{13} = v + w_{123}.$$

In particular, the "additivity" axiom gives

$$\psi_i(w_{12}) + \psi_i(w_{13}) = \psi_i(v) + \psi_i(w_{123}).$$

If $i = 1$, then $1/2 + 1/2 = \psi_1(v) + 1/3$, while, if $i = 3$, then $0 + 1/2 = \psi_3(v) + 1/3$. Hence $\psi_1(v) = 2/3$ and $\psi_2(v) = \psi_3(v) = 1/6$. This means that player I has two-thirds of the arbitration value, while players II and III have one-third between them.

Example: the 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. How much is the position of each player worth in the sense of Shapley value? Note that

$$1 = v_{1234} = v_{24} = v_{34},$$

while $v = 1$ on any 3-tuple, and $v = 0$ in each other case.

We will assume that the value v may be written in the form

$$v = \sum_{S \neq \varnothing} c_S w_S.$$

Later (in the proof of Theorem 4.2.1), 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

$$0 = v_1 = c_1$$

(we write c_1 for $c_{\{1\}}$, and so on). 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, $c_{34} = 1$. We have that

$$1 = v_{123} = c_{123},$$

while

$$1 = v_{124} = c_{24} + c_{124} = 1 + c_{124},$$

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 $c_{234} = -1$. We also have

$$\begin{aligned} 1 = v_{1234} &= c_{24} + c_{34} + c_{123} + c_{124} + c_{134} + c_{234} + c_{1234} \\ &= 1 + 1 + 1 + 0 + 0 - 1 + c_{1234}, \end{aligned}$$

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.

4.2 Probabilistic interpretation of Shapley value

Suppose that the players arrive at the board meeting in a uniform random order. 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 is the probability of that player being the one to make the existing coalition effective. We will now prove this assertion.

Recall that we are given $v(S)$ for all sets $S \subseteq [n] := \{1, \dots, n\}$, with $v(\varnothing) = 0$, and $v(S \cup T) \geq v(S) + v(T)$ if $S, T \subseteq [n]$ are disjoint.

Theorem 4.2.1. *Shapley's four axioms uniquely determine the functions ϕ_i. Moreover, we have the random arrival formula:*

$$\psi_i(v) = \frac{1}{n!} \sum_{k=1}^{n} \sum_{\pi \in S_n : \pi(k) = i} \Big(v\big(\pi(1), \dots, \pi(k)\big) - v\big(\pi(1), \dots, \pi(k-1)\big) \Big)$$

Remark. Note that this formula indeed specifies the probability just mentioned.

Proof. Recall the game for which $w_S(T) = 1$ if $S \subseteq T$, and $w_S(T) = 0$ in the other case. We showed that $\psi_i(w_S) = 1/|S|$ if $i \in S$, and $\psi_i(w_S) = 0$ otherwise. Our aim is, given v, to find coefficients $\big\{c_S\big\}_{S \subseteq [n], S \neq \varnothing}$ such that

$$v = \sum_{\varnothing \neq S \subseteq [n]} c_S w_S. \tag{4.1}$$

Firstly, we will assume (4.1), and determine the values of $\big\{c_S\big\}$. Applying (4.1) to the singleton $\{i\}$:

$$v\big(\{i\}\big) = \sum_{\varnothing \neq S \subseteq [n]} c_S w_S\big(\{i\}\big) = c_{\{i\}} w_i(i) = c_i, \tag{4.2}$$

where we may write c_i in place of $c_{\{i\}}$. More generally, suppose that we have determined c_S for all S with $|S| < \ell$. We want to determine $c_{\tilde{S}}$ for some $\tilde{S}$ with $|\tilde{S}| = \ell$. We have that

$$v(\tilde{S}) = \sum_{\varnothing \neq S \subseteq [n]} c_S w_S\big(\tilde{S}\big) = \sum_{S \subseteq \tilde{S}, |S| < \ell} c_S \, + c_{\tilde{S}}. \tag{4.3}$$

This determines $c_{\tilde{S}}$. Now let us verify that (4.1) does indeed hold. Define the coefficients $\{c_S\}$ via (4.2) and (4.3), inductively for sets $\tilde{S}$ of size $\ell > 1$; that is,

$$c_{\tilde{S}} = v(\tilde{S}) \, - \sum_{S \subseteq \tilde{S} : |S| < \ell} c_S.$$

However, once (4.2) and (4.3) are satisfied, (4.1) also holds (something that should be checked by induction). We now find that

$$\psi_i(v) = \psi_i\Big(\sum_{\varnothing \neq S \subseteq [n]} c_S w_S \Big) = \sum_{\varnothing \neq S \subseteq [n]} \psi_i\big(c_S w_S\big) = \sum_{S \subseteq [n], i \in S} \frac{c_S}{|S|}.$$

This completes the proof of the first statement made in the theorem.

As for the second statement: for each permutation π with $\pi(k) = i$, we define

$$\phi_i(v, \pi) = v\big(\pi(1), \dots, \pi(k)\big) - v\big(\pi(1), \dots, \pi(k-1)\big),$$

and

$$\Psi_i(v) = \frac{1}{n!} \sum_{\{\pi : \pi(k) = i\}} \phi_i(v, \pi).$$

Our goal is to show that $\Psi_i(v)$ satisfies all four axioms.

For a given π, note that $\phi_i(v, \pi)$ satisfies the "dummy" and "efficiency" axioms. It also satisfies the "additivity" axiom, but not the "symmetry" axiom. We now show that averaging produces a new object that is already symmetric — that is, that $\{\Psi_i(v)\}$ satisfies this axiom. To this end, suppose that i and j are such that

$$v\big(S \cup \{i\}\big) = v\big(S \cup \{j\}\big)$$

for all $S \subseteq [n]$ with $S \cap \{i, j\} = \varnothing$. For every permutation π, define π^* that switches the locations of i and j. That is, if $\pi(k) = i$ and $\pi(\ell) = j$, then $\pi^*(k) = j$ and $\pi^*(\ell) = i$, with $\pi^*(r) = \pi(r)$ with $r \neq k, \ell$. We claim that

$$\phi_i(v, \pi) = \phi_j(v, \pi^*).$$

Suppose that $\pi(k) = i$ and $\pi(\ell) = j$. Note that $\phi_i(v, \pi)$ contains the term

$$v\big(\pi(1), \dots, \pi(k)\big) - v\big(\pi(1), \dots, \pi(k-1)\big),$$

whereas $\phi_i(v, \pi^*)$ contains the corresponding term

$$v\big(\pi^*(1), \dots, \pi^*(k)\big) - v\big(\pi^*(1), \dots, \pi^*(k-1)\big).$$

We find that

$$\begin{aligned}
\Psi_i(v) &= \frac{1}{n!} \sum_{\pi \in S_n} \phi_i(v, \pi) = \frac{1}{n!} \sum_{\pi \in S_n} \phi_j(v, \pi^*) \\
&= \frac{1}{n!} \sum_{\pi^* \in S_n} \phi_j(v, \pi^*) = \Psi_j(v),
\end{aligned}$$

where in the second equality, we used the fact that the map $\pi \mapsto \pi^*$ is a one-to-one map from S_n to itself, for which $\pi^{**} = \pi$. Therefore, $\Psi_i(v)$ is indeed the unique Shapley value. $\square$

4.3 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. We find the Shapley value: suppose that the buyer pays $\$x$ for the fish, with $0 < x \leq 10$. Writing S and B for the seller and buyer, we have that $v(S) = 0$, $v(B) = 0$, with $v(S, B) = (10 - x) + x$, so that $\psi_S(v) = \psi_B(v) = 5$.

