

2020

Introduction to Game Theory: A Discovery Approach

Jennifer Firkins Nordstrom
Linfield College

Follow this and additional works at: <https://digitalcommons.linfield.edu/linfauth>



Part of the [Logic and Foundations Commons](#), [Set Theory Commons](#), and the [Statistics and Probability Commons](#)

Recommended Citation

Nordstrom, Jennifer Firkins, "Introduction to Game Theory: A Discovery Approach" (2020). *Linfield Authors Book Gallery*. 83.

<https://digitalcommons.linfield.edu/linfauth/83>

This Book is protected by copyright and/or related rights. It is brought to you for free via open access, courtesy of DigitalCommons@Linfield, with permission from the rights-holder(s). Your use of this Book must comply with the [Terms of Use](#) for material posted in DigitalCommons@Linfield, or with other stated terms (such as a Creative Commons license) indicated in the record and/or on the work itself. For more information, or if you have questions about permitted uses, please contact digitalcommons@linfield.edu.

Introduction to Game Theory

a Discovery Approach

Jennifer Firkins Nordstrom

Linfield College
McMinnville, OR

January 4, 2020

Version 4, January 2020

©2020 Jennifer Firkins Nordstrom

This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

Preface

Many colleges and universities are offering courses in quantitative reasoning for all students. One model for a quantitative reasoning course is to provide students with a single cohesive topic. Ideally, such a topic can pique the curiosity of students with wide ranging academic interests and limited mathematical background. This text is intended for use in such a course. Game theory is an excellent topic for a non-majors quantitative course as it develops mathematical models to understand human behavior in social, political, and economic settings. The variety of applications can appeal to a broad range of students. Additionally, students can learn mathematics through playing games, something many choose to do in their spare time!

This text particularly explores the ideas of game theory through the rich context of popular culture. At the end of each chapter is a section on applications of the concepts to popular culture. It suggests films, television shows, and novels with themes from game theory. The questions in each of these sections are intended to serve as essay prompts for writing assignments.

Course Goals.

- Introduce students to the mathematics of game theory.
- Teach students how to use mathematical models to solve problems in social and economic situations.
- Build students' quantitative intuition.
- Introduce students to the power of mathematics to frame human behavior.
- Provide students an opportunity to use algebraic techniques, such as linear models and systems of equations, in game theoretic applications.
- Provide students an opportunity to use basic ideas of probability, such as expected value, in game theoretic applications.

Course Format. The material is presented in a discovery format, requiring students to make conjectures frequently. Each section is structured as a class activity. Any introduction material can be read by the students, and the numbered problems or questions are to be out in class and as homework, depending on how far a particular student progresses through the section. Most sections require students to attempt to solve the problem *before* they have been provided much framework. The sections then build the necessary tools to solve the problem or understand the key ideas. Being able to compare their original solutions and ideas to the more sophisticated mathematical ones helps build

their mathematical intuition and helps them to understand the power of using mathematics in situations where their intuition falls short.

Suggestions for Use. This text is primarily for use in a college-level quantitative reasoning course. It can also be used for an introductory course in game theory for the social sciences. It approaches the subject matter through an inquiry-based format. Most of the topics can be introduced by providing the students with the activity to work through during class, followed by a discussion. Almost all of the activities are intended to work through the concepts without additional lecture or introduction. Students with even a rudimentary background in algebra will find the material accessible. Any necessary mathematical background can be introduced as needed.

Contents

Preface	v
1 What is Game Theory?	2
1.1 Players and Strategies	2
1.2 Game Matrices and Payoff Vectors	5
2 Two-Person Zero-Sum Games	8
2.1 Introduction to Two-Person Zero-Sum Games	8
2.2 Dominated Strategies	13
2.3 Probability and Expected Value	19
2.4 A Game of Chance	22
2.5 Equilibrium Points	24
2.6 Strategies for Zero-Sum Games and Equilibrium Points	26
2.7 Popular Culture: Rationality and Perfect Information	27
3 Repeated Two-Person Zero-sum Games	29
3.1 Introduction to Repeated Games	29
3.2 Mixed Strategies: Graphical Solution.	32
3.3 Using Sage to Graph Lines and Solve Equations.	39
3.4 Mixed Strategies: Expected Value Solution	43
3.5 Liar's Poker	46
3.6 Augmented Matrices.	48
3.7 Undercut	51
4 Non-Zero-Sum Games	54
4.1 Introduction to Two-Player Non-Zero-Sum Games	54
4.2 Prisoner's Dilemma and Chicken	57
4.3 A Class-Wide Experiment.	61
4.4 What Makes a Prisoner's Dilemma?	63
4.5 Another Multiplayer Experiment	64
4.6 Volunteer's Dilemma	66
4.7 Repeated Prisoner's Dilemma	68
4.8 Popular Culture: Prisoner's Dilemma and Chicken.	71

CONTENTS

viii

References

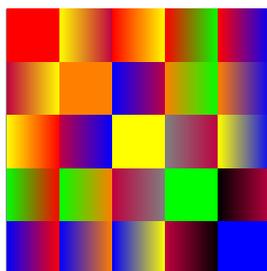
73

Index

74

Chapter 1

What is Game Theory?



“Game Theory is not about ‘playing games.’ It is about conflict resolution among rational but distrustful beings.” Poundstone, *Prisoner’s Dilemma*, p. 39.

Although we will play many games throughout this book, our goal is to understand how rational, distrustful players would play the game. These games are meant to serve as models for situations of conflict. We will explore how to “solve” games under certain assumptions about our players. As with any mathematical model, we will need to make assumptions about how our players will behave, what information they have, and the constraints of the game. For example, we will assume that our players will use all information available to them, and that players will follow the rules of the game.

Games can provide hours of recreational enjoyment, and are worth studying for this alone. However, even simple games can be used to model political, social, and economic interactions. Understanding some foundations of game theory can help us interpret, predict, and respond in competitive situations.

1.1 Players and Strategies

In this book most of the games will be played by two *players*. Each player must decide how he or she will play the game. In order to study games mathematically, we need to make some assumptions about how the players should play the game. This allows us to be able to better predict what our players should do. The following example illustrates the characteristics we will assume about our players.

1.1.1 Example: Cake Division

How can two children fairly divide a cake? One classic solution is to have one child cut the cake and have the other child choose a piece.

Why does this work? In other words, why should both children feel they received a fair share of the cake?

What are the underlying assumptions that make this process work?

1. The goal of each player is to get the largest piece. We can think of this as each player acting in his or her *self-interest*.
2. Both players know that the other player has the same goal, and will act to further this goal. Thus, we know that each player is *rational*. Even more, each player knows the *other* player is rational.

We need both (1) and (2) to reach the solution that the cake is divided evenly and both children receive equal sized pieces.

The idea that players are self-interested is crucial to game theory. There are lots of other ways to play games, and those might be worth exploring. But to get started with game theory, we must make specific assumptions and develop the mathematical context from these assumptions.

Assumption 1: Players are self-interested. The goal is to win the most or lose the least. But what does it mean to win?

A player's *payoff* is the amount (points, money, or anything a player values) a player receives for a particular outcome of a game. We say that the player's goal is to maximize his or her payoff. We should note that the maximum payoff for a player might even be negative, in which case the player wants the least negative (or closest to 0) payoff.

It is important to recognize the difference between having the goal of maximizing the payoff and having the goal of simply winning. Here are some examples.

1. If two players were racing, a player wouldn't just want to finish first, she would want to finish by as large a margin as possible.
2. If two teams were playing basketball, the team wouldn't want to just have the higher score, they would want to win by the largest number of points. In other words, a team would prefer to win by 10 points rather than by 1 point.
3. In an election poll, a candidate doesn't just want to be ahead of her opponent, she wants lead by as large a margin as possible, (especially if she needs to account for error in the polls).

It is important to keep in mind the the goal of each player is to win the most (or lose the least). It will be tempting to look for strategies which simply assure a player of a positive payoff, but we need to make sure a player can't do even better with a different strategy.

Assumption 2: Players are perfectly logical. A Player will always take into account all available information and make the decision which maximizes his payoff. This includes knowing that his opponent is also making the best decision for herself. For example, in the cake cutting game a player wouldn't cut one large piece hoping that his opponent will by chance pick the smaller piece. A player must assume that her opponent will always choose the larger piece.

Now you may be wondering what these assumptions have to do with reality. After all, no one's perfect. But we often study ideal situations (especially in math!). For example, you've all studied geometry. Can anyone here draw a perfectly straight line? Yet, you've all studied such an object!

Our Goal: Develop strategies for our perfectly logical, self-interested players.

1.1.2 Developing Strategies: Tic Tac Toe

Exercise 1.1.1 Play the game. Play several games of Tic Tac Toe with an opponent. Make sure you take turns being the first player and the second player. Develop a strategy for winning Tic Tac Toe. You may have a different strategy for the first player and for the second player. Be as specific as possible. You may need to consider several possibilities which depend on what your opponent does.

- a Who wins? Player 1 or Player 2?
- b What must each player do in order to have the best possible outcome?
- c How did you develop your strategy?
- d How do you know it will always work?

Let us note some characteristics of Tic Tac Toe.

- There are two players.
- Players have *perfect information*. This means each player knows what all of his or her own options are, what all of his or her opponent's options are, and both players know what the outcome of each option is. Additionally, players know that both players have all of this information.
- This game has a *solution*. A solution for a game consists of a strategy for each player and the outcome of the game when each player plays his or her strategy. In Tic Tac Toe, if both players play their best, the game will always end in a tie.
- The game is *finite*. This means the game must end after a finite number of moves or turns. In other words, the game cannot go on forever. A game that is not finite is called *infinite*. Note, an infinite game may end after a finite number of turns, but there is no maximum number of turns or process to ensure the game ends. In Tic Tac Toe, the game must end after 9 or fewer turns.

Exercise 1.1.2 Perfect information, more examples. Can you think of another example of a game with perfect information? What is an example of a game that does not have perfect information?

Exercise 1.1.3 Finite and infinite, more examples. Give some examples of finite games and infinite games.

Definition 1.1.4 A **strategy** for a player is a complete way to play the game regardless of what the other player does. \diamond

The choice of what a player does may depend on the opponent, but that choice is predetermined before game play. For example, in the cake cutting game, it doesn't matter which piece the "chooser" will pick, the "cutter" will always cut evenly. Similarly, it doesn't matter how the cutter cuts, the chooser will always pick the largest piece. In Tic Tac Toe, Player 2's strategy should determine his first move no matter what Player 1 plays first. For example, if Player 1 plays the center square, where should Player 2 play? If Player 1 plays a corner, where should Player 2 play?

Exercise 1.1.5 Describe your favorite game. What is your favorite game?

- a Give a brief description of the game, including what it means to "win" or "lose" the game.
- b How many players do you need?

- c Do the players have perfect information for the game?
- d Is the game finite or can it go on forever?
- e Give some possible strategies for the player(s). Note, depending on the game, these strategies may not always result in a definite win, but they should suggest a way to increase a player's chances of winning (or not losing).

We have established a few assumptions and looked at how to describe strategies in some familiar games. Not all games easily fit into the context we will be using throughout this text. But you might keep in mind some of your favorite games and see how well the strategies and solutions can be applied to them. In the next section we develop some useful notation for describing most of the games we will study.

1.2 Game Matrices and Payoff Vectors

We need a way to describe the possible choices for the players and the outcomes of those choices. For now, we will stick with games that have only two players. We will call them Player 1 and Player 2.

1.2.1 Example: Matching Pennies

Suppose each player has two choices: Heads (H) or Tails (T). If they choose the same letter, then Player 1 wins \$1 from Player 2. If they don't match, then Player 1 loses \$1 to Player 2. We can represent all the possible outcomes of the game with a *matrix*.

Player 1's options will always correspond to the rows of the matrix, and Player 2's options will correspond to the columns. See [Table 1.2.1](#).

Table 1.2.1 A game matrix showing the strategies for each player

		Player 2	
		Head	Tail
Player 1	Head		
	Tail		

Definition 1.2.2 A **payoff** is the amount a player receives for given outcome of the game. ◇

Now we can fill in the matrix with each player's payoff. Since the payoffs to each player are different, we will use ordered pairs where the first number is Player 1's payoff and the second number is Player 2's payoff. The ordered pair is called the *payoff vector*. For example, if both players choose H, then Player 1's payoff is \$1 and Player 2's payoff is -\$1 (since he loses to Player 1). Thus the payoff vector associated with the outcome H, H is $(1, -1)$.

We fill in the matrix with the appropriate payoff vectors in [Table 1.2.3](#)

Table 1.2.3 A game matrix showing the payoff vectors

		Player 2	
		H	T
Player 1	H	$(1, -1)$	$(-1, 1)$
	T	$(-1, 1)$	$(1, -1)$

It is useful to think about different ways to quantify winning and losing. What are some possible measures of value? For example, we could use money, chips, counters, votes, points, amount of cake, etc.

Remember, a player always prefers to win the MOST points (money, chips, votes, cake), not just more than her opponent. If you want to study a game where players simply win or lose (such as Tic Tac Toe), we could just use “1” for a win and “-1” for a loss.

1.2.2 Understanding the Players

Recall that we said there are two major assumptions we must make about our players:

- Our players are *self-interested*. This means they will always prefer the largest possible payoff. They will choose a strategy which maximizes their payoff.
- Our players are *perfectly logical*. This means they will use all the information available and make the wisest choice for themselves.

It is important to note that each player also knows that his or her opponent is also self-interested and perfectly logical!

Exercise 1.2.4 Preferred payoffs.

- Which payoff does the player prefer: 0, 2, or -2?
- Which payoff does the player prefer: -2, -5, or -10?
- Which payoff does the player prefer: -1, -3, or 0?

It may be straightforward to decide the best payoff for a player out of a list of values, and it would be great if a player could just determine the biggest value in the table and choose that strategy. However, when there are two players a player may have to choose a strategy more carefully, since Player 1 can only choose the row, and Player 2 can only choose the column. Thus, the outcome of the game depends on BOTH players.

Example 1.2.5 A 2×2 Game. Suppose two players are playing a game in which they can choose A or B with the payoffs given in the game matrix in [Table 1.2.6](#).

Table 1.2.6 Payoff matrix for Exercise 1.2.7

		Player 2	
		A	B
Player 1	A	(100, -100)	(-10, 10)
	B	(0, 0)	(-1, 11)

In the following exercise, we will try to determine what each player should do. □

Exercise 1.2.7 Finding strategies.

- Just by quickly looking at the matrix, which player appears to be able to win more than the other player? Does one player seem to have an advantage? Explain.
- Determine what each player should do. Explain your answer.
- Compare your answer in (b) to your answer in (a). Did the player you suggested in (a) actually win more than the other player?

- d According to your answer in (b), does Player 1 end up with the largest possible payoff (for Player 1) in the matrix?
- e According to your answer in (b), does Player 2 end up with the largest possible payoff (for Player 2) in the matrix?
- f Do you still think a player has an advantage in this game? Is it the same answer as in (a)?

Example 1.2.8 A 3×3 Game. Suppose there are two players with the game matrix given in [Table 1.2.9](#).

Table 1.2.9 Payoff matrix for Exercise 1.2.10

		Player 2		
		X	Y	Z
Player 1	A	(1000, -1000)	(-5, 5)	(-15, 15)
	B	(200, -200)	(0, 0)	(-5, 5)
	C	(500, -500)	(20, -20)	(-25, 25)

In the following exercise, we will try to determine what each player should do. □

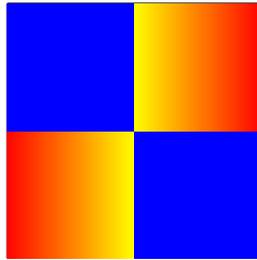
Exercise 1.2.10 More practice finding strategies.

- a Just by quickly looking at the matrix, which player appears to be able to win more than the other player? Does one player seem to have an advantage? Explain.
- b Determine what each player should do. Explain your answer.
- c Compare your answer in (b) to your answer in (a). Did the player you suggested in (a) actually win more than the other player?
- d According to your answer in (b), does Player 1 end up with the largest possible payoff (for Player 1) in the matrix?
- e According to your answer in (b), does Player 2 end up with the largest possible payoff (for Player 2) in the matrix?
- f Do you still think a player has an advantage in this game? Is it the same answer as in (a)?

This chapter has introduced you to who the players are and how to organize strategies and payoffs into a matrix. In the next chapter we will study some methods for how a player can determine his or her best strategy.

Chapter 2

Two-Person Zero-Sum Games



In this chapter we will look at a specific type of two-player game. These are often the first games studied in game theory as they can be straightforward to analyze. All of our games in this chapter will have only two players. We will also focus on games in which one player's win is the other player's loss.

2.1 Introduction to Two-Person Zero-Sum Games

In all of the examples from the last section, whatever one player won, the other player lost.

Definition 2.1.1 A two player game is called a **zero-sum** game if the sum of the payoffs to each player is constant for all possible outcomes of the game. More specifically, the terms (or coordinates) in each payoff vector must add up to the same value for each payoff vector. Such games are sometimes called **constant-sum** games instead. \diamond

We can always think of zero-sum games as being games in which one player's win is the other player's loss.

Example 2.1.2 Zero-sum in Poker. Consider a poker game in which each player comes to the game with \$100. If there are five players, then the sum of money for all five players is always \$500. At any give time during the game, a particular player may have more than \$100, but then another player must have less than \$100. One player's win is another player's loss. \square

Example 2.1.3 Zero-sum in Cake Division. Consider the cake division game. Determine the payoff matrix for this game. It is important to determine what each player's options are first: how can the "cutter" cut the cake? How can the "chooser" pick her piece? The payoff matrix is given in [Table 2.1.4](#).

