

MATH 499, GAME THEORY, SPRING 2025

STEVEN HEILMAN

ABSTRACT. These notes partly follow the book of Karlin and Peres, available [here](#).

CONTENTS

1. Introduction	2
2. Combinatorial Games	3
2.1. Impartial Games	5
2.2. Chomp	7
2.3. Nim	8
2.4. Sprague-Grundy Theorem	10
2.5. Partisan Games	10
2.6. Hex	11
3. Two-person zero-sum games	13
3.1. Von Neumann's Minimax Theorem	16
3.2. Domination	21
4. General Sum Games	22
4.1. Nash equilibria	24
4.2. Correlated equilibria	26
4.3. Fixed Point Theorems	29
4.4. Games with More than Two Players	32
4.5. Nash's Theorem	33
4.6. Evolutionary Game Theory	34
4.7. Signaling and Asymmetric Information	37
5. Price of Anarchy	39
6. Coalitions and Shapley Value	41
7. Mechanism Design	47
7.1. Auctions	47
8. Social Choice	52
8.1. Fourier analysis and influences	52
8.2. Noise stability	54
8.3. Arrow's Impossibility Theorem	56
9. Quantum Games	59
10. Algorithmic Game Theory	62
11. Reinforcement Learning	65
11.1. Bandits	65
11.2. Markov Decision Process	70

12. Appendix: Hoeffding's Inequality	71
13. Appendix: Notation	73

1. INTRODUCTION

Game theory concerns strategies for interactions of multiple players in games. We will never concern ourselves with single-person games. And given a game, we will often ask for the best possible strategy that each player can have. Making precise the meaning of “best” will be a challenge itself.

1.0.1. *Some history.* Most people say that game theory began with von Neumann's 1928 paper in which he proved his Minimax Theorem via fixed point theorems. von Neumann was partially motivated to find the optimal strategy in poker, though he was interested in other applications as well. The subject was reinvigorated with von Neumann and Morgenstern's 1944 book, “Theory of Games and Economic Behavior.” Nash published his first papers on the subject in 1950.

1.0.2. *Some comments about Hollywood depictions of game theory.* In the movie “A Beautiful Mind,” there is one scene where some game theory is demonstrated. However, the game theoretic content of this scene is questionable, and the scene does not accurately depict Nash's insights. The situation is presented as follows. Four men are trying to pick up women at a bar. One woman and four of her friends are also at the bar. The first woman is supposed to be more desirable than the others. If more than one man talks to the first woman, she will decide to talk to no one, and the other women will not want to talk to anyone either. Fictional Nash says that it will be best for all of the men to not talk to the first woman, and instead split up and talk to the other women. This may be beneficial for all of the men, but is it realistic, and is it beneficial to each individual man?

This strategy is not really the best for each man. If every man talks to the four woman friends, then no one is talking to the first woman. So, if one man changes his mind and talks to the first woman, he can improve his outcome. That is, if the men are really not cooperating with each other, it is in their best interest to talk to the first woman. This realization is essential in defining the Nash equilibrium for non-cooperative games. However, we could also argue that all of the men trying to talk to the first woman is not realistic either, due to the “wingman” concept that many people use. So, it seems that both strategies (no man talking to the first woman, versus every man talking to the first woman), although justifiable from a certain perspective, may not be realistic.

Such is the situation with game theory, and more generally with applications of mathematics. We can make some assumptions, and we can prove some theorems. But maybe the assumptions are not quite correct, and maybe the predictions are not that accurate.

For some more discussion and videos, look [here](#).

1.0.3. *Overview.* This course will begin with combinatorial games. In these games, two players take turns making moves until a winning position for one player is reached. In this setting, we will look for **winning strategies**. A winning strategy is a collection of moves for one player, for any possible game position, which guarantees that player wins. Some popular combinatorial games include checkers, chess and Go.

We will then discuss **two-person zero-sum games**, where both players move simultaneously, and each player benefits at the expense of the other. For example, consider rock-paper-scissors. In this case, we will look for an optimal strategy, which will typically be a random choice of available moves. For example, always choosing rock in rock-paper-scissors is not a very good strategy.

We then move on to **general-sum games**. In this setting, an optimal strategy may not exist, but there can still be some logical choice to be made. A **Nash equilibrium** is a set of strategies for each player of the game such that no player can gain something by unilaterally changing their strategy. For example, in the Hollywood depiction discussed above, if the four men do not talk to the first woman, then we are not in a Nash equilibrium, since one man can improve his outcome by talking to the first woman. A Nash equilibrium occurs when all the men talk to the first woman.

In Nash's Theorem, we show that every general-sum game has at least one Nash equilibrium. This Theorem uses another deep theorem known as the **Brouwer fixed point theorem**. The latter theorem is a statement about continuous functions on closed convex sets in Euclidean space.

We will then discuss a bit of **cooperative game theory**, where players form coalitions to achieve a common goal. In the context of the Hollywood depiction above, perhaps this situation is most realistic. It is probably beneficial (and realistic) when some men work as a "wingman" for the others. The notion of **Shapley value** helps us to determine the influence of each player in an optimal outcome.

In **mechanism design**, we switch our perspective from the player of the game, to the designer of the game. We look for a way to design a game such that some optimal outcome occurs when people act selfishly. Auctions provide a nice example of games with **incomplete information**, though a general theory of incomplete information can become quite complicated (the Maschler, Solan and Zamir book has a few chapters on this topic). Or, in **social choice** we try to design a voting method that is fair to everyone, or that is stable to some corruption of the votes.

Finally, we discuss **quantum games**, in which players can share some random (quantum) information.

1.0.4. *Preliminaries.* Since we will use certain examples of games often, the reader is expected to be familiar with: **connect four**, **checkers**, **chess** and **go**.

2. COMBINATORIAL GAMES

Definition 2.1. A **combinatorial game** consists of two players, a set of game positions, and a set of legal moves between game positions. Some game positions are denoted as terminal. A terminal position is denoted as a winning position for exactly one of the two players. The players take turns moving from position to position. The goal for each player is to reach their winning position.

More formally, the set of game positions of a game is a finite set X . We have two players denoted by I and II . The set of legal moves for I is a subset $\Omega_I \subseteq X \times X$, so that $(x, y) \in \Omega_I$ if and only if moving from position x to position y is a legal move for player I . Similarly, The set of legal moves for II is a subset $\Omega_{II} \subseteq X \times X$. The game has some designated opening board position $x_0 \in X$, and player I makes the first legal move. A **(memoryless) strategy** for player I is a function $f: X \rightarrow X$ such that $(x, f(x)) \in \Omega_I$ for all $x \in X$. So, if player I

is presented with game position x , they then use the legal move from x to $f(x)$. Similarly, a (memoryless) strategy for player II is a function $g: X \rightarrow X$ such that $(x, g(x)) \in \Omega_{II}$ for all $x \in X$.

We say that the strategy is memoryless since the strategy only depends on the current game position, rather than all previous game positions. It is also possible to consider strategies which take into account all previous game positions, but we will not need to do so. Unless otherwise stated, all strategies discussed below will be memoryless.

Definition 2.2. A combinatorial game is **impartial** if both players have the same available moves, that is if $\Omega_I = \Omega_{II}$. A combinatorial game is **partisan** if it is not impartial, i.e. if $\Omega_I \neq \Omega_{II}$.

Example 2.3. Connect four, chess, checkers and Go are all partisan, since different players move different game pieces. That is, different players can change the game position in different ways.

Remark 2.4. Some partisan games can end in a **tie**, which is a terminal position in which neither player wins. Also, some combinatorial games can be **drawn** or go on forever. For example, chess can end in a draw if both players only have a king remaining.

Definition 2.5. In a combinatorial game, a **winning strategy** is a strategy for one of the players guaranteeing a win for that player in a finite number of steps.

If player I has a winning strategy for a game, then beginning from the start of the game, no matter what player II does, player I will *always win*. So, having a winning strategy is quite a strong condition.

For combinatorial games, we concern ourselves with the following two questions.

Question 2.6. Can one player always win? That is, does a winning strategy exist for one of the players?

Question 2.7. If a winning strategy exists, can we describe the winning strategy?

Remark 2.8. In a combinatorial game, if one player has a winning strategy, then the other player does not.

Example 2.9. In the game Connect Four, the first player always has a winning strategy. This strategy is fairly complicated to describe, and it is essentially derived from a case-by-case analysis. However, it is known that the best move for the first player is to place their token in the center column.

Example 2.10. For the game checkers, it was shown in 2007 that each of the two players has a strategy guaranteeing at least a draw. A draw occurs in checkers when no piece is captured after forty moves. That is, the first player has a strategy that either wins or achieves a draw. And similarly for the second player. However, actually describing these strategies in simple terms seems difficult if not impossible. This strategy was only found by a brute force enumeration of all possible game positions.

For chess and Go, we still do not know whether or not the first player has a winning strategy. However, we can make a weaker statement:

Theorem 2.11. *In chess, exactly one of the following situations is true:*

- *White has a winning strategy. (That is, for any game of chess, white has at least one strategy guaranteeing a win.)*
- *Black has a winning strategy. (That is, for any game of chess, black has at least one strategy guaranteeing a win.)*
- *Each of the two players has a strategy guaranteeing at least a draw.*

We will eventually derive this result from a more general statement, which is sometimes called Zermelo's Theorem.

Definition 2.12. A combinatorial game with a position set X is said to be **progressively bounded** if, starting from any position $x \in X$, the game will terminate after at most $B(x)$ moves.

Note that the standard Connect Four game is played on a 7×6 board, and each legal move involves placing a disk on the board, so Connect Four will always terminate after at most 42 moves. So, we can choose $B(x) = 42$ for all game positions x , and Connect Four is progressively bounded. For chess, checkers and go it may seem like these games could go on forever, but there are rules in place to ensure that these games are progressively bounded.

2.1. Impartial Games. Recall the following definition.

Definition 2.13. An **impartial combinatorial game** is a combinatorial game with two players who both have the same set of legal moves. A terminal position is a position from which there are no legal moves. Every non-terminal position has at least one legal move. Under **normal play**, the player who moves to the terminal position wins.

We can think of the game positions of an impartial combinatorial game as **nodes**, where a move from position x to position y is a directed **link** from x to y . In this way, the collection of nodes (i.e. **vertices**) together with the links (i.e. **edges**) form a graph. If (x, y) and (y, x) are both legal moves, we can treat the edge (x, y) as an undirected edge. The graph formed in this way can help us visualize the game at hand. A game presented in the form of a graph is sometimes called a game in **extensive form**.

Example 2.14 (A Subtraction Game). Consider the following impartial combinatorial game. Let x_0 be a nonnegative integer. The game begins with a pile of x_0 chips. Each player alternately removes between 1 to 4 chips from the pile. The player who removes the last chip wins.

First, note that since at least one chip is removed from the pile on each turn, the game is progressively bounded with $B(x) = x$ for any game position x . (We label the game positions by the number of chips x in the pile.)

Now, let's try to look for winning strategies in this game. If the game starts with 0 chips, then the first player has lost. If the game starts with 1 to 4 chips, then the first player can immediately win by removing all of the chips. If the game starts with 5 chips, then the first player must leave between 1 to 4 chips remaining, so the second player can win by removing these chips. If the game starts with 6 chips, then the first player can remove 1 chip, leaving 5 chips. So, from the previous case, after the second player takes their turn, the first player can win. Similarly if the game starts with 7, 8 or 9 chips, the first player can leave 5 chips remaining.

To continue our analysis, we define

$$\mathbf{N} = \left\{ \begin{array}{l} \text{positive integers } x \text{ such that the first ("next") player can ensure a win,} \\ \text{if there are } x \text{ chips at the start of the game.} \end{array} \right\}.$$

$$\mathbf{P} = \left\{ \begin{array}{l} \text{positive integers } x \text{ such that the second ("previous") player can ensure a win,} \\ \text{if there are } x \text{ chips at the start of the game.} \end{array} \right\} \cup \{0\}.$$

So, far, we have reasoned that $\{0, 5\} \subseteq \mathbf{P}$ and $\{1, 2, 3, 4, 6, 7, 8, 9\} \subseteq \mathbf{N}$. Continuing our reasoning, we can prove by induction that \mathbf{P} is the set of all nonnegative integers divisible by 5, and \mathbf{N} consists of all other positive integers.

We can extend the definitions of \mathbf{N} and \mathbf{P} to any impartial combinatorial game as follows.

Definition 2.15. For any impartial combinatorial game, let \mathbf{N} (for “next”) be the set of game positions such that the first player to move can guarantee a win. Let \mathbf{P} denote the set of game positions such that *any* legal move leads to a position in \mathbf{N} . We also let \mathbf{P} contain all terminal positions. Note that if a player makes a legal move into a \mathbf{P} position, then they can guarantee a win, since they will encounter an \mathbf{N} position on their next turn.

Since only one player can win the game, \mathbf{N} and \mathbf{P} are always disjoint sets.

In the Subtraction Game discussed above, $\mathbf{N} \cup \mathbf{P}$ is the set of all nonnegative integers. That is, all game positions are either in \mathbf{N} or in \mathbf{P} . This statement holds for any progressively bounded impartial game, as we show below using a method similar to that used above. Starting from the terminal positions, we work our way backwards to find which player has the winning strategy at each position. Such a proof strategy is sometimes called **backward induction**.

Theorem 2.16. *In any progressively bounded impartial game under normal play, any game position x lies in $\mathbf{N} \cup \mathbf{P}$.*

Proof. We induct on $B(x)$, where $B(x)$ is equal to the maximum number of moves it takes to reach a terminal position when starting from the game position x .

We first consider the base case of induction. In the case $B(x) = 0$, the position x is terminal, so that $x \in \mathbf{P}$, by the definition of \mathbf{P} . So, the base case is done.

We now check the inductive step. Let n be a nonnegative integer such that the Theorem holds whenever $B(x) \leq n$. We need to consider the case $B(x) = n + 1$. Note that any legal move from x will satisfy the inductive hypothesis, by the definition of $B(x)$. That is, any legal move from x lies in $\mathbf{N} \cup \mathbf{P}$. We split into two cases.

Case 1. Each move from x leads to a position in \mathbf{N} . In this case, $x \in \mathbf{P}$, by the definition of \mathbf{P} .


Case 2. There is some move from x that leads to a position y such that y is not in \mathbf{N} . Since $B(y) \leq n$, we know that $y \in \mathbf{N} \cup \mathbf{P}$ by the inductive hypothesis. Since $y \notin \mathbf{N}$, we have $y \in \mathbf{P}$. But then by the definition of \mathbf{P} , either y is a terminal position (in which case $x \in \mathbf{N}$), or any move away from y must lie in \mathbf{N} . So, by the definition of \mathbf{N} , we conclude that $x \in \mathbf{N}$. (If we are in game position x , the winning strategy begins by moving to y .)

In any case, $x \in \mathbf{N} \cup \mathbf{P}$. With the induction complete, we conclude that all game positions lie in $\mathbf{N} \cup \mathbf{P}$. \square

2.2. Chomp.

Example 2.17 (Chomp). The game board consists of a finite rectangular grid of squares. Alternately, each player removes a square, and then removes all squares above and to the right of this square. One player wins when the other player removes the bottom left corner of the board. So, the only terminal position occurs when all of the squares are removed.

Below is an example of a sequence of game positions of Chomp played on a 4×4 game board.



Note that this game is progressively bounded, since each move removes at least one square, and the game begins with a finite number of squares. So, from Theorem 2.16, for a given starting position, either player one or player two has a winning strategy. We will in fact show that the first player always has a winning strategy. In fact, if the game position is rectangular at any time, then the next person to move has a winning strategy. We will prove this using a **strategy-stealing** argument.

Proposition 2.18. *Suppose a game of Chomp begins with a finite rectangle of size greater than 1×1 . Then the next player to move has a winning strategy.*

Proof. Denote the rectangle by R . As part of a possible strategy, the next player to move will consider removing the top right corner of the rectangle. Let R_0 denote the resulting arrangement of squares. (Since R is larger than 1×1 , we know R_0 is nonempty.) We split into two cases, which cover all possible cases by Theorem 2.16. (Since at least one square is removed at each turn, the game ends in a finite number of steps, so the game is progressively bounded.)

Case 1. $R_0 \in \mathbf{P}$. In this case, by the definition of \mathbf{N}, \mathbf{P} , we know that $R \in \mathbf{N}$ and part of a winning strategy is to move from R to R_0 .

Case 2. $R_0 \in \mathbf{N}$. If $R_0 \in \mathbf{N}$, then there is some move from R_0 to S , where $S \in \mathbf{P}$. However, any move from R_0 to S could have also been achieved by moving from R to S . Specifically, any move from R_0 to S will start by choosing a square not in the upper right corner. But the upper right corner would have been removed if it were present during the move from R_0 to S . Since $S \in \mathbf{P}$ and since it is possible to move from R to S , we conclude that $R \in \mathbf{N}$.

In any case, $R \in \mathbf{N}$, as desired. □

Remark 2.19. Strictly speaking, Theorem 2.16 does not apply to Chomp, since in Chomp the player who moves to the terminal position loses. However, the conclusion of Theorem 2.16 still holds, with essentially the same proof, for a progressively bounded impartial game where the player who moves to the terminal position loses.

Remark 2.20. This proof has not explicitly described the winning strategy. The proof has only shown the *existence* of a winning strategy. It is an open research problem to explicitly describe the winning strategy for Chomp (for starting game positions of arbitrary size).

2.3. Nim.

Example 2.21 (Nim). The game of Nim is an impartial combinatorial game. In the game of Nim, there are a finite number of piles, each with a finite number of chips. A legal move consists of removing any positive number of chips from a single pile. The two players alternate taking turns. The goal is to remove the last chip. So the only terminal position of the game occurs when no chips are left.

Since the game begins with a finite number of chips, and at least one chip is removed during each turn, the game of Nim is progressively bounded. From Theorem 2.16, all game positions are therefore in \mathbf{N} or in \mathbf{P} . Below, we will describe a way to find who has the winning strategy. For now, let's consider some examples. We denote a game position in nim as (n_1, \dots, n_k) . With this notation, the game position has k piles of chips, where the first pile has n_1 chips, the second pile has n_2 chips, and so on.

If the game begins with exactly one pile with a positive number of chips, then the first player can win by taking all of the chips. If the game position is (n_1, n_2) with $n_1 \neq n_2$, then $(n_1, n_2) \in \mathbf{N}$, since the next player can force both piles to have an equal number of chips. This player can continue to make both piles have an equal number of chips, resulting in a win, no matter what the other player does. So, if $n_1 = n_2$, we see that $(n_1, n_2) \in \mathbf{P}$. Dealing with more than two piles becomes a bit more complicated.

Lemma 2.22. *Let j, k be positive integers. Let $x = (x_1, \dots, x_j)$ and let $y = (y_1, \dots, y_k)$ be game positions in Nim. We denote $(x, y) := (x_1, \dots, x_j, y_1, \dots, y_k)$ as another game position in Nim.*

- (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 if $x \in \mathbf{N}$ and $y \in \mathbf{P}$), then $(x, y) \in \mathbf{N}$.*
- (iii) *If x and y are in \mathbf{N} , then (x, y) can be in either \mathbf{N} or \mathbf{P} .*

Proof. We prove (i) and (ii) simultaneously by induction on the number of chips in (x, y) . If (x, y) has zero chips, then both x and y have zero chips, so x, y and (x, y) are all terminal positions. So, x, y and (x, y) are all in \mathbf{P} by definition of \mathbf{P} .

Now, fix a nonnegative integer n , and assume that (i) and (ii) both hold whenever (x, y) has at most n total chips. For the inductive step, we consider when (x, y) has $n + 1$ total chips.

If $x \in \mathbf{P}$ and $y \in \mathbf{N}$, then we describe the winning strategy for (x, y) . The next player should perform the winning strategy on the piles corresponding to y , resulting in a game position (x, y') , where $y' \in \mathbf{P}$. Then (x, y') has at most n total chips, so by the inductive hypothesis, $(x, y') \in \mathbf{P}$. Since we have found a way to move (x, y) into a position in \mathbf{P} , we conclude that $(x, y) \in \mathbf{N}$.

If $x \in \mathbf{P}$ and $y \in \mathbf{P}$, then the next player to move must take chips from one of the piles. Without loss of generality, the chips are taken from a pile corresponding to y . The resulting game position is then (x, y') where $y' \in \mathbf{N}$, since we removed chips from $y \in \mathbf{P}$. Then (x, y') has at most n total chips, so by the inductive hypothesis, $(x, y') \in \mathbf{N}$. Since any move from (x, y) results in a position in \mathbf{N} , we conclude that $(x, y) \in \mathbf{P}$.

So, items (i) and (ii) are proven by induction on n . To see item (iii), note that any single pile of chips is in \mathbf{N} , whereas $(1, 1) \in \mathbf{P}$ and $(1, 2) \in \mathbf{N}$, as we discussed above. \square

Definition 2.23 (Nim Sum). Let m, n be nonnegative integers. The **nim-sum** of m, n is defined as the following binary number. Write m and n as binary numbers. Sum the digits

modulo 2, producing the digits of a new binary number. The resulting binary number is called the nim-sum of m, n . We denote the nim-sum of m, n by $m \oplus n$. Note that $m \oplus n = n \oplus m$.

Example 2.24. Let's compute the nim-sum of 5 and 9. We have $5 = 1 \cdot 2^2 + 0 \cdot 2^1 + 1 \cdot 2^0$, so that 5 is written as 101 in binary. Also, $9 = 1 \cdot 2^3 + 0 \cdot 2^2 + 0 \cdot 2^1 + 1 \cdot 2^0$, so that 9 is written as 1001 in binary. So, in binary $5 \oplus 9$ is equal to $0101 \oplus 1001 = 1100$, since $0 + 1 = 1$, $1 + 0 = 0$, $0 + 0 = 0$ and $1 + 1 = 0$ modulo 2.

More generally, given any nonnegative integers x_1, \dots, x_j , where $j > 0$, we define their nim sum to be

$$x_1 \oplus \dots \oplus x_j := x_1 \oplus (x_2 \oplus \dots \oplus (x_{j-1} \oplus x_j)).$$

Exercise 2.25. Let a, b be nonnegative integers. Then $a \oplus a = 0$ and $(a \oplus b) \oplus 0 = a \oplus b$.

The following Theorem gives a complete characterization of who has a winning strategy in Nim, according to the nim-sum of the piles.

Theorem 2.26 (Bouton's Theorem). *Let j be a positive integer. A Nim position $x = (x_1, \dots, x_j)$ is in \mathbf{P} if and only if $x_1 \oplus \dots \oplus x_j$ is zero.*

Proof. Let Z denote the set of game positions with zero nim-sum. Suppose $x = (x_1, \dots, x_j) \in Z$, so that $x_1 \oplus \dots \oplus x_j = 0$. If all piles of x are zero, then $x \in \mathbf{P}$. Otherwise, some pile $k \in \{1, \dots, j\}$ has a nonzero number of chips, i.e. $x_k \neq 0$. Suppose we remove some chips from the k^{th} pile leaving $x'_k < x_k$ chips in that pile. Since $x'_k < x_k$, we have $x'_k \oplus x_k \neq 0$. So, using Exercise 2.25, the nim-sum of the resulting piles is then

$$\begin{aligned} x_1 \oplus \dots \oplus x_{k-1} \oplus x'_k \oplus x_{k+1} \oplus \dots \oplus x_j \\ &= [x_1 \oplus \dots \oplus x_{k-1} \oplus x'_k \oplus x_{k+1} \oplus \dots \oplus x_j] \oplus 0 \\ &= [x_1 \oplus \dots \oplus x_{k-1} \oplus x'_k \oplus x_{k+1} \oplus \dots \oplus x_j] \oplus [x_1 \oplus \dots \oplus x_j] = x'_k \oplus x_k \neq 0. \end{aligned}$$

So, any move from a position in Z results in a position not in Z .

Now, suppose $x = (x_1, \dots, x_j) \notin Z$, so that $s := x_1 \oplus \dots \oplus x_j \neq 0$. Since $s \neq 0$, there is a leftmost binary digit of s which is nonzero. We label this digit as digit ℓ . Since this digit is nonzero, there are an odd number of values of $i \in \{1, \dots, j\}$ such that the ℓ^{th} binary digit of x_i is 1. Choose one such value of i . Then the ℓ^{th} binary digit of $s \oplus x_i$ is zero. And since the binary digits $\ell + 1, \ell + 2, \dots$ are all zero for s , the digits $\ell + 1, \ell + 2, \dots$ of $s \oplus x_i$ are identical to those digits of x_i . In summary, $s \oplus x_i < x_i$. Consider the move where the next player removes $x_i - (x_i \oplus s)$ chips from the i^{th} pile, leaving $x_i \oplus s$ chips in the pile. The nim-sum of the resulting position is

$$x_1 \oplus \dots \oplus x_{i-1} \oplus (x_i \oplus s) \oplus x_{i+1} \oplus \dots \oplus x_j = s \oplus s = 0.$$

So, the new game position lies in Z . So, from any position $x \notin Z$, there is a legal move from x to a position in Z .

In summary, from any game position not in Z , the first player can use the strategy of moving to a position in Z . Presented with a position in Z , the second player, if they have any moves, has no choice but to move to a position not in Z , always allowing the first player to make another move. (After the second player moves, the nim sum is nonzero, so there are a nonzero number of chips in the game position.) So, any position not in Z is in \mathbf{N} . And from any position in Z , the next player to go will always move into a position not in Z , i.e. in \mathbf{N} . So, any position in Z is in \mathbf{P} . \square

2.4. Sprague-Grundy Theorem.

Definition 2.27 (Sum of Games). Let G_1, G_2 be two combinatorial games. The **sum** of G_1 and G_2 , denoted $G_1 + G_2$, is another combinatorial game. In this game each player, in their turn, chooses one of the games G_1, G_2 to play, and they make one move in the chosen game. A game position in $G_1 + G_2$ is an ordered pair (x_1, x_2) where x_i is a game position in G_i for each $i \in \{1, 2\}$. The terminal positions in the sum game are (t_1, t_2) , where t_i is a terminal position in G_i for each $i \in \{1, 2\}$.

Example 2.28. If x and y are two game positions of Nim, then the sum of x and y is the same as (x, y) .

Exercise 2.29. Let G_1, G_2 be games. Let x_i be a game position for G_i , and let $\mathbf{N}_{G_i}, \mathbf{P}_{G_i}$ denote \mathbf{N} and \mathbf{P} respectively for the game G_i , for each $i \in \{1, 2\}$.

- (i) If $x_1 \in \mathbf{P}_{G_1}$ and if $x_2 \in \mathbf{P}_{G_2}$, then $(x_1, x_2) \in \mathbf{P}_{G_1+G_2}$.
- (ii) If $x_1 \in \mathbf{P}_{G_1}$ and if $x_2 \in \mathbf{N}_{G_2}$, then $(x_1, x_2) \in \mathbf{N}_{G_1+G_2}$.
- (iii) If $x_1 \in \mathbf{N}_{G_1}$ and if $x_2 \in \mathbf{N}_{G_2}$, then (x_1, x_2) could be in either $\mathbf{N}_{G_1+G_2}$ or $\mathbf{P}_{G_1+G_2}$.

Definition 2.30 (Equivalent Games). Let G_1, G_2 be progressively bounded combinatorial games with starting positions x_1, x_2 , respectively. We say that G_1 and G_2 are **equivalent** if: for any progressively bounded combinatorial game G_3 with position x_3 , the outcome of (x_1, x_3) in $G_1 + G_3$ (i.e. whether it is in \mathbf{N} or \mathbf{P}) is the same as the outcome of (x_2, x_3) in $G_2 + G_3$.

Theorem 2.31 (Sprague-Grundy Theorem). *Let G be a progressively bounded impartial combinatorial game under normal play with starting position x . Then G is equivalent to a single Nim pile of size $g(x) \geq 0$, where $g(x)$ is a function whose domain is the set of game positions of G .*

So, in a certain sense, any progressively bounded impartial combinatorial game under normal play is easy to understand, since it is more or less the same as Nim. Having a good understanding of impartial games, we now move on to partisan games.

2.5. Partisan Games. Recall that a partisan game is a combinatorial game that is not impartial.

As our first task, we will partially extend Theorem 2.16 to the case of impartial games. In fact, the proof below resembles that of Theorem 2.16, though we are forced to use more notation in Theorem 2.32.

Theorem 2.32 (Zermelo's Theorem). *In any progressively bounded combinatorial game with no ties allowed, one of the players has a memoryless winning strategy.*

Proof. Suppose the game has players I and II and a set X of game positions. If x is a game position and the next player to have a turn is $i \in \{I, II\}$, we say that (x, i) is the **state of the game**. We will recursively define a function W which determines the winner of the game, given the state of the game. Let $i \in \{I, II\}$, let $o(i)$ denote the opponent of the player i , and let $B(x, i)$ be the maximum number of moves that it takes to reach a terminal point in the game, starting from the game state (x, i) . We let $S_i(x)$ denote the set of game states that result from a legal move away from (x, i) .

If $B(x, i) = 0$, we let $W(x, i)$ be the player who wins from the terminal position x .

Suppose that we have inductively defined $W(x, i)$, whenever $B(x, i) < k$, where k is a positive integer. Let (x, i) be a game state with $B(x, i) = k$. Let $(y, o(i))$ be any game state resulting from a legal move away from (x, i) . Then $B(y, o(i)) < k$, so that $W(y, o(i))$ is defined. We split into two cases.

Case 1. There is some game state $(y, o(i)) \in S_i(x)$ such that $W(y, o(i)) = i$. In this case, we define $W(x, i) = i$, since player i can move to state y , from which they can win (since $W(y, o(i)) = i$). So, any such state $(y, o(i))$ is a winning position for player i .

Case 2. For all game states $(y, o(i)) \in S_i(x)$, we have $W(y, o(i)) = o(i)$. In this case, we define $W(x, i) = o(i)$. If player i moves from x to any game position y , player $o(i)$ will win (since $W(y, o(i)) = o(i)$ for all $(y, o(i)) \in S_i(x)$).

We have completed the inductive step. So, the function W exists, which exactly specifies which player has a winning strategy at each game state. \square

Exercise 2.33. By imitating the proof of Theorem 2.32, prove Theorem 2.11.