A potential problem with using the Shapley value in this case is the possibility that the buyer underreports his desire for the fish to the party that arbitrates the transaction.

Many right gloves. Find the Shapley values for 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. Note that $V(S)$ is equal to the maximal number of proper and disjoint pairs of gloves. In other words, $v(S)$ is equal to the minimum of the number of left, and of right, gloves held by members of S. Note that $\psi_1(v) = \psi_2(v)$, and $\psi_r(v) = \psi_3(v)$, for each $r \geq 3$. Note also that

$$2\psi_1(v) + r\psi_3(v) = 2,$$

provided that $r \geq 2$. For which permutations does the third player add value to the coalition already formed? The answer is the following orders:

$$13, 23, \{1, 2\}3, \{1, 2, j\}3,$$

where j is any value in $\{4, \ldots, n\}$, and where the curly brackets mean that each of the resulting orders is to be included. The number of permutations corresponding to these possibilities is: $r!$, $r!$, $2(r-1)!$, and $6(r-1) \cdot (r-2)!$. This gives that

$$\psi_3(v) = \frac{2r! + 8(r-1)!}{(r+2)!}.$$

That is,

$$\psi_3(v) = \frac{2r + 8}{(r+2)(r+1)r}.$$

Exercises

4.1 **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 $i = 1,2,\ldots,n$.

5
Mechanism design

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 desireable 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 divy 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.

5.1 Auctions

We will introduce a few of the basic types of auctions. The set-up for game theoretic analysis is as follows: There is a seller, known as the **principal**, some number of buyers, known as the **agents**, and a single item (for simplicity) to be sold, of value v_* to the principal, and of value v_i to agent i. Frequently, the principal has a **reserve price** v_{res}: She will not sell the item, unless the final price is at least the reservation price. The following are some of the basic types of auction:

Definition 5.1.1 (English Auction). In an English auction,

- agents make increasing bids,

- when there are no more bids, the highest bidder gets the item at the price he bids, if that price is at least v_{res}.

Definition 5.1.2 (Dutch Auction). The Dutch auction works in the other direction: in a Dutch auction,

- the principal gives a sequence of decreasing prices,
- the first agent to say "stop" gets the item at the price bid, if this is at least v_{res}.

Definition 5.1.3 (Sealed-bid, First-price Auction). This type of auction is the hardest to analyze. Here,

- buyers bid in sealed envelopes,
- the highest bidder gets the item at the price bid, if this is at least v_{res}.

Definition 5.1.4 (Vickrey Auction). This is a sealed-bid, second-price auction. In the Vickrey auction,

- buyers bid in sealed envelopes,
- the highest bidder gets the item at the *next-highest* bid, if this is at least v_{res}.

Why would a seller *ever* choose to run a Vickrey auction, when they could have a sealed-bid, first-price auction? Intuitively, the rules of the Vickrey auction will encourage the agents to bid higher than they would in a first-price auction. A Vickrey auction has the further theoretical advantage:

Theorem 5.1.1. *In a Vickrey auction, it is a pure Nash equilibrium for each agent to bid his or her value v_i.*

To make sense out of this, we need to specify that if agent i buys the item for ψ_i, the payoff is $v_i - \psi_i$ for agent i, and 0 for all other agents. The role the principal plays is in choosing the rules of the game — she is not a player in the game.

It is clear that in the Vickrey auction, if the agents are following this Nash equilibrium strategy, then the item will sell for the value of the second-highest bidder. This turns out to also be true in the English and the Dutch auctions. In both cases, we need to assume that the bids move in a continuous fashion (or by infinitesimal increments), and that ties are dealt with in a reasonable fashion. In the Dutch auction, we also need to assume that the agents know each other's values.

This implies that in the English and Vickrey auctions, agents can behave optimally knowing only their own valuation, whereas in the Dutch and sealed-bid first-price auctions, they need to guess the others' valuations.

We now prove the theorem.

Proof. To show that agent i bidding their value v_i is a pure Nash equilibrium, we need to show that each agent can't gain by bidding differently. Assume, for simplicity, that there are no ties.

Suppose that agent i changes his bid to $h_i > v_i$. This changes his payoff only if this causes him to get the item, i.e., if there is a $j \neq i$ such that $v_i < v_j < h_i$, and $h_i > v_k$ for all other k. In this case, he pays v_j, his new payoff is $v_i - v_j < 0$, as opposed to the payoff of zero he achieved, before switching.

Now suppose that agent i changes his bid to $\ell_i < v_i$. This changes his payoff only if he was previously going to get the item, and bidding ℓ_i would cause him not to get it, i.e., $v_i > v_k$ for all $k \neq i$, and there exists a v_j such that $\ell_i < v_j < v_i$. In this case, his payoff changes from $v_i - v_j > 0$ to zero. In both cases, he ends up either the same, or worse off. $\square$

Note: Revenue equivalence theorem. More discussion.

Remark. The above pre-supposes that people know their values for the items in the auction. In practice this isn't always the case. The internet auction site eBay uses a second price auction, so that people can bid their "true value". Nonetheless, many people increase their bids when someone else bids higher than their "true value". Having some knowledge of other people's values can influence how much a person values an item. Learning how much other people value an item influences how much a given person values an item, and this is an important consideration when setting up an auction.

5.2 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 $\big\{k/10 : k \in \{0, \ldots, 10\}\big\}$. 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 $\frac{k}{n}$, 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.

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.

There is further refinement in that even if the weatherman doesn't know the average rainfall, he can still do quite well.

Theorem 5.2.1. *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.*

$$\limsup_{n\to\infty} \frac{1}{n} \sum_{k=0}^{10} \left| r_k - \frac{k}{10} n_k \right| \leq \frac{1}{20}.$$

One proof of this can be found in ([**?**]), 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 [**?**]), by considering the situation as a zero-sum game between the weatherman and a certain adversary. In this case the adversary is 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.

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 is 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 [**?**].

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 $g_p : [0,1] \to \mathbb{R}$ should satisfy $g_p(p) > g_p(x)$ for all $x \in [0,1]\backslash\{p\}$.

One good choice is to let $f(x) = \log x$. In this case, the derivative of $g_p(x)$ will be as follows.

$$g_p'(x) = pf'(x) + (1-p)f'(1-x) = \frac{p}{x} - \frac{1-p}{1-x}.$$

The derivative is positive if $x < p$, and negative if $x > p$. So the maximizer of $g_p(x)$ is at $x = p$.

5.3 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, $0 \le S < M = 10^{15}$.

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 5.3.1. Let $A = \{i_1, \ldots, i_\ell\} \subset \{1, \ldots, n\}$ be any subset of size $\ell < n$. We say that a coalition of ℓ people holding a random vector $(X_{i_1}, \ldots, X_{i_\ell})$ has **no information** about a secret S provided $(X_{i_1}, \ldots, X_{i_\ell})$ is a random vector on $\{0, \ldots, M-1\}^\ell$, whose distribution is independent of S, that is

$$\Pr(X_{i_1} = x_1, \ldots, X_{i_\ell} = x_\ell | S = s)$$

does not depend upon s.