Table 2.1.4 Payoff matrix for Cake Cutting game.

		Chooser	
		Larger Piece	Smaller Piece
Cutter	Cut Evenly	(half, half)	(half, half)
	Cut Unvenly	(small piece, large piece)	(large piece, small piece)

In order to better see that this game is zero-sum (or constant-sum), we could give values for the amount of cake each player gets. For example, half the cake would be 50%, a small piece might be 40%. Then we can rewrite the matrix with the percentage values in [Table 2.1.5](#)

Table 2.1.5 Payoff matrix, in percent of cake, for Cake Cutting game.

		Chooser	
		Larger Piece	Smaller Piece
Cutter	Cut Evenly	(50, 50)	(50, 50)
	Cut Unvenly	(40, 60)	(60, 40)

In each outcome, the payoffs to each player add up to 100 (or 100%). In more mathematical terms, the coordinates of each payoff vector add up to 100. Thus the sum is the same, or constant, for each outcome. \square

It is probably simple to see from the matrix in [Table 2.1.5](#) that Player 2 will always choose the large piece, thus Player 1 does best to cut the cake evenly. The outcome of the game is the *strategy pair* denoted [Cut Evenly, Larger Piece], with resulting payoff vector (50, 50).

But why are we going to call these games “zero-sum” rather than “constant-sum”? We can convert any zero-sum game to a game where the payoffs actually sum to zero.

Example 2.1.6 Poker Payoffs Revisited. Consider the above poker game where each player begins the game with \$100. Suppose at some point in the game the five players have the following amounts of money: \$50, \$200, \$140, \$100, \$10. Then we could think of their gain as -\$50, \$100, \$40, \$0, -\$90. What do these five numbers add up to? \square

Example 2.1.7 Convert the cake division payoffs so that the payoff vectors sum to zero (rather than 100).

The solution is given in [Table 2.1.8](#).

Table 2.1.8 Zero-sum payoff matrix for Cake Cutting game.

		Chooser	
		Larger Piece	Smaller Piece
Cutter	Cut Evenly	(0, 0)	(0, 0)
	Cut Unvenly	(-10, 10)	(10, -10)

But let’s make sure we understand what these numbers mean. For example, a payoff of (0, 0) does not mean each player gets no cake, it means they don’t get any more cake than the other player. In this example, each player gets half the cake (50%) plus the payoff. \square

In the form of [Example 2.1.7](#), it is easy to recognize a zero-sum game since each payoff vector has the form $(a, -a)$ (or $(-a, a)$).

2.1.1 Example: An Election Campaign Game

Two candidates, Arnold and Bainbridge, are facing each other in a state election. They have three choices regarding the issue of the speed limit on I-5:

They can support raising the speed limit to 70 MPH, they can support keeping the current speed limit, or they can dodge the issue entirely. The next three examples present three different payoff matrices for Arnold and Bainbridge.

Example 2.1.9 The Speed Limit Issue. The candidates have the information given in [Table 2.1.10](#) about how they would likely fare in the election based on how they stand on the speed limit.

Table 2.1.10 Percentage of the vote for Example 2.1.9.

		Bainbridge		
		Raise Limit	Keep Limit	Dodge
Arnold	Raise Limit	(45, 55)	(50, 50)	(40, 60)
	Keep Limit	(60, 40)	(55, 45)	(50, 50)
	Dodge	(45, 55)	(55, 45)	(40, 60)

□

Exercise 2.1.11 Analysis of the election game. For the following questions, assume Arnold and Bainbridge have the payoff matrix given in [Example 2.1.9](#).

- Explain why [Example 2.1.9](#) is a zero-sum game.
- What should Arnold choose to do? What should Bainbridge choose to do? Be sure to explain each candidate's choice. And remember, a player doesn't just want to win, he wants to get THE MOST votes— for example, you could assume these are polling numbers and that there is some margin of error, thus a candidate prefers to have a larger margin over his opponent!
- What is the outcome of the election?
- Does Arnold need to consider Bainbridge's strategies is in order to decide on his own strategy? Does Bainbridge need to consider Arnold's strategies is in order to decide on his own strategy? Explain your answer.

Example 2.1.12 A New Scenario. Bainbridge's mother is injured in a highway accident caused by speeding. The new payoff matrix is given in [Table 2.1.13](#).

Table 2.1.13 Percentage of the vote for Example 2.1.12.

		Bainbridge		
		Raise Limit	Keep Limit	Dodge
Arnold	Raise Limit	(45, 55)	(10, 90)	(40, 60)
	Keep Limit	(60, 40)	(55, 45)	(50, 50)
	Dodge	(45, 55)	(10, 90)	(40, 60)

□

Exercise 2.1.14 Analysis of the second scenario. For the following questions, assume Arnold and Bainbridge have the payoff matrix given in [Example 2.1.12](#).

- Explain why [Example 2.1.12](#) is a zero-sum game.
- What should Arnold choose to do? What should Bainbridge choose to do? Be sure to explain each candidate's choice.
- What is the outcome of the election?
- Does Arnold need to consider Bainbridge's strategies is in order to decide

on his own strategy? Does Bainbridge need to consider Arnold's strategies is in order to decide on his own strategy? Explain your answer.

Example 2.1.15 A Third Scenario. Bainbridge begins giving election speeches at college campuses and monster truck rallies. The new payoff matrix is given in [Table 2.1.16](#).

Table 2.1.16 Percentage of the vote for Example 2.1.15.

		Bainbridge		
		Raise Limit	Keep Limit	Dodge
Arnold	Raise Limit	(35, 65)	(10, 90)	(60, 40)
	Keep Limit	(45, 55)	(55, 45)	(50, 50)
	Dodge	(40, 60)	(10, 90)	(65, 35)

□

Exercise 2.1.17 Analysis of the third scenario. For the following questions, assume Arnold and Bainbridge have the payoff matrix given in [Example 2.1.15](#).

- a Explain why [Example 2.1.15](#) is a zero-sum game.
- b What should Arnold choose to do? What should Bainbridge choose to do? Be sure to explain each candidate's choice.
- c What is the outcome of the election?
- d Does Arnold need to consider Bainbridge's strategies is in order to decide on his own strategy? Does Bainbridge need to consider Arnold's strategies is in order to decide on his own strategy? Explain your answer.

Exercise 2.1.18 Changing the strategy. In each of the above scenarios, is there any reason for Arnold or Bainbridge to change his strategy? If there is, explain under what circumstances it makes sense to change strategy. If not, explain why it never makes sense to change strategy.

2.1.2 Equilibrium Pairs

Chances are, in each of the exercises above, you were able to determine what each player should do. In particular, if both players play your suggested strategies, there is no reason for either player to change to a different strategy.

Definition 2.1.19 A pair of strategies is an **equilibrium pair** if neither player gains by changing strategies. ◇

For example, consider the game matrix from [Example 1.2.5](#), [Table 1.2.6](#).

Table 2.1.20 Payoff matrix for Example 1.2.5.

		Player 2	
		X	Y
Player 1	A	(100, -100)	(-10, 10)
	B	(0, 0)	(-1, 11)

You determined that Player 2 should choose to play Y, and thus, Player 1 should play B (i.e., we have the strategy pair [B, Y]). Why is this an equilibrium pair? If Player 2 plays Y, does Player 1 have any reason to change to strategy A? No, she would lose 10 instead of 1! If Player 1 plays B, does Player 2 have any reason to change to strategy X? No, she would gain 0 instead of 1! Thus neither player benefits from changing strategy, and so we say [B, Y] is an equilibrium pair.

For now, we can use a “guess and check” method for finding equilibrium pairs. Take each outcome and decide whether either player would prefer to switch. Remember, Player 1 can only choose a different row, and Player 2 can only choose a different column. In our above example there are four outcomes to check: [A, X], [A, Y], [B, X], and [B, Y]. We already know [B, Y] is an equilibrium pair, but let’s check the rest. Suppose the players play [A, X]. Does Player 1 want to switch to B? No, she’d rather get 100 than 0. Does player 2 want to switch to Y? Yes! She’d rather get 10 than -100. So [A, X] is NOT an equilibrium pair since a player wants to switch. Now check that for [A, Y] Player 1 would want to switch, and for [B, X] both players would want to switch. Thus [A, Y] and [B, X] are NOT equilibrium pairs. Now you can try to find equilibrium pairs in any matrix game by just checking each payoff vector to see if one of the players would have wanted to switch to a different strategy.

Exercise 2.1.21 Checking equilibrium pairs. Are the strategy pairs you determined in the three election scenarios equilibrium pairs? In other words, would either player prefer to change strategies? (You don’t need to check whether any other strategies are equilibrium pairs.)

Exercise 2.1.22 Using “guess and check”. Use the “guess and check” method to determine any equilibrium pairs for the following payoff matrices.

a

$$\begin{bmatrix} (2, -2) & (2, -2) \\ (1, -1) & (3, -3) \end{bmatrix}$$

b

$$\begin{bmatrix} (3, -3) & (1, -1) \\ (2, -2) & (4, -4) \end{bmatrix}$$

c

$$\begin{bmatrix} (4, -4) & (5, -5) & (4, -4) \\ (3, -3) & (0, 0) & (1, -1) \end{bmatrix}$$

After trying the above examples, do you think every game has an equilibrium pair? Can games have multiple equilibrium pairs?

Exercise 2.1.23 Existence of equilibrium pairs. Do all games have equilibrium pairs?

Exercise 2.1.24 Multiple equilibrium pairs. Can a game have more than one equilibrium pair?

The last three exercises give you a few more games to practice with.

Exercise 2.1.25 Rock, paper, scissors. Consider the game ROCK, PAPER, SCISSORS (Rock beats Scissors, Scissors beat Paper, Paper beats Rock). Construct the payoff matrix for this game. Does it have an equilibrium pair? Explain your answer.

Exercise 2.1.26 Battle of the networks. Two television networks are battling for viewers for 7 pm Monday night. They each need to decide if they are going to show a sitcom or a sporting event. [Table 2.1.27](#) gives the payoffs as percent of viewers.

Table 2.1.27 Payoff matrix for Battle of the Networks.

		Network 2	
		Sitcom	Sports
Network 1	Sitcom	(55, 45)	(52, 48)
	Sports	(50, 50)	(45, 55)

- a Explain why this is a zero-sum game.
- b Does this game have an equilibrium pair? If so, find it and explain what each network should do.
- c Convert this game to one in which the payoffs actually sum to zero. Hint: if a network wins 60% of the viewers, how much more than 50% of the viewers does it have?

Exercise 2.1.28 Competitive advantage. This game is an example of what economists call *Competitive Advantage*. Two competing firms need to decide whether or not to adopt a new type of technology. The payoff matrix is in [Table 2.1.29](#). The variable a is a positive number representing the economic advantage a firm will gain if it is the first to adopt the new technology.

Table 2.1.29 Payoff matrix for Competitive Advantage.

		Firm A	
		Adopt New Tech	Stay Put
Firm B	Adopt New Tech	(0, 0)	(a, -a)
	Stay Put	(-a, a)	(0, 0)

- a Explain the payoff vector for each strategy pair. For example, why should the pair [Adopt New Tech, Stay Put] have the payoff $(a, -a)$?
- b Explain what each firm should do.
- c Give a real life example of Competitive Advantage.

We've seen how to describe a zero-sum game and how to find equilibrium pairs. We've tried to decide what each player's strategy should be. Each player may need to consider the strategy of the other player in order to determine his or her best strategy. But we need to be careful, although our intuition can be useful in deciding the best strategy, we'd like to be able to be more precise about finding strategies for each player. We'll learn some of these tools in the next section.

2.2 Dominated Strategies

Recall that in a zero-sum game, we know that one player's win is the other player's loss. Furthermore, we know we can rewrite any zero-sum game so that the player's payoffs are in the form $(a, -a)$. Note, this works even if a is negative; in which case, $-a$ is positive.

Example 2.2.1 A Simpler Payoff Matrix. Consider the zero-sum game with payoff matrix in [Table 2.2.2](#). Note that for simplicity our payoff matrix contains only the payoffs and not the strategy names; but Player 1 still chooses a row and Player 2 still chooses a column.

Table 2.2.2 The payoff matrix for Example 2.2.1.

	Player 2	
Player 1	(1, -1)	(-0, 0)
	(-1, 1)	(-2, 2)

If we know we are playing a zero-sum game, then the use of ordered pairs seems somewhat redundant: If Player 1 wins 1, then we know that Player 2 must lose 1 (win -1). Thus, if we KNOW we are playing a zero-sum game, we can simplify our notation by just using Player 1's payoffs. The above matrix in [Table 2.2.2](#) can be simplified as in [Table 2.2.3](#).

Table 2.2.3 The payoff matrix for Example 2.2.1 using only Player 1's payoffs.

	Player 2	
Player 1	1	0
	-1	-2

□

When simplifying, keep a few things in mind:

1. You MUST know that the game is zero-sum.
2. If it is not otherwise specified, the payoffs represent Player 1's payoffs.
3. You can always give a similar matrix representing Player 2's payoffs. However, due to (2), you should indicate that the matrix is for Player 2. For example, Player 2's payoff matrix would be given by [Table 2.2.4](#).

Table 2.2.4 The payoff matrix for Example 2.2.1 using only Player 2's payoffs.

	Player 2	
Player 1	-1	0
	1	2

4. Both players can make strategy decisions by considering only Player 1's payoff matrix. (Why?) Just to test this out, by looking only at the matrix in [Table 2.2.3](#) determine which strategy each player should choose.

Table 2.2.5 Player 1's payoffs in Example 2.2.1.

	Player 2	
Player 1	1	0
	-1	-2

In this last example, it should be clear that Player 1 is looking for rows which give her the largest payoff. This is nothing new. However, Player 2 is now looking for columns which give Player 1 the SMALLEST payoff. (Why?)

Now that we have simplified our notation for zero-sum games, let's try to find a way to determine the best strategy for each player.

Example 2.2.6 A 2×3 Game. Consider the zero-sum game given in [Table 2.2.7](#).

Table 2.2.7 Payoff matrix for Example 2.2.6.

	Player 2		
Player 1	1	0	2
	-1	-2	2

Determine which row Player 1 should choose. Is there any situation in which Player 1 would choose the other row? \square

Example 2.2.8 Another 2×3 Game. Consider the zero-sum game given in Table 2.2.9.

Table 2.2.9 Payoff matrix for Example 2.2.8.

		Player 2		
Player 1	1	0	2	
	-1	-2	3	

Determine which row Player 1 should choose. Is there any situation in which Player 1 would choose the other row? \square

In Example 2.2.6, no matter what Player 2 does, Player 1 would always choose Row 1, since every payoff in Row 1 is greater than or equal to the corresponding payoff in Row 2 ($1 \geq -1$, $0 \geq -2$, $2 \geq 2$). In Example 2.2.8, this is not the case: if Player 2 were to choose Column 3, then Player 1 would prefer Row 2. In Example 2.2.6 we would say that Row 1 *dominates* Row 2.

Definition 2.2.10 A strategy X **dominates** a strategy Y if every entry for X is greater than or equal to the corresponding entry for Y . In this case, we say Y is **dominated by** X .

If strategy X dominates strategy Y , we can write $X \succ Y$. \diamond

In mathematical notation, let a_{ik} be the value in the i^{th} row and k^{th} column. Similarly, a_{jk} is the value in the j^{th} row and k^{th} column. The i^{th} row dominates the j^{th} row if $a_{ik} \geq a_{jk}$ for all k , and $a_{ik} > a_{jk}$ for at least one k .

This definition can also be used for Player 2: we consider columns instead of rows. If we are looking at Player 1's payoffs, then Player 2 prefers smaller payoffs. Thus one column X dominates another column Y if all the entries in X are *smaller than or equal* to the corresponding entries in Y .

Here is the great thing: we can always eliminate dominated strategies! (Why?) Thus, in Example 2.2.6, we can eliminate Row 2, as in Figure 2.2.11.

		Player 2		
Player 1	1	0	2	
	-1	-2	2	

Figure 2.2.11 Row 2 is dominated by Row 1.

Now it is easy to see what Player 2 should do.

In Example 2.2.8, we cannot eliminate Row 2 since it is not dominated by Row 1. However, it should be clear that Column 2 dominates Column 3 (remember, Player 2 prefers SMALLER values). Thus we can eliminate Column 3 as in Figure 2.2.12.

		Player 2		
Player 1	1	0	2	
	-1	-2	3	

Figure 2.2.12 Column 3 is dominated by Column 2.

AFTER eliminating Column 3, Row 1 dominates Row 2. Now, in Figure 2.2.13 we can eliminate Row 2.

		Player 2		
		1	0	2
Player 1	1	0	2	
	-1	-2	3	

Figure 2.2.13 After eliminating Column 3, Row 2 is dominated by Row 1.

Again, now it is easy to determine what each player should do.

Exercise 2.2.14 Check equilibrium pairs. Check that the strategy pairs we determined in [Example 2.2.6](#) and [Example 2.2.8](#) are, in fact, equilibrium pairs.

Now, look back at the election examples from [Subsection 2.1.1](#) and apply the process of eliminating dominated strategies.

Exercise 2.2.15 Eliminating dominated strategies. Use the idea of eliminating dominated strategies to determine what each player should do in the Arnold/ Bainbridge examples in [Table 2.1.10](#), [Table 2.1.13](#), and [Table 2.1.16](#). Do you get the same strategy pairs as you determined in the related exercises ([Exercise 2.1.11](#), [Exercise 2.1.14](#), [Exercise 2.1.17](#))?

The next three exercises provide more practice in using dominated strategies to find equilibrium pairs.