Remark 2.34. Theorem 2.11 also applies to Checkers, Connect Four, and Go. However, note that Theorems 2.16 and 2.11 only give existential statements about some strategies existing. Finding these strategies or describing them in simple terms is perhaps quite difficult. For example, at the moment, computers still cannot play as well as the best humans at Go. Though, computers are roughly at an equal level with the best chess players. And computer programs can run the winning strategy of the first player for connect four. But even for connect four, it is difficult to describe the optimal strategy even in a few pages of explanation.

2.6. Hex.

Example 2.35 (Hex). The game of Hex is a partisan combinatorial game, played on a rhombus-shaped board tiled with a finite number of hexagons, as in Figure 1. One set of parallel sides of the rhombus is colored blue, and the other set of parallel sides is colored yellow. One player is assigned the color blue, and the other player is assigned the color yellow. The players take turns coloring in white hexagons. The goal for each player is to connect their two opposing sides of the board with a path of hexagons of their color. So, the terminal positions for Hex are partial colorings of the game board with at least one path connecting two opposing sides of the board, or full colorings of the board.

Note that the legal moves and the terminal positions are different for both players in Hex. So, Hex is a partisan game.

It is not quite obvious, but it is impossible for Hex to end in a tie. A tie would occur if neither player had a crossing of hexagons of their color. More specifically, let's say we have reached a terminal position in Hex, and we then just proceed to fill in any remaining uncolored hexagons. If the Hex board is completely filled in with yellow and blue hexagons, then there is either a yellow crossing or a blue crossing, but not both. Consequently, exactly one player won the game, and there is no possibility for a tie. Below, we prove the weaker assertion that there is at least one crossing in a filled-in Hex board.

Theorem 2.36. *Consider a standard Hex board that is completely filled in with yellow and blue hexagons. Then there is at least one crossing between a pair of segments of the same color.*

Proof. Consider the opening game position, as in Figure 1. For every hexagon edge in between two hexagons of different colors, label this edge with an arrow. This arrow points in




FIGURE 1. A Starting Position in Hex

a direction such that the blue hexagon lies on the left of the arrow, and the yellow hexagon lies on the right of the arrow. So, two arrows point outside the game board (exit arrows), and two arrows point inside the game board (entry arrows).

Now, suppose the Hex board is filled in with blue and yellow hexagons. Starting from an entry arrow, we inductively create a path of arrows along the edges of the hexagons, such that blue always lies to the left of every arrow, and yellow always lies to the right of every arrow. Given an arrow A such that a blue hexagon lies to its left and a yellow hexagon lies to its right, the next arrow turns either to the right or the left, depending on the hexagon to which A points. If A points to a blue hexagon, the next arrow turns right. If A points to a yellow hexagon, the next arrow turns left. This procedure inductively defines the path of arrows.

This path of arrows cannot intersect itself. If the path did intersect itself, then a portion of the path would, e.g. be entirely surrounded by a path of blue hexagons. If this situation occurred, then one arrow would need to cross the path of blue hexagons. That is, one arrow would pass between two blue hexagons. However, this condition contradicts our definition of the arrows, in which a yellow hexagon must lie to its right. In conclusion, the path of arrows cannot create a loop.

Since there are only finitely many edges on the game board, and the path never intersects itself, the path must exit the game board using one of the two initially specified exit arrows. If the exit arrow and the entry arrow both touch the same yellow strip from the opening game position, then there is a blue crossing (which follows the path of arrows). If the exit arrow and the entry arrow both touch the same blue strip from the opening game position, then there is a yellow crossing (which follows the path of arrows). \square

Theorem 2.37. *On a standard Hex game board, the first player has a winning strategy.*

Proof. As discussed above, the game cannot end in a tie. So, by Theorem 2.32, one player has a (memoryless) winning strategy. We argue by contradiction. Suppose the second player has a winning strategy. Note that the game positions are symmetric with respect to swapping the colors of the game board. So, the first player can use the second player's winning strategy as follows. On the first move, the first player just colors in any hexagon H . On the first

player's next move, they just use the winning strategy of the second player (pretending that hexagon H is white). If the first player is ever required to color the hexagon H , they just color any other hexagon instead, if an uncolored hexagon exists. Having this extra hexagon colored can only benefit the first player. So, the first player is guaranteed to win, and both players will win, a contradiction. We conclude the first player has a winning strategy. \square

3. TWO-PERSON ZERO-SUM GAMES

Previously, we focused on games where players take turns, and the game position is known to each player. Players move deterministically, in that nothing is left to chance. In contrast, we now focus on games where two players move simultaneously. In this case, players with good strategies will rely on chance in their moves. In this chapter, we focus on **two-person zero-sum** games, in which one player loses what the other gains in every outcome. For example, recall that rock-paper-scissors is a two-person zero-sum game. And choosing moves deterministically, e.g. choosing to always play rock, is not a very good strategy.

Example 3.1 (Pick-a-hand, a betting game). Consider the following two-person zero-sum game. Player I is called the chooser, and player II is called the hider. The hider has two gold coins in their back pocket. At the beginning of the turn, they put their hand behind their back and either put one gold coin in their left hand, or two gold coins in their right hand. The chooser then picks a hand and wins any coins the hider has put there. The chooser may get nothing if the hand is empty, they may get one gold coin, or they may get two.

We record all possible outcomes of the game in a **payoff matrix**. The rows of this matrix are indexed by player I 's possible choices, and the columns of this matrix are indexed by player II 's possible choices. The matrix entry a_{ij} is the amount that player II loses to player I , when player I plays $i \in \{L, R\}$ and player II plays $j \in \{L, R\}$. The description of the game as a matrix is known as the **normal form** or **strategic form** of the game.

		hider	
		L	R
chooser	L	1 0	
	R	0 2	

We will now try to find the best strategy for each player. Suppose the hider wants to minimize their losses by always placing the coin in their left hand. This strategy is reasonable since the right hand will have a greater loss if the chooser finds it, so the hider will lose at most 1 coin during the game. But is there a better strategy? If the chooser somehow guesses or finds out about the hider's strategy, they will be able to get the coin. So, the success of the hider's strategy depends on how much information the chooser has.

The chooser could try to maximize their gain by always picking the right hand. But once again, if the hider somehow guesses or finds out about the chooser's strategy, then the hider can choose the left hand, so the chooser will gain nothing at all. Without knowing what the hider will do, the chooser may not gain anything by always choosing R .

In order to find a strategy that does not depend on the information that the players have, we will introduce uncertainty into the players' moves.

Definition 3.2. A **mixed strategy** is a strategy for a player that assigns some fixed probability to each move. A mixed strategy in which one particular move is played with probability one is called a **pure strategy**.

Definition 3.3. For the purposes of this section, a (discrete) **random variable** Y is a function $Y: [0, 1] \rightarrow \mathbb{R}$ taking finitely many values. For some $y \in \mathbb{R}$, we say that $Y = y$ with probability $p \in [0, 1]$ if the length of the set $\{x \in [0, 1]: Y(x) = y\}$ is p . (There is a more general way to define random variables, but this definition is fine for the moment.)

Definition 3.4. Let n be a positive integer. Let p_1, \dots, p_n be nonnegative numbers with $\sum_{i=1}^n p_i = 1$. Suppose a random variable Y has value y_i with probability p_i for each $1 \leq i \leq n$. Then the **expected value** of Y is $\sum_{i=1}^n y_i p_i$.

Definition 3.5. Let Y and Z be random variables. We say that Y and Z are **independent** if, for any fixed real numbers y, z , the probability that $Y = y$ and that $Z = z$ is equal to the probability that $Y = y$, multiplied by the probability that $Z = z$.

Definition 3.6. Let a, b be two real numbers. We denote $\min(a, b)$ as the minimum of a and b , and we denote $\max(a, b)$ as the maximum of a and b .

Exercise 3.7. Let X, Y be random variables. Let p_1, \dots, p_n be nonnegative numbers with $\sum_{i=1}^n p_i = 1$. Let q_1, \dots, q_n be nonnegative numbers with $\sum_{i=1}^n q_i = 1$. Suppose X has value $x_i \in \mathbb{R}$ with probability p_i for each $1 \leq i \leq n$. Suppose Y has value $y_i \in \mathbb{R}$ with probability q_i for each $1 \leq i \leq n$. Show that the expected value of $X + Y$ is equal to the expected value of X , plus the expected value of Y . Or, using the notation EX to denote the expected value of X , show that $E(X + Y) = (EX) + (EY)$. (Note: in this exercise, it is *not* assumed that X and Y are independent.)

Proposition 3.8. *In the pick-a-hand game, the chooser has a mixed strategy with an expected gain of at least $2/3$, and the hider has a mixed strategy with an expected loss of at most $2/3$.*

Proof. Suppose the chooser picks L with probability p and they pick R with probability $1 - p$, and suppose the hider picks L with probability q and they pick R with probability $1 - q$. We assume that the actions of the hider and chooser are independent. Then the probability that both pick L is pq and the probability that both pick R is $(1 - p)(1 - q)$. So, the expected gain of the chooser is

$$1 \cdot pq + 2 \cdot (1 - p)(1 - q) + 0 \cdot p(1 - q) + 0 \cdot q(1 - p) = pq + 2(1 - p)(1 - q) = 3pq - 2p - 2q + 2 =: f(p, q).$$

The gradient of f is $(3q - 2, 3p - 2)$ which is equal to zero exactly when $q = p = 2/3$. So, the point $q = p = 2/3$ is the only critical point of f . Also, from the second derivative test, the point $(2/3, 2/3)$ is a saddle point of f , and $f(2/3, 2/3) = 2/3$.

Suppose both players either guess or figure out the values of p and q . Then they also know the function f . Starting from any point (p, q) , the chooser will try to increase the value of f by changing the value of p . Since $\partial f / \partial p = 3q - 2$, the chooser will increase p to $p = 1$ if $q > 2/3$, so that $f(p, q) = f(1, q) = q > 2/3$; the chooser will decrease p to $p = 0$ if $q < 2/3$, so that $f(p, q) = f(0, q) = 2(1 - q) > 2/3$; and the chooser will not change p if $q = 2/3$, so that $f(p, q) = 2/3$. In any case, the chooser selects p so that f has value $\max(q, 2(1 - q))$. And no matter what the value of q is, the chooser gets an expected gain of at least $2/3$.

Similarly, starting from any point (p, q) , the hider will try to decrease the value of f by changing the value of q . Since $\partial f / \partial q = 3p - 2$, the hider will increase q to $q = 1$ if $p < 2/3$,

so that $f(p, q) = f(p, 1) = p < 2/3$; the hider will decrease q to $q = 0$ if $p > 2/3$, so that $f(p, q) = f(p, 0) = 2(1 - p) < 2/3$; and the hider will not change q if $p = 2/3$, so that $f(p, q) = 2/3$. In any case, the hider selects q so that f has value $\min(p, 2(1 - p))$. So, no matter what the value of p is, the hider gets an expected loss of at most $2/3$. \square

It is no coincidence that both players have the same expected gain from their best mixed strategies. This is a general phenomenon of two-person zero-sum games, which is proven in von Neumann's Minimax Theorem, Theorem 3.29 below.

Definition 3.9. Let m, n be positive integers. Suppose we have an arbitrary two-person zero-sum game, which is described by an $m \times n$ payoff matrix $(a_{ij})_{1 \leq i \leq m, 1 \leq j \leq n}$. The mixed strategies for player I correspond to vectors $x = (x_1, \dots, x_m)$, where x_i is the probability that player I chooses the pure strategy i , where $1 \leq i \leq m$. That is, the set of mixed strategies for player I is denoted by

$$\Delta_m := \{x = (x_1, \dots, x_m) \in \mathbb{R}^m : \sum_{i=1}^m x_i = 1, x_i \geq 0, \forall 1 \leq i \leq m\}.$$

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

$$\Delta_n := \{y = (y_1, \dots, y_n) \in \mathbb{R}^n : \sum_{j=1}^n y_j = 1, y_j \geq 0, \forall 1 \leq j \leq n\}.$$

Unless otherwise stated, a vector in \mathbb{R}^n is treated as a column vector. Also, x^T denotes the transpose of x (so that x^T is a row vector). Lastly, the set Δ_m is called the **m -dimensional simplex**.

Note that, using this notation, pure strategies are exactly the standard basis vectors. Suppose the choices of player I and player II are independent. Then the probability that player I plays i and that player II plays j is $x_i y_j$, for any $1 \leq i \leq m, 1 \leq j \leq n$. So, the expected payoff to player I is

$$\sum_{i=1}^m \sum_{j=1}^n x_i a_{ij} y_j = x^T A y.$$

Definition 3.10. We refer to Ay as the **payoff vector** for player I corresponding to the mixed strategy y for player II . The elements of the vector Ay represent the expected payoffs for player I for each of their pure strategies. Similarly, we refer to $x^T A$ as the **payoff vector** for player II corresponding to the mixed strategy x for player I . The elements of the vector $x^T A$ represent the expected payoffs for player II for each of their pure strategies.

We say that a vector $w \in \mathbb{R}^d$ **dominates** another vector $v \in \mathbb{R}^d$ if $w_i \geq v_i$ for all $1 \leq i \leq d$. In this case, we write $w \geq v$.

We can now make a formal definition of a strategy being optimal in a two-person zero-sum game.

Definition 3.11. A mixed strategy $\tilde{x} \in \Delta_m$ is **optimal for player I** if

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

A mixed strategy $\tilde{y} \in \Delta_n$ is **optimal for player II** if

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

For example, the first definition says that, no matter what strategy player II chooses, the strategy \tilde{x} achieves the best payoff for player I .

Example 3.12. In the pick-a-hand game, we had $n = m = 2$ and $x^T Ay = x_1 y_1 + 2x_2 y_2 = x_1 y_1 + 2(1 - x_1)(1 - y_1)$. Also, in our analysis of the hider's best strategy, we observed that $\min_{y \in \Delta_2} x^T Ay = \min(x_1, 2(1 - x_1))$. And if the hider uses the best strategy, the chooser can still try to maximize their payoff, so that

$$\max_{x \in \Delta_n} \min_{y \in \Delta_n} x^T Ay = \max_{x_1 \in [0,1]} \min(x_1, 2(1 - x_1)) = 2/3.$$

And $\max_{x \in \Delta_2} x^T Ay = \max(y_1, 2(1 - y_1))$, which we saw in our analysis of the chooser's strategy. And if the chooser uses the best strategy, the hider can try to minimize the chooser's payoff, so that

$$\min_{y \in \Delta_n} \max_{x \in \Delta_m} x^T Ay = \min_{y_1 \in [0,1]} \max(y_1, 2(1 - y_1)) = 2/3.$$

In summary, each player can respond to the other player's best strategy, resulting in the same payoff in both cases. This is the general fact expressed in Theorem 3.29 below.

Lastly, we can verify that $x = (2/3, 1/3)$ is the only optimal strategy for Player I , and $y = (2/3, 1/3)$ is the only optimal strategy for Player II . In general, there can be more than one optimal strategy for each player.

3.1. Von Neumann's Minimax Theorem. Let d be a positive integer. In preparation for the proof of the Minimax Theorem, we will discuss some convex geometry.

Definition 3.13. A set $K \subseteq \mathbb{R}^d$ is **convex** if, for any two points $a, b \in K$, the line segment between them

$$\{ta + (1 - t)b : t \in [0, 1]\}$$

also lies in K .

Example 3.14. The unit disc $\{(x_1, x_2) \in \mathbb{R}^2 : x_1^2 + x_2^2 \leq 1\}$ is convex. The unit circle $C = \{(x_1, x_2) \in \mathbb{R}^2 : x_1^2 + x_2^2 = 1\}$ is not convex, since $(1, 0) \in C$, $(-1, 0) \in C$, $(0, 0) \notin C$, while $(0, 0) = (1/2)(1, 0) + (1/2)(-1, 0)$. So, if C were convex, then $(0, 0)$ would have to be in C .

We now recall some facts from Calculus concerning closed sets and continuous functions. Below, we let $0 \in \mathbb{R}^d$ denote the origin of \mathbb{R}^d .

Definition 3.15. Let $x = (x_1, \dots, x_d) \in \mathbb{R}^d$. We define the **length** of x to be

$$\|x\| := (x_1^2 + \dots + x_d^2)^{1/2} = \left(\sum_{i=1}^d x_i^2\right)^{1/2} = (x^T x)^{1/2}.$$

Let $R > 0$. We define the **ball of radius R centered at x** to be the set

$$B_R(x) := \{y \in \mathbb{R}^d : \|y - x\| \leq R\}.$$

Definition 3.16. Recall that a sequence of points $x^{(1)}, x^{(2)}, \dots$ in \mathbb{R}^d **converges** to a point $x \in \mathbb{R}^d$ if $\lim_{j \rightarrow \infty} \|x^{(j)} - x\| = 0$.

Let $K \subseteq \mathbb{R}^d$. We say that K is **closed** if the following condition is satisfied. If $x^{(1)}, x^{(2)}, \dots$ is any sequence of points in K , and if this sequence converges to any point $x \in \mathbb{R}^d$, then we must have $x \in K$.

We say that K is **bounded** if there exists a radius $R > 0$ such that $K \subseteq B_R(0)$.

Exercise 3.17. Using this definition, show that $[0, 1]$ is closed, but $(0, 1)$ is not closed.

Exercise 3.18. Show that the intersection of two closed sets is a closed set.

Exercise 3.19. Let m be a positive integer. Show that Δ_m as defined in Definition 3.9 is a closed and bounded set.

Exercise 3.20. Let $K_1, K_2 \subseteq \mathbb{R}^d$ be closed and bounded sets. Show that $K_1 \times K_2$ is a closed and bounded set.

Definition 3.21. Let $f: \mathbb{R}^d \rightarrow \mathbb{R}$ be a function. We say that f is **continuous** if, for any sequence of points $x^{(1)}, x^{(2)}, \dots$ in \mathbb{R}^d that converges to any point $x \in \mathbb{R}^d$, we have $\lim_{j \rightarrow \infty} |f(x^{(j)}) - f(x)| = 0$.

Exercise 3.22. Define $f: \mathbb{R}^d \rightarrow \mathbb{R}$ by $f(x) := \|x\|$. Show that f is continuous. (Hint: you may need to use the triangle inequality, which says that $\|x + y\| \leq \|x\| + \|y\|$, for any $x, y \in \mathbb{R}^d$.)

Theorem 3.23 (Extreme Value Theorem). Let $K \subseteq \mathbb{R}^d$ be a closed and bounded set. Let $f: \mathbb{R}^d \rightarrow \mathbb{R}$ be a continuous function. Then f achieves its maximum and minimum on K . That is, there exist points $a, b \in K$ such that $f(a) \leq f(x) \leq f(b)$ for all $x \in K$. We denote $f(a) = \min_{x \in K} f(x)$ and $f(b) = \max_{x \in K} f(x)$.

The above Theorem is discussed informally in multivariable calculus, and it is proven formally in a second analysis class. The following Lemmas can also be proven in a second analysis class.

Lemma 3.24. Let $K \subseteq \mathbb{R}^d$ be a closed and bounded set. Let $f: \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be a continuous function. Then the function $g: \mathbb{R}^d \rightarrow \mathbb{R}$ defined by $g(y) := \min_{x \in K} f(x, y)$ is continuous, and the function $h: \mathbb{R}^d \rightarrow \mathbb{R}$ defined by $h(y) := \max_{x \in K} f(x, y)$ is continuous.

Lemma 3.25. Let $U \subseteq \mathbb{R}^d$ be a closed and bounded set, and let Y be a closed set. Let A be an $m \times d$ real matrix. The set of points $\{Ax: x \in U\}$ is closed and bounded. Also, the set of points $U + Y = \{u + y: u \in U, y \in Y\}$ is a closed set.

The following Theorem is our main tool for proving the von Neumann Minimax Theorem.

Theorem 3.26 (The Separating Hyperplane Theorem). Let $K \subseteq \mathbb{R}^d$ be a closed and convex set. Assume that $0 \notin K$. Then, there exists $z \in \mathbb{R}^d$ and $c \in \mathbb{R}$ such that

$$0 < c < z^T x, \quad \text{for all } x \in K.$$

Remark 3.27. A **hyperplane** is a set of the form $\{x \in \mathbb{R}^d: z^T x = c\}$, where $z \in \mathbb{R}^d$ and $c \in \mathbb{R}$. Writing this out in coordinates, so that $x = (x_1, \dots, x_d)$ and $z = (z_1, \dots, z_d)$, we have

$$\{x \in \mathbb{R}^d: z^T x = c\} = \{x \in \mathbb{R}^d: \sum_{i=1}^d x_i z_i = c\}.$$

In particular, when $d = 2$, this equation says $x_1 z_1 + x_2 z_2 = c$, which is exactly an equation for a line in \mathbb{R}^2 . And when $d = 3$, this equation says $x_1 z_1 + x_2 z_2 + x_3 z_3 = c$, which is exactly an equation for a plane in \mathbb{R}^3 . So, a hyperplane is higher-dimensional generalization of a plane.

To explain the name of the theorem, note that the hyperplane $\{x \in \mathbb{R}^d: z^T x = c\}$ separates K from the origin. That is, K lies on one side of the hyperplane, so that $K \subseteq \{x \in \mathbb{R}^d: z^T x >$

$c\}$. And 0 lies on the other side of the hyperplane, so that $0 \in \{x \in \mathbb{R}^d : z^T x < c\}$, since $z^T 0 = 0 < c$.

Proof of Theorem 3.26. Let k be any point in K . Choose $R = \|k\| + 1$. Then $k \in B_R(0)$, so that $B_R(0)$ has nonempty intersection with K . Note that $B_R(0)$ is a closed set, so $K \cap B_R(0)$ is also closed by Exercise 3.18. Also, $K \cap B_R(0)$ is bounded, since $(K \cap B_R(0)) \subseteq B_R(0)$. Define $f : \mathbb{R}^d \rightarrow \mathbb{R}$ by $f(x) := \|x\|$. Then f is continuous by Exercise 3.22. Since $K \cap B_R(0)$ is closed and bounded, f achieves its minimum value on $K \cap B_R(0)$, by Theorem 3.23. In particular, there exists $z \in K$ such that $\|z\| \leq \|x\| \leq R$, for all $x \in K \cap B_R(0)$. Any point $x \in K$ has $\|x\| \leq R$ or $\|x\| > R$. Therefore,

$$\|z\| \leq \|x\|, \quad \forall x \in K \quad (*)$$

Let $x \in K$ and let $0 < t < 1$. Since K is convex, we have $tx + (1-t)z \in K$. So, applying (*), we have

$$\begin{aligned} z^T z &= \|z\|^2 \leq \|tx + (1-t)z\|^2 = (tx + (1-t)z)^T (tx + (1-t)z) \\ &= t^2 x^T x + 2t(1-t)x^T z + (1-t)^2 z^T z. \end{aligned}$$

That is,

$$t^2(x^T z - x^T x - z^T z) \leq 2t(x^T z - z^T z).$$

Dividing by t and letting t go to zero, we get

$$0 \leq x^T z - z^T z = x^T z - \|z\|^2.$$

That is, $\|z\|^2 \leq x^T z$, for any $x \in K$. So, define $c := \|z\|^2/2$. Since $0 \notin K$, $z \neq 0$. So, $\|z\| \neq 0$, and $c > 0$. That is, for every $x \in K$, we have $0 < c < \|z\|^2 \leq x^T z = z^T x$. \square

Lemma 3.28. *Let X and Y be closed and bounded sets in \mathbb{R}^d . Let $f : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be continuous. Then*

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

Proof. Fix some $(x^*, y^*) \in X \times Y$. Then $f(x^*, y^*) \leq \max_{x \in X} f(x, y^*)$, and $\min_{y \in Y} f(x^*, y) \leq f(x^*, y^*)$. (For both inequalities, we are using Theorem 3.23 to assert the existence of the maximum and minimum values of f , respectively.) Combining both inequalities,

$$\min_{y \in Y} f(x^*, y) \leq \max_{x \in X} f(x, y^*). \quad (\ddagger)$$

Since this inequality holds for any x^* , we can take the maximum over x^* of both sides. Similarly, this inequality holds for any y^* , so we can take the minimum over y^* of both sides. Performing these operations preserves the inequality, so we get

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

(In order to take the minimum and maximum in this way, we are using Lemma 3.24 to get continuity of each side of (\ddagger) , and we are then using Theorem 3.23 to assert the existence of these maximum and minimum values.) \square

We use the notation of Definition 3.9.

Theorem 3.29 (von Neumann's Minimax Theorem). Let m, n be positive integers. Let A be an $m \times n$ payoff matrix. Then

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

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

Remark 3.30. Suppose $\tilde{x} \in \Delta_m$ is optimal for Player I and $\tilde{y} \in \Delta_n$ is optimal for Player II. Then Theorem 3.29 implies that

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

That is, $\tilde{x}^T A \tilde{y}$ is equal to the value of the game. For this reason, if \tilde{x} is optimal for Player I, and if \tilde{y} is optimal for player II, we call the strategies \tilde{x}, \tilde{y} **optimal strategies**.

Proof. Let $f: \mathbb{R}^m \times \mathbb{R}^n \rightarrow \mathbb{R}$ so that $f(x, y) := x^T A y$. Then f is continuous (since it is a polynomial of degree at most 2), and Δ_m, Δ_n are closed and bounded sets by Exercise 3.19. So, Lemma 3.28 says that

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

It remains to prove the reverse inequality. We argue by contradiction. Suppose there is a real number λ such that

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

We define a new payoff matrix \hat{A} so that $\hat{a}_{ij} = a_{ij} - \lambda$ for all $1 \leq i \leq m, 1 \leq j \leq n$. Then

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

Let K be the set of vectors that dominate $\hat{A}y$, for any $y \in \Delta_n$. That is,

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

We claim that K is closed and convex. To see the convexity, let $\hat{A}y + v, \hat{A}y' + v' \in K$, where $y, y' \in \Delta_n$ and $v, v' \geq (0, \dots, 0)$. Let $0 \leq t \leq 1$. Since Δ_n is convex, $ty + (1-t)y' \in \Delta_n$. Also, since the components of v, v' are all nonnegative, the components of $tv + (1-t)v'$ are all nonnegative, so that $tv + (1-t)v' \geq (0, \dots, 0)$. Therefore, $\hat{A}[ty + (1-t)y'] + [tv + (1-t)v'] \in K$, i.e. $t(\hat{A}y + v) + (1-t)(\hat{A}y' + v') \in K$, so that K is convex. Also, K is closed by Lemma 3.25. (Lemma 3.25 says that $U = \{\hat{A}y : y \in \Delta_n\}$ is closed and bounded, so if $Y = \{v \in \mathbb{R}^m : v \geq (0, \dots, 0)\}$, then $U + Y = K$ is closed.)

The set K does not contain 0. If $0 \in K$, then there would be some $y \in \Delta_n$ and $v \in \mathbb{R}^m$ with $\hat{A}y = -v \leq (0, \dots, 0)$. Then, for any $x \in \Delta_m$, we would have $x^T \hat{A}y \leq 0$, so that $\max_{x \in \Delta_m} x^T \hat{A}y \leq 0$, contradicting the right side of (*). In summary, $0 \notin K$, and K is closed and convex. So, Theorem 3.26 applies.

That is, there exists $z \in \mathbb{R}^m$ and $c > 0$ such that $0 < c < z^T w$ for all $w \in K$. Using the definition of K , we conclude that

$$z^T(\hat{A}y + v) > c > 0, \quad \forall y \in \Delta_n, \forall v \geq (0, \dots, 0). \quad (**)$$

Suppose $z_j < 0$ for some $1 \leq j \leq m$. Note that $z^T(\hat{A}y + v) = z^T\hat{A}y + \sum_{i=1}^m z_i v_i$. So, if we choose v so that $v_i = 0$ for all $i \neq j$ and let v_j go to $+\infty$, we will find some $v \geq (0, \dots, 0)$ with $z^T(\hat{A}y + v) < 0$, contradicting (**). So, $z \geq (0, \dots, 0)$. Similarly, (**) implies that $z \neq (0, \dots, 0)$. So, if we define $s := \sum_{i=1}^m z_i$, then $s > 0$. Then, if we define $x := z/s$, then $x \in \Delta_m$. So, using (**) with $v = 0$, we have

$$x^T(\hat{A}y + v) = x^T\hat{A}y = z^T(\hat{A}y)/s > c/s > 0, \quad \forall y \in \Delta_n.$$

Taking the minimum over all $y \in \Delta_n$, we then have

$$\min_{y \in \Delta_n} x^T\hat{A}y \geq 0.$$

This inequality contradicts the left side of (*). Having achieved a contradiction, the proof is complete. \square

Remark 3.31. The conclusion of the von Neumann Minimax Theorem does not hold for arbitrary functions. As a related exercise, try to find a function $f: [0, 1] \times [0, 1] \rightarrow \mathbb{R}$ such that $\min_{y \in [0, 1]} \max_{x \in [0, 1]} f(x, y) \neq \max_{x \in [0, 1]} \min_{y \in [0, 1]} f(x, y)$.

Exercise 3.32. Suppose we have a two-person zero-sum game with $(n+1) \times (n+1)$ payoff matrix A such that at least one entry of A is nonzero. Let $x, y \in \Delta_{n+1}$. Write $x = (x_1, \dots, x_n, 1 - \sum_{i=1}^n x_i)$, $y = (x_{n+1}, x_{n+2}, \dots, x_{2n}, 1 - \sum_{i=n+1}^{2n} x_i)$. Consider the function $f: \mathbb{R}^{2n} \rightarrow \mathbb{R}$ defined by $f(x_1, \dots, x_{2n}) = x^T A y$. Show that the determinant of the Hessian of f is always negative. Conclude that any critical point of f is a saddle point. That is, if we find a critical point of f (as we do when we look for the value of the game), then this critical point is a saddle point of f . In this sense, the minimax value occurs at a saddle point of f . (This is something we observed in Proposition 3.8.)

(Hint: Write f in the form $f(x_1, \dots, x_{2n}) = \sum_{i=1}^{2n} b_i x_i + \sum_{\substack{1 \leq i \leq n, \\ n+1 \leq j \leq 2n}} c_{ij} x_i x_j$, where $b_i, c_{ij} \in \mathbb{R}$. From here, it should follow that there exists a nonzero matrix C such that the Hessian of f , i.e. the matrix of second order partial derivatives of f , should be of the form $\begin{pmatrix} 0 & C \\ C^T & 0 \end{pmatrix}$.)