The simplest way to ensure that the distribution of $(X_{i_1}, \ldots, X_{i_\ell})$ 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:

$$\Pr(X = x) = \frac{1}{N} \qquad \forall\, x \in \Omega\,.$$

In the case of an ℓ-dimensional vector with elements in $\{0, \ldots, M-1\}$, we have $\Omega = \{0, \ldots, M-1\}^\ell$, of size $1/M^\ell$.

5.3.1 A simple secret sharing method

The following scheme allows the secret holder to split a secret $S \in \{0, \ldots, M-1\}$ 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 $(X_1, X_2, \ldots, X_{n-1})$, whose distribution is uniform on $\{0, \ldots, M-1\}^{n-1}$. She gives the number X_i to the i^{th} person for $1 \le i \le n-1$, and the number

$$X_n = \left(S - \sum_{i=1}^{n-1} X_i\right) \mod M \tag{5.1}$$

to the last person. Notice that with this definition, X_n is also a uniformly random variable on $\{0, \ldots, M-1\}$, you will prove this in Ex. 5.2.

It is enough to show that any coalition of size $n-1$ has no useful information. For $\{i_1, \ldots, i_{n-1}\} = \{1, \ldots, n-1\}$, 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. 5.1:

Lemma 5.3.1. *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: $A = \{1, \ldots, j-1, j+1, \ldots, n\}$. Let $T_j((X_1, \ldots, X_{n-1})) = (X_1, \ldots, X_{j-1}, X_{j+1}, \ldots, X_n)$, where X_n is defined by Eq. (5.1). This map is one-to-one and onto for each j since we can explicitly define its inverse:

$$T_j^{-1}((Z_1, \ldots, Z_{j-1}, Z_{j+1}, \ldots Z_n)^T) = (Z_1, \ldots, Z_{j-1}, Z_j, Z_{j+1}, \ldots, Z_{n-1})^T,$$

where $Z_j = S - \sum_{1 \le i \ne j \le n-1} Z_i$.

So if a coalition (that does not include all players) puts together all its available information, it 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.

5.3.2 Polynomial method

The following method, devised by Adi Shamir [?], can also be used to split the secret among n players. It has an interesting advantage: using this method we can share a secret between n individual in such a way 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 p be a prime number such that $0 \leq S < p$ and $n < p$. We define a polynomial of order $m-1$:

$$F(z) = \sum_{i=0}^{m-1} A_i z^i \mod p,$$

where $A_0 = S$ and $(A_1, \ldots, A_{m-1})$ is a uniform random vector on $\{0, \ldots, p-1\}^{m-1}$.

Let $z_1, \ldots, z_n$ be distinct numbers in $\{1, \ldots p-1\}$. 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 5.3.1. *A coalition of size m or bigger can reconstruct the secret S, but a coalition of size $\ell < m$ has no useful information:*

$$\Pr(F(z_1) = x_1, \ldots, F(z_\ell) = x_\ell | S) = \frac{1}{p^\ell}, \quad x_i \in \{0, \ldots, p-1\}.$$

Proof. Again it's enough to consider the case $\ell = m-1$. We will show that for any fixed distinct non-zero integers $z_1, \ldots, z_m \in \{0, \ldots, p-1\}$,

$$T((A_0, \ldots, A_{m-1})) = (F(z_1), \ldots, F(z_m))$$

is an invertible linear map on $\{0, \ldots, p-1\}^m$, and hence m people together can recover all the coefficients of F, including $A_0 = S$.

Let's construct these maps explicitly:

$$T \begin{pmatrix} A_0 \\ \vdots \\ A_{m-1} \end{pmatrix} = \begin{pmatrix} \sum_{i=0}^{m-1} A_i z_1^i \mod p \\ \vdots \\ \sum_{i=0}^{m-1} A_i z_m^i \mod p \end{pmatrix}.$$

We see that T is a linear transformation on $\{0, \ldots, p-1\}^m$ 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. 5.3 that

$$\det(M) = \prod_{1 \le i < j \le m} (z_j - z_i).$$

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:

$$F(z) = \sum_{j=1}^{m} F(z_j) \prod_{\substack{1 \le i \le m \\ i \ne j}} \frac{z - z_i}{z_j - z_i} \mod p.$$

Once we expand the right-hand side and bring it to the standard form, $(A_0, \ldots, A_{m-1})$ will appear as the coefficients of the corresponding powers of the indeterminate z. Evaluating at $z = 0$ gives back the secret. $\square$

5.4 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 was no disinterested third party who could be trusted by all the faculty members, they hit upon the following scheme:

All the faculty members gathered around a table. A designated first person picked a very large integer M (which he kept private), added his salary to that number, and passed the result to his neighbor on the right. She, in turn, added her salary and passed the result to her right. The intention was that the total should eventually return to the designated first person, who would then subtract M, compute and reveal the average. 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 5.3.1: 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 $X_{i,1}, \ldots, X_{i,n-1}$, uniformly distributed in the set $\{0, 1, 2, \ldots, M-1\}$, and produces $X_{i,n}$, such that $X_{i,1} + \cdots + X_{i,n} = s_i \bmod M$. 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 could 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. 5.5 that the above scheme leaks no additional information about the salaries.

5.5 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 the cake has no indivisible constituents.

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 5.5.1 (**Algebra of sets**)**.** More formally, we say that a collection $\mathcal{F}$ of subsets of Ω forms an algebra if:

(i) $\varnothing \in \mathcal{F}$;
(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 5.5.2. A non-negative real-valued set function μ defined on $\mathcal{F}$ is called a **finite measure** if:

(i) $\mu(\varnothing) = 0$ and $\mu(\Omega) = M < \infty$;
(ii) if $A, B \in \mathcal{F}$ and $A \cap B = \varnothing$ 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 (we exclude 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 $\{A_1^*, \ldots, A_K^*\}$, 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

$$\mu_j(\tilde{\Omega}) \geq \mu_j(\Omega)\Big(1 - \frac{1}{K}\Big) = \mu_j(\Omega)\frac{K-1}{K}.$$

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.

5.6 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 Alice does not know a 3-coloring, 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 Alice does not put the values of a 3-coloring in the envelopes, then there is 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 about to cheat 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.

5.7 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 4^{th} 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

5.1 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 $Y = T(X)$.

5.2 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

(a) Each X_i is a uniform random variable on $\{0, \dots, M-1\}$.

(b) X_i's are independent random variables.

(c) 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 $\{0, \ldots, M-1\}$.

5.3 Prove that the Vandermonde matrix has the following determinant:

$$\det \begin{bmatrix} 1 & z_1 & \ldots & z_1^{m-1} \\ 1 & z_2 & \ldots & z_2^{m-1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & z_{m-1} & \ldots & z_{m-1}^{m-1} \\ 1 & z_m & \ldots & z_m^{m-1} \end{bmatrix} = \prod_{1 \le i < j \le m} (z_j - z_i).$$

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.

5.4 Evaluate the following determinant, known as a **Cauchy determinant**:

$$\det \begin{bmatrix} \frac{1}{x_1 - y_1} & \frac{1}{x_1 - y_2} & \ldots & \frac{1}{x_1 - y_m} \\ \frac{1}{x_2 - y_1} & \frac{1}{x_2 - y_2} & \ldots & \frac{1}{x_2 - y_m} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{1}{x_m - y_1} & \frac{1}{x_m - y_2} & \ldots & \frac{1}{x_m - y_m} \end{bmatrix}.$$

Hint: find the zeros and poles and the constant factor. It is helpful to consider the limit $x_i \to y_j$.