Exercise 2.2.16 More practice with dominated strategies. Use the idea of eliminating dominated strategies to determine any equilibrium pairs in the zero-sum game given in [Table 2.2.17](#). Note, since it is a zero-sum game we need only show Player 1's payoffs. Explain all the steps in your solution. If you are unable to find an equilibrium pair, explain what goes wrong.

Table 2.2.17 Payoff matrix for Exercise 2.2.16.

		Player 2			
		W	X	Y	Z
Player 1	A	1	0	0	10
	B	-1	0	-2	9
	C	1	1	1	8
	D	-2	0	0	7

Exercise 2.2.18 Determine equilibrium pairs. Determine any equilibrium pairs in the zero-sum game given in [Table 2.2.19](#). Explain all the steps in your solution. If you are unable to find an equilibrium pair, explain what goes wrong.

Table 2.2.19 Payoff matrix for Exercise 2.2.18.

		Player 2			
		W	X	Y	Z
Player 1	A	1	2	3	4
	B	0	-1	-0	5
	C	-1	3	2	4
	D	0	1	-1	1

Exercise 2.2.20 Practice finding equilibrium pairs. Determine any equilibrium pairs in the zero-sum game given in [Table 2.2.21](#). Explain all the steps in your solution. If you are unable to find an equilibrium pair, explain what goes wrong.

Table 2.2.21 Payoff matrix for Exercise 2.2.20.

		Player 2			
		W	X	Y	Z
Player 1	A	-2	0	3	20
	B	1	-2	-3	0
	C	10	-10	-1	1
	D	0	0	10	15

Exercise 2.2.22 A more challenging example. Determine any equilibrium pairs in the zero-sum game given in [Table 2.2.23](#). Explain all the steps in your solution. If you are unable to find an equilibrium pair, explain what goes wrong.

Table 2.2.23 Payoff matrix for Exercise 2.2.22.

		Player 2			
		W	X	Y	Z
Player 1	A	-2	0	3	20
	B	1	-2	-5	-3
	C	10	-10	-1	1
	D	0	0	10	8

Chances are you had trouble determining an equilibrium pair for the game in [Exercise 2.2.22](#). Does this mean there isn't an equilibrium pair? Not necessarily, but we are stuck if we try to use only the idea of eliminating dominated strategies. So we need a new method.

We might think of our next method as the “worst case scenario,” or “extremely defensive play.” The idea is that we want to assume our opponent is the best player to ever live. In fact, we might assume our opponent is telepathic. So no matter what we do, our opponent will always guess what we are going to choose.

Example 2.2.24 Playing Against the Best. Assume you are Player 1, and you are playing against this “infinitely smart” Player 2. Consider the game in [Table 2.2.17](#). If you pick row A, what will Player 2 do? Player 2 will pick column X or Y. Try this for each of the rows. Which row is your best choice? If you pick A, you will get 0; if you pick B, you will get -2 ; if you pick C, you will get 1; and if you pick D you will get -2 . Thus, your best choice is to choose C and get 1. Now assume you are Player 2, and Player 1 is “infinitely smart.” Which column is your best choice? If you pick W, Player 1 will get 1 (you will get -1); if you pick X, Player 1 will get 1; if you pick Y, Player 1 will get 1; and if you pick Z, you will get 10. Thus, you can choose W, X, or Y (since you want Player 1 to win the least) and get -1 . \square

Exercise 2.2.25 A new method. Using the method described in [Example 2.2.24](#), determine what each player should do in the game in [Table 2.2.19](#).

Exercise 2.2.26 More practice with the new method. Using the method described in [Example 2.2.24](#), determine what each player should do in the game in [Table 2.2.21](#).

After working through a few examples can you describe more generally the process used in [Example 2.2.24](#)? What is Player 1 looking for in each row? Then how does she choose which row to play? What is Player 2 looking for in each column? How does he choose which column to play?

Exercise 2.2.27 Generalizing the new method. Generalize the method described in [Example 2.2.24](#). In other words, give a general rule for how Player 1 should determine his or her best move. Do the same for Player 2.

Exercise 2.2.28 The new method and equilibrium points. What do you notice about using this method on the games in [Tables 2.2.17](#), [Table 2.2.19](#), and [Table 2.2.21](#)? Is the solution an equilibrium pair?

Exercise 2.2.29 The new method on the challenging example. Now try this method on the elusive payoff matrix in [Table 2.2.23](#). What should each player do? Do you think we get an equilibrium pair? Explain.

The strategies we found using the above method have a more official name. Player 1's strategy is called the *maximin* strategy. Player 1 is maximizing the minimum values from each row. Player 2's strategy is called the *minimax* strategy. Player 2 is minimizing the maximum values from each column. Notice, we can find the maximin and minimax strategies for any zero-sum game. But do our players always want to use these strategies? Will they always result in an equilibrium pair? The next five exercises explore these questions.

Exercise 2.2.30 Look for dominated strategies. Let's consider another game matrix, given in [Table 2.2.31](#). Explain why you cannot use dominated strategies to find an equilibrium pair. Do you think there is an equilibrium pair for this game (why or why not)?

Table 2.2.31 Payoff matrix for Exercise 2.2.30.

		Player 2			
		W	X	Y	Z
Player 1	A	-2	0	3	20
	B	1	2	-3	0
	C	10	-10	-1	1
	D	0	0	10	15

Exercise 2.2.32 Find the maximin/minimax strategy. If both players use the maximin/ minimax strategy, what is the outcome of the game in [Table 2.2.31](#)?

Exercise 2.2.33 Predicting a maximin strategy. In the game in [Table 2.2.31](#), if Player 1's opponent can guess that Player 1 will choose to use a maximin strategy, is Player 1 better off *not* using the maximin strategy?

Exercise 2.2.34 Deviating from the maximin/minimax strategy. Suppose both players initially decide to use the minimax/ maximin strategy in the game in [Table 2.2.31](#). Is Player 1 better off choosing a different strategy? If Player 2 guesses a change, is Player 2 better off changing strategies? Continue this line of reasoning for several iterations. What strategies do each of the players choose? Is at least one player always better off switching strategies? Can we conclude that the maximin/ minimax strategy does not lead to an equilibrium pair?

Exercise 2.2.35 Comparing examples. Compare your answers in [Exercise 2.2.34](#) to what happens in [Exercise 2.2.16](#), [Exercise 2.2.18](#), and [Exercise 2.2.20](#). Can you identify any key differences between the games in [Exercise 2.2.34](#) and [Exercise 2.2.16](#), [Exercise 2.2.18](#), and [Exercise 2.2.20](#)?

Given a zero-sum matrix game, we can find equilibrium pairs (if they exist) by the "guess and check" method, by eliminating dominated strategies, and by looking for the minimax/maximin strategies. You should be able to apply all three methods and think about which method might be the most appropriate

for a given matrix game. For example, although “guess and check” should always find an equilibrium point if it exists, it may be very tedious to apply to a really large matrix. The maximin/minimax method might be much faster.

2.3 Probability and Expected Value

Many games have an element of chance. In order to model such games and determine strategies, we should understand how mathematicians use probability to represent chance.

2.3.1 Some Basic Probability

You are probably a little bit familiar with the idea of probability. People often talk about the chance of some event happening. For example, a weather forecast might say there is a 20% chance of rain. Now determining the chance of rain can be difficult, so we will stick with some easier examples.

Consider a standard deck of 52 playing cards. What is the chance of drawing a red card? What is the *probability* of drawing a red card? Is there a difference between chance and probability? Yes! The probability of an event has a very specific meaning in mathematics.

The *probability* of an event E is the number of different outcomes resulting in E divided by the total number of equally likely outcomes. In mathematical symbols,

$$P(E) = \frac{\text{number of different outcomes resulting in } E}{\text{total number of equally likely outcomes}}.$$

Notice that the probability of E will always be a number between 0 and 1. An impossible event will have probability 0; an event that always occurs will have probability 1.

Thus, the probability of drawing a red card is $\frac{1}{2}$, not 50%. Although we can convert between probability and percent (since 0.5 converted to percent is 50%), it is important to answer a question about probability with a probability, not a percent.

Example 2.3.1 Drawing a Particular Suit. Given a standard deck of playing cards, what is the probability of drawing a heart?

Answer: You might say since there are four suits, and one of the suits is hearts, you have a probability of $\frac{1}{4}$. You’d be correct, but be careful with this reasoning. This works because each suit has the same number of cards, so each suit is *equally likely*. Another way to calculate the probability is to count the number of hearts (13) divided by the number of cards (52). Thus, we get a probability of $\frac{13}{52} = \frac{1}{4} = 0.25$. \square

Example 2.3.2 A Card is Missing. Now suppose the ace of spades is missing from the deck. What is the probability of drawing a heart?

Answer: As before, there are still four suits in the deck, so it might be tempting to say the probability is still $\frac{1}{4}$. But we’d be wrong! Each suit is no longer equally likely since, it is slightly *less* likely that we draw a spade. Each individual card is still equally likely, though. So now

$$P(\text{drawing a heart}) = \frac{\text{number of hearts}}{\text{number of cards}} = \frac{13}{51} = 0.255.$$

As you can see, it is now slightly more likely that we draw a heart if the ace of spades is removed from the deck. \square

Now try to compute some probabilities on your own.

Exercise 2.3.3 Probability with a single die. Consider rolling a single die. List the possible outcomes. Assuming that it is a fair die, are all the outcomes equally likely? What is the probability of rolling a 2? What is the probability of rolling an even number?

Exercise 2.3.4 Probability with red and green dice. Now consider rolling two fair dice, say a red die and a green die.

- a How many equally likely outcomes are there? List them.
- b What is the probability that you get a two on the red die and a four on the green die?
- c What is the probability that you roll a three on the red die?
- d What is the probability that you roll a two and a four?
- e What is the probability that you roll a three?
- f Compare your answers in (b) and (c) with your answers in (d) and (e). Are they the same or different? Explain.

Exercise 2.3.5 Probability with two of the same dice. Again consider rolling two fair dice, but now we don't care what color they are.

- a Does this change the number of equally likely outcomes from [Exercise 2.3.4](#)? Why or why not? It may be helpful to list the possible outcomes.
- b What is the probability that you get snake eyes (two ones)?
- c What is the probability that you roll a two and a four?
- d What is the probability that you roll a three?
- e What is the probability that you roll a two OR a four?

Exercise 2.3.6 Sums of dice. Suppose we roll two dice and add them.

- a List the possible sums.
- b What is the probability that you get a total of seven on the two dice?
- c What is the probability that you get a total of four when you roll two dice?
- d Are the events of getting a total of seven and getting a total of four equally likely? Explain.

It is important to note that just because you can list all of the possible outcomes, they may not be equally likely. As we see from [Exercise 2.3.6](#), although there are 11 possible sums, the probability of getting any particular sum (such as seven) is *not* $\frac{1}{11}$.

2.3.2 Expected Value

Definition 2.3.7 The **expected value** of a game of chance is the average net gain or loss that we would expect per game if we played the game many times. We compute the expected value by multiplying the value of each outcome by its probability of occurring and then add up all of the products. \diamond

For example, suppose you toss a fair coin: Heads, you win 25 cents, Tails, you lose 25 cents. The probability of getting Heads is $1/2$, as is the probability

of getting Tails. The expected value of the game is

$$\left(\frac{1}{2} \times .25\right) + \left(\frac{1}{2} \times (-.25)\right) = 0.$$

Thus, you would expect an average payoff of \$0, if you were to play the game several times. Note, the expected value is not necessarily the actual value of playing the game.

Exercise 2.3.8 Expected value and a two-coin game. Consider a game where you toss two coins. If you get two Heads, you win \$2. If you get a Head and a Tail, you win \$1, if you get two Tails, you lose \$4. Find the expected value of the game. (Caution: first you need to find the probability of each event— think about “equally likely” events.)

Exercise 2.3.9 Play the two-coin game. Now play the game in [Exercise 2.3.8](#) the indicated number of times. Give your actual payoff and compare it to the expected value.

- a One time.
- b Ten times.
- c Twenty-five times.
- d Is there a single possible outcome where you would actually win or lose the exact amount computed for the expected value? If not, why do we call it the expected value?

Exercise 2.3.10 Expected value of roulette. A standard roulette wheel has 38 numbered slots for a small ball to land in: 36 are marked from 1 to 36, with half of those black and half red; two green slots are numbered 0 and 00. An allowable bet is to bet on either red or black. This bet is an even money bet, which means if you win you receive twice what you bet. Many people think that betting black or red is a fair game. What is the expected value of betting \$1000 on red? Is this a fair game? Explain.

Exercise 2.3.11 Another roulette example. Considering again the roulette wheel, if you bet \$100 on a particular number and the ball lands on that number, you win \$3600. What is the expected value of betting \$100 on Red 4?

After finding the expected value of the games in the above exercises, what do you think the expected value can tell us about a game? Can you use it to decide whether you should play that game of chance or not? When will a game be advantageous for the player? We often care whether a game is “fair.” Can the expected value help you determine if a game is fair?

Exercise 2.3.12 Expected value and fairness. Use the idea of expected value to explain “fairness” in a game of chance.

The last exercise is a good challenge for exploring expected value.

Exercise 2.3.13 A betting game with two dice. You place a bet and roll two fair dice. If you roll a 7 or an 11, you receive your bet back (you break even). If you roll a 2, a 3, or a 12, then you lose your bet. If you roll anything else, you receive half of the sum you rolled in dollars. How much should you bet to make this a fair game?

Hint. It might be helpful to begin with a table showing the possible sums, their probability, and the payoff for each.

In the next section we use the ideas of probability and expected value to understand how set up a payoff matrix for a game of chance.

2.4 A Game of Chance

Now that we have worked with expected value, we can begin to analyze some simple games that involve an element of chance.

Example 2.4.1 One-Card Stud Poker. We begin with a deck of cards in which 50% are Aces (you can use Red cards for Aces) and 50% are Kings (you can use Black cards for Kings). There are two players and one dealer. The play begins by each player putting in the ante (1 chip). Each player is dealt one card face down. **WITHOUT LOOKING AT HIS OR HER CARD**, the players decide to Bet (say, 1 chip) or Fold. Players secretly show the dealer their choice. If one player bet and the other folded, then the player who bet wins. If both bet or both fold, then Ace beats King (or Red beats Black); winner takes the pot (all the chips from the ante and any bets). If there is a tie, they split the pot. \square

Exercise 2.4.2 Play One-Card Stud Poker. Play the game several times with two other people (so you have two players and a dealer), keeping track of the strategy choices of the players and the resulting payoffs.

Exercise 2.4.3 Guess a strategy. Based on playing the game, determine a possible winning strategy.

Exercise 2.4.4 Check if zero-sum. Is this a zero-sum game? Why or why not?

Exercise 2.4.5 Relationship between the deal and the strategy. Does the actual deal affect the choice of strategy?

Exercise 2.4.6 Strategy choices. On any given deal, what strategy choices does a player have?

Before moving on, you should attempt to determine the payoff matrix. The remainder of this section will be more meaningful if you have given some thought to what the payoff matrix should be. It is OK to be wrong at this point, it is not OK to not try.

Exercise 2.4.7 Determine a possible payoff matrix. Write down a possible payoff matrix for this game.

Now let's work through creating the payoff matrix for One-Card Stud Poker.

Exercise 2.4.8 Payoff for [Bet, Fold]. If Player 1 Bets and Player 2 Folds, does it matter which cards were dealt? How much does Player 1 win? How much does Player 2 lose? What is the payoff vector for [Bet, Fold]? (Keep in mind your answer to [Exercise 2.4.4](#).)

Exercise 2.4.9 Payoff for [Fold, Bet]. If Player 1 Folds and Player 2 Bets, does it matter which cards were dealt? What is the payoff vector for [Fold, Bet]?

Exercise 2.4.10 Payoff and the actual deal. If both players Bet, does the payoff depend on which cards were dealt?

To determine the payoff vector for [Bet, Bet] and [Fold, Fold] we will need to consider which cards were dealt. We can use some probability to determine the remaining payoff vectors.

Exercise 2.4.11 Probability of each deal. There are four possible outcomes of the deal— list them. What is the probability that each occurs? (Remember: the probability of an event is a number between 0 and 1.)

Exercise 2.4.12 Payoff for each deal with [Bet, Bet]. Consider the pair of strategies [Bet, Bet]. For each possible deal, determine the payoff vector. For example, if the players are each dealt an Ace (Red), how much does each player win? (Again, keep in mind your answer to [Exercise 2.4.4](#).)

In order to calculate the payoff for [Bet, Bet], we need to take a weighted average of the possible payoff vectors in [Exercise 2.4.12](#). In particular, we will “weight” a payoff by the probability that it occurs. Recall that this is the *expected value*. We will calculate the expected value separately for each player.

Exercise 2.4.13 Player 1’s expected value for [Bet, Bet]. Find the expected value for [Bet, Bet] for Player 1.

Exercise 2.4.14 Player 2’s expected value for [Bet, Bet]. Find the expected value for [Bet, Bet] for Player 2.

The pair of expected values from [Exercise 2.4.13](#) and [Exercise 2.4.14](#) is the payoff vector for [Bet, Bet].

Exercise 2.4.15 Justify using expected value as the payoff. Explain why it should make sense to use the expected values for the payoffs in the matrix for the strategy pair [Bet, Bet].

Hint. Think about what a player needs to know to choose a strategy in a game of chance.

We can use a similar process to find the payoff vector for [Fold, Fold].

Exercise 2.4.16 Repeat for [Fold, Fold]. Repeat [Exercise 2.4.12](#), [Exercise 2.4.13](#), and [Exercise 2.4.14](#) for the pair of strategies [Fold, Fold].

Exercise 2.4.17 Complete payoff matrix. Summarize the above work by giving the completed payoff matrix for One-Card Stud Poker.