Exercise 3.33. Let $x \in \Delta_m$, $y \in \Delta_n$ and let A be an $m \times n$ matrix. Show that

$$\max_{x \in \Delta_m} x^T A y = \max_{i=1, \dots, m} (A y)_i, \quad \min_{y \in \Delta_n} x^T A y = \min_{j=1, \dots, n} (x^T A)_j.$$

Using this fact, show that

$$\min_{y \in \Delta_n} \max_{x \in \Delta_m} x^T A y = \min_{y \in \Delta_n} \max_{i=1, \dots, m} (A y)_i.$$

$$\max_{x \in \Delta_m} \min_{y \in \Delta_n} x^T A y = \max_{x \in \Delta_m} \min_{j=1, \dots, n} (x^T A)_j.$$

Using the second equality, conclude that the value of the game with payoff matrix A can be found via the following Linear Programming problem:

Maximize t subject to the constraints: $\sum_{i=1}^n a_{ij} x_i \geq t$, for all $1 \leq j \leq n$; $\sum_{i=1}^m x_i = 1$; $x \geq (0, \dots, 0)$.

Efficient methods for solving linear programming problems are well-known. However, below we will focus on ways to compute the values of two-person zero-sum games by hand.

		Player II	
		A	B
Player I	C	2	3
	D	1	2

3.2. Domination.

Example 3.34. Consider the following two-person zero-sum game.

We claim that the second row of this matrix can be ignored for the purpose of computing the value of the game. In particular, the second row is **dominated by** the first row, since $a_{11} \geq a_{21}$ and $a_{12} \geq a_{22}$. So, if $y = (y_1, y_2) \in \Delta_2$, since $y \geq (0, 0)$, we have $a_{11}y_1 + a_{12}y_2 \geq a_{21}y_1 + a_{22}y_2$. That is, $(Ay)_1 \geq (Ay)_2$. From Exercise 3.33, the value of the game is equal to

$$\min_{y \in \Delta_2} \max_{x \in \Delta_2} x^T Ay = \min_{y \in \Delta_2} \max_{i=1,2} (Ay)_i = \min_{y \in \Delta_2} (Ay)_1.$$

That is, we can ignore the second row of the matrix A for computing the value of the game. The value of the game is therefore

$$\min_{y \in \Delta_2} (Ay)_1 = \min_{y \in \Delta_2} (2y_1 + 3y_2) = \min_{y_1 \in [0,1]} (2y_1 + 3(1 - y_1)) = \min_{y_1 \in [0,1]} (3 - y_1) = 2.$$

Example 3.35. Consider the two-person zero-sum game defined by the following payoff matrix. Let $x, y \in \Delta_n$. In this example, each row with index at least 4 is dominated by the

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

first row. That is, $(Ay)_1 \geq (Ay)_i$ for any $4 \leq i \leq n$. From Exercise 3.33, the value of the game is equal to

$$\min_{y \in \Delta_n} \max_{x \in \Delta_n} x^T Ay = \min_{y \in \Delta_n} \max_{i=1,\dots,n} (Ay)_i = \min_{y \in \Delta_n} \max_{i=1,2,3} (Ay)_i.$$

That is, for the purpose of computing the value of the game, we can ignore all rows except for the first three. Similarly, each column with index at least 4 dominates the first column. That is, $(x^T A)_1 \leq (x^T A)_j$ for any $4 \leq j \leq n$. From Exercise 3.33, the value of the game is equal to

$$\max_{x \in \Delta_n} \min_{y \in \Delta_n} x^T Ay = \max_{x \in \Delta_n} \min_{j=1,\dots,n} (x^T A)_j = \max_{x \in \Delta_n} \min_{j=1,2,3} (x^T A)_j.$$

That is, for the purpose of computing the value of the game, we can ignore all columns except for the first three.

In summary, the value of the original game has the same value as the following game

		Player II		
		1	2	3
Player I	1	0	-1	2
	2	1	0	-1
	3	-2	1	0

Let A denote this 3×3 payoff matrix. Note that $A = -A^T$. That is, A is antisymmetric. We first claim that the value of the game corresponding to A is zero. To see this, note that the value is

$$\begin{aligned} \min_{y \in \Delta_3} \max_{x \in \Delta_3} x^T A y &= \min_{y \in \Delta_3} \max_{x \in \Delta_3} x^T (-A^T) y = \min_{y \in \Delta_3} \max_{x \in \Delta_3} -(y^T A x)^T = \min_{y \in \Delta_3} [-\min_{x \in \Delta_3} y^T A x] \\ &= -\max_{y \in \Delta_3} \min_{x \in \Delta_3} y^T A x = -\min_{x \in \Delta_3} \max_{y \in \Delta_3} y^T A x = -\min_{y \in \Delta_3} \max_{x \in \Delta_3} x^T A y. \end{aligned}$$

In the penultimate equality, we used von Neumann's Minimax Theorem, Theorem 3.29. So, the value is equal to the negative of itself, so it must be zero. Now, using Exercise 3.33, the value is

$$0 = \max_{x \in \Delta_3} \min_{j=1,2,3} (x^T A)_j = \max_{x \in \Delta_3} \min(x_2 - 2x_3, -x_1 + x_3, 2x_1 - x_2).$$

So, if $x \in \Delta_3$ is the optimal strategy for Player I, we have $\min(x_2 - 2x_3, -x_1 + x_3, 2x_1 - x_2) = 0$, so that $(x_2 - 2x_3, -x_1 + x_3, 2x_1 - x_2) \geq (0, 0, 0)$, i.e. $x_2 \geq 2x_3$, $2x_3 \geq 2x_1$ and $2x_1 \geq x_2$. If any of these inequalities were strict, their combination would say $x_2 > x_2$, which is false. So, each of these inequalities must be an equality. That is, $x_2 = 2x_3 = 2x_1 = x_2$. Since $x_1 + x_2 + x_3 = 1$, we conclude that $x_3 = x_1 = 1/4$ and $x_2 = 1/2$. So, the optimal strategy for player I is $(1/4, 1/2, 1/4)$.

Using similar reasoning for y , the optimal strategy for player II is also $(1/4, 1/2, 1/4)$.

Remark 3.36. A domination argument also implies the following: if there is some $t \in (0, 1)$ and there are three rows u, v, w of a payoff matrix such that $tu + (1-t)v$ dominates w , then we can similarly remove the row w from the payoff matrix, for the purpose of computing the value of the game. Similarly, if there is some $t \in (0, 1)$ and there are three columns u, v, w of a payoff matrix such that w dominates $tu + (1-t)v$, then we can also remove the column w from the payoff matrix, for the purpose of computing the value of the game.

For example, the last row of the following matrix can be ignored in the computation of the minimax, since $(1/2)(10, 0, 0) + (1/2)(0, 0, 10) > (4, 0, 4)$.

		Player II		
		1	2	3
Player I	1	10	0	0
	2	0	0	10
	3	4	0	4

4. GENERAL SUM GAMES

We now discuss **general-sum games**. A general-sum game for two players is given in **strategic form** by two matrices A and B whose entries give the payoffs for each of the players according to the strategies played by the two players. (The zero-sum case corresponds to $A + B = 0$, or $A = -B$.) For general-sum games, there is typically no joint optimal strategy

for both players. Nevertheless, there is a generalization of the von Neumann minimax which is known as the Nash equilibrium. The Nash equilibrium gives strategies that “rational,” “selfish” players would make. However, there are often several Nash equilibria, so if both players want to choose strategies corresponding to the same Nash equilibrium, they may have to cooperate in making such a choice. Also, strategies for the players based on cooperation (as opposed to only selfish actions) may be better than any Nash equilibrium strategy.

Example 4.1 (The Prisoner’s Dilemma). Two suspects of a crime are held by the police in separate cells. The charge is serious, but there is little evidence. If one person confesses while the other remains silent, the confessor is set free, and the silent one is put in jail for ten years. If both confess, then they both spend eight years in prison. If both remain silent, then the sentence is one year each, for some minor crime based on the small amount of evidence. We write a negative payoff for the number of years spent in prison, resulting in the following payoff matrix.

Prisoner <i>I</i>	Prisoner <i>II</i>	
	silent	confess
silent	(−1, −1)	(−10, 0)
confess	(0, −10)	(−8, −8)

The payoff matrix for player *I* is the matrix given by the first entries in each vector in the above matrix. The payoff matrix for player *II* is the matrix given by the second entries in each vector in the above matrix.

Consider the following domination argument for player *I*. The payoff matrix for player *I* has a larger second row than first row. That is, no matter what player *II* does, player *I* will always improve their payoff by choosing to confess. Similarly, the payoff matrix for player two has a larger second column than first column. That is, no matter what player *I* does, player *II* will always improve their payoff by choosing to confess. So, the logical choice for both players is for both of them to confess, resulting in 8 years of jail for both. Note that this outcome is much worse than if both remained silent, in which case both would have 1 year of jail. Put another way, acting selfishly is not in the best interest of both parties. It is better for both parties if they cooperate, and they both agree to stay silent.

The conclusion (that both parties should confess) is valid when the game is played once. However, if the game is played repeatedly a random number of times, then the optimal strategy could change. In particular, when the game is repeated, the players can remember previous rounds of the game. This information can then be used in the players’ strategies.

Remark 4.2. The prisoner’s dilemma appeared in a scene in the movie *The Dark Knight*; however, some doubt about the enforcement of the dilemma causes the Nash equilibrium not to be reached. Also, the game is repeated many times, so that confessing is not necessarily the optimal strategy.

The prisoner’s dilemma can be applied to some real world situations. For example, we can think of the two players as two nations that want to reduce greenhouse gas emissions. Both countries would like to reduce emissions, but if one country reduces emissions while the other does not, the lower emission producer is (arguably) at an economic disadvantage. So, the Nash equilibrium suggests that the two countries do not reduce their emissions. This situation seems to continue to occur in the real world. Other examples of the prisoner’s dilemma include drug addiction, nuclear armament, etc.

Example 4.3 (Battle of the Spouses). Two spouses are choosing what to do on a Saturday night. Spouse *I* wants to go to the opera, and Spouse *II* prefers watching baseball. If they each do different activities, neither is satisfied. The payoff matrices are

		Spouse <i>II</i>	
		opera	baseball
Spouse <i>I</i>	opera	(4, 1)	(0, 0)
	baseball	(0, 0)	(1, 4)

Since there are now two payoff matrices, von Neumann's Minimax Theorem, Theorem 3.29, does not necessarily tell us anything about strategies in the game. However, we could still try to use this Theorem.

For example, suppose Spouse *I* wants to maximize their payoff, regardless of what Spouse *II* does. That is, if A denotes the payoff matrix of Spouse *I*, then they can guarantee a payoff at least

$$\max_{x \in \Delta_2} \min_{y \in \Delta_2} x^T A y = \max_{x \in \Delta_2} \min_{y \in \Delta_2} (4x_1 y_1 + x_2 y_2) = \max_{x \in \Delta_2} \min_{y \in \Delta_2} (4x_1, x_2) = \max_{t \in [0,1]} \min(4t, (1-t)) = 4/5.$$

That is, Spouse *I* can get a payoff of $4/5$ with the mixed strategy $(1/5, 4/5)$. Similarly, Spouse *II* can get a payoff of $4/5$ with the mixed strategy $(4/5, 1/5)$. However, note that these payoffs are strictly lower than the pure strategies, where both Spouses choose baseball, or both Spouses choose opera.

Perhaps the above situation between two Spouses is a bit unrealistic, since we are implicitly assuming that they make their final decisions in separate rooms. In reality, both Spouses would decide what to through negotiation, perhaps. And after several repetitions of the game, the maximum payoff for both players occurs when both Spouses go to the opera half the time, and both Spouses watch baseball half the time (which is again perhaps more realistic). Nevertheless, this example demonstrates that, in general sum games, the minimax approach no longer really tells us about the actions of “rational” players. The Nash equilibrium, which we describe below, attempts to describe the actions of “rational” players in general sum games.

4.1. Nash equilibria.

Definition 4.4 (Nash equilibrium). Let m, n be positive integers. Suppose we have a two-player general sum game with $m \times n$ payoff matrices. Let A be the payoff matrix for player *I* and let B be the payoff matrix for player *II*. A pair of vectors (\tilde{x}, \tilde{y}) with $\tilde{x} \in \Delta_m$ and $\tilde{y} \in \Delta_n$ is a **Nash equilibrium** if no player gains by unilaterally deviating from the strategy. That is,

$$\begin{aligned} \tilde{x}^T A \tilde{y} &\geq x^T A \tilde{y}, \quad \forall x \in \Delta_m, \\ \tilde{x}^T B \tilde{y} &\geq \tilde{x}^T B y, \quad \forall y \in \Delta_n. \end{aligned}$$

Definition 4.5. A two-player general sum game is **symmetric** if $m = n$ and if $A_{ij} = B_{ji}$ for all $i, j \in \{1, \dots, n\}$ (that is, $A = B^T$). A pair (x, y) of strategies is **symmetric** if $m = n$ and if $x_i = y_i$ for all $i \in \{1, \dots, n\}$.

If (x, y) is a Nash equilibrium, and if x has one coordinate equal to 1 with all other coordinates 0, and if y has one coordinate equal to one with all other coordinates zero, then the equilibrium (x, y) is called **pure**. If (x, y) is a Nash equilibrium that is not pure, then it is called **mixed**.

Example 4.6. Recall the Prisoner's Dilemma in Example 4.1. We had $A = \begin{pmatrix} -1 & -10 \\ 0 & -8 \end{pmatrix}$, $B = A^T = \begin{pmatrix} -1 & 0 \\ -10 & -8 \end{pmatrix}$. Suppose $\tilde{x}, \tilde{y} \in \Delta_2$ so that (\tilde{x}, \tilde{y}) is a Nash equilibrium. Then we must have $\tilde{x} = (0, 1)$ and $\tilde{y} = (0, 1)$. (So, the game is symmetric, and its only Nash equilibrium is symmetric and pure.) So, the only Nash equilibrium occurs when both prisoners confess. We can see this from a domination argument. If $y = (y_1, y_2) \in \Delta_2$, note that $-y_1 - 10y_2 < -8y_2$ since $y_1 \geq 0$, $y_2 \geq 0$ and $y_1 + y_2 = 1$. So, if $t \in [0, 1]$, we have $t(-y_1 - 10y_2) + (1 - t)(-8y_2) \leq -8y_2$ with equality only when $t = 0$. Therefore, using Definition 4.4,

$$\tilde{x}^T A \tilde{y} = \max_{x \in \Delta_2} x^T A \tilde{y} = \max_{t \in [0, 1]} t(-\tilde{y}_1 - 10\tilde{y}_2) + (1 - t)(-8\tilde{y}_2) \stackrel{(*)}{=} -8\tilde{y}_2.$$

And the equality $(*)$ is only attained when $t = 0$, i.e. when $\tilde{x} = (0, 1)$. Similarly, using again Definition 4.4,

$$\tilde{x}^T B \tilde{y} = \max_{y \in \Delta_2} \tilde{x}^T B y = \max_{t \in [0, 1]} t(-\tilde{x}_1 - 10\tilde{x}_2) + (1 - t)(-8\tilde{x}_2) \stackrel{(**)}{=} -8\tilde{x}_2.$$

And the equality $(**)$ is only attained if $\tilde{x} = (0, 1)$.

Note that this argument is a bit strange since we assumed at the outset that a Nash equilibrium exists, and then we found this equilibrium. A priori, a Nash equilibrium may not even exist. However, Nash's Theorem (Theorem 4.29 below) guarantees that a Nash equilibrium always exists.

Example 4.7. Recall the Battle of the Spouses in Example 4.3. In this game there are two pure Nash equilibria, where both choose baseball, or both choose opera. There is also a mixed equilibrium of $(4/5, 1/5)$ for player I and $(1/5, 4/5)$ for player II . Recall that we have $A = \begin{pmatrix} 4 & 0 \\ 0 & 1 \end{pmatrix}$, $B = \begin{pmatrix} 1 & 0 \\ 0 & 4 \end{pmatrix}$. Let's first show that $\tilde{x} = (1, 0)$, $\tilde{y} = (1, 0)$ is a Nash equilibrium. Note that

$$\max_{x \in \Delta_2} x^T A \tilde{y} = \max_{x \in \Delta_2} 4x_1 = 4 = \tilde{x}^T A \tilde{y}.$$

$$\max_{y \in \Delta_2} \tilde{x}^T B y = \max_{y \in \Delta_2} y_1 = 1 = \tilde{x}^T B \tilde{y}.$$

So, $\tilde{x} = (1, 0)$, $\tilde{y} = (1, 0)$ is a Nash equilibrium. We can similarly see that $\tilde{x} = (0, 1)$, $\tilde{y} = (0, 1)$ is a Nash equilibrium.

Let's also verify that $\tilde{x} = (4/5, 1/5)$, $\tilde{y} = (1/5, 4/5)$ is a Nash equilibrium. In this case,

$$\max_{x \in \Delta_2} x^T A \tilde{y} = (4/5) \max_{x \in \Delta_2} (x_1 + x_2) = 4/5 = 4(1/5)(4/5) + 1(1/5)(4/5) = \tilde{x}^T A \tilde{y}.$$

$$\max_{y \in \Delta_2} \tilde{x}^T B y = (4/5) \max_{y \in \Delta_2} (y_1 + y_2) = 4/5 = 1(4/5)(1/5) + 4(1/5)(4/5) = \tilde{x}^T B \tilde{y}.$$

The payoffs in the mixed equilibrium are strictly worse than the payoffs for both of the pure equilibria.

4.2. Correlated equilibria. In Examples 4.3 and 4.7, we saw that the Battle of the Spouses has two pure Nash equilibria, corresponding to both Spouses going to the opera, or both watching baseball. The pure equilibria give higher payoffs than the mixed equilibrium, but each of the pure equilibria are unfair to one of the players, since one player gets a higher payoff than the other. (One pure equilibrium has payoffs of 1 and 4, while the other equilibrium has payoffs of 4 and 1.) If we could flip a coin to choose between the two pure equilibria, then the expected payoffs for both players would be $(4 + 1)/2 = 5/2$. This strategy seems both more equitable and more realistic than any of the Nash equilibria. We can add such a “coin flip” to the definition of an equilibrium by defining the correlated equilibrium.

Definition 4.8 (Correlated Equilibrium). Suppose we have a joint distribution of strategies in a two-player general sum game. That is, every choice of strategy has some probability assigned to it. Such a distribution is called a **correlated equilibrium** if no player gains by unilaterally deviating from it.

More formally, suppose A is an $m \times n$ payoff matrix for player I and B is an $m \times n$ payoff matrix for player II . Then a joint distribution of strategies is an $m \times n$ matrix $z = (z_{ij})_{1 \leq i \leq m, 1 \leq j \leq n}$ such that $z_{ij} \geq 0$ for all $i \in \{1, \dots, m\}$, $j \in \{1, \dots, n\}$, and such that

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

So, z_{ij} represents the probability that player I plays i and player II plays j . If the distribution of strategies recommends that player I plays i , then the expected payoff by playing move i is $\sum_{j=1}^n z_{ij} a_{ij}$. If the distribution of strategies recommends that player I plays i , but they play k instead, then the expected payoff by playing move k is $\sum_{j=1}^n z_{ij} a_{kj}$. So, player I does not gain by deviating from the equilibrium strategy when

$$\sum_{j=1}^n z_{ij} a_{ij} \geq \sum_{j=1}^n z_{ij} a_{kj}, \quad \forall i \in \{1, \dots, m\}, \forall k \in \{1, \dots, m\}.$$

Similarly, player II does not gain by deviating from the equilibrium when

$$\sum_{i=1}^m z_{ij} b_{ij} \geq \sum_{i=1}^m z_{ik} b_{ik}, \quad \forall j \in \{1, \dots, n\}, \forall k \in \{1, \dots, n\}.$$

Exercise 4.9. Find all Correlated Equilibria for the Prisoner’s Dilemma.

Example 4.10. We continue Examples 4.3 and 4.7, recalling that $A = \begin{pmatrix} 4 & 0 \\ 0 & 1 \end{pmatrix}$, $B = \begin{pmatrix} 1 & 0 \\ 0 & 4 \end{pmatrix}$. We show that $z = \begin{pmatrix} 1/2 & 0 \\ 0 & 1/2 \end{pmatrix}$ is a correlated equilibrium. In the Battle of the Spouses, this corresponds to the Spouses flipping a single coin, and deciding where to go based on the coin flip.

$$\begin{aligned} \sum_{j=1}^2 z_{1j} a_{1j} &= (1/2)4 > 0 = \sum_{j=1}^2 z_{1j} a_{2j}. \\ \sum_{j=1}^2 z_{2j} a_{2j} &= (1/2)1 > 0 = \sum_{j=1}^2 z_{2j} a_{1j}. \end{aligned}$$

Similarly,

$$\sum_{i=1}^2 z_{i1} b_{i1} > \sum_{i=1}^2 z_{i1} b_{i2}, \quad \sum_{i=1}^2 z_{i2} b_{i2} > \sum_{i=1}^2 z_{i2} b_{i1}.$$

This Correlated equilibrium has payoff $(1/2)4 + (1/2)1 = 5/2$ for each player. This payoff is much higher than the $4/5$ payoff for each player which arose from the mixed Nash equilibrium.

Recall that this game has Nash equilibria $(1, 0)$, $(0, 1)$; $(0, 1)$, $(0, 1)$; and $(4/5, 1/5)$, $(1/5, 4/5)$. All of these equilibria are also Correlated equilibria.

Remark 4.11. Any Nash equilibrium is a Correlated Equilibrium.

Proof. We argue by contradiction. Suppose (\tilde{x}, \tilde{y}) is a Nash equilibrium. Let $z = \tilde{x}\tilde{y}^T$. Suppose for the sake of contradiction that z is not a correlated equilibrium. Then the negation of the definition of correlated equilibrium holds. Without loss of generality, the negated condition applies to player I . That is, there exists $i, k \in \{1, \dots, m\}$ such that

$$\sum_{j=1}^n z_{ij} a_{ij} < \sum_{j=1}^n z_{kj} a_{kj}.$$

That is,

$$\tilde{x}_i \sum_{j=1}^n \tilde{y}_j a_{ij} < \tilde{x}_i \sum_{j=1}^n \tilde{y}_j a_{kj}. \quad (*)$$

This inequality suggests that Player I can benefit by switching from strategy i to strategy k in the mixed strategy \tilde{x} . Let $e_i \in \Delta_m$ denote the vector with a 1 in the i^{th} entry and zeros in all other entries. Define $x \in \Delta_m$ so that $x = \tilde{x} - \tilde{x}_i e_i + \tilde{x}_i e_k$. Observe that

$$x^T A \tilde{y} - \tilde{x}^T A \tilde{y} = (-\tilde{x}_i e_i + \tilde{x}_i e_k)^T A \tilde{y} = -\tilde{x}_i \sum_{j=1}^n a_{ij} \tilde{y}_j + \tilde{x}_i \sum_{j=1}^n a_{kj} \tilde{y}_j \stackrel{(*)}{>} 0.$$

But this inequality contradicts that (\tilde{x}, \tilde{y}) is a Nash equilibrium. \square

Remark 4.12. To state Remark 4.11 in another way, Nash equilibria are exactly Correlated Equilibria that are product distributions. That is, Nash equilibria are Correlated Equilibria z of the form $z = \tilde{x}\tilde{y}^T$, where $\tilde{x} \in \Delta_m$ and $\tilde{y} \in \Delta_n$.

Exercise 4.13. Any convex combination of Nash equilibria is a Correlated Equilibrium. That is, if $z(1), \dots, z(k)$ are Nash Equilibria, and if $t_1, \dots, t_k \in [0, 1]$ satisfy $\sum_{i=1}^k t_i = 1$, then $\sum_{i=1}^k t_i z(i)$ is a Correlated Equilibrium.

Remark 4.14. A Correlated Equilibrium can be found via linear programming, since the conditions for finding the matrix z are all linear in the entries of z .

From Remarks 4.11 and 4.12, a Nash equilibrium is always a Correlated Equilibrium. However, there are Correlated equilibria that are not Nash equilibria, as we demonstrate below. However, already from Example 4.10, we see that the correlated equilibrium z is not a Nash equilibrium, since for any $x, y \in \Delta_2$, $z \neq xy^T$.

Example 4.15. In the **Game of Chicken**, as depicted in the James Dean movie *Rebel Without a Cause*, two cars are speeding towards each other. Each player chooses to chicken out (C) by swerving away, or they can continue driving straight (D). Each player would

prefer to continue driving while the other chickens out. However, if both players choose to continue driving, a catastrophe occurs. The payoffs follow:

		Player II	
		C	D
Player I	C	(6, 6)	(2, 7)
	D	(7, 2)	(0, 0)

Note that (C, D) is a pure Nash equilibrium, as is (D, C) . Also, there is a symmetric mixed Nash equilibrium where player I chooses C with probability $2/3$ and they choose D with probability $1/3$ (and player II acts in the same way). The expected payoff from this strategy for each player is

$$(2/3)^2(6) + (2/3)(1/3)(2) + (1/3)(2/3)(7) + (1/3)^2(0) = \frac{24}{9} + \frac{4}{9} + \frac{14}{9} = \frac{42}{9} = 4 + \frac{2}{3}.$$

Each player can actually ensure a higher expected payoff using a Correlated Equilibrium. This particular Correlated Equilibrium is not a Nash equilibrium. This Correlated Equilibrium corresponds to the matrix

$$z = \begin{pmatrix} 1/3 & 1/3 \\ 1/3 & 0 \end{pmatrix}.$$

Since z cannot be written as xy^T where $x, y \in \Delta_2$, we know that z is not a Nash equilibrium by Remark 4.12. For this equilibrium, the expected payoff for player I is

$$\sum_{i,j=1}^2 z_{ij} a_{ij} = (1/3)(6) + (1/3)(2) + (1/3)(7) = 15/3 = 5.$$

Similarly, player II has an expected payoff of 5. We can think of this Correlated Equilibrium as some mediator suggesting strategies to both players. This mediator recommends each of the strategies (C, C) , (C, D) and (D, C) to the players with equal probability $1/3$.

Let's sketch why this strategy should be a Correlated Equilibrium. Suppose the mediator tells player I to play D . In this case, player I knows that player II must be playing C , so there is no incentive for player I to deviate from the strategy. Now, suppose the mediator tells player I to play C . Given this information, player I knows there is an equal chance for player II to play C or D . That is, player I knows that player II will play C with probability $1/2$ and player II will play D with probability $1/2$. So, if player I follows the suggested strategy C , their expected payoff will be $(1/2)6 + (1/2)2 = 4$. However, if player I changes their strategy to D , their expected payoff will be $(1/2)(7) + (1/2)(0) = 7/2 < 4$. So, it is in the best interest of player I to maintain the suggested strategy of C .

Exercise 4.16. Find all Nash equilibria for the Game of Chicken. Prove that these are the only Nash equilibria. Then, verify that z described above is a Correlated Equilibrium. Can you find a Correlated Equilibrium such that both players have a payoff larger than 5? (Hint: when trying to find such a matrix z , assume that $z_{22} = 0$ and $z_{12} = z_{21}$.)

Exercise 4.17. In the Game of Chicken, you should have found only three Nash equilibria. Recall from Exercise 4.13 that any convex combination of Nash equilibria is a correlated

equilibrium. However, the converse is false in general! We can see this already in the Game of Chicken. Show that the Correlated Equilibrium

$$z = \begin{pmatrix} 1/3 & 1/3 \\ 1/3 & 0 \end{pmatrix}$$

is not a convex combination of the Nash equilibria. Put another way, the payoffs from this Correlated Equilibrium cannot be found by random choosing among the Nash equilibria.

4.3. Fixed Point Theorems. Fixed Point Theorems are used in many mathematical disciplines. They are often used when we want to prove the existence of some object, when perhaps finding or describing that object explicitly can be difficult. For example, certain Fixed Point Theorems give the existence of solutions to several differential equations. Also, Fixed Point Theorems can be used to prove the Inverse Function Theorem. Further below, we will use Brouwer's Fixed Point Theorem so show that at least one Nash equilibrium always exists for any general-sum game.

The following Theorem is discussed informally in calculus class, and it is proven rigorously in a first analysis class.

Theorem 4.18 (Intermediate Value Theorem). *Let $a, b \in \mathbb{R}$ with $a < b$. Let $f: [a, b] \rightarrow \mathbb{R}$ be a continuous function. For any real number y between $f(a)$ and $f(b)$, there exists some $x \in [a, b]$ such that $f(x) = y$.*

Proposition 4.19 (Brouwer Fixed-Point Theorem, One-Dimension). *Let $a, b \in \mathbb{R}$ with $a < b$. Let $f: [a, b] \rightarrow [a, b]$ be a continuous function. Then, there exists some $x \in [a, b]$ such that $f(x) = x$. The point x where $f(x) = x$ is called a **fixed point** of f .*

Proof. Define a function $g: [a, b] \rightarrow \mathbb{R}$ by $g(t) := f(t) - t$, where $t \in [a, b]$. Then g is continuous, since it is the sum of continuous functions. By the definition of f , we have $a \leq f(t) \leq b$ for all $t \in [a, b]$. In particular, $f(a) - a \geq 0$ and $f(b) - b \leq 0$. That is, $g(a) \geq 0$ and $g(b) \leq 0$. By the Intermediate Value Theorem, if we let $y = 0$, then there exists some $x \in [a, b]$ such that $g(x) = y = 0$. That is, $f(x) - x = 0$, i.e. $f(x) = x$. \square

Remark 4.20. Note that the set $[a, b]$ is closed, bounded and convex. These are the only three properties that will be used in the more general theorem. Also, note that the proof of Proposition 4.19 does not explicitly identify where the fixed point occurs. We are only told that some fixed point exists somewhere.

Theorem 4.21 (Brouwer Fixed-Point Theorem). *Let d be a positive integer. Let $K \subseteq \mathbb{R}^d$ be a closed, convex and bounded set. Let $f: K \rightarrow K$ be continuous. Then there exists some $x \in K$ such that $f(x) = x$.*

Remark 4.22. Each of the three hypotheses on K is necessary. In each of the following examples, the function f has no fixed point.