5.5 Show that for the scheme for computing average salary described in section 5.4, 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).

6
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 the U.S. presidential elections. Suppose that a society (group of voters) 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, a simple concept of **majority rule** can be applied to yield an outcome that more than half of the voters find satisfactory. When the vote is evenly split between the two alternatives, an additional tie-breaking mechanism might be necessary. As the number of options increases to three or more, simple majority often becomes inapplicable. In order to find a unique winner, special procedures called **voting mechanisms** are used. The troubling aspect is that the result of the election will frequently depend on the particular mechanism selected.

6.1 Voting mechanisms and fairness criteria

A systematic study of voting mechanisms began in the 18$^{\text{th}}$ century with the stand-off between two members of the French Academy of Sciences — Jean-Charles, Chevalier de Borda and Marie Jean Antoine Nicolas de Caritat, Marquis de Condorcet. Chevalier de Borda observed that the current method in use by the Academy often led to the election of a candidate that was considered less desirable by the majority of the Academicians. He proposed an alternative mechanism (discussed in § 6.2.4) which was soon adapted. However, Marquis de Condorcet immediately demonstrated that the new mechanism itself suffered many undesirable properties and proceeded to invent his own method based on a certain **fairness criterion** now known as the **Condorcet criterion** (discussed in § 6.2.5). Roughly speaking, he postulated that if a candidate can defeat every other in a one-on-one contest, he should be the overall winner. He went further yet and discovered that his method based on pairwise contests also suffered a vulnerability, now known as the **Condorcet paradox** [**?**].

Since then, this pattern continued. A plethora of voting mechanisms have been introduced to satisfy a number of desirable fairness criteria, yet each one has been shown to contain a certain flaw or a paradox. The work of Kenneth Arrow from 1951 has elucidated the problem. He showed that no "fair" procedure can be devised that is free from strategic manipulation [**?**].

6.1.1 Arrow's fairness criteria

The notion of "fairness" used by Arrow requires some elaboration. Consider a finite set $A = \{a, b, c, \ldots\}$, consisting of m alternatives, where $m \geq 3$. For an entity X, **preference** $\succeq_X$ over A is a relationship which specifies, for each of the $\binom{m}{2}$ unordered pairs of alternatives, the one that is preferred by X (with ties allowed). We write $a \succeq b$ whenever a is preferred to b and $a \succ b$ if a is strictly preferred to b. A preference is called **transitive** if $a \succeq b$ and $b \succeq c$ implies that $a \succeq c$. In this case, a preference ($\succeq$) gives a **ranking** p listing the sets of equivalent alternatives from the most to the least preferred ones. Transitivity is not part of the definition of social preference, but will be required below.

Suppose that a society consists of N individuals, each with a transitive preference over A. A constitution is a function that associates to every N-tuple $\pi = (p_1, \ldots, p_N)$, of transitive preferences (called a profile) a social preference ($\succeq_S$).

A "fair" constitution should have the following properties:

(i) **Transitivity** of social preference.

(ii) **Unanimity**: if for every individual preference $a \succ_i b$, then $a \succ_S b$ for the social preference.

(iii) **Independence of irrelevant alternatives**: for any profiles π_1 and π_2 with fixed rankings between a and b, the social ranking of a and b should be the same.

The first two are self-explanatory. The third one is more subtle — this requirement ensures that there can be no "strategic misrepresentation" of individual preferences in order to achieve a desired social preference. Suppose that all individual preferences in π_1 and π_2 have the same rankings of a and b, but c and d are ranked differently, where $|\{c, d\} \cap \{a, b\}| \leq 1$. If π_1

leads to $a \succeq_S b$ and π_2 leads to $b \succeq_S a$, then the independence of irrelevant alternatives is violated, and the group of individuals who prefer b to a have an incentive to conceal their true preferences between c and d in order to achieve a desired social ranking.

A single-winner voting mechanism assigns to each alternative a numerical score. Such a system most commonly produces a special type of constitution — one which distinguishes a single strictly preferred alternative, and ranks all the others as equivalent to one another. In cases when a complete social ranking is extracted, it is again based on numerical scores and must therefore be transitive.

Next we will discuss a few of the most popular voting mechanisms, show how each one comes short of satisfying Arrow's "fairness" criteria, and finally state and prove Arrow's Impossibility theorem.

6.2 Examples of voting mechanisms

Single-winner voting mechanisms can be characterized by the ballot type they use. The **binary** methods use a **simple ballot** where the candidates or the alternatives are listed and each one is either selected or not. The **ranked** methods use a **preference ballot** where each alternative is ranked in the order of preference. Finally, in the **range** or **rated** methods each alternative listed on a ballot is given a score.

6.2.1 Plurality

Probably the most common mechanism is **Plurality**, also know as **First past the Post**. A *simple ballot* is used.

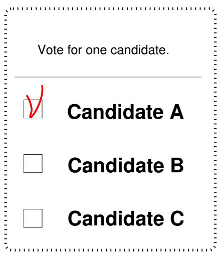

Fig. 6.1. Under *plurality* only one of the alternatives can be selected.

The alternative with the most votes wins. It need not have the majority of the votes. In the U.S., congressional elections within each congressional district are conducted using the plurality system.

This system has several advantages. It is particularly attractive because of its simplicity and transparency. In parliamentary elections, it is often praised for excluding extremists and encouraging political parties to have broader appeal. It also gives a popular independent candidate a chance, since ultimately people and not parties are elected. It may, however, lead to underrepresentation of the minorities and encourage the formation of parties based on geographic or ethnic appeal.

Another undesirable property of plurality is that the candidate that is ultimately elected might be the least favorite for a substantial part of the population. Suppose that there are three candidates A, B, and C, and the voters have three different types of rankings of the candidates.

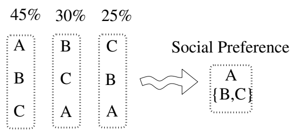

Fig. 6.2. Option A is preferred by 45% of the population, option B by 30% and option C by 25%.

Under simple plurality, A wins the election, in spite of a strong opposition by a majority of voters. If voters who favor C were to cast their votes for B, their second choice, B would win by a 55% majority. This is an example of strategic or **insincere** voting. The strategy of voting for a less desirable but more popular alternative is called **compromising**.

This example shows that plurality violates the *independence of the irrelevant alternatives* criterion, since change in the social preference between B and A can be accomplished without changing any individual A-B preferences. Notice that the full individual rankings are never revealed in the voting.

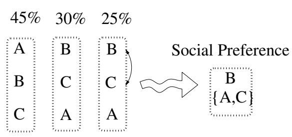

Fig. 6.3. When 25% *insincerely* switch their votes from C to B, social preference between A and B changes.

6.2.2 Runoff elections

Runoff elections are also known as **Plurality with Elimination**. A simple ballot is used, and the voting is carried out in rounds. After each round, if none of the candidates achieves the majority, the candidate (or alternative) with the fewest first place votes is eliminated, and a new round is carried out with the remaining candidates. When only two candidates remain in a round, the one with the most votes wins the election. For an N candidate election, runoff elections require at most $N-1$ rounds.

Notice that runoff voting changes the winner in the example above:

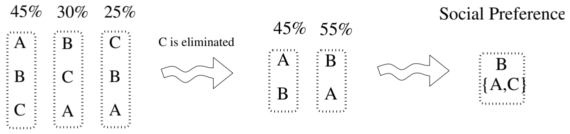

Fig. 6.4. In the first round C is eliminated. When votes are redistributed, B gets the majority. The full voter rankings are not revealed in the process.