Now that you have done all the hard work of finding the payoff matrix for One-Card Stud Poker, we can analyze our two-player zero-sum game using the techniques we learned in the previous sections. It is also important to see how the mathematical solution compares to our conjectured solution from [Exercise 2.4.3](#).

Exercise 2.4.18 Best strategy for One-Card Stud. Use the payoff matrix to determine the best strategy for each player. If each player uses their best strategy, what will be the outcome of the game?

Exercise 2.4.19 Compare strategies. Compare the strategy you found in [Exercise 2.4.18](#) to your suggested strategy in [Exercise 2.4.3](#). In particular, discuss how knowing the payoff matrix might have changed your strategy. Also compare the payoff that results from the strategy in [Exercise 2.4.18](#) to the payoff that results from your original strategy in [Exercise 2.4.3](#).

Since One-Card Stud Poker has an element of chance, we should see what happens if we play the game several times using the strategy from [Exercise 2.4.18](#).

Exercise 2.4.20 Payoff for repeated One-Card Stud. Use the payoff matrix to predict what the payoff to each player would be if the game is played ten times.

Exercise 2.4.21 Play repeated One-Card Stud. Play the game ten times using the best strategy. How much has each player won or lost after ten hands of One-Card Stud Poker? Compare your answer to your prediction in [Exercise 2.4.20](#). Does the actual payoff differ from the theoretical payoff? If so, why do you think this might be?

Exercise 2.4.22 Fair game. Explain why this game is considered fair.

Example 2.4.23 Generalized One-Card Stud Poker. In One-Card Stud Poker we ante one chip and bet one chip. Now, suppose we let players ante a different amount and bet a different amount (although players will still ante and bet the same amount as each other). Suppose a player antes a and bets b . How might this change the game? \square

Exercise 2.4.24 Payoff matrix for Generalized One-Card Stud. Use the method outlined for One-Card Stud Poker to determine the payoff matrix for Generalized One-Card Stud Poker.

Exercise 2.4.25 Strategy for Generalized One-Card Stud. Does the strategy change for the generalized version of the game? Explain.

Now that we have analyzed several zero-sum games, we can see how important it is to find any equilibrium pairs. In the next section we explore equilibrium strategies in more detail.

2.5 Equilibrium Points

In this section, we will try to gain a greater understanding of equilibrium strategies in a game. In general, we call the pair of equilibrium strategies an *equilibrium pair*, while we call the specific payoff vector associated with an equilibrium pair an *equilibrium point*.

Exercise 2.5.1 Find equilibrium points. Determine the equilibrium point(s) for the following games.

$$\text{a } \begin{bmatrix} (2, -2) & (-1, 1) \\ (2, -2) & (-1, 1) \end{bmatrix}$$

$$\text{b } \begin{bmatrix} (0, 0) & (-1, 1) & (0, 0) \\ (-1, 1) & (0, 0) & (-1, 1) \\ (0, 0) & (1, -1) & (0, 0) \end{bmatrix}$$

Exercise 2.5.2 An observation about equilibrium points. What do you notice about the values of the equilibrium points of the games in [Exercise 2.5.1](#)?

The big question we want to answer is “Can two equilibrium points for a two-player zero-sum game have different values?” By experimenting with some examples, try to create an example of a game with two equilibrium points where those points have different values for one of the players. If you can successfully create such an example, you will have answered the question. But just because you can’t find an example, that doesn’t mean one doesn’t exist!

After trying several examples, you might be beginning to believe that the answer to the above question is “no.” Now you are ready to try to prove the following theorem:

Solution Theorem for Zero-Sum Games. Every equilibrium point of a two-person zero-sum game has the same value.

Let’s start with the 2×2 case. We will use a *proof by contradiction*. We will assume the theorem is false and show that we get a logical contradiction. Once we reach a logical contradiction (a statement that is both true and false), we can conclude we were wrong to assume the theorem was false; hence, the theorem must be true. Make sure you are comfortable with the logic of this before moving on.

Since we want to assume the theorem is false, we assume we have a two-player zero-sum game with two different equilibrium values. Since we don’t

have a specific example of such a game, we want to represent the game in a general form. In particular, we can represent the general game

$$\begin{bmatrix} (a, -a) & (c, -c) \\ (d, -d) & (b, -b) \end{bmatrix}.$$

Note that if a is negative, then $-a$ is positive; thus, every possible set of values is represented by this matrix. We want to look at the possible cases for equilibria.

Exercise 2.5.3 Equilibria in Column 1. Explain what goes wrong if $(a, -a)$ and $(d, -d)$ are equilibria with $a \neq d$.

Hint. Think about the different cases, such as $a < d$, $a > d$.

Exercise 2.5.4 Equilibria in the same column/row. Generalize your answer to [Exercise 2.5.3](#) to explain what goes wrong if the two equilibria are in the same column. Similarly, explain what happens if the two equilibria are in the same row.

Exercise 2.5.5 Diagonal equilibria. Does the same explanation hold if the two equilibria are diagonal from each other? (Explain your answer!)

From your last answer, you should see that we need to do more work to figure out what happens if the equilibria are diagonal. So let's assume that the two equilibria are $(a, -a)$ and $(b, -b)$ with $a \neq b$. It might be helpful to draw the payoff matrix and circle the equilibria.

Exercise 2.5.6 A player prefers the value of an equilibrium. Construct a system of inequalities using the fact that a player prefers an equilibrium point to another choice. For example, Player 1 prefers a to d . Thus, $a > d$. List all four inequalities you can get using this fact. You should get two for each player—remember that Player 1 can only compare values in the same column since he has no ability to switch columns. If necessary, convert all inequalities to ones without negatives. (Algebra review: $-5 < -2$ means $5 > 2$!)

Exercise 2.5.7 Create strings of inequalities. Now string your inequalities together. For example, if $a < b$ and $b < c$ then we can write $a < b < c$. (Be careful, the inequalities must face the same way; we cannot write $a > b < c$!)

Exercise 2.5.8 You now have a contradiction. Explain why you now have a contradiction (a statement that *must* be false). We can now conclude that our assumption that $a \neq b$ was wrong.

Exercise 2.5.9 Diagonal case for c and d . Repeat the above argument ([Exercise 2.5.6](#), [Exercise 2.5.7](#), and [Exercise 2.5.8](#)) for the case that the two equilibria are $(d, -d)$ and $(c, -c)$ with $d \neq c$.

Exercise 2.5.10 Summarize conclusion. Explain why you can conclude that all equilibria in a 2×2 two-player zero-sum game have the same value.

We just worked through the proof, step-by-step, but now you need to put all the ideas together for yourself.

Exercise 2.5.11 The complete proof. Write up the complete proof for the 2×2 case in your own words.

Exercise 2.5.12 Generalizing to a larger game. Can you see how you might generalize to a larger game matrix? You do not need to write up a proof of the general case, just explain how the key ideas from the 2×2 case would apply to a bigger game matrix.

Hint. Think about equilibria in (a) the same row, (b) in the same column, or (c) in a different row and column.

We've seen that any two equilibrium points must have the same value. However, it is important to note that just because an outcome has the same value as an equilibrium point, that does not mean it is also an equilibrium point.

Exercise 2.5.13 Equal values may not be equilibria. Give a specific example of a game matrix with two outcomes that are $(0, 0)$, where one is an equilibrium point and the other is not.

Working through the steps of a mathematical proof can be challenging. As you think about what we did in this section, first make sure you understand the argument for each step. Then work on understanding how the steps fit together to create the larger argument.

The next section summarizes all our work with finding equilibrium points for zero-sum games.

2.6 Strategies for Zero-Sum Games and Equilibrium Points

Throughout this chapter, we have been trying to find solutions for two player zero-sum games by deciding what two rational players should do. In this section, we will try to understand where we are with solving two-player zero-sum games. The exercises in this section are intended to review the concepts of dominated strategies, equilibrium points, and the maximin/minimax strategies. By working through your own examples, we hope to tie these concepts together and ask some bigger questions about equilibrium points. For example, should a player always play an equilibrium strategy? Will the maximin/minimax strategy always find an equilibrium point if one exists? What should a player do if no equilibrium exists? Although the formal answers to some of these questions are outside the scope of this book, you should be able to make some good conjectures about equilibrium points and rational solutions to two-player zero-sum games.

Exercise 2.6.1 Random 2×2 matrix. Write down a random payoff (zero-sum) matrix with two strategy choices for each player.

Exercise 2.6.2 Random 3×3 matrix. Write down a random payoff (zero-sum) matrix with three strategy choices for each player.

Exercise 2.6.3 Random 4×4 matrix. Write down a random payoff (zero-sum) matrix with four strategy choices for each player.

Exercise 2.6.4 Analyze several examples. Exchange your list of matrices with another student in the class. For each matrix you have been given

- a try to determine any dominated strategies, if they exist.
- b try to determine any equilibrium points, if they exist.
- c determine the maximin and minimax strategies for Player 1 and Player 2, respectively. Can you always find these?

Exercise 2.6.5 Classify examples. Now combine all the examples of payoff matrices in a group of 3 or 4 students. Make a list of the examples with equilibrium points and a list of examples without equilibrium points. If you have only one list, try creating examples for the other list. Based on your lists, do you think random payoff matrices are likely to have equilibrium points?

We want to use lists of matrices as experimental examples to try to answer some of the remaining questions we have about finding rational solutions for

games and equilibrium points. If you don't feel you have enough examples, you are welcome to create more or gather more from your classmates.

Exercise 2.6.6 Playing an equilibrium strategy. If a matrix has an equilibrium point, can a player ever do better to *not* play an equilibrium strategy? Explain.

Exercise 2.6.7 Equilibria and maximin/minimax. If a matrix has an equilibrium point, does the maximin/minimax strategy always find it? Explain.

Exercise 2.6.8 No equilibria and maximin/minimax. If a matrix does NOT have an equilibrium point, should a player always play the maximin/minimax strategy? Explain.

Exercise 2.6.9 Strategy and games with no equilibria. If a matrix does NOT have an equilibrium point is there an ideal strategy for each player? Explain.

Exercise 2.6.10 Summarize the connections. Write a brief summary of the connections you have observed between finding a rational solution for a game and equilibrium points.

Now you should have an understanding of how to find equilibrium strategies in two-player zero-sum games. The main advantage of equilibrium strategies is that if both players play them, neither player would have gained by playing a different strategy. Thus, we can think of the equilibrium strategies as the solution to the game for two rational players. But what should our players do if the game has no equilibrium point? We will look more closely at games with no equilibrium point in the next chapter.

2.7 Popular Culture: Rationality and Perfect Information

In this section, we will look at applications of rationality and perfect information in popular culture. We present films with connections to game theory and suggest some related questions for essays or class discussion.

The movie *Dr. Strangelove or: How I Learned to Stop Worrying and Love the Bomb* (1964) depicts the cold war era with the USA and the Soviet Union on the brink of atomic war.

Question 2.7.1 Society would generally consider General Jack Ripper to be irrational. What is his goal or his optimal payoff? Give evidence that he is, in fact, acting rationally in light of his goal.

Question 2.7.2 Explain how the Soviet's Doomsday Machine is supposed to be the ultimate deterrent to nuclear war. How does the lack of perfect information impact the effectiveness of the Doomsday Machine?

In the movie *The Princess Bride* (1987) the Dread Pirate Roberts and kidnapper Vizzini engage in a battle of wits in which Vizzini is to deduce which goblet contain a lethal poison.

Question 2.7.3 In your own words describe how the poison scene demonstrates that **knowing that both players have the same knowledge** can help one deduce additional information. In other words, just knowing that one player has all the information is not enough; that player, must also know that the other player has the same knowledge.

Question 2.7.4 Of course, in the poison scene, both players break the rules. How does this impact the players' ability to use perfect information?

Now try to apply the ideas of rationality and perfect information to your own popular culture examples.

Question 2.7.5 Consider a game-theoretic scenario in a novel or film of your choice. Are the players rational? What are the players' goals, and are they making choices which will maximize their payoff? Explain. \square

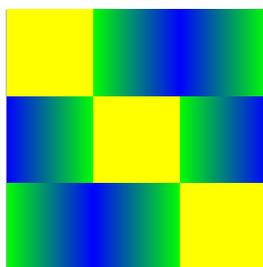
Question 2.7.6 Consider the statement "One of the main differences between horror films and suspense films is that in horror films characters behave irrationally while in suspense films they behave rationally." Do you agree or disagree with this statement? Give an example of a suspense film and a horror film with evidence from the films that supports your position. \square

Question 2.7.7 Think of other films where two characters engage in a "game." What are the assumptions of the players? Do they have perfect information? Does the amount of information a player has give him or her an advantage? Explain. \square

Question 2.7.8 Give an example from a film, current events, or your own life where if one player "breaks the rules," while the other player assumes perfect knowledge (both players know the possible strategies and outcomes), it will change the outcome of the "game." \square

Chapter 3

Repeated Two-Person Zero-sum Games



If we are presented with a two-person zero-sum game we know that our first step is to look for an equilibrium point. If a game has an equilibrium point, then we know that our players should play the corresponding strategy pair. In this case the equilibrium pair and its payoff vector is the “solution” to the game. In this chapter we will explore games that do not necessarily have an equilibrium point. We will also try to determine what a player should do if they play the game repeatedly.

3.1 Introduction to Repeated Games

Now that we are experts at finding equilibrium pairs, what happens when a game doesn't have any equilibrium pairs? What should our players do?

Example 3.1.1 A 2×2 Repeated Game. Consider the following zero-sum game

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

Does this game have an equilibrium pair? Play this game with an opponent 10 times. Tally your wins and losses. Describe how you chose which strategy to play. Describe how your opponent chose which strategy to play. \square

When playing the game several times, does it make sense for either player to play the same strategy all the time? Why or why not?

Although we use the term “strategy” to mean which row (or column) a player chooses to play, we will also refer to how a player plays a repeated game as the player's strategy. In order to avoid confusion, in repeated games we will define some specific strategies.

Definition 3.1.2 In a repeated game, if a player always plays the same row (or column), we say that she is playing a **pure strategy**. \diamond

For example, if Player 1 always plays Row A, we say she is playing *pure strategy A*.

Definition 3.1.3 If a player varies which row (or column) he plays, then we say he is playing a **mixed strategy**. \diamond

For example, if a player plays Row A 40% of the time and Row B 60% of the time, we will say he is playing a (.4, .6) strategy, as we generally use the probability rather than the percent. The probabilities of each strategy will be listed in the same order as the strategies in the matrix.

It is not enough just to determine how often to play a strategy. Suppose Player 1 just alternates rows in [Example 3.1.1](#). Can Player 2 “out-guess” Player 1? What might be a better way for Player 1 to play?

We’d really like to find a way to determine the best mixed strategy for each player in repeated games. Let’s start with what we already know: games with equilibrium points. If a game has an equilibrium pair, would a player want to play a mixed strategy? Recall that a strategy pair is an equilibrium pair if neither player gains by switching strategy.

Example 3.1.4 Repeating a Game with an Equilibrium. Consider the following zero-sum game

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

This game has an equilibrium pair. Convince yourself that if this game is played repeatedly, each player should choose to play a pure strategy. \square

Thus, if the game has an equilibrium we know that players will play the pure strategies determined by the equilibrium pairs. So let’s get back to thinking about games without equilibrium pairs. If we play such a game once, can we predict the outcome? What about if we repeat the game several times— can we predict the outcome? Think about tossing a coin. If you toss it once, can you predict the outcome? What if you toss it 100 times— can you predict the outcome? Not exactly, but we can say what we *expect*: if we toss a coin 100 times we expect to have half of the coins turn up heads and half turn up tails. This may not be the *actual* outcome, but it is a reasonable prediction. Now is a good time to remind yourself about finding the *expected value*!!

Recall the familiar game *Rock-Paper-Scissors*: ROCK beats SCISSORS, SCISSORS beat PAPER, and PAPER beats ROCK. Using the payoff matrix and experimentation, we will try to determine the best strategy for this game.

Exercise 3.1.5 RPS payoff matrix. Construct a game matrix for Rock-Paper-Scissors.

Exercise 3.1.6 RPS and equilibrium points. Is Rock-Paper-Scissors a zero-sum game? Does it have an equilibrium point? Explain.

Exercise 3.1.7 Play RPS. We want to look at what happens if we repeat Rock-Paper-Scissors. Play the game ten times with an opponent. Record the results (list strategy pairs and payoffs for each player).

Exercise 3.1.8 Conjecture a strategy. Describe any strategy you used in [Exercise 3.1.7](#).

Exercise 3.1.9 Strengths and weaknesses of your strategy. Reflect on your chosen strategy. Does it guarantee you a “win”? What should it mean to “win” in a repeated game? What are the strengths and weaknesses of your strategy?

Exercise 3.1.10 Share your strategy. Discuss your strategy with someone else in the class (it can be your opponent). After sharing your ideas for a strategy, can you improve your previous strategy?

Although you may have come up with a good strategy, let's see if we can't decide what the "best" strategy should be for Rock-Paper-Scissors. Let's assume we are playing Rock-Paper-Scissors against the smartest player to ever live. We will call such an opponent the "perfect" player.

Exercise 3.1.11 The weakness of a pure strategy. Explain why it is not a good idea to play a pure strategy; i.e., to play only ROCK, only PAPER, or only SCISSORS.

Exercise 3.1.12 An uneven strategy. Does it make sense to play one option more often than another (for example, ROCK more often than PAPER)? Explain.

Exercise 3.1.13 Frequency of R, P, S. How often should you play each option?

Exercise 3.1.14 Playing a pattern. Do you want to play in a predictable pattern or randomly? What are some advantages and disadvantages of a pattern? What are some advantages and disadvantages of a random strategy?