- Let $K = \mathbb{R}$ with $f(x) = x + 1$. Then K is closed and convex, but not bounded.
- Let $K = (0, 1)$ with $f(x) = x/2$. Then K is convex and bounded, but not closed.
- Let $K = \{(x_1, x_2) \in \mathbb{R}^2: x_1^2 + x_2^2 = 1\}$ with $f(x) = -x$. Then K is bounded and closed, but not convex, as we showed in Example 3.14.

There are a few different ways to prove Theorem 4.21. Perhaps the most direct proof uses Sperner's Lemma, which is a purely combinatorial statement.

Lemma 4.23 (Sperner's Lemma). (Case $d = 1$). Suppose the unit interval $[0, 1]$ is partitioned such that $0 = t_0 < t_1 < \dots < t_n = 1$, where each t_i is marked with a 1 or 2 whenever $0 < i < n$, t_0 is marked 1 and t_n is marked 2. Then the number of ordered pairs (t_i, t_{i+1}) , $0 \leq i < n$ with different markings is odd.

(Case $d = 2$). Suppose we divide a large triangle into smaller triangles, such that the intersection of any two adjacent triangles is a common edge of both. All vertices of the smaller triangles are labelled 1, 2 or 3. The three vertices of the large triangle are labelled 1, 2 and 3. Vertices of small triangles that lie on an edge of the large triangle must receive a label of one of the endpoints of that edge. Given such a labeling, the number of small triangles with three differently labeled vertices is odd; in particular, this number is nonzero.

(General case). Suppose we divide the simplex Δ_{d+1} into smaller simplices such that the intersection of any of the smaller adjacent simplices is a common face of both. All vertices of the smaller simplices are labelled $1, \dots, d+1$. The $d+1$ vertices of the large simplex are labelled $1, \dots, d$ and $d+1$. For any $0 < k < d+1$, vertices of small simplices that lie on any k -dimensional face must receive a label of one of the vertices of that face. Given such a labeling, the number of small simplices with $d+1$ differently labeled vertices is odd; in particular, this number is nonzero.

Proof of the case $d = 1$. Let $f(t_i)$ be the marking of t_i , for any $0 \leq i < n$. Using a telescoping sum, $1 = 1 - 0 = f(t_n) - f(t_0) = \sum_{i=0}^{n-1} [f(t_{i+1}) - f(t_i)]$. So, we have a sum of elements of $\{-1, 0, 1\}$ which add to 1. The total number of appearances of 1 and -1 must therefore be odd. The total number of appearances of 1 and -1 is equal to the number of ordered pairs (t_i, t_{i+1}) , $0 \leq i < n$, with different markings. \square

Proof of the case $d = 2$. Let M be the number of ordered pairs of small triangles, together with an edge on the small triangle with the labeling 1, 2. Let A be the number of small triangle edges with labeling 1, 2 that lie inside the boundary of the big triangle. Let B be the number of small triangle edges with labeling 1, 2 that lie in the interior of the big triangle. Then $M = A + 2B$, since M counts edges from A once, but it counts edges from B twice, since an edge inside the large triangle belongs to two small triangles. (But an edge contained in the boundary of the big triangle belongs to only one small triangle.)

Let N be the number of small triangles which use all labels 1, 2 and 3. Let N' be the number of small triangles which use the labels 1 and 2, but not 3. Then $M = N + 2N'$, since a triangle with labeling 1, 2, 3 has one edge with labeling 1, 2; a triangle with labeling 1, 1, 2 has two edges with labeling 1, 2; and a triangle with labeling 1, 2, 2 has two edges with labeling 1, 2. In summary,

$$N + 2N' = A + 2B.$$

From the case $d = 1$, we know that A is odd. Therefore, N is odd. \square

Proof of the general case. Let M be the number of ordered pairs of small simplices, together with a face with labeling $1, 2, \dots, d$. Let A be the number of small simplex faces with labeling $1, 2, \dots, d$ that lie inside the boundary of the big simplex. Let B be the number of small simplex faces with labeling $1, 2, \dots, d$ that lie in the interior of the big simplex. Then $M = A + 2B$, since M counts faces from A once, but it counts faces from B twice, since a face inside the large simplex belongs to two small simplices.

Let N be the number of small simplices with all labels $1, \dots, d+1$. Let N' be the number of small simplices that have all labels $1, \dots, d$ except $d+1$. So, N' is the number of small

simplices with exactly one label appearing twice. Then $M = N + 2N'$, since a simplex with labeling $1, 2, \dots, d+1$ has one face with labeling $1, 2, \dots, d$; while a simplex with exactly no label of $d+1$ and with one label appearing twice has two faces with labeling $1, 2, \dots, d$. In summary,

$$N + 2N' = A + 2B.$$

From the inductive hypothesis (inducting on the dimension d), A is odd. Therefore, N is odd, as desired. \square

Definition 4.24. Let $K \subseteq \mathbb{R}^d$ be a closed and bounded set. We define the **diameter** of K , denoted by $\text{diam}(K)$, to be the maximum of $\|x - y\|$ over all $x, y \in K$. This maximum exists by the Extreme Value Theorem, Theorem 3.23 applied to the continuous function $f(x, y) = \|x - y\|$ over the closed and bounded set $K \times K$. (The set $K \times K$ is closed and bounded by Exercise 3.20)

The following Lemma is typically proven in a second analysis class.

Lemma 4.25. *Let $K_1 \supseteq K_2 \supseteq K_3 \supseteq \dots$ be a sequence of nonempty closed and bounded sets in \mathbb{R}^d such that $\text{diam}(K_i) \rightarrow 0$ as $i \rightarrow \infty$. Then there exists a unique point $x \in \mathbb{R}^d$ such that $x \in K_i$ for all $i \geq 1$. In fact, if $x^{(i)} \in K_i$ for all $i \geq 1$, then $\lim_{i \rightarrow \infty} x^{(i)} = x$.*

Sperner's Lemma allows us to prove Brouwer's fixed-point theorem.

Proof of Theorem 4.21. We first assume that $K = \Delta_{d+1}$. We argue by contradiction. Suppose $f: K \rightarrow K$ is continuous and $f(x) \neq x$ for all $x \in K$.

Let S be any subdivision of K into smaller simplices such that the diameter of the smaller simplices is at most half the diameter of the large simplex. We also assume that if any of the small simplices intersect each other, then they share a common face.

We now create a labeling. Fix $i \geq 1$. For each vertex $x = (x_1, \dots, x_{d+1})$ of a small simplex, give x any label $i \in \{1, \dots, d+1\}$ such that $(f(x))_i < x_i$. Since $f(x) \neq x$ and $\sum_{i=1}^{d+1} x_i = \sum_{i=1}^{d+1} (f(x))_i = 1$, there must be some $i \in \{1, \dots, d+1\}$ such that $(f(x))_i < x_i$, otherwise we would have $f(x) = x$. This labeling satisfies the hypothesis of Sperner's Lemma. If x is a vertex of K , e.g. $x = (0, \dots, 0, 1, 0, \dots, 0)$, with a 1 in the i^{th} coordinate, then x must be labelled with i , since i is the only coordinate where we could possibly have $(f(x))_i < x_i$. Similarly, if x lies in any k -dimensional face of K with $0 < k < d$, then the only nonzero coordinates of x are those corresponding to the vertices of K lying in that face. So, x must receive a labeling among the vertices in that face.

Applying Sperner's Lemma, Lemma 4.23, there exists a small simplex that receives all labels $1, \dots, d+1$. We now iterate this process. Starting from this small simplex, subdivide it into even smaller simplices such that the diameter of the even smaller simplices is at most half the diameter of the small simplex. We can then create a labeling as before, and find an even smaller simplex that receives all labels $1, \dots, d+1$. In this way, we have a sequence of simplices K_1, K_2, K_3, \dots such that $K_1 \supseteq K_2 \supseteq K_3 \supseteq \dots$ and such that $\text{diam}(K_{i+1}) \leq \text{diam}(K_i)/2$ for all $i \geq 1$. Also, for each $i \geq 1$, the vertices of K_i receive all labels $1, \dots, d+1$.

For each $i \geq 1$, let $x^{(ij)} \in K_i$ denote the vertex of K_i that receives the label $j \in \{1, \dots, d+1\}$. That is, $(f(x^{(ij)}))_j < x_j^{(ij)}$. By Lemma 4.25, there exists a point $x \in \Delta_{d+1}$ such that $x \in K_i$ for all $i \geq 1$. Also, $\lim_{i \rightarrow \infty} x^{(ij)} = x$, for every $j \in \{1, \dots, d+1\}$. Since f is

continuous, since limits preserve non-strict inequalities and since $(f(x^{(ij)}))_j < x_j^{(ij)}$, we get

$$(f(x))_j = \lim_{i \rightarrow \infty} f(x^{(ij)})_j \leq \lim_{i \rightarrow \infty} x_j^{(ij)} = x_j.$$

That is, $(f(x))_j \leq x_j$ for all $j \in \{1, \dots, d+1\}$. Since $\sum_{i=1}^{d+1} x_i = \sum_{i=1}^{d+1} (f(x))_i = 1$, we conclude that $x = f(x)$, a contradiction. Having achieved a contradiction, we have proven the Theorem when $K = \Delta_{d+1}$.

To extend the Theorem to any K , let Δ be a copy of a simplex such that $\Delta \subseteq K$ and such that any simplex of dimension larger than Δ cannot be contained in K . Let x be an interior point of Δ . Translating Δ and K if necessary, we may assume $x = 0$. We first define a function $H: \Delta \rightarrow K$. Define $H(0) = 0$. Let $y \in \Delta$ with $y \neq 0$. Let a be the intersection of $\partial\Delta$ with the line segment starting at x and going through y , and let b be the intersection of ∂K with the line segment starting at x and going through y . Define $H(y) = y \|b\| / \|a\|$. Note that $H(a) = b$, since a, b lie on the same line segment. In words, H “stretches out” Δ onto K . Then $H: \Delta \rightarrow K$ is continuous with continuous inverse. And if $f: K \rightarrow K$ has a fixed point x , then so does $H^{-1} \circ f \circ H$. (Namely, $H^{-1}(x)$ is a fixed point of the map $H^{-1} \circ f \circ H$.) But $H^{-1} \circ f \circ H: \Delta \rightarrow \Delta$ is continuous, so this map has no fixed point. We conclude that f also does not have a fixed point, as desired. \square

4.4. Games with More than Two Players.

Definition 4.26 (General Sum Game, k players). Let $k \geq 2$ be an integer. A general sum game with k players is defined as follows. We label each of the k players as integers in $\{1, \dots, k\}$. Each player $i \in \{1, \dots, k\}$ has a set $S_i = \{1, 2, \dots, |S_i|\}$ of pure strategies and a payoff function $F_i: S_1 \times S_2 \times \dots \times S_k \rightarrow \mathbb{R}$. That is, if for each $i \in \{1, \dots, k\}$, player i uses the pure strategy $\ell_i \in S_i$, then player j has payoff $F_j(\ell_1, \dots, \ell_k)$, for every $j \in \{1, \dots, k\}$.

Definition 4.27 (Symmetric Game). Let $k \geq 2$ be an integer. A general sum game with k players is **symmetric** if $S_1 = \dots = S_k$, and for every $i', j' \in \{1, \dots, k\}$, there is a permutation $\pi: \{1, \dots, k\} \rightarrow \{1, \dots, k\}$ such that $\pi(i') = j'$ and such that

$$F_{\pi(i)}(\ell_{\pi(1)}, \dots, \ell_{\pi(k)}) = F_i(\ell_1, \dots, \ell_k), \quad \forall \ell_1, \dots, \ell_k \in S_1, \quad \forall i \in \{1, \dots, k\}.$$

Definition 4.28 (Nash equilibrium). A **pure Nash equilibrium** for a k -player game is a set of pure strategies for each set of players

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

such that, for every $j \in \{1, \dots, k\}$ and for every $\ell_j \in S_j$,

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

A **mixed strategy** for a k -player game is a set of vectors $(x^{(1)}, \dots, x^{(k)})$ such that $x^{(i)} \in \Delta_{|S_i|}$. The expected payoff for player j for the mixed strategy $(x^{(1)}, \dots, x^{(k)})$ is

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

We say a mixed strategy $(x^{(1)}, \dots, x^{(k)})$ is a **Nash equilibrium** if: for any $j \in \{1, \dots, k\}$, and for any $x \in \Delta_{|S_j|}$,

$$F_j(x^{(1)}, \dots, x^{(j-1)}, x, x^{(j+1)}, \dots, x^{(k)}) \leq F_j(x^{(1)}, \dots, x^{(k)}).$$

4.5. Nash's Theorem. Implicit in von Neumann's Minimax Theorem, Theorem 3.29, is that the optimal strategies exist for two-person zero-sum games. That is, there is some way that "rational" players should act. Unfortunately, the Minimax Theorem no longer applies to general-sum games, as we have seen on the homework. So, the minimax does not help us to find the strategies of "rational" players in general-sum games. Giving a rigorous definition of a "rational" strategy for players in general sum games was a major open problem, which Nash solved. He solved this problem by defining the Nash equilibrium.

However, at the moment, it may be possible that a Nash equilibrium does not exist at all. If no Nash equilibrium exists in general, then the definition of the Nash equilibrium may not be so useful for investigating the strategies of "rational" players. Thankfully, as Nash showed in Theorem 4.29, at least one Nash equilibrium always exists. The proof below, which uses Brouwer's Fixed Point Theorem, is more or less identical to Nash's original proof. The original proof of von Neumann's Minimax Theorem also used a (different) fixed point theorem, so the use of such a tool in this subject was not unprecedented.

Theorem 4.29 (Nash's Theorem). *For any general-sum game with $k \geq 2$ players, there exists at least one Nash equilibrium.*

Proof for $k = 2$. Let m, n be positive integers. Suppose the players have $m \times n$ payoff matrices A and B . Let $A_{(i)}$ denote the i^{th} row of A , and let $B^{(j)}$ denote the j^{th} column of B . Let $K = \Delta_m \times \Delta_n$. Note that K is closed and bounded by Exercises 3.19 and 3.20. We define a map $F: K \rightarrow K$. For any $(x, y) \in K$, and for any $1 \leq i \leq m$, $1 \leq j \leq n$, define

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

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

Then c_i represents the payoff of player I switching from strategy x to the pure strategy i . Similarly, d_j represents the payoff of player II switching from strategy y to the pure strategy j . We can now define F . For any $(x, y) \in K$, define $F(x, y) := (\tilde{x}, \tilde{y}) \in K$, where, for any $1 \leq i \leq m$, $1 \leq j \leq n$,

$$\tilde{x}_i = \frac{x_i + c_i}{1 + \sum_{k=1}^m c_k}, \quad \tilde{y}_j = \frac{y_j + d_j}{1 + \sum_{k=1}^n d_k}.$$

(The denominators are chosen so that $\tilde{x} \in \Delta_m$ and $\tilde{y} \in \Delta_n$.) Note that F is continuous, since the numerator of each component of (\tilde{x}, \tilde{y}) is continuous, and each denominator is also continuous and at least equal to 1. So, we can apply Brouwer's Fixed-Point Theorem, Theorem 4.21 to find some $(x, y) \in K$ where $(x, y) = F(x, y) = (\tilde{x}, \tilde{y})$. We claim $c_i = d_j = 0$ for all $1 \leq i \leq m$, $1 \leq j \leq n$. To see this, suppose for example that $c_1 > 0$, so $A_{(1)}y - x^T A y > 0$. Let $\ell \in \{1, \dots, m\}$ so that $A_{(\ell)}y = \min\{A_{(i')}y : x_{i'} > 0, 1 \leq i' \leq m\}$. Then

$$A_{(1)}y > x^T A y = \sum_{i=1}^m x_i A_{(i)}y \geq \sum_{i=1}^m x_i A_{(\ell)}y = A_{(\ell)}y.$$

So, by definition of c_ℓ , we have $c_\ell = 0$. Using $(x, y) = F(x, y)$, $c_1 > 0$ and $x_\ell > 0$, we have

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

a contradiction. We conclude that $c_1 = 0$. Similarly, all other c_i and d_j are zero, for all $1 \leq i \leq m$, $1 \leq j \leq n$.

In summary, for all $1 \leq i \leq m$, $1 \leq j \leq n$, we have

$$A_{(i)}y \leq x^T Ay, \quad x^T B^{(j)} \leq x^T By$$

So, for any $x' \in \Delta_m$,

$$x'^T Ay = \sum_{i=1}^m x'_i A_{(i)}y \leq \sum_{i=1}^m x'_i x^T Ay = x^T Ay.$$

Similarly, for any $y' \in \Delta_n$,

$$x^T By' = \sum_{j=1}^n x^T B^{(j)} y'_j \leq \sum_{j=1}^n (x^T By) y'_j = x^T By.$$

That is, (x, y) is a Nash equilibrium. \square

Proof sketch for general k . The same proof works for general k with minor modifications. Suppose we have $k > 2$ players. For each $1 \leq j \leq k$, let $n(j)$ be the number of strategies for player j , so that player j has mixed strategies in $\Delta_{n(j)}$. Let $x^{(j)} \in \Delta_{n(j)}$. For each $1 \leq i, j \leq k$, define a function $c_i^{(j)}(x^{(1)}, \dots, x^{(k)})$, so that $c_i^{(j)}$ is the maximum of zero and the gain when player j switches from strategy $x^{(j)}$ to the pure strategy i . The remaining parts of the proof are essentially the same, where now $K = \Delta_{n(1)} \times \dots \times \Delta_{n(k)}$, and $F: K \rightarrow K$. \square

Corollary 4.30. *For any symmetric general-sum game with $k \geq 2$ players, there exists at least one symmetric Nash equilibrium. (At this equilibrium, the payoffs are equal for all players.)*

Proof. We repeat the proof above. Since the game is symmetric, we have $n(1) = \dots = n(k)$. However, instead of defining F on K , we define F on the smaller diagonal set $D \subseteq K$, where

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

Defining F as before, we now have $F: D \rightarrow D$, since, in a symmetric game, $c_i^{(1)}(x, \dots, x) = \dots = c_i^{(k)}(x, \dots, x)$ for all $1 \leq i \leq k$ and for all $x \in \Delta_n$. Brouwer's Fixed-Point Theorem then gives a point $(x, \dots, x) \in D$ such that $F(x, \dots, x) = (x, \dots, x)$, as desired. \square

From Remark 4.11, we deduce the following Corollary of Theorem 4.29.

Corollary 4.31. *For any general-sum game with $k \geq 2$ players, there exists at least one Correlated Equilibrium.*

Remark 4.32. Note that Nash's Theorem uses the Brouwer Fixed Point Theorem, and Remark 4.20 says that Brouwer's Theorem gives no explicit way of finding the fixed point of a function. So, the proof of Nash's Theorem similarly does not give an explicit way of finding a Nash equilibrium. We are only guaranteed that this equilibrium occurs somewhere. In fact, the Nash equilibrium can be difficult to find explicitly in general.

4.6. Evolutionary Game Theory.

Example 4.33 (Hawks and Doves). The game of Hawks and Doves is a variant of the Game of Chicken. The present game models two different behaviors in a single species population. Among the single species, an individual can behave aggressively (like a hawk, (H)) or an individual can behave passively (like a dove, (D)). Let v, c be constants. The payoff matrices follow:

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

Individuals in the population who succeed in the game produce more offspring. We then interpret the Nash equilibrium as specifying the proportion of hawks versus doves in a population of the species, after many generations exist. (Since the game is symmetric, at least one symmetric Nash equilibrium exists, and we can interpret this equilibrium as a proportion of hawks versus doves.)

We split into two cases, depending on the values of v, c . If $c < v/2$, then Player I always improves their payoff by choosing H , no matter what Player II does. Similarly, Player II always prefers playing H , no matter what Player I does. So, (H, H) is a pure Nash equilibrium. It can also be shown that there are no other Nash equilibria. In fact, when $c < v/2$, Hawks and Doves is essentially the same as the Prisoner's Dilemma.

We now consider the case $c > v/2$. In this case, Hawks and Doves is essentially the same as the Game of Chicken. There are two pure Nash equilibria: (H, D) and (D, H) . Since the game is symmetric, there exists a symmetric mixed Nash equilibrium. The symmetric equilibrium has equal payoffs for both players. To find this equilibrium, suppose player I plays H with probability p where $0 < p < 1$. Then player I plays D with probability $1 - p$, and $(p, 1 - p), (p, 1 - p)$ is a Nash equilibrium. That is, if A denotes the payoff matrix for player I ,

$$(p, 1 - p)A \begin{pmatrix} p \\ 1 - p \end{pmatrix} = \max_{t \in [0, 1]} (t, 1 - t)A \begin{pmatrix} p \\ 1 - p \end{pmatrix}.$$

The function $f(t) = (t, 1 - t)A(p, 1 - p)^T$ is linear in t , where $f: [0, 1] \rightarrow \mathbb{R}$. Since $f(t)$ achieves its maximum on $[0, 1]$ at $p \in (0, 1)$, and since f is linear, we conclude that $f(0) = f(1) = f(p)$. That is, the payoff for player I playing H is the same as their payoff for playing D . Writing out $f(0) = f(1)$ gives

$$p(\frac{v}{2} - c) + (1 - p)v = (1 - p)\frac{v}{2}. \quad (*)$$

Solving for p , we get $p = v/(2c)$. Since $c > v/2$, we have $p < 1$. Since this equilibrium is symmetric, player II has the same strategy as player I .

4.6.1. Population Dynamics. We now re-interpret the constant p . Suppose we have a large population of a single species. Let p be the fraction of Hawks in the population. If the population is very large, and interactions among the population are random, then the Law of Large Numbers from Probability Theory says: the total payoff of Hawks is proportional to the expected payoff of a single Hawk playing against an opponent whose mixed strategy is $(p, 1 - p)$. Similarly, the total payoff of Doves is proportional to the expected payoff of a single Dove playing against an opponent whose mixed strategy is $(p, 1 - p)$. More specifically, the total payoff of Hawks divided by the total payoff of Doves is equal to the left side of $(*)$, divided by the right side of $(*)$.

Now, if $p < v/(2c)$, then the left side of $(*)$ exceeds the right side. That is, the total payoff of Hawks exceeds the total payoff of Doves. So, the number of Hawks will increase, i.e. p will increase.

If $p > v/(2c)$, then the left side of (*) is smaller than the right side. That is, the total payoff of Hawks is less than the total payoff of Doves. So, the number of Hawks will decrease, i.e. p will decrease. This result might be surprising, since there are a lot of Hawks, but the proportion of Doves in the population is increasing. We could explain this observation by noting that, when there are many Hawks, they will injure each other a lot.

In any case, note that p seems to want to converge to $p = v/(2c)$. Let's try to capture this behavior with the following definition.

Definition 4.34 (Evolutionarily Stable Strategy). Suppose we have a two-player symmetric game (so that the payoff matrix for player I is A , the payoff matrix for player II is B , and with $A = B^T$). Assume that A, B are $n \times n$ matrices. A mixed strategy $x \in \Delta_n$ is said to be an **evolutionarily stable strategy** if, for any pure strategy w (which is called a “mutant” strategy), we have

$$\begin{aligned} w^T A x &\leq x^T A x, \\ \text{If } w^T A x &= x^T A x, \text{ then } w^T A w < x^T A w. \end{aligned}$$

To motivate the definition, suppose x is the evolutionarily stable strategy, and w is any pure strategy. Suppose the population with strategy x is invaded by a small population with strategy w . Let $\varepsilon > 0$. The new population then has strategy $\varepsilon w + (1 - \varepsilon)x$. And the new payoffs are

$$\begin{cases} \varepsilon x^T A w + (1 - \varepsilon)x^T A x & , \text{ for individuals with strategy } x \\ \varepsilon w^T A w + (1 - \varepsilon)w^T A x & , \text{ for individuals with strategy } w \end{cases}$$

Now if $\varepsilon > 0$ is small, then the definition of evolutionarily stable strategy shows that the payoff for individuals with strategy x exceeds the payoff for individuals with strategy w . So, the population w will die out.

Remark 4.35. The condition $w^T A x \leq x^T A x$ implies that x is a Nash equilibrium.

Remark 4.36. Strictly speaking, by definition, if x is a pure Nash equilibrium, then x is not evolutionarily stable, since we can choose $w = x$ and then the second condition in the definition does not hold.

Example 4.37. Consider again Hawks and Doves where $c > v/2$. Let's show that the mixed strategy $x = (v/(2c), 1 - v/(2c))$ is evolutionarily stable. Let $w \in \Delta_2$ be a pure strategy. The first condition $x^T A x \geq w^T A x$ follows since x is a Nash equilibrium. We now check the second condition. It suffices to show that $w^T A w < x^T A w$ for any pure strategy w . Recall that $p = v/(2c)$ and $0 < p < 1$.

If $w = (1, 0)$, then $w^T A w = (v/2) - c < p(v/2 - c) + (1 - p)(0) = x^T A w$, using $p \in (0, 1)$ and $v/2 - c < 0$.

If $w = (0, 1)$, then $w^T A w = v/2 < pv + (1 - p)(v/2) = x^T A w$, using $p > 0$.

Exercise 4.38. Recall that the game of Rock-Paper-Scissors is defined by the payoff matrices

$$A = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix}, \quad B = A^T.$$

Then the game is symmetric. (And also, note that $A + B = 0$, so that the game is a zero-sum game.)

Show that $(1/3, 1/3, 1/3)$ is the unique Nash equilibrium. Then, show that this Nash equilibrium is **not** evolutionarily stable.

This observation leads to interesting behaviors in population dynamics. A certain type of lizard has three kinds of sub-species whose interactions resemble the Rock-Paper-Scissors game. The dynamics of the population cycle between large, dominant sub-populations of each of the three sub-species. That is, first the “Rock” lizards are a majority of the population, then the “Paper” lizards become the majority, then the “Scissors” lizards become the majority, and then the “Rock” lizards become the majority, and so on.

4.7. Signaling and Asymmetric Information. So far, we have considered games where both parties know the rules of the game. However, in politics, poker, etc. this assumption is not really true. Even simplified versions of poker are a bit difficult to analyze, so for now we only discuss a simplified example of signaling from biology.

In the wild, antelopes are sometimes seen jumping very high. One theory postulates that this behavior shows predators that the antelope is physically fit. If the antelope is physically fit, it is not worthwhile to try to chase the antelope. That is, the antelope is **signaling** a trait to others, for their own benefit.

Example 4.39 (Lions and Antelopes). Suppose an antelope sees a lion in the distance. Assume there are two kinds of antelopes: healthy ones (H) and weak ones (W). If the lion chases a healthy antelope, it will not catch the antelope, and both will expend a lot of energy. If the lion chases a weak antelope, it will catch the antelope. The antelope knows whether or not it is healthy, but the lion cannot know. The lion can choose to chase the antelope, or ignore it. If the antelope is healthy, the payoff matrices are described by A^H below; if the antelope is weak, the payoff matrices are described by A^W below. The lion does not know which game it is playing, but it does know it will be playing with one of the two payoff matrices:

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

 $A^H =$

		antelope	
		chase	run-if-chased
lion	chase	$(5, -1000)$	
	ignore	$(0, 0)$	

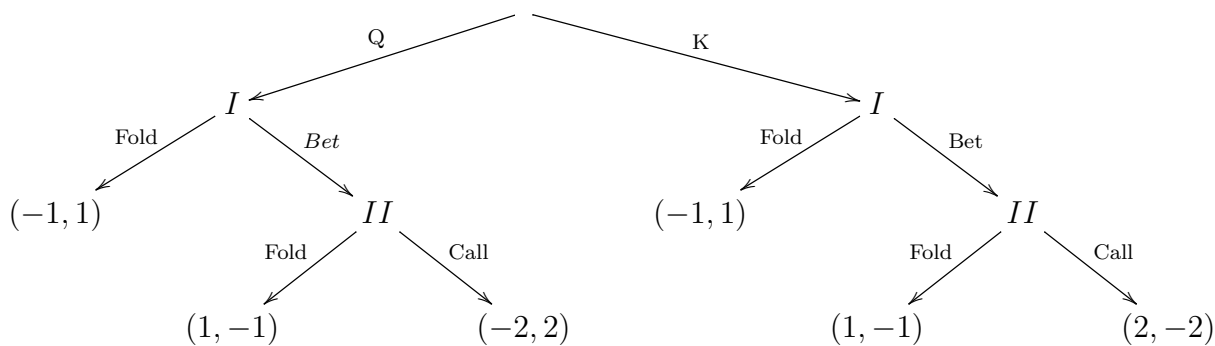
 $A^W =$

Suppose 20% of antelopes are weak. By choosing to chase, the lion can expect a payoff of $(.8)(-1) + (.2)(5) = .2$. So, the lion will have a positive payoff by chasing every antelope it sees. So it should be inclined to always chase instead of ignore, since ignoring always has a payoff of 0. But the lion cannot really know its true payoff ahead of time, since it does not know whether the antelope is healthy or weak. However, the antelope does know whether or not it is weak. For this reason, the game has **asymmetric information**. If a healthy antelope conveys its health by jumping very high, then the lion knows it should not in fact chase this antelope. Given that the antelope conveys their health, both parties benefit more when the lion chooses to ignore. So, when the healthy antelope jumps, it changes the optimal behavior of the lion.

Example 4.40 (A Simplified Game of Poker). We consider a simplified version of poker. Even in this simplified game, we will see that bluffing (i.e. risking money in an unfavorable game position) is part of an optimal strategy. Novice poker players tend to avoid bluffing

since it seems like it adds unnecessary risk. However, as we will see, the best strategy should really incorporate bluffing.

In the following simplified game of poker between two players, we have a deck of 2 cards, consisting of one queen (Q) and one king (K). Both players begin by putting one dollar each into the bank. Player *I* takes a single card from the deck, and looks at their card (so that Player *II* does not know which card has been drawn). Player *I* can now either bet or fold. If Player *I* folds, then the game ends and Player *II* wins the \$2 in the bank. If Player *I* bets, then Player *I* puts an additional dollar into the bank. Player *II* can now call or fold. If Player *II* folds, then the game ends and Player *I* wins the \$3 in the bank. If Player *II* calls, then Player *II* puts an additional dollar in the bank and the game ends. If Player *I* has a king, then Player *I* wins the \$4 in the bank. If Player *I* has a queen, then Player *II* wins the \$4 in the bank. The tree depicting all game scenarios appears below, together with payoffs. This tree is called the **Kuhn Tree** of the game.