This method is rarely used in its full form because of the additional costs and a lower voter turn-out associated with the multiple rounds. The most widely used version is the **top-two runoff** election. When a clear winner is not determined in the first round, a single runoff round is carried out between the top two candidates. In the U.S., runoff elections are often used in party primary elections and various local elections.

6.2.3 Instant runoff

With the use of the *preference ballot*, runoff elections can be accomplished in one round.

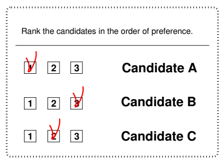

Fig. 6.5. In the *instant runoff*, voters specify their rankings of the candidates.

The candidate with the least number of first votes is eliminated from consideration, and the votes of those who put him first are redistributed to their second favorite candidate.

This method is cheaper than a proper runoff. It also encourages voter turn-outs since there is only one round. Yet this method suffers from the same weakness as plurality — it's open to strategic manipulations. Consider the following scenario:

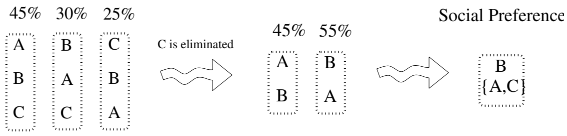

Fig. 6.6. After C is eliminated, B gets the majority of votes.

If voters in the first group knew the distribution of preferences, they could ensure a victory for A by getting 10% of their constituents to conceal their true preference and insincerely move C from the bottom to the top of their rankings. In the first round, B would be eliminated. Subsequently A would win against C. This strategy is called **push-over**.

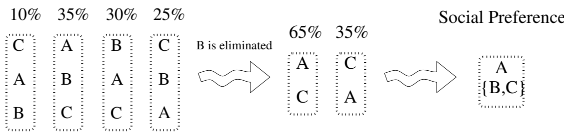

Fig. 6.7. A small group misrepresents their true preferences, ensuring that B is eliminated. As a result, A wins the election.

This example shows that instant runoff violates the independence of irrelevant alternatives criterion, since it allows for the social preference between A and B to be switched without changing any of the individual A-B preferences.

Instant runoff is used in Australia for elections to the Federal House of Representatives, in Fisi for the Fijian House of Representatives, to elect the President of Ireland, and for various municipal elections in Australia, the United States, and New Zealand.

6.2.4 Borda count

Borda count also uses the *preference ballot.* Under this method, given a numbered list of N alternatives, each voter assigns to it a permutation of $\{1, 2, \ldots, N\}$, where N corresponds to the most and 1 to the least desirable alternative. The candidate with the largest point total wins the election. Chevalier de Borda proposed this method in 1770 when he discovered that the *plurality* method then used by the French Academy of Sciences suffered from the paradox that we have described. The Borda method was subsequently used by the Academy for the next two decades.

Donald G. Saari showed that *Borda count* is in some sense the least problematic of all single winner mechanisms [?],[?]. Yet it is not free from the same flaw that plagues all other single-winner methods: it too can be manipulated by strategic voting.

Consider the following example:

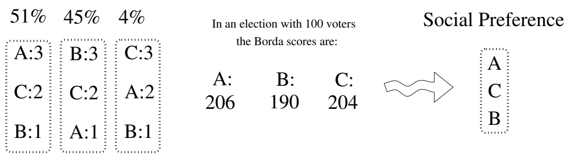

Fig. 6.8. Alternative A has the overall majority and is the winner under *Borda count.*

In this case, A has an unambiguous majority of votes and is also the winner under the *Borda count.* However, if supporters of C were to insincerely rank B above A, they could ensure a victory for C. This strategy is called **burying**.

This is again a violation of the independence of the irrelevant alternatives, since none of the individual A-C preferences had been changed.

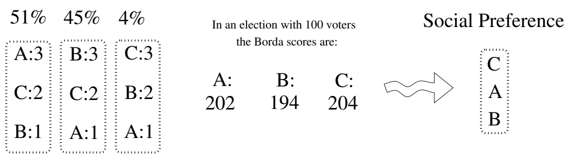

Fig. 6.9. Supporters of C can *bury* A by moving B up in their rankings.

6.2.5 Pairwise contests

Pairwise contests, also known as **Condorcet methods**, are a family of methods in which each alternative is matched one-on-one with each of the others. A one-on-one win brings 1 point, and a tie brings half a point. The alternative with the most total points wins. With N alternatives, this procedure requires $\binom{N}{2}$ stages. It can be accomplished in a single stage with the use of the *preference ballot.*

Marquis de Condorcet advanced this method after he demonstrated weaknesses in the *Borda count.* He then proceeded to show a vulnerability in his own method — a tie in the presence of a *preference cycle* [?].

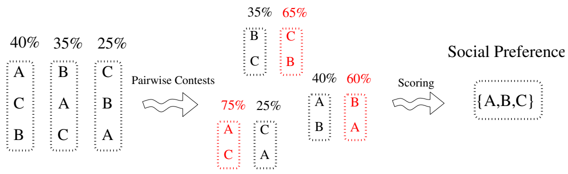

Fig. 6.10. Preference cycle: in one-on-one contests A defeats C, C defeats B, and B defeats A.

To resolve such a situation, a variety of tie breaking mechanisms exist. **Black's method**, for instance, uses *Borda count*, while more sophisticated methods run *Instant Runoffs* on a certain subset of the candidates.

In addition to frequently producing ties, **Condorcet methods** are in turn vulnerable to strategic voting. In the following example, supporters of C use *compromising* to resolve a cycle in favor of their second favorite alternative B.

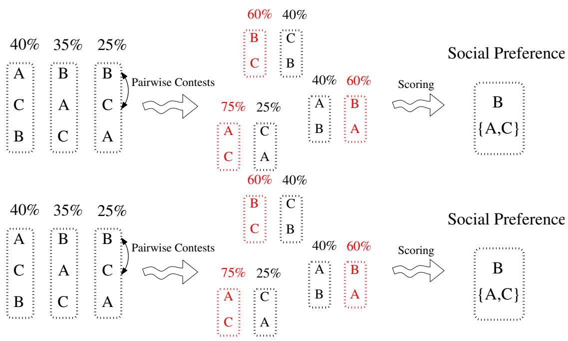

Fig. 6.11. Preference cycle (above panel): in one-on-one contests A defeats C, C defeats B, and B defeats A. After the third group of voters compromises and places B ahead of C (lower panel), B defeats C as well as A, so B is the overall winner.

This again violates the independence of irrelevant alternatives criterion, since B moves up relative to A in the social preference, while all individual A-B preferences remain constant.

6.2.6 Approval voting

Recently **approval voting** has become very popular in certain professional organizations. This is a procedure in which voters can vote for, or approve of, as many candidates as they wish. An **approval ballot** is used where

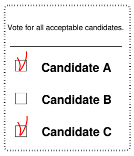

Fig. 6.12. Candidate A and C will receive one vote.

each approved candidate is marked off. Each approved candidate receives one vote, and the one with the most votes wins.

It should not come as a surprise that this method is also vulnerable to strategic voting. We give an example where a strategy equivalent to *compromising* allows supporters of C to get their second preferred candidate elected.