Hopefully, you concluded that the best strategy against our perfect player would be to play ROCK, PAPER, SCISSORS $1/3$ of the time each, and to play randomly. We can say that our strategy is to play each option randomly with a probability of $1/3$, and call this the $\text{Random}(1/3, 1/3, 1/3)$ strategy.

Exercise 3.1.15 Long-term payoff. Using this "best" strategy, what do you predict the long term payoff will be for Player 1? For Player 2?

Exercise 3.1.16 Testing the $\text{Random}(1/3, 1/3, 1/3)$ strategy. Let's check our prediction. Using a die, let 1 and 2 represent ROCK, 3 and 4 represent PAPER, and 5 and 6 represent SCISSORS. Play the game 20 times with someone in class where each player rolls to determine the choice of ROCK, PAPER, or SCISSORS. Keep track of the strategy pairs and payoffs. What was the total payoff for each player? (At this point, if you still feel that you have a better strategy, try your strategy against the random one— see what happens!)

Exercise 3.1.17 Compare to your prediction. How did the actual outcome compare to your predicted outcome? What do you expect would happen if you play the game 100 times? (Or more?)

Using ideas about probability and expected value we can more precisely answer [Exercise 3.1.17](#).

Exercise 3.1.18 Probabilities when both players play the random strategy. Assume both players are using the $\text{Random}(1/3, 1/3, 1/3)$ strategy. List all the possible outcomes for a single game (recall the outcome is the strategy pair and the payoff, for example [R, P], $(-1, 1)$). What is the probability that any particular pair of strategies will be played? Are the strategy pairs equally likely?

Exercise 3.1.19 Expected value. Using the probabilities and payoffs from [Exercise 3.1.18](#) calculate the expected value of the game for each player.

Exercise 3.1.20 Strategy for the repeated 2×2 game. Now consider the matrix from [Example 3.1.1](#) above:

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

See if you can determine how often Player 1 should play each row, and how often Player 2 should play each column. Try testing your proposed strategy (you may be able to use a variation on the dice as we saw in [Exercise 3.1.16](#)). Write up any conjectured strategies and the results from playing the game with your strategy. Do you think you have come up with the best strategy? Explain.

You may have had an idea about the best way to play Rock-Paper-Scissors before working through this section, but how can we find solutions to other games, such as the one in [Exercise 3.1.20](#)? We don't want to just use a "guess and check" method. Especially since there are infinitely many possible mixed strategies to try! The rest of the chapter will develop mathematical methods for solving repeated games with no equilibrium point.

3.2 Mixed Strategies: Graphical Solution

In this section we will learn a method for finding the maximin solution for a repeated game using a graph.

Let's continue to consider the game given in [Example 3.1.1](#) by

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

In order to make our analysis easier, let's name the row and column strategies as in [Table 3.2.1](#).

Table 3.2.1 Example 3.1.1 with named strategies.

	C	D
A	1	0
B	-1	2

We want to determine how often Player 1 should play A and how often she should play B.

Exercise 3.2.2 Conjecture a strategy. First it is good to test your instinct. Do you think she should play one of the strategies more often than the other? If so, which strategy should she play the most?

What we are really trying to find is the probability with which Player 1 plays A (or B). Since we know that the probabilities sum to one, if we can find one probability, then we know the other.

Here is one way to do this. Let p be the probability that Player 1 plays B. Let m be the payoff to Player 1. Since we are trying to find a mixed strategy for Player 1, we will pick a strategy for Player 2 and try to determine the possible payoffs for Player 1.

Let us determine some pairs (p, m) .

Step 1. Assume *Player 2* plays pure strategy *C*.

Step 1a. Find the probability (p) and payoff (m) if Player 1 always plays A. If Player 1 plays pure strategy A, then she never plays B. Thus the probability she plays B is 0. Hence, $p = 0$. In the case where Player 1 plays A and Player 2 plays C, what is the payoff to Player 1? This is m , so $m = 1$. Thus, for the strategy pair $[A, C]$ we get $(0, 1)$ for (p, m) . It is important to note that $(0, 1)$ is *not* a payoff vector. This is common notation for any ordered pair. With payoff vectors, the ordered pair represents the payoff to each player. Here the ordered pair represents a probability of playing B and the payoff to Player 1.

Step 1b. Find the probability (p) and payoff (m) if Player 1 always plays B . If Player 1 plays pure strategy B , then what is the probability that she plays B ? Since she always plays B , $p = 1$. In the case where Player 1 plays B and Player 2 plays C , what is the payoff to Player 1? $m = -1$. Thus, for the strategy pair $[B, C]$ we get $(1, -1)$ for (p, m) .

Step 1c. Now we want to know what Player 1's payoff will be as she varies the probability, p , with which she plays B . We can draw a graph where the x -axis represents to probability with which she plays B (p) and the y -axis represents the expected payoff (m). See [Figure 3.2.3](#).

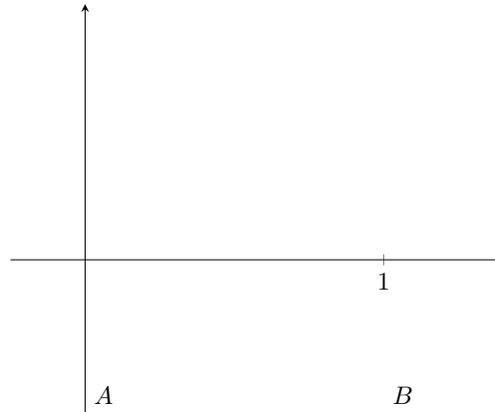


Figure 3.2.3 Labeled axes.

Thus, when Player 1 plays only A , she is playing B with probability 0; when Player 1 plays only B , she is playing B with probability 1. It might be easier to remember if you label your graph as in [Figure 3.2.3](#).

Step 1d. Now we can plot the points we determined in Step 1a and Step 1b. We will connect them with a line representing Player 2's pure strategy C . See [Figure 3.2.4](#).

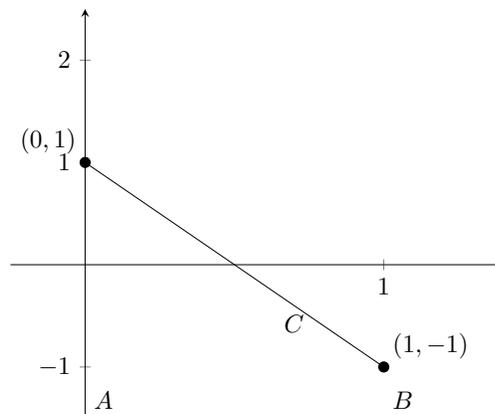


Figure 3.2.4 Player 2's strategy C .

Before moving on, let's make sure we understand what this line represents. Any point on it represents the expected payoff to Player 1 as she varies her strategy, *assuming Player 2 only plays C*. In this case, we can see that as she plays B more often, her expected payoff goes down.

Now let's do the same thing, assuming Player 2 plays only D .

Step 2. Assume *Player 2* plays pure strategy *D*.

Step 2a. Find the probability (p) and payoff (m) if *Player 1* always plays *A*. If *Player 1* plays pure strategy *A*, then what is the probability that she plays *B*? $p = 0$. What is the payoff to *Player 1*? $m = 0$. Thus, for the strategy pair $[A, D]$ we get $(0, 0)$ for (p, m) .

Step 2b. Find the probability (p) and payoff (m) if *Player 1* always plays *B*. If *Player 1* plays pure strategy *B*, then what is the probability that she plays *B*? $p = 1$. What is the payoff to *Player 1*? $m = 2$. Thus, for the strategy pair $[B, D]$ we get $(1, 2)$ for (p, m) .

Step 2c. Now, on our same graph from Step 1, we can plot the points we determined in Step 2a and Step 2b. We will connect them with a line representing *Player 2*'s pure strategy *D*. See [Figure 3.2.5](#).

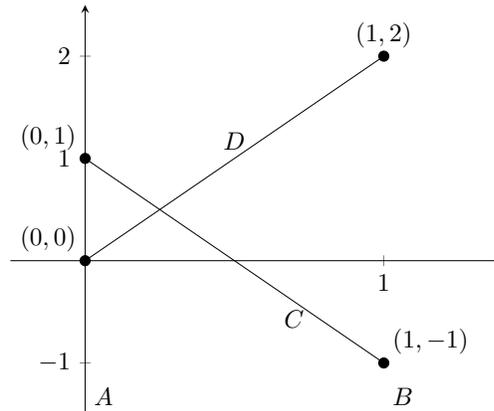


Figure 3.2.5 *Player 2*'s strategy *D*.

Now we can see that if *Player 2* plays only *D*, then *Player 1* does best by playing only *B*.

So we have this nice graph, but what does it really tell us? Although we drew lines representing each of *Player 2*'s pure strategies, *Player 1* doesn't know what *Player 2* will do. Suppose *Player 1* only played *A*, while *Player 2* plays an unknown mixed strategy. Then the possible payoffs for *Player 1* are 1 or 0. The more often *Player 2* plays *C*, the more often *Player 1* gets 1. So the *expected payoff* per game for a repeated game varies between 0 and 1. We can see the possible expected values as the red line on the graph in [Figure 3.2.6](#).

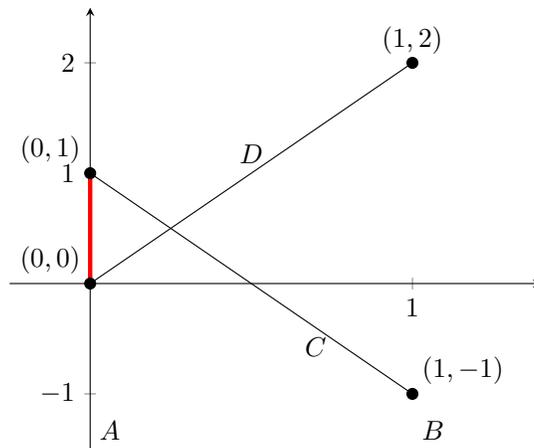


Figure 3.2.6 Figure 3.2.6 of the expected payoffs for Player 1 playing only A.

Since we want to understand mixed strategies for Player 1, what would happen if Player 1 played A half the time and B half the time? In other words, what happens if $p = 1/2$? Although we may not easily be able to see the exact values, we can represent the possible expected values on the graph in Figure 3.2.7.

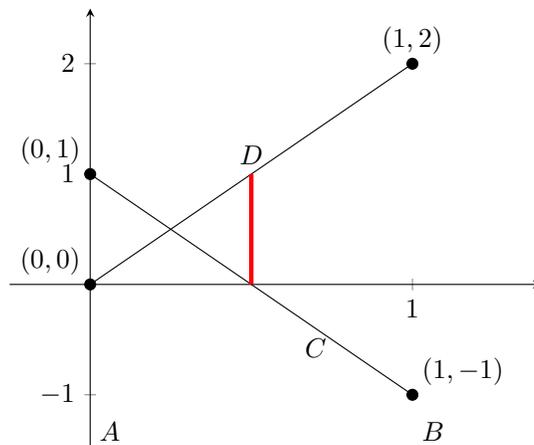


Figure 3.2.7 The expected payoffs for Player 1 playing B half the time.

Hopefully, you've begun to see that for each choice of p , the top line represents the highest expected value for Player 1; the bottom line represents the lowest expected value for Player 1; the area between the lines represents the possible expected values for Player 1. As we did with non-repeated games, let's look at the "worst case scenario" for Player 1. In other words, let's assume that Player 2 can figure out Player 1's strategy. Then Player 1 would want to *maximize the minimum expected value*. Aha! This is just looking for the *maximin* strategy!

Now the minimum expected value for each choice of p is given by the bottom lines on the graph, shown in red in Figure 3.2.8.

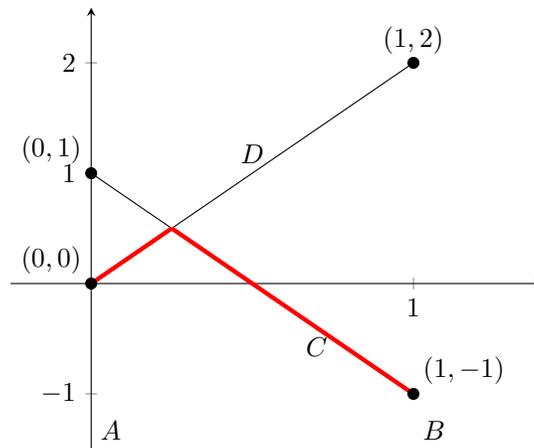


Figure 3.2.8 The minimum expected payoffs for Player 1.

It should be easy to see that the *maximum* of the minimum expected payoffs occurs at the intersection of the two lines.

Step 3. Find the intersection of the two lines.

Step 3a. Find the equation for Line C. This is the line passing through the points $(0, 1)$ and $(1, -1)$. It has slope -2 and y -intercept 1 . Thus, it has equation $y = -2x + 1$. [Although the x -axis represents probability p and the y -axis represents expected payoff m , you are probably more comfortable solving equations—at least for the moment—in x and y .]

Step 3b. Find the equation for Line D. This is the line passing through the points $(0, 0)$ and $(1, 2)$. It has slope 2 and y -intercept 0 . Thus, it has equation $y = 2x$.

Step 3c. Use substitution to find the point of intersection.

$$\begin{aligned} 2x &= & -2x + 1 \\ 4x &= & 1 \\ x &= & \frac{1}{4} \end{aligned}$$

Substituting $x = \frac{1}{4}$ back in to either original equation, say $y = 2x$, gives us $y = \frac{1}{2}$. Thus, the point of intersection is $(1/4, 1/2)$.

Step 4. Determine Player 1's maximin mixed strategy. Recalling that the first coordinate is p , the probability that Player 1 plays B , we know that Player 1 will play B with probability $1/4$, and thus, play A with probability $3/4$ [$1 - 1/4 = 3/4$]. The expected payoff for Player 1 is $1/2$. It is important to check back to your original intuition about the game from [Exercise 3.2.2](#). Did it seem as though Player 1 should play A more often than B ?

Let's make a few important observations. First, it should be clear from the graph that Player 1 expects a payoff of $1/2$ NO MATTER WHAT PLAYER 2 DOES. Furthermore, since this is a zero-sum game, we know that Player 2's expected payoff is $-1/2$. It is important to note that this graph does not give us any information about an optimal strategy for Player 2. We will see how to find a strategy for Player 2 in the following exercises. Can you think of how you might do this?

We can use the graphical method to find the maximin and minimax mixed strategies for repeated two-person zero-sum games.

Using the same game matrix as above:

$$\begin{bmatrix} 1 & 0 \\ -1 & 2 \end{bmatrix},$$

we will continue to label Player 1's strategies by A and B , and Player 2's strategies by C and D . We now want to determine the minimax strategy for Player 2. Keep in mind the payoffs are still the payoffs to Player 1, so Player 2 wants the payoff to be as small as possible.

Exercise 3.2.9 Sketch the graph. Sketch the graph for Player 1 that we drew above. Be sure to label the endpoints of each line. Also label each line according to which strategy they represent.

Exercise 3.2.10 Indicate the minimax strategy. Describe the minimax strategy and show it on the graph. (You do not need to find the actual mixed strategy for Player 2.)

Exercise 3.2.11 Payoff vectors. Are the payoff vectors for the maximin and minimax strategies the same?

For non-repeated games we have seen that if the maximin value is the same as the minimax value, then the game has a pure strategy equilibrium. The same idea applies to mixed strategy games. If the value of the maximin strategy is the same as the value of the minimax strategy, then the corresponding mixed strategies will be an equilibrium point. Thus, your answer to [Exercise 3.2.11](#) should tell you this game has a mixed strategy equilibrium point consisting of the maximin/ minimax strategy.

We now know that Player 2 wants to play the minimax strategy in response to Player 1's maximin strategy, so we need to find the actual mixed strategy for Player 2 to employ. Since we are minimizing Player 1's maximum expected payoff, we will continue to use the matrix representing Player 1's payoff. We will repeat the process we used for Player 1, except the x -axis now represents the probability that Player 2 will play D , and the lines will represent Player 1's strategies A and B . The y -axis continues to represent Player 1's payoff.

Exercise 3.2.12 Draw the axes for Player 2's strategy. First sketch the axes. (Recall, the x -axis only goes from 0 to 1.)

Exercise 3.2.13 Player 1 plays A . Assume Player 1 only plays A .

- If Player 2 only plays C , what is the payoff to Player 1? Recall we called this m . What is the probability that Player 2 plays D ? Recall we called this p . On your graph, plot the point (p, m) .
- If Player 2 plays only D , find m and p . Plot (p, m) on the graph.
- Now sketch the line through your two points. This line represents Player 1's pure strategy A and the expected payoff (to Player 1) for Player 2's mixed strategies. Label it A .

Exercise 3.2.14 Player 1 plays B . Now assume Player 1 plays only B . Repeat the steps in [Exercise 3.2.13](#), using B instead of A , to find the line representing Player 1's pure strategy B . (Label it!)

Exercise 3.2.15 The graph for Player 2. It is important to keep in mind that although the x -axis refers to how often Player 2 will play C and D , the y -axis represents the payoff to *Player 1*.

- Explain why we are looking for the *minimax* strategy for Player 2.

- b Show on the graph the *maximum* payoff that Player 1 can expect for each of Player 2's possible mixed strategies.
- c Show the point on the graph that represents the minimax strategy.

Exercise 3.2.16 Equations for the lines. Find the equations of the lines you drew in [Exercise 3.2.13](#) and [Exercise 3.2.14](#).

Exercise 3.2.17 The point of intersection. Using the equations from [Exercise 3.2.16](#), find the point of intersection of the two lines.

Exercise 3.2.18 Player 2's mixed strategy. How often should Player 2 play *C*? How often should he play *D*? What is Player 1's expected payoff? And hence, what is Player 2's expected payoff?