For example, if Player *I* draws a Q and then Folds, then the payoff is $(-1, 1)$. That is, Player *I* has a payoff of -1 (since they lost one dollar), and Player *II* has a payoff of 1 (since they gained one dollar).

Alternatively, let's describe all possible strategies with payoff matrices. We label strategies of Player *II* as Call or Fold. We also label strategies of Player *I* by two letters; the first letter denotes their action given that they draw Q, and the second letter denotes their action given that they draw K. So, the strategy BF says that Player *I* will Bet if they draw Q, and they will Fold if they draw K.

We then write the expected payoffs for each pair of strategies in the following matrix.

		Player <i>II</i>	
		Call	Fold
Player <i>I</i>	BB	$(0, 0)$	$(1, -1)$
	FB	$(1/2, -1/2)$	$(0, 0)$
	BF	$(-3/2, 3/2)$	$(0, 0)$
	FF	$(-1, 1)$	$(-1, 1)$

For example, suppose Player *I* chooses FB and Player *II* chooses to Call. With probability $1/2$, Player *I* has Q, so that Player *I* folds and the payoffs are $(-1, 1)$. And with probability $1/2$, Player *I* has K, so that Player *I* wins and the payoffs are $(2, -2)$. So, the expected payoff of FB is $(1/2)(-1, 1) + (1/2)(2, -2) = (1/2, -1/2)$. The other payoffs are calculated similarly.

Domination for Player I 's payoffs shows that the strategies FF and BF are not optimal. So, Player I should only consider player strategies BB and FB, and the payoff matrix reduces to the following.

		Player II	
		Call	Fold
Player I	BB	$(0, 0)$	$(1, -1)$
	FB	$(1/2, -1/2)$	$(0, 0)$

By inspection, pure strategies do not give a Nash equilibrium. So, a “rational” strategy is a mixed Nash equilibrium. Suppose Player I plays BB with probability $p \in [0, 1]$ and Player II Calls with probability $q \in [0, 1]$. If A denotes the payoff matrix for player I ,

$$(p, 1-p)A \begin{pmatrix} q \\ 1-q \end{pmatrix} = \max_{t \in [0,1]} (t, 1-t)A \begin{pmatrix} q \\ 1-q \end{pmatrix}.$$

The function $f(t) = (t, 1-t)A \begin{pmatrix} q \\ 1-q \end{pmatrix}$ is linear in t , where $f: [0, 1] \rightarrow \mathbb{R}$. Since $f(t)$ achieves its maximum on $[0, 1]$ at $p \in (0, 1)$, and since f is linear, we conclude that $f(0) = f(1) = f(p)$. That is, the payoff for Player I playing BB is the same as their payoff for playing FB . Writing out $f(0) = f(1)$ gives $q/2 = 1 - q$, so that $q = 2/3$. Similarly, $p = 1/3$.

In conclusion, the optimal strategy is the mixed Nash equilibrium given by $(1/3, 2/3)$ and $(2/3, 1/3)$. In particular, Player I should play the bluffing strategy BB $1/3$ of the time.

Games with Asymmetric Information can be quite complicated in general. We will eventually cover more complicated examples. There are some general theories of asymmetric information that are quite complicated, so we will not discuss them. The theory of auctions, though limited in scope, is perhaps the most complete theory of games with asymmetric information.

5. PRICE OF ANARCHY

As we have seen above, even simple games such as the Game of Chicken can have multiple Nash equilibria, so it then becomes unclear what “rational” players would do. One way to sidestep this issue is to examine all possible payoffs from all Nash equilibria, and check how much they vary. If all Nash equilibria have similar payoffs, then the exact choice of Nash equilibria is not so significant. It is then natural to study the ratios of payoffs of different equilibria, as a measure of “inefficiency.”

Definition 5.1. Suppose we have a general sum game with k players. Denote the set of all of possible mixed strategies of all players as S , and denote the set of Nash equilibria as NE . Let $f: S \rightarrow \mathbb{R}$ be continuous. The **price of anarchy** of the game with respect to f is the ratio

$$\frac{\max_{s \in S} f(s)}{\min_{s \in NE} f(s)}.$$

Since $S \subseteq NE$, the price of anarchy is at least 1.


Typically, we take f to be the payoff of a fixed player in the game, or the sum of payoffs of all players in the game, so the price of anarchy is the ratio of the worst possible payoff to the best possible payoff. We also typically choose f to be a positive function. In this case,

the price of anarchy is ideally close to one, since then selfish behavior of players is close to being societally optimal.

The related **price of stability** is

$$\frac{\max_{s \in S} f(s)}{\max_{s \in \text{NE}} f(s)}.$$

Example 5.2 (Pigou's Example). Consider a network with two vertices v_1, v_2 (cities) and two edges (roads) both going from v_1 to v_2 . Each edge has a cost which describes the time it takes for a driver to traverse that road. The first edge has a cost of 1, which represents a long road without any traffic. The second edge has a cost $t \geq 0$ where t denotes the amount of traffic on that road. The second edge represents a short but possibly faster highway.




Suppose there is one unit of traffic, representing a large number of players. Each player wants to go from v_1 to v_2 . Each player acts independently of each other player. And each player wants to minimize their travel time. Under the above assumptions, each player will travel on the second edge, as a Nash equilibrium, since it will always have a faster travel time. But then each player's travel time is $t = 1$. Put another way, each player's travel time would be the same if the second highway did not exist at all.

But is there a better strategy that is societally optimal? Suppose each player independently chooses the top road with probability $p \in [0, 1]$, and they choose the bottom road with probability $1 - p$. Then on average there is $t = 1 - p$ units of traffic on the bottom road, so the average travel time of each player is $1 \cdot p + (1 - p)^2$, which is minimized when $1 - 2(1 - p) = 0$, i.e. when $1 - p = 1/2$, i.e. $p = 1/2$. So, the average travel time of each player is $1/2 + 1/4 = 3/4$. Since we found only one Nash equilibrium, the price of anarchy and price of stability are both equal to $1/(3/4) = 4/3$.

In fact, in any routing game where the edge costs are affine functions, the price of anarchy is at most $4/3$. However, even in this Pigou example, if we change the second edge cost to t^q and let $q \rightarrow \infty$, then the price of anarchy becomes unbounded.


Example 5.3 (Braess Paradox). Pigou's example shows that adding an extra road might not improve the flow of traffic. In fact there is an example showing that adding an extra right can increase traffic!

Consider a network with four vertices v_1, v_2, v_3, v_4 (cities) and four edges (roads): $(v_1, v_3), (v_3, v_2), (v_1, v_4), (v_4, v_2)$. Suppose the edges have costs $t, 1, 1, t$, respectively.



Suppose there is one unit of traffic, representing a large number of players. Each player wants to go from v_1 to v_2 . Each player acts independently of each other player. And each

player wants to minimize their travel time. Under the above assumptions, the unique Nash equilibrium occurs when each player chooses the top or bottom path with probability $1/2$ each, for a mean travel time of $1/2 + 1 = 3/2$.



Suppose now we add a short and fast (one way) highway from v_3 to v_4 with zero cost. Any driver taking the v_1 to v_3 route will use the new highway to then proceed from v_4 to v_2 , since this is strictly better for such a driver, giving them a travel time of 2. In fact, the unique Nash equilibrium occurs when each player chooses this path, since any player starting on the v_1 to v_4 path would rather take the previously mentioned path, as its travel time is shorter. Since we found only one Nash equilibrium, the price of anarchy and price of stability are both equal to $2/(3/2) = 4/3$.

Example 5.4 (The Prisoner’s Dilemma). Recall Example 4.1 which had payoffs

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

Recall that domination implies that the bottom right corner is the unique Nash equilibrium. Let $f: S \rightarrow \mathbb{R}$ be the sum of payoffs of both players plus 18. (It makes more intuitive sense to look at the maxima/minima of positive functions f when discussing the price of anarchy.) Then a domination argument shows that $\max_{s \in S} f(x) = -2 + 18 = 16$, i.e. it is societally best for both players to stay silent. Since there is only one Nash equilibrium where f takes value $-16 + 18 = 2$, both the price of anarchy and price of stability are equal to $16/2 = 8$.

Different choices of f can lead to different ratios here, but the point is that the Nash equilibrium is not an optimal

6. COALITIONS AND SHAPLEY VALUE

In coalitional games, the players can form coalitions in order to achieve each of their goals. A central question is: what kind of “power” does each individual player have? For example, the US and China have a great deal of bargaining power at the UN, but how can we quantify this power?

Example 6.1 (A Glove Market). Suppose we have a three-player game, where Player I has a left-handed glove, and Players II and III each have a right-handed glove. Each of the Players is selling their glove. A customer comes to the players, and they would like to pay \$100 total for a pair of gloves. The Players negotiate with each other before the purchase. How much money can each Player rationally ask of the customer?

We first define a characteristic function v , which is supposed to measure the value that a subset of players has when they form a coalition.

Definition 6.2 (Characteristic Function). Let n be a positive integer. Suppose we have a game with n players. A **characteristic function** is a function $v: 2^{\{1, \dots, n\}} \rightarrow \mathbb{R}$ satisfying $v(\emptyset) = 0$. That is, v is a function on the subsets of the set $\{1, \dots, n\}$.

In the glove market example, v takes the values 0 and 1, with a domain consisting of all subsets of players. Formally, we write $v: 2^{\{1, 2, 3\}} \rightarrow \{0, 1\}$, so that $2^{\{1, 2, 3\}}$ denotes the set of (eight) subsets of $\{1, 2, 3\}$. The function v takes the value 1 exactly for subsets of players who can accomplish their goal. With the game defined above, this means v is 1 only when the subset of players has both a right-handed and a left-handed glove. That is,

$$1 = v(\{123\}) = v(\{12\}) = v(\{13\}).$$

And v takes value 0 on all other subsets of $\{1, 2, 3\}$. In particular, $v(\emptyset) = 0$.

Note that v is also monotone. That is, if $S, T \subseteq \{1, 2, 3\}$ satisfy $S \subseteq T$, then $v(S) \leq v(T)$. So, the value of a coalition increases when the size of the coalition increases.

Definition 6.3 (Shapley value). Suppose we have a game with n players together with a characteristic function $v: 2^{\{1, \dots, n\}} \rightarrow \mathbb{R}$. For each $i \in \{1, \dots, n\}$, we define the **arbitration value** or **Shapley value** $\phi_i(v) \in \mathbb{R}$ to be any set of real numbers satisfying the following four axioms:

- (i) **(Symmetry)** If for some $i, j \in \{1, \dots, n\}$ we have $v(S \cup \{i\}) = v(S \cup \{j\})$ for all $S \subseteq \{1, \dots, n\}$ with $i, j \notin S$, then $\phi_i(v) = \phi_j(v)$.
- (ii) **(No power/ no value)** If for some $i \in \{1, \dots, n\}$ we have $v(S \cup \{i\}) = v(S)$ for all $S \subseteq \{1, \dots, n\}$, then $\phi_i(v) = 0$.
- (iii) **(Additivity)** If u is any other characteristic function, then $\phi_i(v + u) = \phi_i(v) + \phi_i(u)$, for all $i \in \{1, \dots, n\}$.
- (iv) **(Efficiency)** $\sum_{i=1}^n \phi_i(v) = v(\{1, \dots, n\})$.

The quantity $\phi_i(v)$ is supposed to measure the “bargaining power” of player $i \in \{1, \dots, n\}$. From the definition of Shapley values, it is not clear whether or not the Shapley values exist. And even if they do exist, they may not be uniquely defined. Thankfully, the Shapley values exist and are unique.

Theorem 6.4 (Shapley). *There exists a unique solution $\phi = (\phi_1, \dots, \phi_n)$ satisfying the four axioms in Definition 6.3.*

Example 6.5 (A veto game). Let n be a positive integer. Fix a nonempty subset $T \subseteq \{1, \dots, n\}$. We define the **T -veto game** where effective coalitions are subsets that contain T . That is, we define a characteristic function $u_T: 2^{\{1, \dots, n\}} \rightarrow \mathbb{R}$ by

$$u_T(S) = \begin{cases} 1 & , \text{ if } T \subseteq S \\ 0 & , \text{ otherwise.} \end{cases}$$

(Note that $u_\emptyset(S) = 1$ for all $S \subseteq \{1, \dots, n\}$; in particular $u_\emptyset(\emptyset) = 1$ so u_\emptyset is not a characteristic function.)

Let's compute the Shapley values of u_T . Fix some $i \in \{1, \dots, n\}$ with $i \notin T$. If $S \subseteq \{1, \dots, n\}$, then $u_T(S \cup \{i\}) = u_T(S)$. (If $T \subseteq S$, then $T \subseteq S \cup \{i\}$; if T is not a subset of

S , then T is not a subset of $S \cup \{i\}$, since $i \notin T$.) So, Axiom (ii) implies that

$$\phi_i(u_T) = 0 \quad \text{if } i \notin T.$$

Now, let $i, j \in T$ with $i \neq j$ and let $S \subseteq \{1, \dots, n\}$ with $i, j \notin S$. Then $u_T(S \cup \{i\}) = u_T(S \cup \{j\}) = 0$. (Since $i, j \in T$ with $i, j \notin S$ and $i \neq j$, we know that T is not a subset of $S \cup \{i\}$, and T is not a subset of $S \cup \{j\}$.) So, Axiom (i) implies that

$$\phi_i(u_T) = \phi_j(u_T) \quad \text{if } i, j \in T.$$

By Axiom (iv), $\sum_{i=1}^n \phi_i(u_T) = u_T(\{1, \dots, n\}) = 1$, as long as T is nonempty. So, if $j \in T$, we have

$$1 = \sum_{i=1}^n \phi_i(u_T) = \sum_{i \in T} \phi_i(u_T) = |T| \phi_j(u_T).$$

In conclusion, if $i \in \{1, \dots, n\}$, we have

$$\phi_i(u_T) = \begin{cases} \frac{1}{|T|} & , \text{ if } i \in T \\ 0 & , \text{ if } i \notin T. \end{cases}$$

Remark 6.6. We note in passing that if $c \in \mathbb{R}$, then $\phi_i(cu_T) = c\phi_i(u_T)$.

Remark 6.7. For any $T \subseteq \{1, \dots, n\}$ we have used the notation $|T|$ to denote the number of elements of $\{1, \dots, n\}$ in T .

Example 6.8. Let's return to the glove market example. We begin with the equality

$$u_{\{1,2\}} + u_{\{1,3\}} = v + u_{\{1,2,3\}}.$$

Let $i \in \{1, 2, 3\}$. Using Axiom (iii), we have

$$\phi_i(u_{\{1,2\}}) + \phi_i(u_{\{1,3\}}) = \phi_i(v) + \phi_i(u_{\{1,2,3\}}).$$

In the case $i = 1$, this equality says $1/2 + 1/2 = \phi_1(v) + 1/3$, so that $\phi_1(v) = 2/3$. In the case $i = 2$, this equality says $1/2 + 0 = \phi_2(v) + 1/3$, so that $\phi_2(v) = 1/6$. Similarly, $\phi_3(v) = 1/6$.

Recalling that a person wanted to spend \$100 to buy a pair of gloves, we interpret the Shapley values (arbitration values) as follows. Player *I* can reasonably expect to earn 2/3 of the \$100, or around \$66; Player *II* can reasonably expect to earn 1/6 of the \$100, or around \$17, and Player *III* can reasonably expect to earn 1/6 of the \$100, or around \$17.

Theorem 6.4 follows from the following statement.

Theorem 6.9. *Let n be a positive integer. Suppose we have a game with n players together with a characteristic function $v: 2^{\{1, \dots, n\}} \rightarrow \mathbb{R}$. Let S_n denote the set of permutations on n elements (so $\pi \in S_n$ is a bijective function $\pi: \{1, \dots, n\} \rightarrow \{1, \dots, n\}$). Then, $\forall i \in \{1, \dots, n\}$,*

$$\phi_i(v) = \frac{1}{n!} \sum_{k=1}^n \sum_{\pi \in S_n: \pi(k)=i} [v(\{\pi(1), \dots, \pi(k)\}) - v(\{\pi(1), \dots, \pi(k-1)\})].$$

Remark 6.10. In the case that v is monotone and $\{0, 1\}$ -valued, either $v(\{\pi(1), \dots, \pi(k)\}) - v(\{\pi(1), \dots, \pi(k-1)\}) = 0$, or $v(\{\pi(1), \dots, \pi(k)\}) - v(\{\pi(1), \dots, \pi(k-1)\}) = 1$. The latter case only occurs when $v(\{\pi(1), \dots, \pi(k)\}) = 1$ and $v(\{\pi(1), \dots, \pi(k-1)\}) = 0$, i.e. when adding person $\pi(k) = i$ to the coalition changes the value of v from 0 to 1. Since the number

of elements of S_n is $n!$, we conclude that $\phi_i(v)$ denotes the probability that person i changes the value of v from 0 to 1, averaged over all possible orderings of coalitions.

Remark 6.11. Fix $i \in \{1, \dots, n\}$. Let $S \subseteq \{1, \dots, n\}$ with $i \notin S$. The number of possible orderings of elements of S in an ordered list is $|S|!$. And the number of possible orderings of elements of $\{1, \dots, n\} \setminus (S \cup \{i\})$ is $(n - |S| - 1)!$. So, the term $v(S \cup \{i\}) - v(S)$ is repeated in our formula for $\phi_i(v)$ a total of $|S|!(n - |S| - 1)!$ times. That is,

$$\phi_i(v) = \sum_{S \subseteq \{1, \dots, n\}: i \notin S} \frac{|S|!(n - |S| - 1)!}{n!} (v(S \cup \{i\}) - v(S)).$$

Proof of Theorem 6.9. In Example 6.5, we defined a set of functions $\{u_T\}_{T \subseteq \{1, \dots, n\}}$, where $u_T: 2^{\{1, \dots, n\}} \rightarrow \mathbb{R}$. Recall that $u_T(S) = 1$ if $T \subseteq S \subseteq \{1, \dots, n\}$, and $u_T(S) = 0$ otherwise. The set of functions $\{u_T\}_{T \subseteq \{1, \dots, n\}}$ is a basis for the set of functions from $2^{\{1, \dots, n\}}$ to \mathbb{R} (we will prove this below). So, there exist constants $\{c_T\}_{T \subseteq \{1, \dots, n\}}$ such that

$$v(S) = \sum_{T \subseteq \{1, \dots, n\}} c_T u_T(S), \quad \forall S \subseteq \{1, \dots, n\}. \quad (*)$$

Let's find the coefficients c_T in terms of values v . First, note that $v(\emptyset) = 0 = c_\emptyset$ from (*). Now, let $i \in \{1, \dots, n\}$. From (*),

$$v(\{i\}) = \sum_{T \subseteq \{1, \dots, n\}} c_T u_T(\{i\}) = c_{\{i\}} u_{\{i\}}(\{i\}) = c_{\{i\}}. \quad (**)$$

We can now find the coefficients c_T inductively. Let $\ell \geq 0$. Suppose we have already found the coefficients c_T for all $T \subseteq \{1, \dots, n\}$ with $|T| < \ell$. Now let $S \subseteq \{1, \dots, n\}$ with $|S| = \ell$. From (*),

$$v(S) = \sum_{T \subseteq \{1, \dots, n\}} c_T u_T(S) = c_S + \sum_{T \subseteq S: |T| < \ell} c_T.$$

That is, we can solve for c_S , so that c_S is determined uniquely by v :

$$c_S = v(S) - \sum_{T \subseteq S: |T| < \ell} c_T. \quad (***)$$

We now show that (*) holds, avoiding circular reasoning. Define coefficients $\{c_T\}_{T \subseteq \{1, \dots, n\}}$ inductively so that $c_\emptyset = 0$, $c_{\{i\}} = v(\{i\})$ for each $i \in \{1, \dots, n\}$, and so on, using (**). We claim that (*) must now hold. We show this by induction on the size of S . We already know that (*) holds when $S = \emptyset$ and when $|S| = 1$, as in (**). So, fix $\ell \geq 0$ and suppose (*) holds when $|S| < \ell$. Now, consider $S \subseteq \{1, \dots, n\}$ with $|S| = \ell$. Then using (**), we see that

$$\sum_{T \subseteq \{1, \dots, n\}} c_T u_T(S) = c_S + \sum_{T \subseteq S: |T| < \ell} c_T = v(S).$$

That is, (*) holds for this particular S . Having completed the inductive step, we conclude that (*) holds. Finally, using (*), Axiom (iii), Remark 4.12, and Example 6.5,

$$\phi_i(v) = \phi_i \left(\sum_{T \subseteq \{1, \dots, n\}} c_T u_T \right) = \sum_{T \subseteq \{1, \dots, n\}} \phi_i(c_T u_T) = \sum_{T \subseteq \{1, \dots, n\}} c_T \phi_i(u_T) = \sum_{T \subseteq \{1, \dots, n\}: i \in T} c_T \frac{1}{|T|}.$$

That is, $\phi_i(v)$ is uniquely determined by the values of v . So, the uniqueness part of the Theorem is complete. It remains to prove existence.

To prove existence, we show the explicit formula for $\phi_i(v)$ involving permutations holds. It then suffices to show the formula from Remark 6.11 satisfies all four of Shapley's Axioms. That is, we need to verify that the following functions (f_1, \dots, f_n) verify Shapley's Axioms:

$$f_i(v) = \sum_{S \subseteq \{1, \dots, n\}: i \notin S} \frac{|S|!(n - |S| - 1)!}{n!} (v(S \cup \{i\}) - v(S)), \quad \forall i \in \{1, \dots, n\}.$$

Axioms (ii), (iii) follow immediately. To verify Axiom (i), let $i \neq j$, $i, j \in \{1, \dots, n\}$, and suppose $v(S \cup \{i\}) = v(S \cup \{j\})$ for all $S \subseteq \{1, \dots, n\}$ such that $i, j \notin S$. Note the following cancellation occurs.

$$\begin{aligned} f_i(v) - f_j(v) &= \sum_{S \subseteq \{1, \dots, n\}: i \notin S, j \in S} \frac{|S|!(n - |S| - 1)!}{n!} (v(S \cup \{i\}) - v(S)) \\ &\quad - \sum_{S \subseteq \{1, \dots, n\}: j \notin S, i \in S} \frac{|S|!(n - |S| - 1)!}{n!} (v(S \cup \{j\}) - v(S)) \\ &= \sum_{S \subseteq \{1, \dots, n\}: i, j \notin S} \frac{(|S| + 1)!(n - |S| - 2)!}{n!} (v(S \cup \{i\} \cup \{j\}) - v(S \cup \{j\})) \\ &\quad - \sum_{S \subseteq \{1, \dots, n\}: i, j \notin S} \frac{(|S| + 1)!(n - |S| - 2)!}{n!} (v(S \cup \{i\} \cup \{j\}) - v(S \cup \{i\})) \\ &= \sum_{S \subseteq \{1, \dots, n\}: i, j \notin S} \frac{(|S| + 1)!(n - |S| - 2)!}{n!} (v(S \cup \{i\}) - v(S \cup \{j\})) = 0. \end{aligned}$$

That is, (f_1, \dots, f_n) satisfy Axiom (i). Finally, we verify Axiom (iv). Fix a bijection $\pi: \{1, \dots, n\} \rightarrow \{1, \dots, n\}$. Let $i \in \{1, \dots, n\}$. Let $k = \pi^{-1}(i)$. Define a function $g_{i,\pi}(v) := v(\{\pi(1), \dots, \pi(k)\}) - v(\{\pi(1), \dots, \pi(k-1)\})$. Then we have a telescoping sum

$$\begin{aligned} \sum_{i=1}^n g_{i,\pi}(v) &= \sum_{i=1}^n (v(\{\pi(1), \dots, \pi(\pi^{-1}(i))\}) - v(\{\pi(1), \dots, \pi(\pi^{-1}(i)) - 1\})) \\ &= v(\{\pi(1), \dots, \pi(n)\}) - v(\emptyset) = v(\{1, \dots, n\}). \end{aligned}$$

That is, the functions (g_1, \dots, g_n) satisfy Axiom (iv). Using Remark 6.11, each of the functions f_1, \dots, f_n is an average of the functions g_1, \dots, g_n . More specifically,

$$f_i = \frac{1}{n!} \sum_{\pi \in S_n} g_{i,\pi}.$$

Since the functions (g_1, \dots, g_n) satisfy Axiom (iv), so do the functions (f_1, \dots, f_n) :

$$\sum_{i=1}^n f_i(v) = \frac{1}{n!} \sum_{\pi \in S_n} \sum_{i=1}^n g_{i,\pi}(v) = \frac{1}{n!} \sum_{\pi \in S_n} v(\{1, \dots, n\}) = v(\{1, \dots, n\}).$$

□

Exercise 6.12. Let n be a positive integer. Let $v: 2^{\{1, \dots, n\}} \rightarrow \{0, 1\}$ be a characteristic function that only takes values 0 and 1. Assume also that v is monotonic. That is, if

$S, T \subseteq \{1, \dots, n\}$ with $S \subseteq T$, then $v(S) \leq v(T)$. The **Shapley-Shubik power index** of each player is defined to be their Shapley value.

By monotonicity of v , we have $v(S \cup \{i\}) \geq v(S)$ for all $S \subseteq \{1, \dots, n\}$ and for all $i \in \{1, \dots, n\}$. Also, since v only takes values 0 and 1, we have

$$v(S \cup \{i\}) - v(S) = \begin{cases} 1 & \text{when } v(S \cup \{i\}) > v(S) \\ 0 & \text{when } v(S \cup \{i\}) = v(S) \end{cases}.$$

Consequently, we have the following simplified formula for the Shapley-Shubik power index of player $i \in \{1, \dots, n\}$:

$$\phi_i(v) = \sum_{S \subseteq \{1, \dots, n\}: v(S \cup \{i\})=1 \text{ and } v(S)=0} \frac{|S|!(n - |S| - 1)!}{n!}.$$

Compute the Shapley-Shubik power indices for all players on the UN security council, with pre-1965 and post-1965 structure. Which structure is better for nonpermanent members?

In pre-1965 rules, the UN security council had five permanent members, and six nonpermanent members. A resolution passes only if all five permanent members want it to pass, and at least two nonpermanent members want it to pass. So, we can model this voting method, by letting $\{1, 2, \dots, 11\}$ denote the council, and letting $\{1, 2, 3, 4, 5\}$ denote the permanent members. Then we use the characteristic function $v: 2^{\{1, \dots, 11\}} \rightarrow \{0, 1\}$ so that, for any $S \subseteq \{1, \dots, 11\}$, $v(S) = 1$ if $\{1, 2, 3, 4, 5\} \subseteq S$ and if $|S| \geq 7$. And $v(S) = 0$ otherwise.

This voting method was called unfair, so it was restructured in 1965. After the restructuring, the council had the following form (which is still used today). The UN security council has five permanent members, and now ten nonpermanent members. A resolution passes only if all five permanent members want it to pass, and at least four nonpermanent members want it to pass. So, we can model this voting method, by letting $\{1, \dots, 15\}$ denote the council, and letting $\{1, 2, 3, 4, 5\}$ denote the permanent members. Then we use the characteristic function $v: 2^{\{1, \dots, 15\}} \rightarrow \{0, 1\}$ so that, for any $S \subseteq \{1, \dots, 15\}$, $v(S) = 1$ if $\{1, 2, 3, 4, 5\} \subseteq S$ and if $|S| \geq 9$. And $v(S) = 0$ otherwise.

Exercise 6.13. Let n be a positive integer. Let $v: 2^{\{1, \dots, n\}} \rightarrow \{0, 1\}$ be a characteristic function that only takes values 0 and 1. Assume also that v is monotonic and $v(\{1, \dots, n\}) = 1$. For each $i \in \{1, \dots, n\}$, let B_i be the number of subsets $S \subseteq \{1, \dots, n\}$ such that $v(S) = 0$ and $v(S \cup \{i\}) = 1$. The **Banzhaf power index** of player i is defined to be

$$\frac{B_i}{\sum_{j=1}^n B_j}.$$

Like the Shapley-Shubik power index, the Banzhaf power index is another way to measure the relative power of each player.

Compute the Banzhaf power indices for all players for the glove market example, Example 6.1.

Then, compute the Banzhaf power indices for all players on the UN security council, with pre-1965 and post-1965 structure. Which structure is better for nonpermanent members?

7. MECHANISM DESIGN

In Mechanism Design, we wish to modify the rules of a game so that rational players exhibit some desirable property. For example, in our first topic of auctions, the auctioneer would like to design their auction in a way such that the auction participants pay the most amount of money.

7.1. Auctions. The following are common examples of auctions with n bidders, where n is a positive integer.

Example 7.1 (English Auction (Open-bid, ascending auction)). In an English auction, the auctioneer first declares a (typically low) reserve price for an item to be auctioned. The n bidders take turns making increasing bids. When there are no more bids, the highest bidder gets the item for the price they last bid.

Example 7.2 (Dutch Auction (Open-bid, descending auction)). In a Dutch auction, the auctioneer first declares a (typically high) reserve price for an item to be auctioned. Starting from the reserve price, the auctioneer then states a sequence of decreasing prices. The first bidder to say “stop” receives the item for the last price that was stated.

Example 7.3 (Sealed Bid, First-Price Auction). Every buyer submits a sealed envelope with their desired bid for the item. The buyer who has submitted the highest bid receives the item for the price of their bid.

Example 7.4 (Vickrey Auction (Sealed Bid, Second-Price Auction)). Every buyer submits a sealed envelope with their desired bid for the item. The buyer who has submitted the highest bid receives the item for the *second-highest bid*.

Some auctions behave much differently than others, as the following example demonstrates. In class, we will simulate the auction of a \$20 bill with ascending \$1 increments. (Microbidding websites such as QuiBids use the auction method below to make money.)