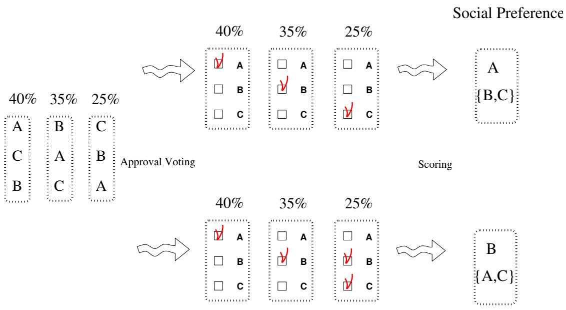

Fig. 6.13. When supporters of C also mark B as approved, the social preference between A and B changes, while all the individual AB preferences persist.

This shows that *approval voting* also violates the independence of the irrelevant alternative criterion.

6.3 Arrow's impossibility theorem

In 1951 Kenneth Arrow formulated and proved his famous *Impossibility Theorem*. He showed that the only constitution that is transitive, respects unanimity, and is invulnerable to strategic voting is a **dictatorship**. A constitution is called a **dictatorship** by an individual D if for any pair of alternatives α and β,

$$\alpha \succ_S \beta \iff \alpha \succ_D \beta$$

(see [**?**]). In essence, the preference of the dictator determines the social preference.

Theorem 6.3.1. *[Arrow's Theorem] Any constitution that respects transitivity, unanimity, and independence of irrelevant alternatives is a dictatorship.*

We present here a simplified proof of Arrow's theorem that is due to Geanakoplos [**?**]. The proof requires that we consider extremal alternatives — those at the very top or bottom of the rankings. Fix an individual X and an alternative β. Given a profile π, define two new profiles:

$$\pi^+(X,\beta) \text{ such that } \beta \succ_X \alpha \text{ for all } \alpha \neq \beta, \text{ all other preferences are as in } \pi$$
$$\pi^-(X,\beta), \text{ such that} \beta \prec_X \alpha \text{ for all } \alpha \neq \beta, \text{ all other preferences are as in } \pi.$$

Definition 6.3.1. X is said to be **extremely pivotal** for an alternative β at the profile π if $\pi^+(X,\beta)$ leads to a social ranking where $\beta \succ_S \alpha$ for each $\alpha \neq \beta$ in A, while $\pi^-(X,\beta)$ leads to a social ranking where $\beta \prec_S \alpha$ for each $\alpha \neq \beta$ in A.

Such an individual can move an alternative β from the very bottom of the social preference to the very top. We will show that there is an extremely pivotal individual X who must be a genuine dictator.

Lemma 6.3.1 (Extremal Lemma)**.** *Let alternative b be chosen arbitrarily from A. For a profile where every individual preference has the alternative b in an extremal position (at the very top or bottom), b must occupy an extremal position in the social preference, as well.*

Proof. Suppose, toward a contradiction, that for such a profile and distinct a, b, c, the social preference puts $a \succeq_s b$ and $b \succeq_s c$. Consider a new profile where every individual moved c strictly above a in their ranking. None of the ab or bc rankings changes since b is in an extremal location. Hence, by the independence of the irrelevant alternatives, in such a profile $a \succeq b$ and $b \succeq c$, still. By transitivity then, $a \succeq c$, while unanimity implies $c \succ a$, a contradiction. □

Next we argue that there is a voter $X = X(b)$ who is extremely pivotal for b at a certain profile π_1. Consider a profile such that each individual preference has b at the very bottom of the ranking. By unanimity the social ranking does the same. Now let the individuals successively 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 individuals put b at the top of their rankings, so must the society. Hence, there must be the first individual X

(*a priori*, his identity seems to depend on the order of preference switching), whose change in preference precipitates the change in the social ranking of b.

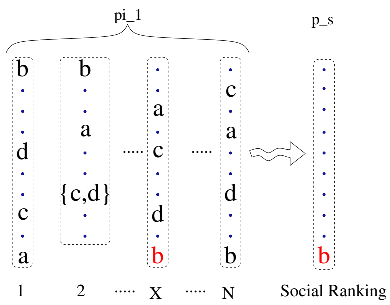

Fig. 6.14.

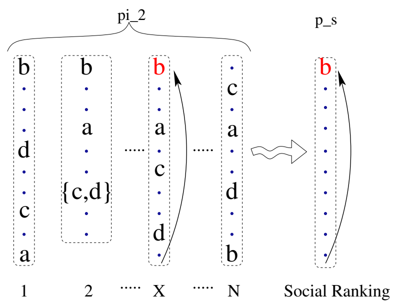

Fig. 6.15.

Denote by π_1 the profile just before X has switched b to the top, and by π_2 the profile immediately after the switch.

We argue that X must be a *limited dictator* over any pair ac not involving b. An individual X is called a **limited dictator** over ac, denoted by $D(ac)$, if $a \succ_S c$ whenever $a \succ_X c$ and $c \succ_S a$ whenever $c \succ_X a$. Let's choose one element, say a, from the pair ac. Construct a profile π_3 from π_2 by letting X move a above b so that $a \succ_X b \succ_X c$, and letting all the other individuals arbitrarily rearrange their relative rankings of a and c. By independence of irrelevant alternatives the social ranking corresponding to π_3 would necessarily put $a \succ_S b$, since all the individual ab preferences are as in profile π_1 where X put b at the bottom. Since all the individual bc ratings are as in profile π_2, where X puts b at the top, in the social ranking we must have $b \succ_S c$. Hence, by transitivity, $a \succ_S c$. This means, due to the independence of irrelevant alternatives, that the social ranking between ac necessarily coincides with that of individual X. Hence $X = D(ac)$ — a dictator over ac. It remains to show that X is also a dictator over ab.

We pick another distinct alternative d and construct an extremely pivotal voter $X(d)$. From the argument above, such a person is a dictator over any pair $\alpha\beta$ not involving d, for instance ab. But X can affect the social ranking of ab at profiles π_1 and π_2, hence $X = X(d) = D(ab)$ and is thus the ab dictator in question. This completes the proof.

Notice that the theorem does not say that the specific voting mechanism doesn't matter. The theorem merely asserts that dictatorship by an individual is the only mechanism that is free from strategic manipulation. A group in search of an acceptable voting mechanism should keep this in mind.

There are many other “fairness” criteria that could and should be used to select a “good” mechanism.

Exercises

6.1 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.

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. A matching ψ is **unstable** if there exist one man and one woman who are not matched to each other in ψ, but prefer each other to their partners in ψ. Otherwise, the matching is called **stable**.

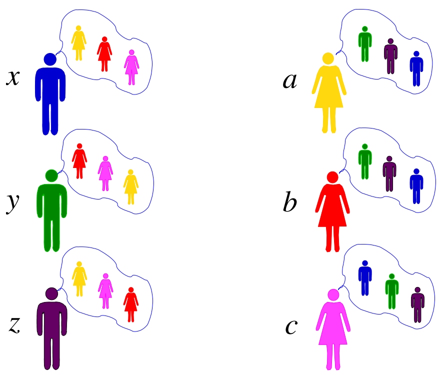

Fig. 7.1.

Let's see an example. Suppose we have 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** is introduced by Gale and Shapley.

(i) Each man proposes to his most preferred woman.
(ii) Each woman evaluates her proposers and rejects all but the most preferred one. She does not accept her preferred suitor at this stage, but puts him on hold.
(iii) Each rejected man proposes to his next preferred woman.
(iv) Repeat step (ii) and (iii) until each woman has one proposing man. At that point each woman accepts her proposer.