Exercise 3.2.19 Equilibrium strategies. Explain why each player should play the maximin/ minimax mixed strategy. In other words, explain why neither player benefits by changing his or her strategy.

Hint. Think about playing defensively and assuming the other player is the "perfect" player.

Now it may have occurred to you that since this is a zero-sum game, we could have just converted our matrix to the payoff matrix for Player 2, and found Player 2's maximin strategy. But it is important to understand the relationship between the maximin and the minimax strategies. So for the sake of practice and a little more insight, find Player 2's maximin strategy by writing the payoff matrix for Player 2 and repeating the process that we did for Player 1. Keep in mind that Player 2 is finding the probability of playing *C* and *D* rather than *A* and *B*.

Exercise 3.2.20 Finding the maximin using Player 2's payoffs. Convert the payoff matrix above into the payoff matrix for Player 2. Find the maximin strategy for Player 2 using the graphical method. Be sure to include a sketch of the graph (labeled!!), the equations for the lines, the probability that Player 2 will play *C* and *D*, and the expected payoff for Player 2.

Exercise 3.2.21 Compare the solutions. Compare your answer in [Exercise 3.2.20](#) to your answer in [Exercise 3.2.18](#).

Exercise 3.2.22 Fairness. Is this game fair? Explain.

Exercise 3.2.23 Expected payoff. We saw above that the expected payoff for Player 1 was $1/2$. Explain what this means for a repeated game.

Hint. Is it actually possible for a player to win $1/2$ in a given game?

Now you are ready to try to analyze some more games!

Exercise 3.2.24 Practice finding mixed strategies. Consider the zero-sum game $\begin{bmatrix} -4 & 4 \\ 2 & -2 \end{bmatrix}$.

- a Does this game have a pure strategy equilibrium?
- b Just by looking at the matrix, do you think this game will be fair? (Would you rather be Player 1 or Player 2?)
- c Sketch (and label!) the appropriate graph for this game.
- d Use your graph to determine if there is a mixed strategy equilibrium point. If there is, how often should Player 1 play each strategy? What is the expected payoff to each player?
- e Is this game fair? Explain. Compare your answer to (b).

Exercise 3.2.25 More practice finding mixed strategies. Consider the zero-sum game $\begin{bmatrix} 0 & 1 \\ 1 & -10 \end{bmatrix}$.

- Does this game have a pure strategy equilibrium?
- Just by looking at the matrix, do you think this game will be fair? (Would you rather be Player 1 or Player 2?)
- Sketch (and label!) the appropriate graph for this game.
- Use your graph to determine if there is a mixed strategy equilibrium point. If there is, how often should Player 1 play each strategy? What is the expected payoff to each player?
- Is this game fair? Explain. Compare your answer to (b).

Although it is worth working through examples by hand in order to understand the algebraic process, in the next section we will see how technology can help us solve systems of equations.

3.3 Using Sage to Graph Lines and Solve Equations

In this section we will use technology to graph lines and solve for the intersection point. In particular, we will use an open online resource called *Sage*.

Let's continue to consider the game given in [Table 3.2.1](#) by

Table 3.3.1 Game from Example 3.1.1.

	C	D
A	1	0
B	-1	2

Recall, our goal is to determine how often Player 1 should play A and how often she should play B.

We will follow the same steps as in [Section 3.2](#). Let p be the probability that Player 1 plays B. Let m be the payoff to Player 1. Since we are trying to find a mixed strategy for Player 1, we will pick a strategy for Player 2 and try to determine the possible payoffs for Player 1.

Let us determine some pairs (p, m) .

Step 1. Assume *Player 2* plays pure strategy *C*.

Step 1a. If Player 1 always plays *A*, then we are considering the strategy pair $[A, C]$. Since Player 1 never plays *B*, $p = 0$. The payoff to Player 1 for $[A, C]$ is $m = 1$. Thus, for the strategy pair $[A, C]$ we get $(p, m) = (0, 1)$.

Step 1b. If Player 1 always plays *B*, then we are considering the strategy pair $[B, C]$. Since Player 1 always plays *B*, $p = 1$. The payoff to Player 1 for $[B, C]$ is $m = -1$. Thus, for the strategy pair $[B, C]$ we get $(p, m) = (1, -1)$.

Step 1c. Now we want to know what Player 1's payoff will be as she varies the probability, p , with which she plays *B*. We can draw a graph where the x -axis represents to probability with which she plays B (p) and the y -axis represents the expected payoff (m). Thus, when Player 1 plays only *A*, she is playing *B* with probability 0; when Player 1 plays only B,

she is playing B with probability 1. It might be easier to remember if you label your graph as in [Figure 3.2.3](#).

Step 1d. Now we can use *Sage* to plot the points we determined in Step 1a and Step 1b and the line between them. This line represents Player 2's pure strategy *C*. See [Figure 3.2.4](#). Click on the "Evaluate (*Sage*)" button to plot the line between the points $(0, 1)$ and $(1, -1)$.

```
u=(0,1);
v=(1,-1);
show(line([u, v], thickness=2, color='blue')+point(u,
    color='blue', pointsize=70)+point(v,
    color='blue', pointsize=70), figsize=4)
```

Before moving on, let's again, make sure we understand what this line represents. Any point on it represents the expected payoff to Player 1 as she varies her strategy, *assuming Player 2 only plays C*. In this case, we can see that as she plays *B* more often, her expected payoff goes down. You can now use this *Sage* cell to plot any line for Player 2's pure strategy *C*. Just edit the values for the points *u* and *v*. Go ahead and try it! (Don't worry the original values will reset when you refresh the page.)

Now let's do the same thing, assuming Player 2 plays only *D*.

Step 2. Assume *Player 2* plays pure strategy *D*.

Step 2a. If Player 1 always plays *A*, then we are considering the strategy pair $[A, D]$. Since Player 1 never plays *B*, $p = 0$. The payoff to Player 1 for $[A, D]$ is $m = 0$. Thus, for the strategy pair $[A, D]$ we get $(p, m) = (0, 0)$.

Step 2b. If Player 1 always plays *B*, then we are considering the strategy pair $[B, D]$. Since Player 1 always plays *B*, $p = 1$. The payoff to Player 1 for $[B, D]$ is $m = 2$. Thus, for the strategy pair $[B, D]$ we get $(p, m) = (1, 2)$.

Step 2c. Now, on our same graph from Step 1, we can plot the points we determined in Step 2a and Step 2b. We will connect them with a line representing Player 2's pure strategy *D*. See [Figure 3.2.5](#).

```
u=(0,1);
v=(1,-1);
w=(0,0);
z=(1, 2);
show(line([u, v], thickness=2, color='blue')+point(u,
    color='blue', pointsize=70)
+point(v, color='blue', pointsize=70)
+line([w, z], thickness=2, color='red'))
+point(w, color='red', pointsize=70)+point(z,
    color='red', pointsize=70), figsize=4)
```

Now we can see that if Player 2 plays only *D*, then Player 1 does best by playing only *B*. Again, you can use this *Sage* cell to plot both Player 2's pure strategies. Points *u* and *v* are for strategy *C*, while points *w* and *z* are for strategy *D*.

As we saw in [Section 3.2](#), for each choice of p , the top line represents the highest expected value for Player 1; the bottom line represents the lowest expected value for Player 1; the area between the lines represents the possible

expected values for Player 1. Thus, Player 1 wants to maximize the minimum expected value, which means she wants to find the maximin strategy. And, as we saw in [Section 3.2](#), the maximin strategy occurs at the intersection of the two lines.

Step 3. Find the intersection of the two lines.

Step 3a. Find the equation for Line C. This is the line passing through the points $(0, 1)$ and $(1, -1)$. It has slope -2 and y -intercept 1 . Thus, it has equation $m = -2p + 1$. (Recall the x -axis represents probability p and the y -axis represents expected payoff m .)

Step 3b. Find the equation for Line D. This is the line passing through the points $(0, 0)$ and $(1, 2)$. It has slope 2 and y -intercept 0 . Thus, it has equation $m = 2p$.

Step 3c. Now we can use technology to help us solve the equations.

```
p, m = var('p, m')
solve([m == -2*p + 1, m == 2*p], p, m)[0]
```

```
[p == (1/4), m == (1/2)]
```

The solution for Player 1 is (p, m) . Where p is the probability Player 1 plays B, and m is the expected payoff to Player 1.

We can use this Sage cell to solve for p and m for any 2×2 game by editing the equations $m == -2 * p + 1, m == 2 * p$.

Step 4. Determine Player 1's maximin mixed strategy. Recalling that p is the probability that Player 1 plays B, we know that Player 1 will play B with probability $1/4$, and thus, play A with probability $3/4$. The expected payoff for Player 1, m , is $1/2$. It is important to check the algebraic solution with where the intersection point appears on the graph. Although we are using technology to help us graph and solve for the intersection point, we need to be able to catch any errors we make entering the information into Sage.

We have seen that we can use the same matrix with Player 1's payoffs to find the strategy for Player 2. Using the same game matrix as above:

$$\begin{bmatrix} 1 & 0 \\ -1 & 2 \end{bmatrix},$$

and continuing to label Player 1's strategies by A and B , and Player 2's strategies by C and D , we can graph lines for Player 1's pure strategies A and B . We now let the x -axis represent the probability that Player 2 plays D . In the Sage applet below, for u and v enter the coordinates of two points that determine the line for when Player 1 plays A , then the two points for w and z that determine the line for when Player 1 plays B . We will then have Sage graph the lines. You can enter new values for u, v, w, z if you would like to draw the graph for a different matrix.

```
@interact(layout=dict(top=[[ 'u', 'v'], [ 'w', 'z' ]]))
def endpoints(u=vector((0, 1.0)), v=vector((1, 0.0)),
             w=vector((0, -1.0)), z=vector((1, 2.0))):
    L1 = line([u, v], thickness=2, color=('blue'))
    L2 = line([w, z], thickness=2, color=('red'))
```

```

P1 = point(u, color=('blue'), pointsize=70)
P2 = point(v, color=('blue'), pointsize=70)
P3 = point(w, color=('red'), pointsize=70)
P4 = point(z, color=('red'), pointsize=70)
pretty_print(html("Enter the coordinates of the endpoints
for the two lines you'd like to graph. Note that u and
v are for one line, w and z for the other."))
show(L1+L2+P1+P2+P3+P4, figsize=4)

```

Now determine and enter the equations of the two lines and have Sage solve for the intersection point.

```

p, m = var('p, m')
@interact
def
    intersection(Slope1=-1, Intercept1=1, Slope2=3, Intercept2=-1):
    Eq1 = m==Slope1*p+Intercept1
    Eq2 = m==Slope2*p+Intercept2
    S = solve([Eq1, Eq2], p, m)[0]
    pretty_print(html("The intersection point is %s$,
    %s$. "%(latex(S[0]), latex(S[1]))))

```

You can now use these last two Sage cells to solve any 2×2 game with a mixed strategy equilibrium. You can also take some time to experiment with what happens if the game has a pure strategy equilibrium.

Now use the Sage cells to help you analyze some more games!

Exercise 3.3.2 Use Sage to find a mixed strategy equilibrium. Con-

sider the zero-sum game $\begin{bmatrix} -4 & 4 \\ 2 & -2 \end{bmatrix}$.

- Does this game have a pure strategy equilibrium?
- Just by looking at the matrix, do you think this game will be fair? (Would you rather be Player 1 or Player 2?)
- Use Sage to draw the appropriate graph for this game. Which pure strategy does the blue line represent? Which pure strategy does the red line represent?
- Use your graph to determine if there is a mixed strategy equilibrium point. If there is, use Sage to determine how often Player 1 should play each strategy. What is the expected payoff to each player?
- Is this game fair? Explain. Compare your answer to (b).

Exercise 3.3.3 More practice with Sage. Consider the zero-sum game

$\begin{bmatrix} 0 & 1 \\ 1 & -10 \end{bmatrix}$.

- Does this game have a pure strategy equilibrium?
- Just by looking at the matrix, do you think this game will be fair? (Would you rather be Player 1 or Player 2?)
- Use Sage to draw the appropriate graph for this game. Which pure strategy does the blue line represent? Which pure strategy does the red line represent?
- Use your graph to determine if there is a mixed strategy equilibrium

point. If there is, use Sage to determine how often Player 1 should play each strategy. What is the expected payoff to each player?

e Is this game fair? Explain. Compare your answer to (b).

3.4 Mixed Strategies: Expected Value Solution

In this section, we will use the idea of expected value to find the equilibrium mixed strategies for repeated two-person zero-sum games.

One of the significant drawbacks of the graphical solution from the previous sections is that it can only solve 2 X 2 matrix games. If each player has 3 options, we would need to graph in three dimensions. Technically this is possible, but rather complicated. If each player has more than 3 options, since we can't graph in four or more dimensions, we are at a complete loss. So we need to think about an alternate way to solve for the mixed strategies. Although we will begin with 2 X 2 games, this method will easily generalize to larger games.

Example 3.4.1 Matching Pennies Game. Consider the game in which each player can choose HEADS (H) or TAILS (T); if the two players match, Player 1 wins; if the two players differ, Player 2 wins. What strategy should each player play? \square

Exercise 3.4.2 Payoff matrix. Determine the payoff matrix for the Matching Pennies game.

Exercise 3.4.3 Pure strategy equilibria. Explain why the Matching Pennies game has no pure strategy equilibrium point.

Exercise 3.4.4 Conjecture a mixed strategy. Since we know that there is no pure strategy equilibrium point, we need to look for a mixed strategy equilibrium point. Just by looking at the payoff matrix for Matching Pennies, what do you think an ideal strategy for each player would be? Explain your choice.

Exercise 3.4.5 Expected value of conjecture. Suppose both players play your ideal strategy in the Matching Pennies game, what should the expected value of the game be?

We could use our previous graphical method to determine the expected value of the game (you might quickly try this just to verify your prediction). However, as we have noted, a major drawback of the graphical solution is that if our players have 3 (or more) options, then we would need to graph an equation in 3 (or more!) variables; which, I hope you agree, we don't want to do. Although we will continue to focus on 2×2 games, we will develop a new method which can more easily be used to solve to larger games.

We will need a little notation. Let

$$\begin{aligned} P_1(H) &= \text{the probability that Player 1 plays H;} \\ P_1(T) &= \text{the probability that Player 1 plays T;} \\ P_2(H) &= \text{the probability that Player 2 plays H;} \\ P_2(T) &= \text{the probability that Player 2 plays T.} \end{aligned}$$

Also, we will let $E_1(H)$ be the expected value for Player 1 playing pure strategy H against a given strategy for Player 2. Similarly, $E_2(H)$ will be Player 2's expected value for playing pure strategy H.

Exercise 3.4.6 The (60, 40) strategy for Player 2. Suppose Player 2 plays H 60% of the time and T 40% of the time.

- What are $P_2(H)$ and $P_2(T)$?
- What do you think Player 1 should do? Does this differ from your ideal mixed strategy in [Exercise 3.4.4](#)? Explain.
- We can use expected value to compute what Player 1 should do in response to Player 2's 60/40 strategy. First, consider a pure strategy for Player 1. Compute the expected value for Player 1 if she only plays H while Player 2 plays H with probability .6 and T with probability .4. This expected value is $E_1(H)$, above.
- Similarly, compute the expected value for Player 1 if she plays only T (call it $E_1(T)$).
- Which pure strategy has a higher expected value for Player 1? If Player 1 plays this pure strategy, will she do better than your predicted value of the game?

Exercise 3.4.7 The (60, 40) strategy is not ideal for Player 2. Hopefully, you concluded that in [Exercise 3.4.6](#) a pure strategy is good for Player 1. Explain why this means the 60/40 strategy is bad for Player 2.

Exercise 3.4.8 When to play H. For any particular mixed (or pure) strategy of Player 2, we could find $E_1(T)$ and $E_1(H)$. Explain why if $E_1(H) > E_1(T)$, Player 1 should always play H.

Exercise 3.4.9 When to play T. Similarly, explain why if $E_1(H) < E_1(T)$, Player 1 should always play T.

Exercise 3.4.10 Player 2 is unhappy when Player 1's expected values are unequal. Explain why the situations in [Exercise 3.4.8](#) and [Exercise 3.4.9](#) are bad for Player 2.

Exercise 3.4.11 Equal expected values are better. Use your answers from [Exercise 3.4.8](#), [Exercise 3.4.9](#), and [Exercise 3.4.10](#) to explain why the situation in which $E_1(H) = E_1(T)$ is the best for Player 2.

From [Exercise 3.4.11](#) we now know that Player 2 wants $E_1(H) = E_1(T)$, we can use a little algebra to compute the best defensive strategy for Player 2. Remember, we want to assume that Player 1 will always be able to choose the strategy that will be best for her, and thus Player 2 wants to protect himself. Let's find the probabilities with which Player 2 should play H and T.

Exercise 3.4.12 Equations for Player 1's expected values. Let $P_2(H)$ and $P_2(T)$ be the probabilities that Player 2 plays H and T respectively. Find equations for $E_1(H)$ and $E_1(T)$ in terms of $P_2(H)$ and $P_2(T)$ for the game of Matching Pennies. Since we want $E_1(H) = E_1(T)$, set your two equations equal to each other. This gives you one equation in terms of $P_2(H)$ and $P_2(T)$.

Exercise 3.4.13 The sum equation. Explain why we must also have $P_2(H) + P_2(T) = 1$.

So in general, to solve for Player 2's strategy, we want to write an equation for each of Player 1's pure strategy expected values in terms of Player 2's probabilities. For example, $E_1(H)$ and $E_1(T)$ in terms of variables $P_2(H)$ and $P_2(T)$. We then set the expected values equal to each other. We now have an equation just in terms of Player 2's probabilities.