Example 7.5 (Shubik’s Evil Dollar Auction). The auctioneer first declares a reserve price of \$1 for the auction of a \$20 bill. The n bidders take turns making increasing bids. When there are no more bids, the highest bidder gets the item for the price they last bid. However, the *second-highest bidder also pays their bid, and they receive nothing*.

We now restate our definition of independence of random variables.

Definition 7.6. Let Y and Z be random variables. We say that Y and Z are **independent** if, for any fixed real numbers y, z , the probability that $Y \leq y$ and that $Z \leq z$ is equal to the probability that $Y \leq y$, multiplied by the probability that $Z \leq z$. We note this statement by $\mathbf{P}(Y \leq y, Z \leq z) = \mathbf{P}(Y \leq y)\mathbf{P}(Z \leq z)$. More generally, if X_1, \dots, X_n are random variables, then X_1, \dots, X_n are independent if $\mathbf{P}(X_1 \leq x_1, \dots, X_n \leq x_n) = \mathbf{P}(X_1 \leq x_1) \cdots \mathbf{P}(X_n \leq x_n)$ for all $x_1, \dots, x_n \in \mathbb{R}$.

Definition 7.7 (A Symmetric Auction). A single object is for sale at an auction. The seller is willing to sell the object at any nonnegative price. There are n buyers, which we identify with the set $\{1, 2, \dots, n\}$. All buyers have some set of **private values** in $[0, 1]$. We denote the private value of buyer $i \in \{1, \dots, n\}$ by V_i , so that V_i is a random variable that takes nonnegative real values. We assume that all of the random variables V_1, \dots, V_n

are independent. We also assume that V_1, \dots, V_n are identically distributed, with a continuous density function. That is, there exists some continuous function $f: \mathbb{R} \rightarrow [0, \infty)$ with $\int_{-\infty}^{\infty} f(x)dx = 1$ such that: for each $i \in \{1, \dots, n\}$, for each $t \in \mathbb{R}$, the probability that $V_i \leq t$ is equal to $\int_{-\infty}^t f(x)dx$. We define the **expected value** of V_1 to be $\int_{-\infty}^{\infty} xf(x)dx$. We also assume that all buyers are **risk-neutral**, so that each buyer seeks to maximize their expected profits.

Buyer $i \in \{1, \dots, n\}$ knows the private value V_i , but they do not know the other private values. So, a symmetric auction is a game with imperfect information. Also, all buyers know the distribution of V_1 .

Finally, we assume that all of the above assumptions are **common knowledge**. That is, every player knows the above assumptions; every player knows that every player knows the above assumptions; every player knows that every player knows that every player knows the above assumptions; etc.

Exercise 7.8. Let V_1, V_2, V_3 be independent random variables that are uniformly distributed in $[0, 1]$. So, for example, for any $0 \leq a < b \leq 1$, the probability that $a \leq V_1 \leq b$ is $b - a$.

Compute the expected value of V_1 . Compute the expected value of $V_1 + V_2 + V_3$. Then compute the expected value of $\max(V_1, V_2, V_3)$. Finally, compute the expected value of V_1^5 .

The assumption that the private values are identically distributed means that each buyer, in planning their strategy, assumes that all other buyers are similar to each other and similar to them.

Under the above assumptions, a **pure strategy** for Player $i \in \{1, \dots, n\}$ is a function $\beta_i: [0, 1] \rightarrow [0, \infty)$. So, if Player i has a private value of V_i , they will make a bid of $\beta_i(V_i)$ in the auction. (We will not discuss mixed strategies in auctions.)

Given the strategies $\beta = (\beta_1, \dots, \beta_n)$, and given any $v \in [0, 1]$, Player i has expected profit $P_i(\beta, v)$, if their private value is v . (If buyer i wins the auction, and if buyer i has private value v and bid b , then the profit of buyer i is $v - b$.) We say that a strategy β is a **(Bayes-Nash) equilibrium** if, given any $v \in [0, 1]$, any $b \geq 0$, and any $i \in \{1, \dots, n\}$,

$$P_i(\beta, v) \geq P_i((\beta_1, \dots, \beta_{i-1}, b, \beta_{i+1}, \dots, \beta_n), v).$$

That is, any player with private value v cannot profit by deviating from the equilibrium bid $\beta_i(v)$ to another bid b .

Remark 7.9. The Dutch auction is equivalent to the sealed-bid first-price auction. In particular, the equilibrium strategies are the same for both auctions. Also, this equivalence holds regardless of assumptions on the buyers, such as independence of private values.

Remark 7.10. In the Vickrey auction, it is an equilibrium when each buyer bids their private value. Moreover, this strategy dominates all other strategies. That is, each buyer with this strategy maximizes their profit, *regardless of what the other players do*. This follows since a buyer with private value v can make a profit at most $\max(v - m, 0)$ where m is the maximum of all other bids, and this profit is achieved when a buyer bids their private value. Note that this conclusion does not require any assumptions on the buyers, such as independence of private values.

In the English auction, it is an equilibrium when each buyer continues to bid up until their private value is reached. Moreover, this strategy dominates all other strategies. (Technically

we will not define an equilibrium strategy for the English auction, so the previous two sentences are informal.)

In both auctions, at these equilibria, the winner of the auction is the buyer with the highest private value, and the selling price is the second highest private value among the buyers. So, at the respective equilibria, the payments of the bidders is the same, and the winner of the auction is the same.

Example 7.11. Suppose we have two buyers, and $f(v) = 1$ for any $v \in [0, 1]$ in a sealed-bid first price auction. That is, the private values V_1 and V_2 are uniformly distributed in the interval $[0, 1]$ (and independent). We will show that an equilibrium strategy is $\beta_1(v) = v/2$, $\beta_2(v) = v/2$, for every $v \in [0, 1]$. That is, each player will bid half of their private value.

Suppose buyer 2 uses this strategy, and suppose buyer 1 has private value $v \in [0, 1]$. Suppose buyer 1 submits the bid $b \in [0, 1]$. Since buyer 2 will bid $V_2/2$, buyer 1 wins the auction only when $b > V_2/2$. Since V_2 is uniformly distributed in $[0, 1]$, we have $b > V_2/2$, i.e. we have $2b > V_2$ with probability $\int_0^{\min(1, 2b)} dx = \min(2b, 1)$.

The expected profit of buyer 1 is $v - b$ multiplied by the probability that buyer 1 wins the auction. That is, the expected profit of buyer 1 is

$$(v - b) \min(2b, 1).$$

So, buyer 1 maximizes their profit by choosing $b = v/2$ (since the function $b \mapsto (v - b) \min(2b, 1)$ is maximized at $b = v/2$.) That is, we must have $\beta_1(v) = v/2$. Using a similar argument, if $\beta_1(v) = v/2$, then buyer 2 maximizes their profit by choosing $\beta_2(v) = v/2$. That is, the strategy $\beta_1(v) = v/2$, $\beta_2(v) = v/2$, $\forall v \in [0, 1]$ is a symmetric equilibrium.

Using a similar argument for a sealed-bid second-price auction, we can find that a symmetric equilibrium occurs when $\beta_1(v) = v$, $\beta_2(v) = v$, $\forall v \in [0, 1]$.

We now ask: *which auction method should the auctioneer prefer?*

We first consider the sealed bid first-price auction. Let $Y = \max(V_1, V_2)$. Then $Y \leq y$ if and only if $V_1 \leq y$ and $V_2 \leq y$. Since V_1 and V_2 are independent,

$$\mathbf{P}(Y \leq y) = \mathbf{P}(V_1 \leq y)\mathbf{P}(V_2 \leq y) = \mathbf{P}(V_1 \leq y)^2 = \left(\int_0^y dv \right)^2 = y^2.$$

So, Y has PDF $f_Y(y) = 2y$ when $0 \leq y \leq 1$. So, the expected value of Y is $\int_0^1 v 2v dv = 2/3$. At equilibrium, buyer 1 bids $V_1/2$ and buyer 2 bids $V_2/2$, the expected profit of the auctioneer is the expected value of $\max(V_1/2, V_2/2) = (1/2) \max(V_1, V_2)$. So, the auctioneer has expected profit $(1/2)(2/3) = 1/3$.

We now consider the sealed bid second-price auction. Let $Y = \min(V_1, V_2)$. At equilibrium, buyer 1 bids V_1 and buyer 2 bids V_2 . And in the second-price auction, the expected profit of the auctioneer is the expected value of $\min(V_1, V_2)$. Using the identity $\max(V_1, V_2) + \min(V_1, V_2) = V_1 + V_2$, the expected profit of the auctioneer is $\int_0^1 v dv + \int_0^1 v dv - 2/3 = 1/3$.

In summary, in the sealed bid first-price auction, and in the sealed bid second-price auction, at each equilibrium, the auctioneer makes the same revenue! Combining this observation with Remarks 7.9 and 7.10, the auctioneer makes the same revenue also for English and Dutch auctions! That is, all four auction methods give the auctioneer the same expected revenue.

Exercise 7.12. Suppose we have two buyers, and $f(v) = 1$ for any $v \in [0, 1]$ in a sealed-bid first price auction. That is, V_1 and V_2 are uniformly distributed in the interval $[0, 1]$. Show

that an equilibrium strategy is $\beta_1(v) = v$, $\beta_2(v) = v$, $\forall v \in [0, 1]$. That is, each player will bid exactly their private value.

Let n be a positive integer. We return now to auctions with n buyers. Let $Y = \max(V_2, V_3, \dots, V_n)$. For any $t \in [0, 1]$, suppose Y has PDF $g: \mathbb{R} \rightarrow [0, \infty)$. The random variable Y represents the price at which buyer 1 can win the auction, since Y is the largest bid of the other buyers.

Theorem 7.13. *In a symmetric sealed-bid first-price auction, the following strategy is a symmetric equilibrium:*

$$\beta_i(t) = \frac{\int_0^t xg(x)dx}{\int_0^t g(x)dx}, \quad \forall i \in \{1, \dots, n\}, \quad \forall t > 0, \quad \beta_i(0) = 0.$$

(That is, $\beta_i(t) = \mathbf{E}(Y | Y \leq t)$.) In a symmetric sealed-bid second-price auction, the following strategy is a symmetric equilibrium: $\beta_i(t) = t$ for all $t \in [0, 1]$, for all $i \in \{1, \dots, n\}$.

Proof Sketch. We first observe from integration by parts that for all $t > 0$,

$$F_Y(t)\mathbf{E}(Y | Y \leq t) = \int_0^t y f_Y(y) dy = \int_0^t y \frac{d}{dy} \mathbf{P}(Y \leq y) dy = tF_Y(t) - \int_0^t F_Y(y) dy. \quad (*)$$

We consider the sealed-bid first-price auction. It can be shown that $\beta(t) := \frac{\int_0^t xg(x)dx}{\int_0^t g(x)dx}$ is strictly increasing and continuous for all $t \geq 0$. So, β has an inverse β^{-1} . When buyer 1 bids $b > 0$ with private value $v > 0$, their expected profit is

$$\begin{aligned} \mathbf{P}(\beta(Y) < b)(v - b) &= \mathbf{P}(Y < \beta^{-1}(b))(v - b) = F_Y(\beta^{-1}(b))(v - \beta(\beta^{-1}(b))) \\ &= F_Y(\beta^{-1}(b))(v - \mathbf{E}(Y | Y \leq \beta^{-1}(b))) \\ &= F_Y(\beta^{-1}(b))(v - \beta^{-1}(b) + \beta^{-1}(b) - \mathbf{E}(Y | Y \leq \beta^{-1}(b))) \\ &\stackrel{(*)}{=} F_Y(\beta^{-1}(b))(v - \beta^{-1}(b)) + \int_0^{\beta^{-1}(b)} F_Y(y) dy. \end{aligned}$$

Denote $z := \beta^{-1}(b)$. Differentiating with respect to z , we get

$$f_Y(z)(v - z) - F_Y(z) + F_Y(z) = f_Y(z)(v - z),$$

which is positive when $z < v$ and zero when $z = v$. Hence the maximum occurs when $z = v$, i.e. at equilibrium, buyer 1 should bid $b = \beta(z) = \beta(v)$. \square

Remark 7.14. Extending Example 7.11 to n players, we get $\mathbf{P}(\max(V_2, \dots, V_n) \leq t) = [\mathbf{P}(V_2 \leq t)]^{n-1} = t^{n-1}$ if $t \in [0, 1]$, so Y has PDF $g(t) = (d/dt)t^{n-1} = (n-1)t^{n-2}$ when $0 \leq t \leq 1$, and a symmetric equilibrium in the first-price auction is

$$\beta_i(t) = \frac{\int_0^t x(n-1)x^{n-2}dx}{\int_0^t (n-1)x^{n-2}dx} = \frac{\frac{n-1}{n}t^n}{t^{n-1}} = \frac{n-1}{n}t, \quad \forall i \in \{1, \dots, n\}, \quad \forall t \in [0, 1].$$

Theorem 7.15 (Revenue Equivalence Theorem). *Suppose we have any symmetric sealed-bid auction. Let β be a symmetric strictly increasing equilibrium such that:*

- The winner of the auction is the buyer with the highest private value.
- The expected payment of each buyer with private value 0 is 0.

Then, the seller's expected revenue is

$$n \int_0^\infty \left(\int_0^t xg(x)dx \right) f(t)dt = n \int_{-\infty}^\infty F_Y(t)\mathbf{E}(Y | Y \leq t)f(t)dt.$$

In particular, the seller's expected revenue does **not** depend on the form of the auction (other than the assumptions specified above).

Proof Sketch. Suppose buyer 1 has private value $v \geq 0$. Let $p(v)$ be the expected payment of buyer 1 when the other bids are $\beta_2(V_2), \dots, \beta_n(V_2)$. Let $a(v)$ denote the probability that buyer 1 wins the auction. Then at the symmetric equilibrium, since β is increasing,

$$a(v) = \mathbf{P}(\beta(v) > \max_{j=2, \dots, n} \beta(V_j)) = \mathbf{P}(v > \max_{j=2, \dots, n} V_j) = F_Y(v) = \mathbf{P}(V_1 \leq v)^{n-1}.$$

Suppose buyer 1 deviates from the equilibrium strategy and bids as if their private value is $w \geq 0$. Then buyer 1 has expected profit

$$va(w) - p(w).$$

If a, p are differentiable, at equilibrium, the derivative of this quantity in w vanishes when $w = v$, i.e. $va'(v) - p'(v) = 0$, so that $p'(v) = va'(v)$. So, if $p(0) = 0$ we have by integrating then integration by parts

$$p(v) = \int_0^v wa'(w)dw = va(v) - \int_0^v a(w)dw$$

From our expression above for $a(v)$, we conclude that $p(v)$ does not depend on the auction method. So in particular, the expected payment is the same as a sealed-bid second-price auction.

In a sealed-bid second-price auction, at the symmetric equilibrium, the expected payment of buyer 1 with private value t is $F_Y(t)\mathbf{E}(Y | Y \leq t)$. So, the expected payment of buyer 1 is $\int_{-\infty}^\infty F_Y(t)\mathbf{E}(Y | Y \leq t)f(t)dt$. The seller's expected revenue is the sum of payments of all n buyers, hence we multiply this quantity by n . \square

Exercise 7.16 (Muddy Children Puzzle/ Blue-Eyed Islanders Puzzle). This exercise is meant to test our understanding of common knowledge.

Situation 1. There are 100 children playing in the mud. All of the children have muddy foreheads, but any single child cannot tell whether or not their own forehead is muddy. Any child can also see all of the other 99 children. The children do not communicate with each other in any way, there are no mirrors or recording devices, etc. so that no child can see their own forehead. The teacher now says, "stand up if you know your forehead is muddy." No one stands up, because no one can see their own forehead. The teacher asks again. "Knowing that no one stood up the last time, stand up now if you know your forehead is muddy." Still no one stands up. No matter how many times the teacher repeats this statement, no child stands up.

Situation 2. After Situation 1, the teacher now says, "I announce that at least one of you has a muddy forehead." The teacher then says, "stand up if you know your forehead is muddy." No one stands up. The teacher pauses then repeats, "stand up if you know your forehead is muddy." Again, no one stands up. The teacher continues making this statement. The hundredth time that they make this statement, all the children suddenly stand up.

Explain why all of the children stand up in Situation 2, but they do not stand up in Situation 1. Pay close attention to what is common knowledge in each situation.

Exercise 7.17. There are five pirates on a ship. It is also common knowledge that every pirate prefers to maximize their amount of gold. There are 100 gold pieces to be split amongst the pirates. The game begins when the first pirate proposes how they think the gold should be split amongst the five pirates. All five pirates vote whether or not to accept the proposal, by a majority vote. If the proposal is accepted, the game ends. If the proposal is not accepted, the first pirate is thrown overboard, and the game continues. The second pirate now proposes how they think the gold should be split amongst the four remaining pirates. All four pirates vote whether or not to accept the proposal, by a majority vote (the second pirate breaks a tie). If the proposal is accepted, the game ends. If the proposal is not accepted, the second pirate is thrown overboard, and the game continues, etc. (At each stage of the game, the pirate that could be thrown overboard next can break the tie in the majority vote.) What is the largest amount of gold that the first pirate can obtain in the game?

8. SOCIAL CHOICE

In social choice theory, we try to design voting methods that elect officials in desirable ways. We consider n voters, where n is a positive integer. At the outset, we consider these n voters casting votes between two candidates. Once the votes are cast, we need to decide who has won the election. We could take the majority vote; or we could just use the vote of some pre-selected person to determine the winner of the election; there are many possible ways to determine the winner of the election. The latter voting method would be called a dictatorship. There are also many ways to measure how well the electorate's desires are reflected in the election's outcome.

Below, we interpret an election of n voters for 2 candidates by labeling the candidates as 1 and -1 , respectively. Then, given the votes $x = (x_1, \dots, x_n) \in \{-1, 1\}^n$, the winner of the election is $f(x)$, where $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ is the voting method. Here, person $i \in \{1, \dots, n\}$ votes for candidate $x_i \in \{-1, 1\}$, so that x is the set of all votes of the n voters.

8.1. Fourier analysis and influences. Let n be a positive integer. Let $x = (x_1, \dots, x_n) \in \{-1, 1\}^n$. Let $f, g: \{-1, 1\}^n \rightarrow \mathbb{R}$. For any subset $S \subseteq \{1, \dots, n\}$, define a function $W_S: \{-1, 1\}^n \rightarrow \mathbb{R}$ by

$$W_S(x) := \prod_{i \in S} x_i.$$

Define also the inner product

$$\langle f, g \rangle := 2^{-n} \sum_{x \in \{-1, 1\}^n} f(x)g(x).$$

Lemma 8.1. *The set of functions $\{W_S\}_{S \subseteq \{1, \dots, n\}}$ is an orthonormal basis for the space of functions from $\{-1, 1\}^n \rightarrow \mathbb{R}$, with respect to this inner product. (When we write $S \subseteq \{1, \dots, n\}$, we include the empty set \emptyset as a subset of $\{1, \dots, n\}$.)*

Exercise 8.2. Prove Lemma 8.1.

By Lemma 8.1, any $f: \{-1, 1\}^n \rightarrow \mathbb{R}$ can be expressed as

$$f(x) = \sum_{S \subseteq \{1, \dots, n\}} \langle f, W_S \rangle W_S(x), \quad \forall x \in \{-1, 1\}^n.$$

For any $S \subseteq \{1, \dots, n\}$, if we denote the **Fourier coefficient** $\widehat{f}(S)$ of f associated to S by

$$\widehat{f}(S) := \langle f, W_S \rangle = 2^{-n} \sum_{y \in \{-1, 1\}^n} f(y) W_S(y),$$

then we have

$$f(x) = \sum_{S \subseteq \{1, \dots, n\}} \widehat{f}(S) W_S(x), \quad \forall x \in \{-1, 1\}^n.$$

The expression on the right of the equals sign is called the **Fourier expansion** of f . Using the orthonormality of the functions $\{W_S\}_{S \subseteq \{1, \dots, n\}}$, we have **Plancherel's theorem**:

$$\begin{aligned} \langle f, g \rangle &= \left\langle \sum_{S \subseteq \{1, \dots, n\}} \widehat{f}(S) W_S, \sum_{S' \subseteq \{1, \dots, n\}} \widehat{g}(S') W_{S'} \right\rangle = \sum_{S, S' \subseteq \{1, \dots, n\}} \widehat{f}(S) \widehat{g}(S') \langle W_S, W_{S'} \rangle \\ &= \sum_{S \subseteq \{1, \dots, n\}} \widehat{f}(S) \widehat{g}(S). \end{aligned}$$

In particular, using $f = g$, we get

$$\langle f, f \rangle = \sum_{S \subseteq \{1, \dots, n\}} |\widehat{f}(S)|^2. \quad (1)$$

Exercise 8.3. Let $f: \{-1, 1\}^2 \rightarrow \{-1, 1\}$ such that $f(x) = 1$ for all $x \in \{-1, 1\}^2$. Compute $\widehat{f}(S)$ for all $S \subseteq \{1, 2\}$.

Let $f: \{-1, 1\}^3 \rightarrow \{-1, 1\}$ such that $f(x_1, x_2, x_3) = \text{sign}(x_1 + x_2 + x_3)$ for all $(x_1, x_2, x_3) \in \{-1, 1\}^3$. Compute $\widehat{f}(S)$ for all $S \subseteq \{1, 2, 3\}$. The function f is called a **majority function**.

Remark 8.4 (Voting Interpretation). Let $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$. We think of $x = (x_1, \dots, x_n) \in \{-1, 1\}^n$, where x_i is the vote of person $i \in \{1, \dots, n\}$ for the candidate $x_i \in \{-1, 1\}$. Then $f(x)$ is the winner of the election, given the votes x .

Definition 8.5 (Influences). Let $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ and let $i \in \{1, \dots, n\}$. Define the i^{th} partial derivative $\partial_i f: \{-1, 1\}^n \rightarrow \{-1, 0, 1\}$ of f by

$$\partial_i f(x) := \frac{f(x_1, \dots, x_n) - f(x_1, \dots, x_{i-1}, -x_i, x_{i+1}, \dots, x_n)}{2}, \quad \forall x \in \{-1, 1\}^n.$$

Define the i^{th} influence $I_i f \in \mathbb{R}$ of a function $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ by

$$I_i f := 2^{-n} \sum_{x \in \{-1, 1\}^n} |\partial_i f(x)|^2.$$

Equivalently,

$$I_i f = \langle \partial_i f, \partial_i f \rangle = \sum_{S \subseteq \{1, \dots, n\}: i \in S} (\widehat{f}(S))^2.$$

To see this, note that if $i \in S$, then $\partial_i W_S = W_S$, and if $i \notin S$, then $\partial_i W_S = 0$. So, $\partial_i f = \sum_{S \subseteq \{1, \dots, n\}: i \in S} \widehat{f}(S) W_S$. Then (1) gives the required equality.

Remark 8.6. $I_i f$ is the probability that person $i \in \{1, \dots, n\}$ can change the outcome of the election. To see this, note that if $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$, then $|\partial_i f(x)|$ is 0 or 1. And $|\partial_i f(x)| = 1$ only when changing the i^{th} person's vote changes the outcome of the election. Also, $|\partial_i f(x)| = |\partial_i f(x)|^2$, so $I_i f = 2^{-n} \sum_{x \in \{-1, 1\}^n} |\partial_i f(x)|$. The latter quantity is

exactly an average over all votes x of the quantity $|\partial_i f(x)|$. In fact, this interpretation of the influence of a function shows that $I_i(f)$ is nearly identical to the Banzhaf power index of voter $i \in \{1, \dots, n\}$. (In particular, if B_i is the Banzhaf power index of voter $i \in \{1, \dots, n\}$, then $B_i = I_i(f)/(\sum_{j=1}^n I_j(f))$. Note that the sum of all the Banzhaf power indices is 1, whereas the sum of all of the influences of a function can be much larger than 1.)

Exercise 8.7. Let n be a positive integer. Show that there is a one-to-one correspondence (or a bijection) between the set of functions f where $f: \{-1, 1\}^n \rightarrow \mathbb{R}$, and the set of functions g where $g: 2^{\{1, 2, \dots, n\}} \rightarrow \mathbb{R}$. For example, you could identify a subset $S \subseteq \{1, \dots, n\}$ with the element $x = (x_1, \dots, x_n) \in \{-1, 1\}^n$ where, for all $i \in \{1, \dots, n\}$, we have $x_i = 1$ if $i \in S$, and $x_i = -1$ if $i \notin S$.

Let $i, j \in \{1, \dots, n\}$ and let $x \in \{-1, 1\}^n$. Let $S(x) = \{j \in \{1, \dots, n\} : x_j = 1\}$. Using this one-to-one correspondence, show that the i^{th} Shapley value of $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ can be written as

$$\phi_i(f) = \sum_{x \in \{-1, 1\}^n : x_i = -1} \frac{|S(x)|(n - |S(x)| - 1)!}{n!} (f(x_1, \dots, x_{i-1}, 1, x_{i+1}, \dots, x_n) - f(x)).$$

So, $\phi_i(f)$ is similar to, but distinct from, $I_i(f)$.

Example 8.8. The dictator function $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ is defined by $f(x_1, \dots, x_n) = x_1$. So, $\hat{f}(\{1\}) = 1$, and $\hat{f}(S) = 0$ for every other $S \subseteq \{1, \dots, n\}$ with $S \neq \{1\}$. Note that $\partial_1 f = f$ and $\partial_i f = 0$ for every $i > 1$. Consequently, $I_1(f) = (\hat{f}(\{1\}))^2 = 1$, and $I_i(f) = 0$ for every $i > 1$. That is, only one person can influence the election, while everyone else cannot.

Example 8.9. The majority function is defined for $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ (with n odd) by the formula

$$f(x_1, \dots, x_n) = \text{sign}(x_1 + \dots + x_n).$$

Let $i \in \{1, \dots, n\}$. Using binomial coefficients and Stirling's formula, we compute

$$\begin{aligned} I_i f &= \binom{n-1}{(n-1)/2} 2^{-n} = \frac{(n-1)!}{((n-1)/2)!^2} 2^{-n} \\ &\sim \frac{\sqrt{2\pi(n-1)}((n-1)/e)^{n-1}}{2\pi((n-1)/2)((n-1)/(2e))^{n-1}} 2^{-n} \sim \frac{1}{\sqrt{2\pi}} \frac{1}{\sqrt{n-1}}. \end{aligned}$$

In particular, all influences are equal to each other. And $\lim_{n \rightarrow \infty} \sqrt{2\pi(n-1)} I_i f = 1$.

So your influence in a majority election is much larger than you think!

Exercise 8.10. Let X, Y be independent random variables. Let n be a positive integer. Let $p_1, \dots, p_n, q_1, \dots, q_n \geq 0$ with $\sum_{i=1}^n p_i = \sum_{i=1}^n q_i = 1$. Let $x_1, \dots, x_n, y_1, \dots, y_n \in \mathbb{R}$. Assume that $X = x_i$ with probability p_i and $Y = y_i$ with probability q_i for all $i \in \{1, \dots, n\}$. Show that the expected value of XY is equal to the expected value of X , multiplied by the expected value of Y .

8.2. Noise stability. Recall that any function $f: \{-1, 1\}^n \rightarrow \mathbb{R}$ can be written in its Fourier series as

$$f(x) = \sum_{S \subseteq \{1, \dots, n\}} \hat{f}(S) W_S(x), \quad \forall x \in \{-1, 1\}^n.$$

Given any $S \subseteq \{1, \dots, n\}$, we define $|S|$ to be the number of elements of S . For example, $|\{1, 2, 5, 6\}| = 4$ and $|\emptyset| = 0$.

Definition 8.11 (Noise stability). Let $\rho \in (0, 1)$ and define $T_\rho f: \{-1, 1\}^n \rightarrow \mathbb{R}$ by

$$T_\rho f(x) := \sum_{S \subseteq \{1, \dots, n\}} \rho^{|S|} \widehat{f}(S) W_S(x), \quad \forall x \in \{-1, 1\}^n.$$

Define the **noise stability** of f with parameter ρ to be

$$\langle f, T_\rho f \rangle = \sum_{S \subseteq \{1, \dots, n\}} \rho^{|S|} |\widehat{f}(S)|^2.$$

Remark 8.12. If $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$. We will now show that the noise stability of f measures whether or not the outcome of the election changes when the votes are corrupted, say, by erroneous tabulation, etc.

Let $\varepsilon \in (0, 1)$. Recall that $x = (x_1, \dots, x_n) \in \{-1, 1\}^n$ is interpreted as a list of votes of n people for two candidates. Suppose each vote is corrupted independently with some probability ε . That is, for any fixed $i \in \{1, \dots, n\}$, x_i is kept the same with probability $1 - \varepsilon$, and x_i is changed to $-x_i$ with probability ε . And if $i \neq j$ with $i, j \in \{1, \dots, n\}$, any change made to x_i is independent of any change made to x_j . Let $y = (y_1, \dots, y_n)$ be the corrupted vote. So, the expected value of y_i is $(1 - \varepsilon)x_i + \varepsilon(-x_i) = (1 - 2\varepsilon)x_i$. And the expected value of $y_i y_j$ is $(1 - 2\varepsilon)x_i(1 - 2\varepsilon)x_j$, by Exercise 8.10. More generally, the expected value of $W_S(y) = \prod_{i \in S} y_i$ is

$$(1 - 2\varepsilon)^{|S|} \prod_{i \in S} x_i = (1 - 2\varepsilon)^{|S|} W_S(x).$$

Summing over S , the expected value of $f(y)$ is then

$$\sum_{S \subseteq \{1, \dots, n\}} (1 - 2\varepsilon)^{|S|} \widehat{f}(S) W_S(x) = T_{(1-2\varepsilon)} f(x).$$

Now, if the corruption of votes did not change the outcome of the election, then we would have $f(y) = f(x) \in \{-1, 1\}$. That is, $f(x)f(y) = 1$. However, if the corruption of votes did in fact change the outcome of the election, then we would have $f(y) = -f(x)$, so that $f(x)f(y) = -1$. The expected value of $f(x)f(y)$, i.e. the quantity $f(x)T_{(1-2\varepsilon)}f(x)$ therefore measures how much the outcome of the election changes as a result of vote corruption. And the quantity $2^{-n} \sum_{x \in \{-1, 1\}^n} f(x)T_{(1-2\varepsilon)}f(x) = \langle f, T_{(1-2\varepsilon)}f \rangle$ is then an average of $f(x)T_{(1-2\varepsilon)}f(x)$.