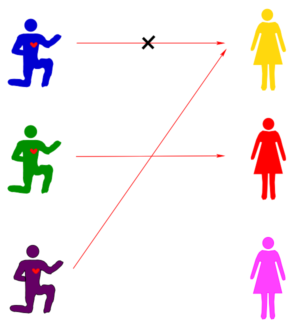

Fig. 7.2. Arrows indicate proposals, cross indicates rejection.

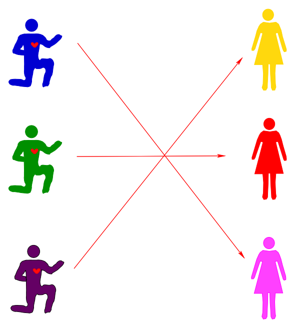

Fig. 7.3. Stable matching is achieved in the second stage.

Similarly, we could define a women-proposing algorithm.

Theorem 7.2.1. *The men-proposing algorithm yields a stable matching.*

Proof. First, the algorithm must terminate because when a man is rejected, he never proposes to the same woman again. Thus, a trivial upper bound for the number of rounds of the algorithm is n^2.

Next we are going to show that this algorithm stops only when each woman has exactly one proposer. Otherwise, the algorithm would stop when one man has all n proposals rejected. But this cannot happen because if a man j has $n-1$ proposals rejected, then these women all have proposers waiting for them. Hence the n^{th} proposal of man j cannot be rejected.

The argument above shows that we get a matching ψ by the algorithm. Now we prove that ψ is a stable matching. Consider a pair Bob and Alice with $\psi(\text{Bob}) \neq \text{Alice}$. If Bob prefers Alice to $\psi(\text{Bob})$, then Bob must have proposed to Alice earlier and was rejected. That means Alice got a better proposal. Hence $\psi^{-1}(\text{Alice})$ is a man she prefers to Bob. This proves that ψ is a stable matching. $\square$

7.3 Properties of stable matchings

We say a woman a is **attainable** for a man x if there exists a stable matching ϕ with $\phi(x) = a$.

Theorem 7.3.1. *Let ψ be the stable matching produced by Gale-Shapley men-proposing algorithm. Then,*

(a) For every man i, $\psi(i)$ is the most preferred attainable woman for i.
(b) For every woman j, $\psi^{-1}(j)$ is the least preferred attainable man for j.

Proof. Suppose ϕ is another stable matching. We prove by induction on the round of the men-proposing algorithm producing ψ, that every man k cannot be rejected by $\phi(k)$. So that $\psi(k)$ is preferred by k than $\phi(k)$.

In the first round, if a man k proposes to $\phi(k)$ and is rejected, then $\phi(k)$ has a better proposal in the first round, say, ℓ. Since $\phi(k)$ is the most preferred woman of ℓ the pair $(\ell, \phi(k))$ is unstable for ϕ, which is a contradiction.

Suppose we have proved the argument for round $1, 2, \ldots, r-1$. Consider round r. Suppose by contradiction that k proposes to $\phi(k)$ and rejected. Then, in this round $\phi(k)$ has better proposal, say, ℓ. By induction hypothesis, ℓ would have never been rejected by $\phi(\ell)$ in the earlier rounds. This means ℓ prefers $\phi(k)$ to $\phi(\ell)$. So $(\ell, \phi(k))$ is unstable for ϕ, which is a contradiction.

Thus we proved (a). For part (b), we could use the same induction. The detailed proof is left to the reader as an exercise. $\square$

Corollary 7.3.1. *If Alice is assigned to the same man in both of 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 $a_{i,j} \neq a_{i,j'}$ when $j \neq j'$, and $a_{i,j} \neq a_{i'j}$ when $i \neq i'$. 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: $j_1 > j_2 > \cdots > j_n$. By the same way, if in the j^{th} column, we have

$$a_{i_1j} < a_{i_2j} < \cdots < a_{i_nj}$$

Then, the preference order of woman j is: $i_1 > i_2 > \cdots > i_n$.

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 minimizes $\sum_i a_{i,\phi(i)}$ among all the stable matchings ϕ. Moreover, this stable matching reaches the unique minimum of $\sum_i a_{i,\phi(i)}$. Meanwhile, the women-proposing algorithm produces a stable matching which minimizes $\sum_j a_{\psi^{-1}(j),j}$ among all the stable matchings ψ, and reaches the unique minimum. Thus the stable matchings produced by the two algorithms are exactly the same. By Corollary 7.3.1, there exists a unique stable matching. □

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):

$$\begin{array}{llll} \text{for } a: & x > y > z & \text{for } x: & c > b > a \\ \text{for } b: & y > x > z & \text{for } y: & a > b > c \\ \text{for } c: & y > x > z & \text{for } z: & c > a > b \end{array}$$

Find the stable matchings that will be produced by the men-proposing and by the women-proposing Gale-Shapley algorithm.

8

Random-turn and auctioned-turn games

In Chapter 1 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 [**?**] and [**?**].

8.1 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 $x_1, \ldots, x_k$.

Claim:

$$u(x) = \frac{1}{2}\Big(\max_i \{u(x_i)\} + \min_j \{u(x_j)\}\Big). \tag{8.1}$$

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

$$u_{\mathrm{I}}(x) = \begin{cases} \sup_{S_{\mathrm{I}}} \{\inf_{S_{\mathrm{II}}}\{\mathbb{E} f(X_\tau)\}\}, & \text{if } \tau < \infty \\ -\infty, & \text{if } \tau = \infty. \end{cases}$$

Likewise, let

$$u_{\mathrm{II}}(x) = \begin{cases} \inf S_{\mathrm{II}} \left\{\sup_{S_{\mathrm{I}}}\{\mathbb{E} f(X_\tau)\}\right\}, & \text{if } \tau < \infty \\ +\infty, & \text{if } \tau = \infty. \end{cases}$$

Then both u_{I} and u_{II} satisfy (8.1).

We call functions satisfying (8.1) "infinity-harmonic". In the original paper by Lazarus, Loeb, Propp, and Ullman, [**?**] they were called "Richman functions".

8.2 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:

8.2.1 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.

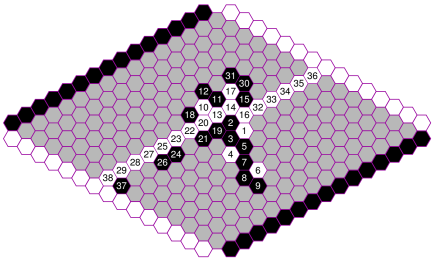

Fig. 8.1. A game between a human player and a program by David Wilson on a 15×15 board.

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.

8.2.2 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. 8.2. The game of random-turn Bridgit 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.

8.2.3 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 $f^*(S_1) = \text{sign}(f(S_1))$.)

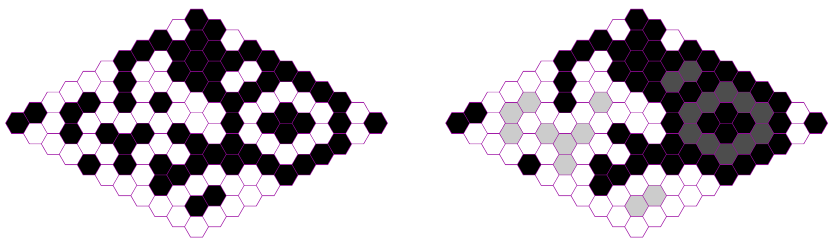

Fig. 8.3. A completed game of Surround before recoloring surrounded territory (on left), and after recoloring (on right). 10 black spaces were recolored white, and 12 white spaces were recolored black, so $f(S_1) = 2$.