In order to solve for the probabilities, we also need to use the fact that Player 2's probabilities sum to 1. For example, $P_2(H) + P_2(T) = 1$. For a 2×2

game, you now have 2 equations with 2 unknowns ($P_2(H)$ and $P_2(T)$). Use an algebraic method such as substitution or elimination to solve the system of equations.

Exercise 3.4.14 Solve for Player 2's probabilities. Using the equations from [Exercise 3.4.12](#) and [Exercise 3.4.13](#), solve for $P_2(H)$ and $P_2(T)$. You now have the equilibrium mixed strategy for Player 2. Does this match the mixed strategy you determined in [Exercise 3.4.4](#)?

Now can you use a similar process to find Player 1's strategy? Whose expected values should you use? Whose probabilities?

Exercise 3.4.15 Find Player 1's probabilities. Set up and solve the analogous equations from [Exercise 3.4.12](#) and [Exercise 3.4.13](#) for Player 1. Does this match the mixed strategy from [Exercise 3.4.4](#)?

Hint. We should have an equation for $E_2(H)$ and one for $E_2(T)$. Since we are looking for the probabilities of each of Player 1's options, the equations should involve $P_1(H)$ and $P_1(T)$.

We now have a new method for finding the best mixed strategies for Players 1 and 2, assuming that each player is playing defensively against an ideal player. But how can we find the value of the game? For Player 2, we assumed $E_1(H) = E_1(T)$. In other words, we found the situation in which Player 1's expected value is the same no matter which pure strategy she plays. Thus, Player 1 is *indifferent* to which pure strategy she plays. If she were not indifferent, then she would play the strategy with a higher expected value. But, as we saw, this would be bad for Player 2. So Player 1 can assume that Player 2 will play the equilibrium mixed strategy. Thus, as long as Player 1 plays a mixed strategy, it doesn't matter whether at any given time, she plays H or T (this is the idea that she is indifferent to them). Hence, the expected value of the game for Player 1 is the same as $E_1(H)$, which is the same as $E_1(T)$. Similarly, we find that the expected value of the game for Player 2 is $E_2(H)$ (or $E_2(T)$). This is a pretty complicated idea. You may need to think about it for a while. In the meantime, use the probabilities you found for each player and the equations for $E_1(H)$ and $E_2(H)$ to find the value of the game for each player.

Exercise 3.4.16 Find Player 1's expected value of the game. Use the probabilities you calculated in [Exercise 3.4.14](#) to find $E_1(H)$, and hence the expected value of the game for Player 1. How does this compare to your prediction for the value of the game that you gave in [Exercise 3.4.5](#)?

Exercise 3.4.17 Find Player 2's expected value of the game. Use the probabilities you calculated in [Exercise 3.4.15](#) to find $E_2(H)$, and hence the expected value of the game for Player 2. How does this compare to your prediction for the value of the game that you gave in [Exercise 3.4.5](#)?

Keep practicing with the expected value method on some other games!

Exercise 3.4.18 Solve a 2×2 repeated game using expected values. Apply this method of using expected value to [Example 3.1.1](#) (which we solved using the graphical method in the previous section) to find the equilibrium mixed strategies for each player and the value of the game for each player:

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

Exercise 3.4.19 Expected value solution for Rock-Paper-Scissors. As we noted in this section, this method can be used to solve bigger games. We just have more equations. Use the expected value method to find the equilibrium

mixed strategy for Rock-Paper-Scissors for Player 2.

Hint. You will need to set $E_1(R) = E_1(P)$ and $E_1(P) = E_1(S)$; solve for $P_2(R), P_2(P), P_2(S)$; etc. It should be very similar to what you did for Matching Pennies, but there will be three equations and three unknowns.

If you found this last exercise to be algebraically arduous, don't worry, we will eventually use technology to help us solve problems with several variables.

3.5 Liar's Poker

In this section, we will continue to explore the ideas of a mixed strategy equilibrium. We have seen several examples of finding an equilibrium. We began with games which had a pure strategy equilibrium and then moved to games with mixed strategy equilibrium. We saw two different methods for finding an equilibrium. The first employed graphs in order to understand and find the maximin and minimax strategies, and hence the equilibrium mixed strategy. The second method employed the ideas of expected value to find the equilibrium strategy. We will continue to solidify these ideas with another game, a simplified variation of poker.

Example 3.5.1 Liar's Poker. We begin with a deck of cards which has 50% aces (A) and 50% kings (K). Aces rank higher than kings. Player 1 is dealt one card, face down. Player 1 can look at the card, but does not show the card to Player 2. Player 1 then says "ace" or "king" depending on what his card is. Player 1 can either tell the truth and say what the card is (T), or he can bluff and say that he has a higher ranking card (B). Note that if Player 1 has an ace, he must tell the truth since there are no higher ranking cards. However, if he is dealt a king, he can bluff, by saying he has an ace. If Player 1 says "king" the game ends and both players break even. If Player 1 says "ace" then Player 2 can either call (C) or fold (F). If Player 2 folds, then Player 1 wins \$0.50. If Player 2 calls and Player 1 does not have an ace, then Player 2 wins \$1. If Player 2 calls and Player 1 does have an ace, then Player 1 wins \$1. \square

Exercise 3.5.2 Play Liar's Poker. Choose an opponent and play Liar's Poker several times. Be sure to play the game as Player 1 and as Player 2. This is important for understanding the game. Keep track of the outcomes.

Exercise 3.5.3 Conjecture a strategy. Just from playing Liar's Poker several times, can you suggest a strategy for Player 1? What about for Player 2? Does this game seem fair, or does one of the players seem to have an advantage? Explain your answers.

Exercise 3.5.4 Try to find the payoff matrix. In order to formally analyze Liar's Poker, we should find the payoff matrix. Do your best to find the payoff matrix. In a single hand of Liar's Poker, what are the possible strategies for Player 1? What are the possible strategies for Player 2? Determine any payoffs that you can.

Finding the payoff matrix in [Exercise 3.5.4](#) is probably more challenging than it appears. Eventually we want to employ the same method for finding the payoff matrix that we used in One-Card Stud Poker from [Example 2.4.1](#) in Chapter 2, but first we need to understand each player's strategies, and the resulting payoffs. We begin with the fact that Player 1 can be dealt an ace or a king.

Exercise 3.5.5 Player 1 has an Ace. Assume Player 1 is dealt an ace. What can Player 1 do? What can Player 2 do? What is the payoff for each situation?

Exercise 3.5.6 Player 1 has a King. Assume Player 1 is dealt a king. What can Player 1 do? What can Player 2 do? What is the payoff for each situation?

Since Player 1 must say “ace” when dealt an ace, he only has a choice of strategy when dealt a king. So we can define his strategy independent of the deal. One strategy is to say “ace” when dealt an ace and say “ace” when dealt a king; call this the **bluffing strategy** (B). The other strategy is to say “ace” when dealt an ace and say “king” when dealt a king; call this the **truth strategy** (T). Recall that the only time Player 2 has a choice is when Player 1 says “ace.” In this case Player 2 can *call* (C) or *fold* (F). Since we need to determine the payoff matrix, we first need to determine the payoffs for pure strategies. This is similar to what we did for the One-Card Stud game.

Exercise 3.5.7 Expected value of [B, C]. Consider Player 1’s pure strategy of always bluffing when dealt a king (B) and Player 2’s pure strategy of always calling (C). Determine the expected value for Player 1. What is Player 2’s expected value?

Hint. You need to consider each possible deal.

Exercise 3.5.8 Expected value of [B, F]. Similarly, determine the expected value for Player 1 for the pure strategy pair [B, F]. What is Player 2’s expected value?

Exercise 3.5.9 Expected value of [T, C]. Determine the expected value for Player 1 for the pure strategy pair [T, C]. What is Player 2’s expected value?

Exercise 3.5.10 Expected value of [T, F]. Determine the expected value for Player 1 for the pure strategy pair [T, F]. What is Player 2’s expected value?

Exercise 3.5.11 Payoff matrix for Liar’s Poker. Using the expected values you calculated in [Exercise 3.5.7](#), [Exercise 3.5.8](#), [Exercise 3.5.9](#), and [Exercise 3.5.10](#), set up the 2×2 payoff matrix for Liar’s Poker.

Once you have determined the payoff matrix for Liar’s Poker, you can use either the graphical method or expected value method to solve the game. But before using either of these methods always check for a pure strategy equilibrium!

Exercise 3.5.12 Pure strategy equilibrium. Using the payoff matrix you found in [Exercise 3.5.11](#), does Liar’s Poker have a pure strategy equilibrium? Explain.

Exercise 3.5.13 Mixed strategy equilibrium. Use any method you have learned to find a mixed strategy equilibrium for Liar’s Poker. Give the mixed strategy for Player 1 and the mixed strategy for Player 2.

Exercise 3.5.14 Compare strategies. Compare your solution from [Exercise 3.5.13](#) to your conjectured strategy from [Exercise 3.5.3](#).

Exercise 3.5.15 Expected value of the game. What is the expected value of the game for each player? How much would Player 1 expect to win if she played 15 games using the equilibrium mixed strategy?

Exercise 3.5.16 Fairness. Is this game fair? Explain. Again, compare your answer to your conjecture in [Exercise 3.5.3](#).

Congratulations! You can now set up matrices for simple games of chance and solve for a mixed strategy equilibrium. Before solving a more complicated game, let's get the help of technology for solving larger matrix games.

3.6 Augmented Matrices

In this section, we will see how to use matrices to solve systems of equations. In both the graphical method and the expected value method, you have had to solve a system of equations. In the graphical method you had systems consisting of two lines such as [Example 3.6.1](#).

Example 3.6.1 Two Equations. An example of a system of two lines:

$$\begin{aligned}y &= \frac{3}{5}x - \frac{1}{5} \\y &= -x + 3.\end{aligned}$$

□

In the expected value method we had systems of three equations such as [Example 3.6.2](#).

Example 3.6.2 Three Equations. An example of a system of three equations where the variables are $E_1(A)$, $E_1(B)$, $P_2(C)$, $P_2(D)$:

$$\begin{aligned}E_1(A) &= P_2(C) - P_2(D) \\E_1(B) &= 2P_2(D) \\1 &= P_2(C) + P_2(D).\end{aligned}$$

□

In [Example 3.6.2](#), even after setting $E_1(A) = E_1(B)$ so that there were only 2 variables, the algebra began to get cumbersome. What if we wanted to solve a much larger game, such as a 5 X 5 game?

We've used matrices to represent our games, but now we want to use them as a mathematical tool to help us solve systems of equations. In order to use matrices to solve our systems of equations, we want to write all our equations in the same form: we will have all the variable terms on the left of the equals sign and all constants on the right.

Thus, in [Example 3.6.1](#) we can write the system as

$$\begin{aligned}\frac{3}{5}x - y &= \frac{1}{5} \\x + y &= 3.\end{aligned}$$

In fact, we can simplify the first equation by multiplying both sides by 5:

$$\begin{aligned}3x - 5y &= 1 \\x + y &= 3.\end{aligned}$$

We can use the coefficients and constants to create a matrix:

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

In this matrix you have a column for the coefficients of each variable. So the coefficients of x are in the first column, the coefficients of y are in the second. The constant terms are always in the last column. Each row corresponds to

one equation. This matrix is called an *augmented matrix*. It is really just a matrix, but we call it *augmented* if we include information from both sides of the equation (the coefficients and the constants).

The algebraic method for solving the system of equations (finding the x and y values that satisfy both equations) is called *row reduction*. It is based on the *elimination method* that you may have seen in a precalculus or college algebra course. We won't go through the algebra here, as we really don't need it. Since our goal is to be able to easily solve larger systems of equations, we will rely on technology to do the algebra.

Computer algebra systems such as *Sage*, *Mathematica*, and *Maple*, as well as graphing calculators, can easily do the row reduction for us.

On a graphing calculator, use the matrix function to enter the matrix. Then use the "rref" function (this stands for "reduced row echelon form"). The result will be the following matrix:

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

We can solve this even more directly using *Sage*.

```
A=matrix([[3, -5, 1], [1, 1, 3]]);
show(A.rref())
```

```
matrix([[1, 0, 2], [0, 1, 1]])
```

Recall that the first column is for x and the second is for y . Each row represents an equation. So translating each row back to equations, we have the following system of equations:

$$\begin{aligned} x + 0y &= 2 \\ 0x + y &= 1 \end{aligned}$$

or

$$\begin{aligned} x &= 2 \\ y &= 1. \end{aligned}$$

By plugging these values back into the original equations, you can verify that this is in fact the solution.

Since the technology does all the algebra for us, our job is to translate the equations into an appropriate matrix and then translate the final matrix back into the solution to the system of equations. Remember, when using a matrix to solve a game, the matrix is only a tool, it is not the solution to the game.

Now, let's try the equations for the expected value method using [Example 3.6.2](#). As presented, how many variables does the system have?

$$\begin{aligned} E_1(A) &= P_2(C) - P_2(D) \\ E_1(B) &= 2P_2(D) \\ 1 &= P_2(C) + P_2(D) \end{aligned}$$

It has 4: $E_1(A)$, $E_1(B)$, $P_2(C)$ and $P_2(D)$. But when we solved these equations, we set the expected values equal. This gives us the two equations

$$\begin{aligned} P_2(C) - P_2(D) &= 2P_2(D) \\ 1 &= P_2(C) + P_2(D). \end{aligned}$$

Now if we put these into the form with all variables on the left and constants on the right, we get

$$\begin{aligned}P_2(C) - 3P_2(D) &= 0 \\P_2(C) + P_2(D) &= 1.\end{aligned}$$

Putting these equations into an augmented matrix, gives us

$$\begin{bmatrix}1 & -3 & 0 \\1 & 1 & 1\end{bmatrix},$$

where the first column corresponds to $P_2(C)$ and the second column corresponds to $P_2(D)$.

```
A=matrix([[1, -3, 0], [1, 1, 1]]);
show(A.rref())
```

```
matrix([[1, 0, 3/4], [0, 1, 1/4]])
```

Using row reduction, we get

$$\begin{bmatrix}1 & 0 & 3/4 \\0 & 1 & 1/4\end{bmatrix}.$$

Thus, recalling Column 1 is for $P_2(C)$ and Column 2 is for $P_2(D)$, our solution is $P_2(C) = 3/4$, and $P_2(D) = 1/4$.

Here are some more systems of equations to practice solving using augmented matrices. If you want to use the above Sage cells just edit the values for each row in the cell.

Exercise 3.6.3 Solve a system of 2 equations. Solve the system of equations.

$$\begin{aligned}2x - 2y &= 6 \\x + 3y &= 7\end{aligned}$$

Exercise 3.6.4 Solve another system of 2 equations. Solve the system of equations.

$$\begin{aligned}4p_1 - 2p_2 &= 0 \\p_1 + p_2 &= 1\end{aligned}$$

For larger matrices, you can edit the Sage cell by adding additional terms in each row, and adding more rows. For example, you can replace $[3, -5, 1]$, $[1, 1, 3]$ with $[4, 8, -4, 4]$, $[3, 8, 5, -11]$, $[-2, 1, 12, -17]$.

Exercise 3.6.5 Solve a system of 3 equations. Consider the system of equations

$$\begin{aligned}4x + 8y - 4z &= 4 \\3x + 8y + 5z &= -11 \\-2x + y + 12z &= -17.\end{aligned}$$

a Set up the augmented matrix for this system.

b Use row reduction to find the solution.

Exercise 3.6.6 Solve another system of 3 equations. Consider the system of equations

$$\begin{aligned} 2x + y - 4z &= 10 \\ 3x + 5z &= -5 \\ y + 2z &= 7. \end{aligned}$$

a Set up the augmented matrix for this system.

b Use row reduction to find the solution.

Exercise 3.6.7 Even more practice with 3 equations. Consider the system of equations

$$\begin{aligned} a + b - 5c &= 0 \\ -4a - b + 6c &= 0 \\ a + b + c &= 1. \end{aligned}$$

a Set up the augmented matrix for this system.

b Use row reduction to find the solution.

Exercise 3.6.8 Now, a sytem with 5 equations. Consider the system of equations

$$\begin{aligned} 3x + 2y - w - v &= 0 \\ 2x - y + 3z + w + 5v &= 0 \\ x + 2y + 6z - w &= 0 \\ -y + z - 3w + v &= 0 \\ x + y + z + w + v &= 1. \end{aligned}$$

a Set up the augmented matrix for this system.

b Use row reduction to find the solution.

Now we are ready to apply everything we have learned about solving repeated zero-sum games to a much more challenging game in the next section!

3.7 Undercut

This section requires you to be able to solve “large” systems of equations. You will be using the matrix techniques from [Section 3.6](#). You are encouraged to use technology such as a graphing calculator or Sage.

As we saw in [Section 3.5](#), an important part of game theory is the process of translating a game to a form that we can analyze.

Example 3.7.1 Undercut. Each player chooses a number 1-5. If the two numbers don’t differ by 1, then each player adds their own number to their score. If the two numbers differ by 1, then the player with the *lower* number adds *both* numbers to his or her score; the player with the higher number gets nothing. (From Douglas Hofstadter’s *Metamagical Themas*.)

For example, suppose in round one Player 1 chooses 4; Player 2 chooses 4. Each player keeps their own number. The score is now 4-4. In the next round, Player 1 chooses 2, Player 2 chooses 5. The score would now be 6-9. In the third round Player 1 chooses 4, Player 2 chooses 5. Now Player 1 gets both numbers, making the score 15-9. \square

Exercise 3.7.2 Play Undercut. Choose an opponent and play Undercut several times. Keep track of the outcomes.

After playing Undercut with an opponent, try to devise a good strategy.