It is therefore desirable to have a voting method with a noise stability that is as large as possible. In other words, we would like to maximize $\langle f, T_\rho f \rangle$ when $\rho \in (0, 1)$ among all functions $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$.

Proposition 8.13. Let $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ and let $0 < \rho < 1$. Then

$$\langle f, T_\rho f \rangle \leq 1,$$

with equality only if f is a constant function. If additionally $\widehat{f}(\emptyset) = 0$, i.e. if we assume that $\sum_{x \in \{-1, 1\}^n} f(x) = 0$, then

$$\langle f, T_\rho f \rangle \leq \rho,$$

with equality only if f is a dictator or anti-dictator. (There exists $i \in \{1, \dots, n\}$ such that $f(x) = f(x_1, \dots, x_n) = x_i$ for all $x \in \{-1, 1\}^n$, or $f(x) = -x_i$ for all $x \in \{-1, 1\}^n$.)

Proof. Since $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$, then $(f(x))^2 = 1$ for all $x \in \{-1, 1\}^n$. So, by Plancherel's Theorem (1), $1 = \langle f, f \rangle = \sum_{S \subseteq \{1, \dots, n\}} |\hat{f}(S)|^2$. Therefore, using $0 < \rho < 1$,

$$\langle f, T_\rho f \rangle = \sum_{S \subseteq \{1, \dots, n\}} \rho^{|S|} |\hat{f}(S)|^2 \leq \sum_{S \subseteq \{1, \dots, n\}} |\hat{f}(S)|^2 = 1.$$

And equality only occurs when $\hat{f}(S) = 0$ for all $S \subseteq \{1, \dots, n\}$ with $|S| \geq 1$. That is, equality only occurs when f is a constant function.

With the additional assumption that $\hat{f}(\emptyset) = 0$, we similarly have

$$\langle f, T_\rho f \rangle = \sum_{S \subseteq \{1, \dots, n\}} \rho^{|S|} |\hat{f}(S)|^2 = \sum_{S \subseteq \{1, \dots, n\}: |S| \geq 1} \rho^{|S|} |\hat{f}(S)|^2 \leq \sum_{S \subseteq \{1, \dots, n\}: |S| \geq 1} |\hat{f}(S)|^2 = 1.$$

And equality only occurs when $\hat{f}(S) = 0$ for all $S \subseteq \{1, \dots, n\}$ with $|S| \neq 1$. So, $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ and $\hat{f}(S) = 0$ for all $S \subseteq \{1, \dots, n\}$ with $|S| \neq 1$. The following Exercise then completes the proof. \square

Exercise 8.14. Let $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$. Assume that $\hat{f}(S) = 0$ whenever $S \subseteq \{1, \dots, n\}$ and $|S| \neq 1$. Then there exists $i \in \{1, \dots, n\}$ such that $f(x) = f(x_1, \dots, x_n) = x_i$ for all $x \in \{-1, 1\}^n$, or $f(x) = -x_i$ for all $x \in \{-1, 1\}^n$.

So, the function f with largest noise stability is a constant function. However, an election where the outcome does not depend on the votes is not a very desirable thing! So, in order to find a sensible voting method that is noise stable, we need to impose some restrictions on the function f . For example, let's assume that each candidate has an equal chance of winning the election. This means that half of the elements $x \in \{-1, 1\}^n$ satisfy $f(x) = 1$ and the other half of the elements of $x \in \{-1, 1\}^n$ satisfy $f(x) = -1$. Equivalently, $\hat{f}(\emptyset) = 2^{-n} \sum_{x \in \{-1, 1\}^n} f(x) = 0$. This assumption eliminates the constant function. But still, the dictator function from Example 8.8 maximizes $\langle f, T_\rho f \rangle$ when $\rho \in (0, 1)$ among all $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ with $\hat{f}(\emptyset) = 0$. So, when only one person decides the election, the result of the election is fairly stable. But this result is kind of an artifact of there being only one influential voter. So, let's impose an additional constraint that no one person has too much influence on the outcome of the election. With this additional assumption, democracy prevails. That is, the majority function is the most noise stable.

Theorem 8.15 (Majority is Stablest). Let $\rho \in (0, 1)$, $n \geq 1$, $\varepsilon > 0$. Let $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ with $\sum_{x \in \{-1, 1\}^n} f(x) = 0$. Let $g_n(x_1, \dots, x_n) = \text{sign}(x_1 + \dots + x_n)$ be the majority function. Then there exists $\tau > 0$ such that, if $I_i f \leq \tau$ for all $i \in \{1, \dots, n\}$, then

$$\langle f, T_\rho f \rangle \leq \lim_{n \rightarrow \infty} \langle g_n, T_\rho g_n \rangle + \varepsilon.$$

8.3. Arrow's Impossibility Theorem. So far, we discussed how a function $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ can be thought of as a voting method where n people vote for either of two candidates. There is also a way to use these functions in an election between more than two candidates. The method we describe now is the **Condorcet voting method**. Suppose we have three candidates, which we label as a, b, c . Every voter now ranks the three candidates. For

example, a voter can rank b first, a second, and c third. This voter prefers candidate b over the other candidates, they prefer a over c , and they do not prefer c at all. We can think of the election between the three candidates as three separate elections between two candidates. That is, once all of the voters rank the candidates, we can just look at every pair of comparisons of the voters. For example, if we just look at the preferences between a and c , then who is more preferred, among all voters?

Consider the following ranking of three candidates a, b, c between three voters 1, 2, 3.

Voter	Rank 1	Rank 2	Rank 3
1	a	b	c
2	b	c	a
3	c	a	b

If we ignore candidate b , then voters 2 and 3 prefer c over a , while voter 1 prefers a over c . So, using a majority rule for these preferences, the voters prefer c over a . If we ignore candidate c , then voters 1 and 3 prefer a over b , while voter 2 prefers b over a . So, using a majority rule again, the voters prefer a over b . Finally, if we ignore candidate a , then voters 1 and 2 prefer b over c , while voter 3 prefers c over b . So, using a majority rule, the voters prefer b over c .

In conclusion, given the above ranking, if we use a majority rule for every comparison between two candidates, the voters prefer a over b , they prefer b over c , and they prefer c over a . So, no one has won the election! We have shown:

Theorem 8.16 (Condorcet Paradox). *In a Condorcet Voting method where the ranking of any two candidates is determined by a majority rule, it is possible that no one wins the election.*

It is desirable to have an actual winner of an election. To do this, we could try to change the voting method by using something other than the majority function. To see how to do this, let's establish some notation. Suppose we are considering an election between three candidates a, b, c and we have n voters. Each voter ranks the three candidates. A voter can rank the candidates by exhibiting preferences between any two candidates. Let $x = (x_1, \dots, x_n), y = (y_1, \dots, y_n), z = (z_1, \dots, z_n) \in \{-1, 1\}^n$. Then $x = (x_1, \dots, x_n)$, and for any $i \in \{1, \dots, n\}$, we think of x, y, z as specifying the votes as follows.

- If $x_i = 1$, then voter i prefers a over b , and if $x_i = -1$, then voter i prefers b over a .
- If $y_i = 1$, then voter i prefers b over c , and if $y_i = -1$, then voter i prefers c over b .
- If $z_i = 1$, then voter i prefers c over a , and if $z_i = -1$, then voter i prefers a over c .

Lemma 8.17. *In order for the three candidates to be ranked in a list by voter $i \in \{1, \dots, n\}$, the numbers x_i, y_i, z_i must be not all equal. Equivalently, the three candidates are ranked in a list by voter i if and only if $(3 - x_i y_i - x_i z_i - y_i z_i)/4 = 1$. If the three candidates are not ranked in a list, then $(3 - x_i y_i - x_i z_i - y_i z_i)/4 = 0$.*

Proof. There are $2^3 = 8$ possible ordered choices of $x_i, y_i, z_i \in \{-1, 1\}$. Let's just check every possible choice. If $x_i = y_i = z_i = 1$, then voter i prefers a over b , they prefer b over c and they prefer c over a . So, this choice of x_i, y_i, z_i is not a ranked list of candidates. Also, we observe that $(3 - x_i y_i - x_i z_i - y_i z_i)/4 = 0$.

If $x_i = y_i = 1$ and $z_i = -1$, then voter i prefers a over b , they prefer b over c and they prefer a over c . That is, they rank a first, b second, and c third. In this case, the voters are in a ranked list. Also, we observe that $(3 - x_i y_i - x_i z_i - y_i z_i)/4 = 1$. The remaining cases are treated similarly. \square

Now, fix $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$. Given the votes x, y, z , the function f determines the societal preference as follows:

- If $f(x) = 1$, the voters prefer a over b , and if $f(x) = -1$, the voters prefer b over a .
- If $f(y) = 1$, the voters prefer b over c , and if $f(y) = -1$, the voters prefer c over b .
- If $f(z) = 1$, the voters prefer c over a , and if $f(z) = -1$, the voters prefer a over c .

In the Condorcet Paradox, we used $f: \{-1, 1\}^3 \rightarrow \{-1, 1\}$ where f was the majority function; $f(x_1, x_2, x_3) = \text{sign}(x_1 + x_2 + x_3)$. But now we are considering any function f .

Definition 8.18 (Condorcet Winner, Three Candidates). Let $x, y, z \in \{-1, 1\}^n$, and assume $(3 - x_i y_i - x_i z_i - y_i z_i)/4 = 1$ for all $i \in \{1, \dots, n\}$. Let $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$. Given the votes x, y, z , we say that a **Condorcet winner** of the election exists if one candidate is preferred over the other two. By Lemma 8.17, a Condorcet winner exists when $(3 - f(x)f(y) - f(x)f(z) - f(y)f(z))/4 = 1$. And a Condorcet winner does not exist when $(3 - f(x)f(y) - f(x)f(z) - f(y)f(z))/4 = 0$.

In Theorem 8.16, we saw that a Condorcet winner does not always exist if $f: \{-1, 1\}^3 \rightarrow \{-1, 1\}$ is the majority function. If we maybe change the function f , then maybe it is possible to always have a Condorcet winner. However, as Arrow's Theorem shows, guaranteeing a Condorcet winner is a very special assumption.

Theorem 8.19 (Arrow's Impossibility Theorem). Let $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$. Let $x, y, z \in \{-1, 1\}^n$, and assume $(3 - x_i y_i - x_i z_i - y_i z_i)/4 = 1$ for all $i \in \{1, \dots, n\}$. The votes then describe the preferences of voters between three candidates. For any such votes x, y, z , assume that a Condorcet winner exists. Then f or $-f$ must be a dictatorship.

Proof. We will average the value of $(3 - f(x)f(y) - f(x)f(z) - f(y)f(z))/4$ over all possible votes $x, y, z \in \{-1, 1\}^n$ where $(3 - x_i y_i - x_i z_i - y_i z_i)/4 = 1$ for all $i \in \{1, \dots, n\}$. The vector (x_i, y_i, z_i) is then interpreted as a random vector, which is independent of (x_j, y_j, z_j) for any $j \neq i$ where $i, j \in \{1, \dots, n\}$. So, for example, x_1 is a random variable independent of the random variable x_2 , x_1 is independent of y_3 , and so on.

For any fixed $i \in \{1, \dots, n\}$, the vector (x_i, y_i, z_i) can be equal to any of the following six vectors:

$$\{(1, 1, -1), (1, -1, 1), (-1, 1, 1), (1, -1, -1), (-1, 1, -1), (-1, -1, 1)\}.$$

Averaging over all six possibilities, the expected value of x_i is $(1/6)(3 - 3) = 0$; similarly, the expected value of y_i is zero. Also, averaging over all six possibilities, the expected value of $x_i y_i$ is

$$\frac{1}{6}(1 - 1 - 1 - 1 - 1 + 1) = -\frac{1}{3}.$$

And if $i \neq j$, x_i, y_i are independent of x_j, y_j . So, if $S, S' \subseteq \{1, \dots, n\}$, the expected value of $W_S(x)W_{S'}(y) = \prod_{i \in S} x_i \prod_{j \in S'} y_j$ is zero by Exercise 8.10, unless $S = S'$. In the case

$S = S'$, the expected value of $W_S(x)W_S(y)$ is then $(-1/3)^{|S|}$. So, summing over all Fourier coefficients of f , the expected value of $f(x)f(y)$ is

$$\sum_{S \subseteq \{1, \dots, n\}} |\hat{f}(S)|^2 (-1/3)^{|S|} = \langle f, T_{(-1/3)} f \rangle.$$

This equality similarly holds for $f(x)f(z)$ and $f(y)f(z)$. In conclusion, the expected value of $(3 - f(x)f(y) - f(x)f(z) - f(y)f(z))/4$ is

$$\frac{3}{4}(1 - \langle f, T_{-1/3} f \rangle).$$

By assumption, a Condorcet winner exists for any votes x, y, z . That is, $(3 - f(x)f(y) - f(x)f(z) - f(y)f(z))/4 = 1$ for all possible votes x, y, z . So,

$$1 = \frac{3}{4}(1 - \langle f, T_{-1/3} f \rangle) = \frac{3}{4} \left(1 - \sum_{S \subseteq \{1, \dots, n\}} |\hat{f}(S)|^2 (-1/3)^{|S|} \right).$$

We now claim that $\hat{f}(S) = 0$ unless $|S| = 1$. To prove this claim, we argue by contradiction. Assume $\hat{f}(S) \neq 0$ for some $S \subseteq \{1, \dots, n\}$ with $|S| \neq 1$. For any nonnegative integer k , we have $(-1/3)^k \geq -1/3$, with equality only when $k = 1$. That is, $-(-1/3)^k \leq 1/3$, with equality only when $k = 1$. Since $\hat{f}(S) \neq 0$ for some $|S| \neq 1$, the following inequality is therefore strict:

$$1 < \frac{3}{4} \left(1 + \frac{1}{3} \sum_{S \subseteq \{1, \dots, n\}} |\hat{f}(S)|^2 \right).$$

Recalling that $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ satisfies $1 = \langle f, f \rangle = \sum_{S \subseteq \{1, \dots, n\}} |\hat{f}(S)|^2$, we get

$$1 < \frac{3}{4}(1 + 1/3) = (3/4)(4/3) = 1.$$

So, $1 < 1$, a contradiction. We conclude that $\hat{f}(S) = 0$ unless $|S| = 1$. The Theorem now follows from Exercise 8.14. \square

9. QUANTUM GAMES

Definition 9.1. Let n be a positive integer. A **classical XOR game** G of size n involves two cooperating players, one referee, and an $n \times n$ matrix $C = (C_{ij})$ where $C_{ij} \in \{-1, 1\}$ for all $i, j \in \{1, \dots, n\}$. Before the game begins, the two players can communicate and decide on a strategy, but once the game starts, they can no longer communicate. To start the game, the referee randomly chooses a pair of integers $i, j \in \{1, \dots, n\}$, so that (i, j) is chosen with probability $p_{ij} \geq 0$, where $\sum_{i,j=1}^n p_{ij} = 1$. The referee sends i to player I and j to player II . Upon receiving this information, the first player responds with $x_i \in \{-1, 1\}$ and the second player responds with $y_j \in \{-1, 1\}$. The payoff is then $x_i y_j C_{ij}$. The goal for the two players is to maximize the expected payoff.

Example 9.2 (The CHSH game). Let $n = 2$ and let $C_{11} = C_{12} = C_{21} = 1$ with $C_{22} = -1$ and let $p_{ij} = 1/4$ for all $i, j \in \{1, 2\}$. The maximum expected payoff is then

$$\begin{aligned} \max_{x_1, x_2, y_1, y_2 \in \{-1, 1\}} \sum_{i,j=1}^2 p_{ij} C_{ij} x_i y_j &= \max_{x_1, x_2, y_1, y_2 \in \{-1, 1\}} \frac{1}{4} (x_1 y_1 + x_1 y_2 + x_2 y_1 - x_2 y_2) \\ &= \max_{x_1, x_2, y_1, y_2 \in \{-1, 1\}} \frac{1}{4} (x_1 (y_1 + y_2) + x_2 (y_1 - y_2)). \end{aligned}$$

Note that, $y_1 + y_2$ is equal to 2, 0 or -2 , and in each such case $y_1 - y_2$ is equal to 0, ± 2 , or 0. So, by the triangle inequality, $|x_1 (y_1 + y_2) + x_2 (y_1 - y_2)| \leq |y_1 + y_2| + |y_1 - y_2| \leq 2$. In fact, choosing $y_1 = y_2 = 1$ and $x_1 = x_2 = 1$, shows that $x_1 (y_1 + y_2) + x_2 (y_1 - y_2) = 2$. Therefore,

$$\max_{x_1, x_2, y_1, y_2 \in \{-1, 1\}} \sum_{i,j=1}^2 p_{ij} C_{ij} x_i y_j = \frac{2}{4} = \frac{1}{2}.$$

Note that if the two players could coordinate with each other exactly, then their payoff would be 1. So, the best they can do is to get half of the payoff of the case of perfect coordination. And to get such a payoff, they could use the pure strategy where $x_1 = x_2 = y_1 = y_2 = 1$.

Could the players get a higher payoff if they use mixed strategies, or even if they use shared randomness somehow? It turns out that the answer is no.

Lemma 9.3. Let $q = (q_1, \dots, q_n)$, where $q_i \geq 0$ for all $i \geq 0$ and $\sum_{i=1}^n q_i = 1$. Let X be a random variable that takes the value z_i with probability q_i for each $1 \leq i \leq n$. Recall the expected value of X is $\sum_{i=1}^n q_i z_i$. Let z such that $z_i \leq z$ for all $i \in \{1, \dots, n\}$. Then $\sum_{i=1}^n q_i z_i \leq z$.

Proof. $\sum_{i=1}^n q_i z_i \leq \sum_{i=1}^n q_i z = z$. □

Suppose the two players use randomness in their strategy, and we can even allow shared randomness (which is independent of the randomness that the referee uses). For example, the players could flip any number of random coins before the start of the game, and they could use this information to make their choices for x_1, x_2, y_1, y_2 . Suppose the outcome of these coin flips is t , so that the players choose $x_i(t), y_j(t)$ if they are shown $i, j \in \{1, 2\}$. Then for this outcome t , the payoff is

$$\frac{1}{4} (x_1(t) y_1(t) + x_1(t) y_2(t) + x_2(t) y_1(t) - x_2(t) y_2(t)) \leq \frac{1}{2}.$$

So, for all outcomes t , the payoff is at most $1/2$. Then by Lemma 9.3, the expected payoff is at most $1/2$. That is, using shared randomness does not improve the expected payoff of the game.

However, using shared quantum randomness can increase the payoff! The axioms of quantum mechanics allow the players to share a quantum state, using entanglement. We can think of this quantum state as a random vector, which is shared by both players. Each player can measure the quantum state using a 2×2 real unitary matrix. Recall that if A is a real unitary matrix, then $A = A^T$ and A^2 is equal to the identity matrix. Recall that the trace of A is the sum of the diagonal elements of A , so that $\text{Tr}(A) = A_{11} + A_{22}$. Suppose player I uses the matrix A and player II uses the matrix B . Then the measurement of

player I is x and the measurement of player II is y , where $x, y \in \{-1, 1\}$. And the expected value of xy is

$$\frac{1}{2}\text{Tr}(AB).$$

So, the players can choose their strategies by selecting unitary matrices to maximize their payoff. Define

$$X_1 := \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad X_2 := \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y_1 := \begin{pmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ 1/\sqrt{2} & -1/\sqrt{2} \end{pmatrix}, \quad Y_2 := \begin{pmatrix} 1/\sqrt{2} & -1/\sqrt{2} \\ -1/\sqrt{2} & -1/\sqrt{2} \end{pmatrix}.$$

Note that all of these matrices are unitary. Consider the following strategy of the game. When player I sees $i \in \{1, 2\}$, they use the measurement X_i producing outcome x_i , and when player II sees $j \in \{1, 2\}$, they use the measurement Y_j producing outcome y_j . Then the quantity

$$\frac{1}{4}(x_1y_1 + x_1y_2 + x_2y_1 - x_2y_2)$$

has expected value

$$\begin{aligned} & \frac{1}{4} \left(\frac{1}{2}\text{Tr}(X_1Y_1) + \frac{1}{2}\text{Tr}(X_1Y_2) + \frac{1}{2}\text{Tr}(X_2Y_1) - \frac{1}{2}\text{Tr}(X_2Y_2) \right) \\ &= \frac{1}{4} \left(\frac{\sqrt{2}}{2} + \frac{\sqrt{2}}{2} + \frac{\sqrt{2}}{2} + \frac{\sqrt{2}}{2} \right) = \frac{\sqrt{2}}{2}. \end{aligned}$$

So, this **quantum strategy** outperforms the classical strategy by a multiplicative factor of $\sqrt{2}$.

In fact, this strategy is the best possible.

Lemma 9.4. *Let a, b be nonnegative real numbers. Then $2ab \leq a^2 + b^2$.*

Proof. Since $(a - b)^2 \geq 0$, we have $a^2 + b^2 - 2ab \geq 0$. □

Recall that a vector $X \in \mathbb{R}^n$ has length $\|X\| := (X^T X)^{1/2}$. The following Lemma was proven in your Linear Algebra class, 115A.

Lemma 9.5 (Cauchy-Schwarz Inequality). *If X, Y are vectors in \mathbb{R}^n , then*

$$X^T Y \leq \|X\| \cdot \|Y\|$$

Proof Sketch. If $Y = 0$, then the inequality holds since both sides are zero. If $Y \neq 0$, multiply out the inequality $\|X - Y(X^T Y)/(Y^T Y)\|^2 \geq 0$. □

Lemma 9.6. *Let $X_1, Y_1, X_2, Y_2 \in \mathbb{R}^n$ be vectors of unit length. Then*

$$X_1^T Y_1 + X_1^T Y_2 + X_2^T Y_1 - X_2^T Y_2 \leq 2\sqrt{2}.$$

Proof.

$$\begin{aligned}
[X_1^T Y_1 + X_1^T Y_2 + X_2^T Y_1 - X_2^T Y_2]^2 &= [X_1^T (Y_1 + Y_2) + X_2^T (Y_1 - Y_2)]^2 \\
&\leq [\|Y_1 + Y_2\| + \|Y_1 - Y_2\|]^2, \quad \text{by Lemma 9.5} \\
&= \|Y_1 + Y_2\|^2 + 2\|Y_1 + Y_2\| \cdot \|Y_1 - Y_2\| + \|Y_1 - Y_2\|^2 \\
&\leq 2(\|Y_1 + Y_2\|^2 + \|Y_1 - Y_2\|^2), \quad \text{by Lemma 9.4} \\
&= 2(Y_1^T Y_1 + Y_2^T Y_2 + Y_1^T Y_1 + Y_2^T Y_2 + 2Y_1^T Y_2 - 2Y_1^T Y_2) \\
&= 2(Y_1^T Y_1 + Y_2^T Y_2 + Y_1^T Y_1 + Y_2^T Y_2) = 8.
\end{aligned}$$

□

Corollary 9.7 (CHSH inequality). *Let $d \geq 2$. Let X_1, X_2, Y_1, Y_2 be $d \times d$ unitary matrices. Then*

$$\frac{1}{4} \left(\frac{1}{d} \text{Tr}(X_1 Y_1) + \frac{1}{d} \text{Tr}(X_1 Y_2) + \frac{1}{d} \text{Tr}(X_2 Y_1) - \frac{1}{d} \text{Tr}(X_2 Y_2) \right) \leq \frac{\sqrt{2}}{2}.$$

Proof. If X, Y are $d \times d$ real matrices, we can interpret these matrices as d^2 -dimensional real vectors, with the inner product $\langle X, Y \rangle = \frac{1}{d} \text{Tr}(XY^T)$. If X is unitary, then $\langle X, X \rangle = \frac{1}{d} \text{Tr}(XX^T) = 1$. The assertion then follows from Lemma 9.6. □

Corollary 9.8 (Bell's inequality). *Let $d \geq 2$. Let $C_{11} = C_{12} = C_{21} = 1$ with $C_{22} = -1$ and let $p_{ij} = 1/4$ for all $i, j \in \{1, 2\}$. Then*

$$\max_{x_1, x_2, y_1, y_2 \in \{-1, 1\}} \sum_{i,j=1}^2 p_{ij} C_{ij} x_i y_j = \frac{1}{2} < \frac{\sqrt{2}}{2} = \max_{\substack{X_1, X_2, Y_1, Y_2 \in \mathbb{R}^{d \times d}: \\ X_i^2 = Y_j^2 = 1, X_i = X_i^T, Y_j = Y_j^T \\ \forall i, j \in \{1, 2\}}} \sum_{i,j=1}^2 p_{ij} C_{ij} \frac{1}{d} \text{Tr}(X_i Y_j)$$

Proof. The leftmost equality was shown in Example 9.2. The rightmost equality follows from Corollary 9.7, and by choosing

$$X_1 := \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad X_2 := \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y_1 := \begin{pmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ 1/\sqrt{2} & -1/\sqrt{2} \end{pmatrix}, \quad Y_2 := \begin{pmatrix} 1/\sqrt{2} & -1/\sqrt{2} \\ -1/\sqrt{2} & -1/\sqrt{2} \end{pmatrix}.$$

□

10. ALGORITHMIC GAME THEORY

Recall from Exercise 3.33 that we can find optimal strategies for two-person zero-sum games efficiently (that is, in polynomial time in the size of the game, using linear programming). A similar result for general-sum games should not hold, since Nash equilibria are computationally hard to find. Below we will provide more detail for this assertion.

Recall the definition of Nash equilibrium in Definition 4.4. Let $\tilde{x} \in \Delta_m$ and let $\tilde{y} \in \Delta_n$ be a Nash equilibrium, i.e.

$$\begin{aligned}
\tilde{x}^T A \tilde{y} &= \max_{x \in \Delta_m} x^T A \tilde{y}. \\
\tilde{x}^T B \tilde{y} &= \max_{y \in \Delta_n} \tilde{x}^T B y.
\end{aligned}$$

From Exercise 3.33, we can rewrite this as

$$\tilde{x}^T A \tilde{y} = \max_{i=1, \dots, m} (A \tilde{y})_i.$$

$$\tilde{x}^T B \tilde{y} = \max_{j=1, \dots, n} (\tilde{x}^T B)_j.$$

Let

$$I := \{1 \leq i \leq m: \tilde{x}_i > 0\}, \quad J := \{1 \leq j \leq n: \tilde{y}_j > 0\}. \quad (*)$$

We claim that

$$\begin{aligned} \max_{i=1, \dots, m} (A\tilde{y})_i &= (A\tilde{y})_i, & \forall i \in I. \\ \max_{j=1, \dots, n} (\tilde{x}^T B)_j &= (\tilde{x}^T B)_j, & \forall j \in J. \end{aligned}$$

To see why, suppose not. Then $(A\tilde{y})_i > (A\tilde{y})_{i'}$ for some $i, i' \in I$. But then $x \in \Delta_m$ defined by $x_i := \tilde{x}_i + \tilde{x}_{i'}$, $x_{i'} := 0$, and $x_k := \tilde{x}_k$ for all other $1 \leq k \leq m$ satisfies $x^T A\tilde{y} > \tilde{x}^T A\tilde{y}$, contradicting the Nash equilibrium property. It then follows from the claim that

$$\begin{aligned} \tilde{x}^T A\tilde{y} &= (A\tilde{y})_i, & \forall i \in I. \\ \tilde{x}^T B\tilde{y} &= (\tilde{x}^T B)_j, & \forall j \in J. \end{aligned}$$

Algorithm 10.1. Input: $m \times n$ matrix A, B , $I \subseteq \{1, \dots, m\}$, $J \subseteq \{1, \dots, n\}$

Assumption: there exists a Nash equilibrium $\tilde{x} \in \Delta_m$, $\tilde{y} \in \Delta_n$ with I, J satisfying $(*)$.

Output: \tilde{x}, \tilde{y} .

- Let \tilde{x}_i , $i \in I$ and let \tilde{y}_j , $j \in J$ be the $|I| + |J|$ unknown variables.
- Consider $|I| - 1$ equations

$$(A\tilde{y})_i = (A\tilde{y})_{i'} \quad \forall i, i' \in I.$$

- Consider $|J| - 1$ equations

$$(\tilde{x}^T B)_j = (\tilde{x}^T B)_{j'} \quad \forall j, j' \in J.$$

- Together with the equations $\sum_{i \in I} \tilde{x}_i = 1$, $\sum_{j \in J} \tilde{y}_j = 1$, we have $|I| + |J|$ linear equations in $|I| + |J|$ equations. By assumption, there exists at least one positive solution to these equations. Output such a solution, if one can be found where \tilde{x}, \tilde{y} have nonnegative values and they also maximize the expected payoffs of the players.

Algorithm 10.2 (Finding a Nash Equilibrium). Input: $m \times n$ matrix A, B . Output: \tilde{x}, \tilde{y} a Nash equilibrium

- Run Algorithm 10.1 for every possible pair of subsets $I \subseteq \{1, \dots, m\}$, $J \subseteq \{1, \dots, n\}$.
- From Nash's Theorem, there exists at least one pair I, J where Algorithm 10.1 succeeds

Remark 10.3. Algorithm 10.2 has run time at least 2^{m+n} , since we might need to check this many pairs of I, J . Also, Algorithm 10.2 indicates that finding a Nash equilibrium is a combinatorial problem.

Remark 10.4. Algorithm 10.1 also works for games with more than two players, though the equations to solve are no longer linear.

Although you might hope to improve on Algorithm 10.2's exponential run time, it seems unlikely much improvement could be made. We might guess: is finding a Nash equilibrium NP-complete? Since a Nash equilibrium exists by Nash's Theorem, this question is not so natural. However, if we modify the question a bit, NP-completeness results can be proven.

(Recall that NP is the set of decision problems, i.e. computational problems with a yes or no answer, where "yes" instances can be verified in polynomial time. A computational

problem A is NP-complete if it is in NP and every other problem in NP has a polynomial time reduction to A . A well-known example of an NP-complete problem is 3-SAT. An instance of 3-SAT is a 3CNF formula, i.e. a logical AND of 3-term ORs. A 3-term OR is a clause of three variables or their negations (e.g. x_3 or x_5 or not x_2 .)