8.2.4 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 $S \backslash S_1$. This is different from ordinary tic-tac-toe in that the game does not end after the first line is completed.

8.2.5 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

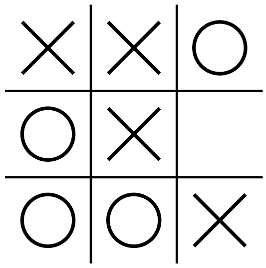

Fig. 8.4. Random-turn tic-tac-toe played out until no new rows can be constructed. $f(S_1) = 1$.

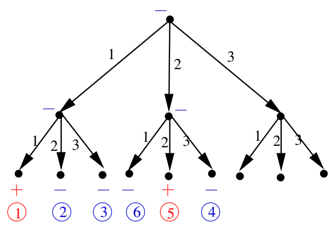

Fig. 8.5. Here player II wins; the circled numbers give the order of the moves.

$S_1(j+1)$. The payoff function $f(S_1)$ for the recursive three-fold majority is -1 if $S_1(0) = \varnothing$ and $+1$ if $S_1(0) = \{\text{root}\}$.

8.2.6 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.

8.3 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 $T_3 = S \backslash (T_1 \cup T_2)$ 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

$$E(T_1, T_2) = \frac{1}{2}(\max_{s \in T_3} E(T_1 \cup \{s\}, T_2) + \min_{t \in T_3} E(T_1, T_2 \cup \{t\}).$$

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 $\tilde{T} \cap T_2 = \varnothing$ we have that $f(\tilde{T}) = C$, then $E(T_1, T_2) = C$.

Theorem 8.3.1. *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 $\mathbb{E}[f(T)] = E(\varnothing, \varnothing)$. 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 $M(\varnothing, \varnothing)$ 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. □

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 8.3.2. *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 8.3.1 and Theorem 8.3.2 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 8.3.1 and Theorem 8.3.2 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.

8.4 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 8.4.1. *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 [**?**]:

$$\begin{aligned}\sum_i I_i(f) &= \mathbb{E}\left[\sum_i f(\vec{x})x_i\right] = \mathbb{E}\left[f(\vec{x})\sum_i x_i 1_{x_i \text{ examined}}\right] \\ &\le \text{(by Cauchy-Schwarz)}\sqrt{\mathbb{E}[f(\vec{x})^2]\,\mathbb{E}\left[\left(\sum_{i:\ x_i \text{ examined}} x_i\right)^2\right]} \\ &= \sqrt{\mathbb{E}\left[\left(\sum_{i:\ x_i \text{ examined}} x_i\right)^2\right]} = \sqrt{\mathbb{E}[\#\text{ bits examined}]}\,. \end{aligned} \tag{8.2}$$

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 (8.2) we have

$$\mathbb{E}[\#\text{ turns}] \ge \left[\sum_i I_i(f)\right]^2.$$

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.

8.4.1 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.

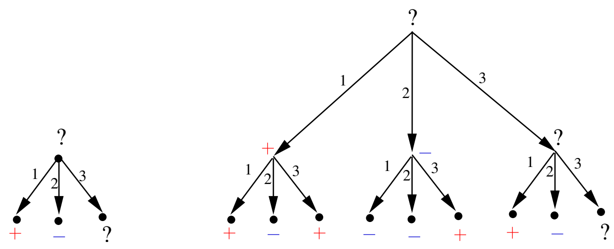

Fig. 8.6.

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 $I_i(f) = \left(\frac{1}{2}\right)^h$. By symmetry this is the same for every leaf.

We now use (8.2) to produce the bound:

$$\mathbb{E}[\#\text{ turns}] \geq \left[\sum_i I_i(f)\right]^2 = \left(\frac{3}{2}\right)^{2h}.$$

8.5 Richman games

Richman games were suggested by the mathematician David Richman, and analyzed by Lazarus, Loeb, Propp, and Ullman in 1995 [**?**]. 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.

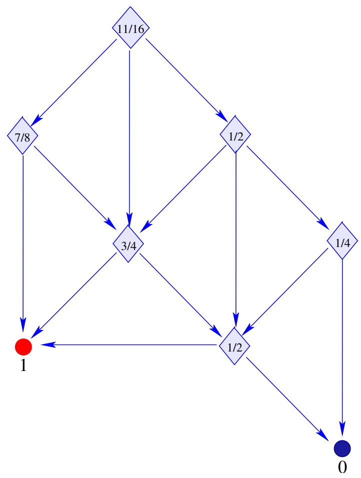

Fig. 8.7.

Theorem 8.5.1. *Suppose Blue has \$$x$, Red has \$$y$, and the current position is v. If*

$$\frac{x}{x+y} > R(v) \tag{8.3}$$

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 (8.3) 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, $R(w) \geq R(z)$. $\square$

Corollary 8.5.1. *If* (8.3) *holds at the beginning of the game, Blue has a winning strategy.*

Proof. When Blue loses, $R(v) = 1$, but $\frac{x}{x+y} \leq 1$. $\square$

Corollary 8.5.2. *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 (8.3) holds.

Exercises

8.1 Generalize the proofs of Theorem 8.3.1 and Theorem 8.3.2 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.

8.2 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 deapth h. For each leaf, determine the probability that flipping the sign of that leaf would change the overall result.

8.3 Even if y is unknown, but (8.3) holds, Blue still has a winning strategy, which is to bid

$$\left(1 - \frac{R(u)}{R(v)}\right) x.$$

Prove this.

8.6 Additional notes on random-turn Hex

8.6.1 *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.

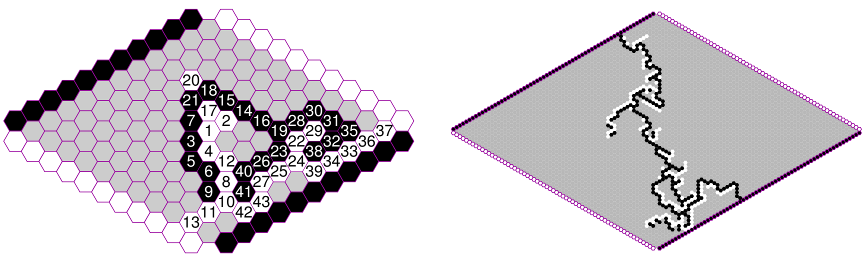

Fig. 8.8. Random-turn Hex on boards of size 11×11 and 63×63 under (near) optimal play.

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 8.8 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 $L^{1.1306\pm0.0003}$ [**?**], where the exponent is known to be strictly larger than 1. We give a stronger lower bound:

Theorem 8.6.1. *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 (8.2), 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 [**?**] (as predicted earlier in [**?**]), 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 $\mathbb{E}[\#\text{ turns}] \geq L^{3/2+o(1)}$. □

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 8.9 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.

8.7 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

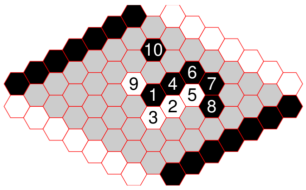

Fig. 8.9. A rare occurrence — a game of random-turn Hex under (near) optimal play with a disconnected play.

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 [**?**], 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.