Theorem 10.5 (Gilboa and Zemel, 1989). *The following problems are NP-complete, even for symmetric games. Given a two-player game in strategic form, decide whether or not:*

- *The game has at least two Nash equilibria.*
- *The game has a Nash equilibrium where player 1 has a payoff larger than a given amount.*
- *The game has a Nash equilibrium where player 1 and player 2 have a combined payoff larger than a given amount.*
- *The game has a Nash equilibrium with support size (i.e. number of pure strategies with positive probability of being played) larger than a given number.*
- *The game has a Nash equilibrium with support contains a given pure strategy.*
- *The game has a Nash equilibrium with support not containing a given pure strategy.*

One might be tempted to go back into the proof of Nash's theorem and try to derive an efficient algorithm for finding Nash equilibria from the proof. However, the key part of the proof used Brouwer's fixed point theorem, which itself only states the existence of a fixed point (equilibrium), without giving an algorithm for finding the fixed point. And indeed, it is known that find a Brouwer fixed point is just as hard as finding a Nash equilibrium. (Both are PPAD-complete problems. If $P = NP$ then $P = PPAD$, since verifying that you have a Nash equilibrium can be done efficiently. However, it could occur that $P = PPAD \neq NP$.)

Even if it seems hard to exactly find Nash equilibria, we might hope to efficiently find approximate equilibria. An ε -approximate Nash equilibrium is a strategy for the players such that, if any player deviates from the strategy, their payoff can increase by at most ε . In this definition, we assume all payoffs of all players lie in the interval $[0, 1]$. Note that an approximate equilibrium is not guaranteed to be close to an actual Nash equilibrium.

For fixed large ε (say around $1/2$ or $1/3$) it would not be surprising to find an ε -approximate Nash equilibrium. The most interesting case is then when $\varepsilon > 0$ is small.

Theorem 10.6 (Lipton, Markakis, Mehta, 2003, Corollary 1). *For any 2-person game with n strategies for each player, for any $\varepsilon > 0$, there is an algorithm with run time $n^{O(\log n)}$ that finds an ε -approximate Nash equilibrium.*

Proof Sketch. Let $k := 12 \log n / \varepsilon^2$. A mixed strategy is called k -uniform if it is the uniform distribution on a multiset of k elements of $\{1, \dots, n\}$. (So, for example, $x \in \Delta_m$ defined by $x_1 = 2/3$, $x_2 = 1/3$ and $x_i = 0$ for all $i > 2$ is 3-uniform on the multiset $\{1, 1, 2\}$.) It is shown using the probabilistic method that there exists a k -uniform ε -approximate Nash equilibrium. So, the algorithm will just exhaustively search over all possible k -uniform strategies to find such an approximate equilibrium. The running time of the algorithm is then proportional to the number of pairs of k -uniform distributions (found by "stars and bars"):

$$\binom{n+k-1}{k}^2 = \left(\frac{(n+k-1)(n+k-2) \cdots (n+1)n}{k!} \right)^2 \leq 4n^{2k} = 4n^{24 \log n / \varepsilon^2}.$$

□

However, this algorithm does not have a polynomial run time. And evidence suggests this run time cannot be improved.

Theorem 10.7 (Rubinstein, 2016). *Assume an Exponential Time Hypothesis for PPAD. Then there exists $\varepsilon > 0$ such that, finding an ε -approximate Nash equilibrium for a 2 player $n \times n$ game requires a run time $n^{(\log(n))^{1-o(1)}}$.*

Also, finding an ε -approximate Nash equilibrium is PPAD complete for an n player game where each player has two strategies to use [Rubinstein, 2015].

An ε -relative approximate Nash equilibrium is a strategy for the players such that, if any player deviates from the strategy, their payoff can increase by at most ε times the absolute value of the new payoff. In the two player case, this says $\tilde{x}^T A \tilde{y} \geq x^T A \tilde{y} - \varepsilon |x^T A \tilde{y}|$ for all $x \in \Delta_m$, and similarly $\tilde{x}^T B \tilde{y} \geq \tilde{x}^T B y - \varepsilon |\tilde{x}^T B y|$ for all $y \in \Delta_n$. In this definition, we do not need to assume that payoffs are restricted to lie in the interval $[0, 1]$.

A result of Daskalakis from 2011 says: for any $\varepsilon > 0$, it is PPAD-complete to find an ε -relative approximate Nash equilibrium. Consequently, the relative approximation of equilibria is harder than the additive approximation.

Lastly, recall that an evolutionary stable strategy is not guaranteed to exist, and so we might expect its computational complexity to be at least as hard as NP. And indeed this is the case, as shown by Etessami and Lochbihler: deciding if an ESS exists is both NP-hard and co-NP hard. In fact, it is S_2^P complete, a results of Conitzer. (Recall that S_2^P contains NP.)

Remark 10.8. The hardness results mentioned above are worst-case in nature. That is, they state there exist sequences of complicated games where e.g. finding approximate equilibria encodes finding the fixed point of some other complicated function (as in Brouwer’s fixed point theorem). One might then ask if it is easier to find or approximately find Nash equilibria in some “average case” of games, rather than in the “worst case.” One such effort appears [here](#).

11. REINFORCEMENT LEARNING

11.1. Bandits.

Definition 11.1. Let $\mathcal{A} = \{1, \dots, k\}$ be a finite set. (k is the number of “arms” of the multi-armed bandit.) Let $(P_a)_{a \in \mathcal{A}}$ be a set of probability distributions. Let H be a positive integer or ∞ . A **bandit** problem is a two-player game with incomplete information played over H rounds. The two players are called the learner and the environment, and H is called the horizon. For each round $0 \leq t \leq H$ of the game, the learner chooses a strategy $A_t \in \mathcal{A}$ (action). The learner then obtains a reward R_t , where R_t is sampled from P_{A_t} . The learner would like to maximize their cumulative reward $\sum_{t=0}^H R_t$. (The environment is not really executing a strategy. The environment can be thought to flip some coins or roll some dice to present the reward to the learner.)

Since the game occurs over many rounds, A_t is a function of the history $A_1, \dots, A_{t-1}, R(A_1), \dots, R(A_{t-1})$ at time t . A **policy** π is a function whose input is the history (at any time t) and whose output is an action in \mathcal{A} .

Often the environment will be assumed to lie in a certain class (e.g. the reward distribution for each arm might be assumed to be Bernoulli, i.e. all rewards are valued in $\{0, 1\}$.)

The game has incomplete information since the learner does not know the distributions $(P_a)_{a \in \mathcal{A}}$.

For any $a \in \mathcal{A}$, let μ_a be the expected value of a random variable with distribution P_a (i.e. we could write $\mu_a = \mathbf{E}_{R \sim P_a} R$.) The **expected regret** r_H at time H is

$$H \max_{a \in \mathcal{A}} \mu_a - \sum_{t=0}^{H-1} \mathbf{E} R_t.$$

The regret of the learner relative to a policy π' is the total expected reward using policy π' for H rounds of the game, minus the total expected reward obtained by the learner in H rounds of the game. The regret relative to a set of policies Π is the maximum regret of the learner relative to all policies $\pi' \in \Pi$.

Remark 11.2. In a bandit problem, the actions of the learner at time t do not affect the rewards at future times $t+1, t+2, \dots$. However, in general reinforcement learning problems, the actions of the learner at time t *do* affect the state of the system, and the rewards at future times.

Example 11.3. Suppose $k = 2$, P_1 corresponds to a random variable taking the value $1/2$ with probability 1 , and P_2 corresponds to a random variable taking value 0 with probability $1/4$ and value 1 with probability $3/4$. Suppose the learner first plays action 1, then plays action 2, then chooses the action with largest reward for all future times. Then with probability $1/4$, the learner chooses strategy 1 at all future times, receiving an expected reward at time H of at most $(1/2)H$. However, $\max_{a \in \mathcal{A}} \mu_a = 3/4$, so the regret at time H is

$$H(3/4) - [(H-1)(1/2) + (3/4)] = (1/4)(H-1).$$

Example 11.3 shows that an insufficient amount of exploration can lead to poor regret. We would like to have the smallest regret possible. It should be possible to have a regret asymptotically smaller than H as $H \rightarrow \infty$. The following algorithm has this property for appropriately chosen j .

Algorithm 11.4 (Explore then Commit). Input: $j \geq 1$ and a bandit problem with k arms

Output: a policy for the bandit problem

- For each action $a \in \{1, \dots, k\}$, play action a , j times, for a total of jk rounds of the game.
- For each action, let M_a be the sample mean of the rewards from action a .
- In the remaining $H - jk$ rounds, play an action that maximizes m_a over all $a \in \mathcal{A}$.

Proposition 11.5. Suppose $0 \leq R_t \leq 1$ for all $t \geq 0$ and $k \leq H^3$. Then Algorithm 11.4 with $j := (H/k)^{2/3}(\log(H))^{1/3}$ has regret bounded by

$$6H^{2/3}(k \log H)^{1/3}$$

for a k -arm bandit problem with horizon H .

Proof. It follows from Hoeffding's inequality (Theorem 12.1 below) that, for any fixed $a \in \mathcal{A}$,

$$\mathbf{P}(|M_a - \mu_a| \geq \sqrt{(2/j) \log H}) \leq 2e^{-2j\sqrt{(2/j) \log H}^2} = 2H^{-4}.$$

(We emphasize that M_a is random, since it is a function of the random rewards R_0, R_1, \dots . So \mathbf{P} denotes the probability with respect to the randomness of M_a , while μ_a is not random.) Let C denote the event $\{|M_a - \mu_a| < \sqrt{(2/j) \log H}, \forall a \in \mathcal{A}\}$. From the union bound,

$$\mathbf{P}(C^c) \leq \sum_{a=1}^k \mathbf{P}(|m_a - \mu_a| > \sqrt{(2/j) \log H}) \leq 2kH^{-4}.$$

That is,

$$\mathbf{P}(C) \geq 1 - 2k/H^4.$$

Let a_1 maximize M_a over all $a \in \mathcal{A}$ and let a_2 maximize μ_a over all $a \in \mathcal{A}$. (Note that a_1 is random, since it is a function of M_1, \dots, M_k , which is itself random.) When C occurs, we have

$$\sqrt{(2/j) \log H} + \mu_{a_1} \geq M_{a_1} \geq M_{a_2} \geq \mu_{a_2} - \sqrt{(2/j) \log H},$$

that is

$$0 \leq \mu_{a_2} - \mu_{a_1} \leq 2\sqrt{(2/j) \log H}.$$

So, under the event C , the regret at time H is bounded by

$$jk + (H - jk)2\sqrt{(2/j) \log H}.$$

And under the event C^c , we have the regret bound of H at time H . Combining these,

$$\begin{aligned} r_H &\leq H\mathbf{P}(C^c) + [jk + (H - jk)2\sqrt{(2/j) \log H}]\mathbf{P}(C) \\ &\leq H(2k/H^4) + [jk + (H - jk)2\sqrt{(2/j) \log H}]. \end{aligned}$$

Choosing $j = (H/k)^{2/3}(\log(H))^{1/3}$ concludes the proof. \square

Remark 11.6. The ε -Greedy algorithm has a similar regret bound and a similar idea to Algorithm 11.4. At each step of this algorithm, with probability $1 - \varepsilon$ we play an action a with largest value of m_a , and with probability ε we play a uniformly random action in \mathcal{A} . Choosing $\varepsilon \approx (k \log(kH/\delta)/H)^{1/3}$ then obtains a regret bound as in Proposition 11.5, with probability at least $1 - \delta$.

Remark 11.7. The assumption $k \leq H^3$ is sensible; in fact $k \leq H$ is also sensible, since $k > H$ implies that there are more arms than chances to explore them.

There is an algorithm with a regret bound strictly better than Proposition 11.5. In Algorithm 11.4, we explored all arms the same number of times. But perhaps a really bad arm should be explored less than others.

Continuing the above notation, let $N_t(a)$ be the number of times action $a \in \mathcal{A}$ has been played at time t , let $M_{t,a}$ be the sample mean of the rewards obtained at time t from action a being played. For any $t \geq 0$ where $N_t(a) \geq 1$, define

$$\begin{aligned} UCB_t(a) &:= M_{t,a} + \sqrt{\frac{2 \log H}{N_t(a)}}. \\ LCB_t(a) &:= M_{t,a} - \sqrt{\frac{2 \log H}{N_t(a)}}. \end{aligned}$$

Algorithm 11.8 (Successive Elimination). Input: A bandit problem with k arms
Output: a policy for the bandit problem

- Initialize all actions $a \in \{1, \dots, k\}$ (or arms) as “active.”
- While there is more than one active arm, do the following.
 - Try all active arms once (so that $N_t(a) \geq 1$ for all $t \geq k$).
 - For each active arm a , if t denotes the current time, check if there exists $a' \in \mathcal{A}$ such that $UCB_t(a) < LCB_t(a')$.
 - If such an a' exists, deactivate arm a .

The intuition of Algorithm 11.8 is that the interval $(LCB_t(a), UCB_t(a))$ is an approximate confidence interval for the value of μ_a , so if one of these intervals for some $a \in \mathcal{A}$ is disjoint from (and smaller than) another one of these intervals for some $a' \in \mathcal{A}$, then μ_a is probably not very large and a should be eliminated.

Theorem 11.9. Suppose $0 \leq R_t \leq 1$ for all $t \geq 0$ and $k \leq H$. Then Algorithm 11.8 has regret bounded by

$$8(kH \log H)^{1/2}$$

for a k -arm bandit problem with horizon H .

Proof. For any $j \geq 1$, let $S_{j,a}$ be the sample mean of the rewards obtained from the first j times that action a was played. From Hoeffding’s inequality, (Theorem 12.1 below), for any fixed $a \in \mathcal{A}$,

$$\mathbf{P}(|S_{j,a} - \mu_a| \geq \sqrt{(2/j) \log H}) \leq 2e^{-2j\sqrt{(2/j) \log H}^2} = 2H^{-4}.$$

(We would like to apply Hoeffding’s inequality to $M_{t,a}$ instead, but $M_{t,a}$ is a sum of a random number of terms, so this cannot be done directly.) Taking the union bound over all $1 \leq j \leq H$ and $a \in \mathcal{A}$,

$$\mathbf{P}(\exists a \in \mathcal{A}, \exists 1 \leq j \leq H: |S_{j,a} - \mu_a| \geq \sqrt{(2/j) \log H}) \leq 2kH^{-3}.$$

That is,

$$\mathbf{P}(C) := \mathbf{P}(\forall a \in \mathcal{A}, \forall 1 \leq j \leq H: |S_{j,a} - \mu_a| \leq \sqrt{(2/j) \log H}) \geq 1 - 2kH^{-3}.$$

Let $\tilde{a} \in \mathcal{A}$ maximize μ_a over all $a \in \mathcal{A}$. The event C implies that $|M_{t,a} - \mu_a| \leq \sqrt{(2/N_t(a)) \log H}$ for all $t \geq k$ and for all $a \in \mathcal{A}$. So, if the event C occurs, then \tilde{a} is never deactivated in Algorithm 11.8, since for any $a \in \mathcal{A}$,

$$M_{t,\tilde{a}} + \sqrt{\frac{2 \log H}{N_t(\tilde{a})}} \geq \mu_{\tilde{a}} \geq \mu_a \geq M_{t,a} - \sqrt{\frac{2 \log H}{N_t(a)}}$$

That is, $UCB_t(\tilde{a}) \geq LCB_t(a)$ for all $a \in \mathcal{A}$ and $t \geq k$. For any $a \in \mathcal{A}$, let $t = t(a)$ be the last time where a is active in Algorithm 11.8 (within the loop that compares UCB and LCB values). By definition of Algorithm 11.8, since a is not eliminated at time t

$$\delta(a) := \mu_{\tilde{a}} - \mu_a \leq 2 \left(\sqrt{\frac{2 \log H}{N_t(\tilde{a})}} + \sqrt{\frac{2 \log H}{N_t(a)}} \right) = 4 \sqrt{\frac{2 \log H}{N_t(a)}},$$

Since $N_t(a) = N_t(\tilde{a})$, by definition of Algorithm 11.8. Since a will be played only one more time, we have $N_H(a) = N_t(a) + 1$, so that

$$\delta(a) \leq 8\sqrt{\frac{\log H}{N_H(a)}}, \quad \forall a \in \mathcal{A} \text{ with } \mu_a < \mu_{\tilde{a}}.$$

Consequently, an arm played many times (which has large $N_H(a)$ value) will have an expected reward close to the best one.

The regret at time H from playing action a is equal to

$$N_H(a)\delta(a) \leq N_H(a)8\sqrt{\frac{\log H}{N_H(a)}} = 8\sqrt{N_H(a)\log H}.$$

Summing over $a \in \mathcal{A}$, we get total regret bounded by

$$\begin{aligned} 8\sqrt{\log H} \sum_{a \in \mathcal{A}} \sqrt{N_H(a)} &= 8\sqrt{\log H} k \frac{1}{k} \sum_{a \in \mathcal{A}} \sqrt{N_H(a)} \leq 8\sqrt{\log H} k \sqrt{\sum_{a \in \mathcal{A}} \frac{N_H(a)}{k}} \\ &= 8\sqrt{k \log H} \sqrt{\sum_{a \in \mathcal{A}} N_H(a)} = 8\sqrt{kH \log H}. \end{aligned}$$

The last inequality used Jensen's inequality, since $f(x) = \sqrt{x}$ is concave for $x \geq 0$, and $\sum_{a \in \mathcal{A}} N_H(a) = H$. \square

Algorithm 11.10 (Upper Confidence Bound). Input: A bandit problem with k arms
Output: a policy for the bandit problem

- For each action $a \in \{1, \dots, k\}$, play action a , one time.
- For each action, let M_a be the sample mean of the rewards from action a , and let $N_t(a)$ denote the number of times action a has been played at time $t > k$.
- At any remaining time $t > H - k$, play an action $a \in \mathcal{A}$ that maximizes $UCB_t(a)$.

Theorem 11.11. Suppose $0 \leq R_t \leq 1$ for all $t \geq 0$ and $k \leq H$. Then Algorithm 11.10 has regret bounded by

$$8(kH \log H)^{1/2}$$

for a k -arm bandit problem with horizon H .

Proof. As in Theorem 11.9, we assume for the moment that C occurs. Let $a_t \in \mathcal{A}$ denote the action chosen in round $t > H - k$. Then $UCB_t(a_t) \geq UCB_t(\tilde{a})$ and C implies that $M_{ta_t} \leq \mu_{a_t} + \sqrt{\frac{2 \log H}{N_t(a_t)}}$ and that $UCB_t(\tilde{a}) \geq \mu_{\tilde{a}}$, so

$$\mu_{a_t} + 2\sqrt{\frac{2 \log H}{N_t(a_t)}} \geq M_{t,a_t} + \sqrt{\frac{2 \log H}{N_t(a_t)}} = UCB_t(a_t) \geq UCB_t(\tilde{a}) \geq \mu_{\tilde{a}}.$$

Therefore,

$$\delta(a_t) := \mu_{\tilde{a}} - \mu_{a_t} \leq 2\sqrt{\frac{2 \log H}{N_t(a_t)}}$$

We can now conclude as in Theorem 11.9. \square

Algorithms 11.8 and 11.10 give essentially the best possible expected regret (in the worst case) in the following sense.

Theorem 11.12. *There exists a constant $c > 0$ such that the following holds. Suppose we have a k -arm bandit problem with horizon H . Then, there exists rewards $0 \leq R_t \leq 1$ for all $t \geq 0$ and $k \leq H$ with expected regret at least*

$$c(kH)^{1/2}.$$

11.2. Markov Decision Process. Let $\Delta([0, 1])$ denote the set of all probability distributions on the interval $[0, 1]$.

Definition 11.13. A deterministic **Markov Decision Process** is a tuple $M = (\mathcal{S}, \mathcal{A}, R, p)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $R: \mathcal{S} \times \mathcal{A} \rightarrow \Delta([0, 1])$ is the stochastic reward function, and $p: \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ is the deterministic transition function.

And MDP M

12. APPENDIX: Hoeffding's Inequality

Theorem 12.1 (Hoeffding Inequality/ Large Deviation Estimate). *For all $i \geq 1$, let $a_i < b_i$ be real numbers. Let X_1, X_2, \dots be independent random variables with $\mathbf{P}(X_i \in [a_i, b_i]) = 1$. Then, for any $n \geq 1$,*

$$\mathbf{P}\left(\frac{1}{n} \sum_{i=1}^n X_i - \mathbf{E}\left(\frac{1}{n} \sum_{j=1}^n X_j\right) \geq t\right) \leq e^{-\frac{2t^2 n^2}{\sum_{i=1}^n (b_i - a_i)^2}}, \quad \forall t \geq 0.$$

Consequently,

$$\mathbf{P}\left(\left|\frac{1}{n} \sum_{i=1}^n X_i - \mathbf{E}\left(\frac{1}{n} \sum_{j=1}^n X_j\right)\right| \geq t\right) \leq 2e^{-\frac{2t^2 n^2}{\sum_{i=1}^n (b_i - a_i)^2}}, \quad \forall t \geq 0.$$

Proof. Let $\alpha > 0$. Using Markov's inequality, and $\alpha t \geq 0$,

$$\begin{aligned} \mathbf{P}\left(\sum_{i=1}^n [X_i - \mathbf{E}X_i] \geq t\right) &= \mathbf{P}(e^{\alpha \sum_{i=1}^n [X_i - \mathbf{E}X_i]} \geq e^{\alpha t}) \\ &\leq e^{-\alpha t} \mathbf{E}e^{\alpha \sum_{i=1}^n [X_i - \mathbf{E}X_i]} = e^{-\alpha t} \prod_{i=1}^n \mathbf{E}e^{\alpha(X_i - \mathbf{E}X_i)}. \end{aligned}$$

Substituting Lemma 12.2 below, for any $t \geq 0$

$$\mathbf{P}\left(\sum_{i=1}^n (X_i - \mathbf{E}X_i) \geq t\right) \leq e^{-\alpha t} e^{\alpha^2 \sum_{i=1}^n (b_i - a_i)^2 / 8} = e^{-\alpha t + \alpha^2 \sum_{i=1}^n (b_i - a_i)^2 / 8}.$$

Since $\alpha > 0$ is arbitrary, we choose α to minimize the right side. This minimum occurs when $\alpha = 4t / \sum_{i=1}^n (b_i - a_i)^2$, so that $-\alpha t + \alpha^2 \sum_{i=1}^n (b_i - a_i)^2 / 8 = -2t^2 / \sum_{i=1}^n (b_i - a_i)^2$, giving the first desired bound. The final bound follows by writing $\mathbf{P}(|\sum_{i=1}^n [X_i - \mathbf{E}X_i]| \geq t) = \mathbf{P}(\sum_{i=1}^n [X_i - \mathbf{E}X_i] \geq t) + \mathbf{P}(-\sum_{i=1}^n [X_i - \mathbf{E}X_i] \geq t)$ and then applying the first inequality twice. \square

Lemma 12.2 (Hoeffding's Lemma). *Let $a < b$ be real numbers. Let X be a random variable with $\mathbf{P}(X \in [a, b]) = 1$. Then for any $\alpha \in \mathbb{R}$,*

$$\mathbf{E}e^{\alpha(X - \mathbf{E}X)} \leq e^{\frac{1}{8}\alpha^2(b-a)^2}.$$

Proof. By replacing X with $X - \mathbf{E}X$, we may assume that $\mathbf{E}X = 0$ and instead prove that

$$\mathbf{E}e^{\alpha X} \leq e^{\frac{1}{8}\alpha^2(b-a)^2}.$$

Since $a \leq X \leq b$ with probability one, there is a random $Y \in [0, 1]$ such that $X =: aY + b(1 - Y)$. That is, this equality holds when $Y := (X - b)/(a - b)$. By convexity of $x \mapsto e^{\alpha x}$,

$$e^{\alpha X} = e^{\alpha(aY + b(1-Y))} \leq Y e^{\alpha a} + (1 - Y) e^{\alpha b} = \frac{X - b}{a - b} e^{\alpha a} + \frac{a - X}{a - b} e^{\alpha b}.$$

Let $\gamma := -b/(a - b)$, $c := \alpha(a - b)$, $f(c) := -\gamma c + \log(1 - \gamma + \gamma e^c)$. Taking expectations of both sides and using $\mathbf{E}X = 0$ we get

$$\mathbf{E}e^{\alpha X} \leq -\frac{b}{a - b} e^{\alpha a} + \frac{a}{a - b} e^{\alpha b} = e^{f(c)}.$$

Note that $f(0) = 0$, $f'(c) = -\gamma + \frac{\gamma e^c}{1-\gamma+\gamma e^c}$, $f'(0) = 0$, and

$$f''(c) = \frac{\gamma e^c}{1-\gamma+\gamma e^c} - \frac{\gamma^2 e^{2c}}{(1-\gamma+\gamma e^c)^2} = \frac{\gamma e^c}{1-\gamma+\gamma e^c} \left(1 - \frac{\gamma e^c}{1-\gamma+\gamma e^c}\right)$$

So, if $s := \frac{\gamma e^c}{1-\gamma+\gamma e^c}$, we have $f''(c) = s(1-s) \leq 1/4$. So, by Taylor's Theorem (with error term), for any $c \in \mathbb{R}$, there exists c_0 between 0 and c such that

$$f(c) = f(0) + cf'(0) + \frac{c^2}{2}f''(c_0) = \frac{c^2}{2}f''(c_0) \leq \frac{c^2}{8}.$$

In conclusion $\mathbf{E}e^{\alpha X} \leq e^{c^2/8}$. □

13. APPENDIX: NOTATION

Let n, m be a positive integers. Let Y, Z be sets contained in a space X .

\mathbb{R} denotes the set of real numbers

\in means “is an element of.” For example, $2 \in \mathbb{R}$ is read as “2 is an element of \mathbb{R} .”

\forall means “for all”

\exists means “there exists”

$\mathbb{R}^n = \{(x_1, x_2, \dots, x_n) : x_i \in \mathbb{R}, \forall 1 \leq i \leq n\}$

$\Delta_m := \{x = (x_1, \dots, x_m) \in \mathbb{R}^m : \sum_{i=1}^m x_i = 1, x_i \geq 0, \forall 1 \leq i \leq m\}$

$f: Y \rightarrow Z$ means f is a function with domain Y and range Z . For example,

$f: \mathbb{R}^2 \rightarrow \mathbb{R}$ means that f is a function with domain \mathbb{R}^2 and range \mathbb{R}

\emptyset denotes the empty set

$Y \subseteq Z$ means $\forall y \in Y$, we have $y \in Z$, so Y is contained in Z

$Y \setminus Z := \{x \in Y : x \notin Z\}$

$Y \cap Z$ denotes the intersection of Y and Z

$Y \cup Z$ denotes the union of Y and Z

S_n denotes the set of permutations on n elements

2^Y denotes the set of all subsets of Y . Equivalently, 2^Y is the set of all functions $f: Y \rightarrow \{0, 1\}$. More generally, Z^Y denotes the set of all functions $f: Y \rightarrow Z$.

Suppose we have an impartial combinatorial game under normal play with some set of game positions. Then

N denotes the set of game positions such that the first player to move can guarantee a win.

P denotes the set of terminal game positions together with the set of all game positions such that any legal move leads to a position in **N**.

Let a, b be real numbers. Let A be a matrix. Let d be a positive integer. Let K be a closed and bounded subset of \mathbb{R}^d . Let $f: \mathbb{R}^d \rightarrow \mathbb{R}$ be a continuous function.

$\min(a, b)$ denotes the minimum of a and b .

$\max(a, b)$ denotes the maximum of a and b .

A^T denotes the transpose of A .

$\min_{x \in K} f(x)$ denotes the minimum value of f on K .

$\max_{x \in K} f(x)$ denotes the maximum value of f on K .

Let $R > 0$, and let $x = (x_1, \dots, x_d) \in \mathbb{R}^d$.

$$\|x\| := \left(\sum_{i=1}^d x_i^2 \right)^{1/2}, \text{ the 2-norm on } \mathbb{R}^d.$$

$B_R(x) := \{y \in \mathbb{R}^d : \|y - x\| \leq R\}$, the ball of radius R centered at x .

∂K denotes the boundary of K , i.e. the set of all points $x \in \mathbb{R}^d$ such that, for any $R > 0$, $B_R(x)$ contains at least one point in K and at least one point not in K .

Let n be a positive integer. Let $f, g: \{-1, 1\}^n \rightarrow \mathbb{R}$. Let $S \subseteq \{1, \dots, n\}$. Let $|S|$ denote the number of elements of S . Let $x = (x_1, \dots, x_n) \in \{-1, 1\}^n$. Let $\rho \in (0, 1)$.

$$W_S(x) = \prod_{i \in S} x_i$$

$$\langle f, g \rangle = 2^{-n} \sum_{x \in \{-1, 1\}^n} f(x)g(x) \quad , \text{ an inner product on the vector space}$$

of functions from $\{-1, 1\}^n$ to \mathbb{R} .

$$\widehat{f}(S) = \langle f, W_S \rangle = 2^{-n} \sum_{x \in \{-1, 1\}^n} f(x)W_S(x) \quad , \text{ the Fourier coefficient of } f \text{ associated to } S.$$

$$T_\rho f(x) = \sum_{S \subseteq \{1, \dots, n\}} \widehat{f}(S) \rho^{|S|} W_S(x)$$

$\langle f, T_\rho f \rangle$ denotes the noise stability of f with parameter ρ .

Let $t \in \mathbb{R}$. We define $\text{sign}: \mathbb{R} \rightarrow \{-1, 0, 1\}$ by

$$\text{sign}(t) = \begin{cases} 1 & , \text{ if } t > 0 \\ -1 & , \text{ if } t < 0 \\ 0 & , \text{ if } t = 0. \end{cases}$$

Let X, Y be sets, and let $f: X \rightarrow Y$ be a function. The function $f: X \rightarrow Y$ is said to be a **one-to-one correspondence** if and only if: for every $y \in Y$, there exists exactly one $x \in X$ such that $f(x) = y$.

USC MATHEMATICS, LOS ANGELES, CA
Email address: stevenmheilman@gmail.com