Exercise 3.7.3 Conjecture a strategy. Just from playing Undercut several times, can you suggest a strategy for Player 1? What about for Player 2? For example, what number(s) should you play most often/ least often, or does it matter? Are there numbers you should never play? Does this game seem fair, or does one of the players seem to have an advantage? Explain your answers.

As we've seen before, a payoff matrix can help with analyzing a game.

Exercise 3.7.4 Payoff matrix. Create a payoff matrix for Undercut. Note that your payoffs should have a score for each player.

Exercise 3.7.5 Zero-sum. Is this a zero-sum game? Explain.

Exercise 3.7.6 Pure strategy equilibrium. Does there appear to be a pure strategy equilibrium for this game? Explain.

Let's assume we are going to play Undercut repeatedly. By the time you and your opponent are done playing, what should it mean to win the game?

Exercise 3.7.7 Long-run winner. How might we determine a "winner" for Undercut after playing several times?

Most likely, you said that someone will win the game if they have the most points. In fact, we probably don't care if the final score is 10-12 or 110-112. In either case, Player 2 wins. Since we will play this game several times, we do care about the point difference. For example, a score of 5-1 would be better for Player 1 than 5-3. So let's think about the game in terms of the point difference between the players in a given game. This is called the *net gain*. For example, with score of 5-1, Player 1 would have a net gain of 4.

Exercise 3.7.8 Net gain. Calculate the net gain for Player 1 for each of the three rounds in [Example 3.7.1](#) in the beginning of this section.

Exercise 3.7.9 Net gain payoff matrix. Create a new payoff matrix for Undercut which uses the players' net gain for the payoff vectors.

Exercise 3.7.10 Zero-sum. Is this now a zero-sum game? Explain.

The method of using net gain to describe the payoffs to each player should be familiar from some of the really early examples where we turned constant-sum payoff vectors into zero-sum vectors. But note that the original form of this game wasn't even a constant-sum game! What we are really doing here is thinking about our payoffs not as points, but a win or loss relative to our opponent. Now that we have reframed Undercut as a zero-sum game, we can apply our methods for solving the game that we have seen in this chapter.

Exercise 3.7.11 Pure strategy equilibrium. Is there a pure strategy equilibrium for this game? Explain.

Hint. Rather than looking at each option, you could compare the values for the pure maximin/ minimax strategies.

This game has one additional property that will help simplify our analysis. This game is *symmetric*, meaning the game looks the same to Players 1 and 2.

Exercise 3.7.12 Symmetric games. Give an example of another game which is symmetric. Give an example of a game which is *not* symmetric.

Exercise 3.7.13 Expected payoff for a symmetric game. What is the expected payoff for a symmetric game? Explain your answer.

Hint. You might think about whether it is possible for a player to have an advantage in a symmetric game.

Hopefully, you determined that there is not a pure strategy equilibrium for Undercut. Thus, we would like to find a mixed strategy equilibrium. Since this is a 5×5 game, we cannot use our graphical solution. We will need to

rely on our expected value solution. We want to decide with what probability we should play each number. Let a, b, c, d, e be the probabilities with which Player 2 plays 1-5, respectively. For example, if Player 1 plays a pure strategy of 2, then the expected value for Player 1, $E_1(2)$, is $-3a + 0b + 5c - 2d - 3e$.

Exercise 3.7.14 Equations for Player 1's expected value. Write down the five equations that give Player 1's expected value for each of Player 1's pure strategies.

Exercise 3.7.15 Expected value of a symmetric game. In [Exercise 3.7.13](#), you should have determined that since this is a symmetric game, the expected value for each Player should be 0. Modify your equations to include this piece of information. It is important to recognize that this step greatly simplifies our work for the expected value method since we don't need to set the expected values equal to each other. HOWEVER, we can ONLY do this since we know the game is symmetric!

If we use that the game is symmetric, and hence the expected value of the game for each player must be 0 since neither player can have an advantage over the other, we do not need to set the equations equal to each other. We could not use this method earlier since we had no way of knowing the expected value of a general game.

We now have five equations and five unknowns. There is a sixth equation: we know that the probabilities must add up to 1. We can now solve for the equilibrium strategy!

Exercise 3.7.16 Solve the system of equations. Use matrices to solve the resulting system of six equations. Give the mixed strategy equilibrium for Player 2. What is the mixed strategy for Player 1?

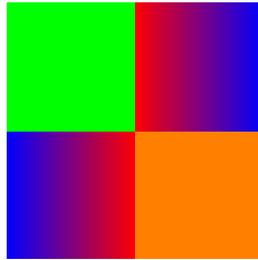
Hint. Should it be different than the strategy for Player 2?

Exercise 3.7.17 Summary. Based on your answer to [Exercise 3.7.16](#), which number(s) should you play the most often? Which should you play the least? Are there any numbers that you should never play? Compare the mathematical solution to your conjectured solution for [Exercise 3.7.3](#). Is there an advantage to knowing the mathematical solution?

You have now solved a rather complex two-person game. Try playing it with your friends and family. It may be difficult (or even impossible) to play randomly with the exact probabilities. It is also unlikely that your opponent will also be playing the equilibrium strategy, but can you use the solution to assure an advantage, or at least assure that your opponent doesn't have an advantage? Can you see the difference between an exact theoretical solution to a game, and a practical strategy for playing the game? In the next chapter we will see even more differences between theoretical and practical solutions to a game.

Chapter 4

Non-Zero-Sum Games



In the previous chapters we concentrated on zero-sum games. We know how to solve any zero-sum game. If it has a pure strategy equilibrium, then we know the players should play the equilibrium strategies. If it doesn't have an equilibrium point, then we have seen methods for finding a mixed strategy equilibrium. Assuming both players are our model rational players, then we know they should always play an equilibrium strategy.

In this chapter we turn our attention to non-zero-sum games.

4.1 Introduction to Two-Player Non-Zero-Sum Games

In this section we introduce non-zero-sum games. In a *non-zero-sum game* the players' payoffs no longer need to sum to a constant value. Now it is possible for both players to gain or both players to lose.

Exercise 4.1.1 Compare properties. What are some properties of a zero-sum game that may no longer hold for a non-zero-sum game? For example, in a zero-sum game is it ever advantageous to inform your opponent of your strategy? Is it advantageous to communicate prior to game play and determine a joint plan of action? Would you still want to minimize your opponents payoff?

Let's build an understanding of non-zero-sum games by looking at an example.

Example 4.1.2 Battle of the Sexes. Alice and Bob want to go out to a movie. Bob wants to see an action movie, Alice wants to see a comedy. Both prefer to go to a movie together rather than to go alone. We can represent the situation with the payoff matrix in [Table 4.1.3](#):

Table 4.1.3 Battle of the Sexes

		Alice	
		Action	Comedy
Bob	Action	(2, 1)	(-1, -1)
	Comedy	(-1, -1)	(1, 2)

□

Exercise 4.1.4 Not zero-sum. Explain why this is not a zero-sum game.

In zero-sum games it is never advantageous to let your opponent know your strategy. Does that property still apply for games like Battle of the Sexes?

Exercise 4.1.5 Announcing a strategy. Could it be advantageous for a player to announce his or her strategy? Could it be harmful to announce his or her strategy? If possible, describe a scenario in which it would be advantageous to announce a strategy. If possible, describe a scenario in which it would be harmful to announce a strategy.

We might first try to analyze Battle of the Sexes using the same techniques as we used for zero-sum games. For example, we might start as we would in zero-sum games by looking for any equilibrium points.

Exercise 4.1.6 Equilibrium points. Since our main goal in analyzing games has been to find equilibrium points, let's find any equilibrium points for Battle of the Sexes. Are there any strategy pairs where players would not want to switch?

Hint. There are two!

Exercise 4.1.7 Values of equilibrium points. Are the equilibrium points the same (in other words, does each player get the same payoff at each equilibrium point)? Compare this to what must happen for zero-sum games.

Now that we know Battle of the Sexes has two equilibrium points, we should try to find actual strategies for Alice and Bob. Is there a good strategy for each if they play the game only once? What if they repeat the game? Recall that with zero-sum games, if there was an equilibrium, rational players always want to play it, even if the game is repeated. Does that still seem to work here? Also, how might the ability to communicate change what the players do?

Exercise 4.1.8 Repeating the game. Suppose the game is played repeatedly. For example, Alice and Bob have movie night once a month. Suggest a strategy for Alice and for Bob. Feel free to play the game with someone from class. Try playing several times *without communicating* about your strategy before each game.

Exercise 4.1.9 Communication. How could communication affect the choice of strategy? Now play several times where you are allowed to communicate about your strategy. Does this change your strategy?

Exercise 4.1.10 Mixed strategy equilibrium. In either case, communicating and not communicating, do you think your strategies for Alice and Bob represent a mixed strategy equilibrium?

Exercise 4.1.11 Compare to zero-sum. In a zero-sum game, if there is a pure strategy equilibrium, then what happens when you repeat a game? If you repeat Battle of the Sexes, does the game always result in an equilibrium pair?

Hopefully, you are beginning to see some of the challenges for analyzing non-zero-sum games. We know there are equilibrium points in Battle of the Sexes,

but even rational play may not result in an equilibrium. For the remainder of this section, let's assume that players are *not* allowed to communicate about strategy prior to play. Such games are called *non-cooperative* games. Before moving on, let's try to find the maximin strategies for our players using the graphical method, as we did with zero-sum games.

Exercise 4.1.12 Bob's payoff matrix. Consider Battle of the Sexes from Bob's point of view. We know that Bob wants to maximize his payoff (that has not changed). So Bob doesn't care what Alice's payoff's are. Write down Bob's payoff matrix.

Exercise 4.1.13 Graphical method on Bob's matrix. Recall that the graphical method represents Bob's expected payoff depending on how often he plays each of his options. Sketch the graph associated with Bob's payoff matrix.

Exercise 4.1.14 Bob's maximin mixed strategy. The area between the two lines still represents the possible expected values for Bob, depending on how often Alice plays each of her strategies. So as before, the bottom lines represent the *least* Bob can expect as he varies his strategy. Thus, the point of intersection will represent the biggest of these smallest values (the maximin strategy). Find this point of intersection. How often should Bob play each option? What is his expected payoff?

So no matter what Alice does, Bob can expect that over the long run he wins at least the value you found in [Exercise 4.1.14](#). Make sure you understand this before moving on.

Exercise 4.1.15 Alice's maximin mixed strategy. Now consider Battle of the Sexes from Alice's point of view. Write down her payoff matrix and use the graphical method to find the probability with which she should play each option and her expected payoff.

Now, from [Exercise 4.1.14](#) and [Exercise 4.1.15](#) you have the minimum payoff each player should expect. Note that since this is not a zero-sum game, both players can expect a positive payoff. But now we want to see how this pair of mixed strategies really works for the players.

Exercise 4.1.16 Alice's expected value when Bob plays his maximin strategy. Assume Bob plays the mixed strategy from [Exercise 4.1.14](#). Calculate Alice's expected value for each of her *pure* strategies ($E_2(\text{Comedy})$ and $E_2(\text{Action})$).

Exercise 4.1.17 Alice's preference. Are Alice's expected values equal? If not, which strategy does she prefer when Bob plays the mixed strategy from [Exercise 4.1.14](#)? Does Alice want to deviate from her mixed strategy?

Exercise 4.1.18 Mixed strategy equilibrium. Now if Alice plays a pure strategy, does Bob want to change his strategy? So is the mixed strategy pair for Bob and Alice from [Exercise 4.1.14](#) and [Exercise 4.1.15](#) an equilibrium? In fact, if Bob changes his strategy, how does his expected value compare to the expected value for his mixed strategy?

Exercise 4.1.19 Downside of the graphical method. What goes wrong with the graphical method in the case of a non-zero-sum game?

Hint. Is it important for Alice to consider the minimax strategy? Does Alice have any reason to minimize Bob's payoff?

As we can see by the above example, the methods used to analyze zero-sum games may not be too helpful for non-zero-sum games. Part of the problem is that in solving zero-sum games we take into consideration that one player

wants to *minimize* the payoff to the other player. This is no longer the case. Changing strategies may allow BOTH players to do better. A player no longer has any reason to prevent the other player from doing better.

Exercise 4.1.20 Response to the mixed strategy. So we know the mixed strategy won't give us an equilibrium. But suppose we start there. In other words, suppose Bob plans to play the mixed strategy from [Exercise 4.1.14](#). Which pure strategy should Alice play? In response, which pure strategy should Bob play? What is the outcome? Do you end up at an equilibrium?

Exercise 4.1.21 Bos's expected value when Alice plays her maximin strategy. Now suppose Alice plans to play the mixed strategy from [Exercise 4.1.15](#). Calculate the expected value for Bob for each of his pure strategies. Which pure strategy does Bob prefer to play? How will Alice respond to Bob's pure strategy? What is the outcome? Do you end up at an equilibrium?

Exercise 4.1.22 Out-guessing the mixed strategy. Suppose Bob thinks Alice will try the mixed strategy and Alice thinks Bob will try the mixed strategy. Which pure strategy will each play? What is the outcome? Do you end up at an equilibrium?

Exercise 4.1.23 Playing the maximin mixed strategy. Considering [Exercise 4.1.20](#), [Exercise 4.1.21](#), and [Exercise 4.1.22](#), is it in a player's best interest to even consider playing the mixed strategy from [Exercise 4.1.14](#) or [Exercise 4.1.15](#)?

We've seen the limitations of the graphical method, but what about the expected value method?

Exercise 4.1.24 Expected value solution. Try applying the expected value method to Battle of the Sexes. What mixed strategies do you get? Explain why this method will also not result in an equilibrium. You might want to consider whether it is important for one player to minimize the expected value for the other player.

Now that we have seen how the methods that allowed us to solve zero-sum games don't work for non-zero-sum games, we will need to find new ways to approach non-zero-sum games.

4.2 Prisoner's Dilemma and Chicken

Before getting any further into non-zero-sum games, let's recall some key ideas about zero-sum games.

- If a zero-sum game has an equilibrium point, then repeating the game does not affect how the players will play.
- If a zero-sum game has more than one equilibrium point then the values of the equilibrium points are the same.
- In a zero-sum game, we can find mixed strategy equilibrium points using the graphical method or the expected value method.
- In a zero-sum game, a player never benefits from communicating her strategy to her opponent.

Hopefully, in the last section you saw that non-zero-sum games can differ on all of the above!

Example 4.2.1 A 2×2 Non-Zero Sum Game. Let's consider the game given by Table 4.2.2.

Table 4.2.2 A non-zero sum example

	C	D
A	(0, 0)	(10, 5)
B	(5, 10)	(0, 0)

□

Exercise 4.2.3 Not zero-sum. Check that this is NOT a zero-sum game.

Even with non-zero-sum games, it is helpful to start by finding any equilibrium points.

Exercise 4.2.4 Equilibrium points. Using the “guess and check” method for finding equilibria, you should be able to determine that Table 4.2.2 has two equilibrium points. What are they?

Exercise 4.2.5 Preference between equilibria. As we saw in Section 4.1, the equilibrium points in non-zero-sum games need not have the same values. Does Player 1 prefer one of the equilibria from Exercise 4.2.4 over the other?

Since it is now possible for BOTH players to benefit at the same time, it might be a good idea for players to communicate with each other. For example, if Player 1 says that she will choose A no matter what, then it is in Player 2's best interest to choose D. If communication is allowed in the game, then we say the non-zero-sum game is *cooperative*. If no communication is allowed, we say it is *non-cooperative*.

We saw in Section 4.1, that our methods for analyzing zero-sum games do not work very well on non-zero-sum games. Let's look a little closer at this.

If we apply the graphical method for Player 1 to the above game, we get that Player 1 should play a $(1/3, 2/3)$ mixed strategy for an expected payoff of $10/3$. Similarly we can determine that Player 2 should play a $(2/3, 1/3)$ mixed strategy for an expected payoff of $10/3$. Recall we developed this strategy as a “super defensive” strategy. But are our players motivated to play as defensively in a non-zero-sum game? Not necessarily! It is no longer true that Player 2 needs to keep Player 1 from gaining!

Now suppose, Player 1 plays the $(1/3, 2/3)$ strategy. Then the expected payoff to Player 2 for playing pure strategy C, $E_2(C)$, is $20/3$; and the expected payoff to Player 2 for playing pure strategy D, $E_2(D)$, is $5/3$. Thus Player 2 prefers C over D. But if Player 2 plays only C, then Player 1 should abandon her $(1/3, 2/3)$ strategy and just play B! This results in the payoff vector $(5, 10)$. Notice, that now the expected value for Player 1 is 5, which is better than $10/3$! Again, since Player 2 is not trying to keep Player 1 from gaining, there is no reason to apply the maximin strategy to non-zero-sum games. Similarly, we don't want to apply the expected value solution since Player 1 does not care if Player 2's expected values are equal. Each player only cares about his or her own payoff, not the payoff of the other player.

OK, so now, how do we analyze these games?

Exercise 4.2.6 Conjecture a strategy. What are some possible strategies for each player? Might some strategies depend on what a player knows about her opponent?

Can you see that some of the analysis might be better understood with psychology than with mathematics?

In order to better understand non-zero-sum games we look at two classic games.

Example 4.2.7 Prisoner’s Dilemma. Two partners in crime are arrested for burglary and sent to separate rooms. They are each offered a deal: if they confess and rat on their partner, they will receive a reduced sentence. So if one confesses and the other doesn’t, the confessor only gets 3 months in prison, while the partner serves 10 years. If both confess, then they each get 8 years. However, if neither confess, there isn’t enough evidence, and each gets just one year. We can represent the situation with the matrix in [Table 4.2.8](#).

Table 4.2.8 The Prisoner’s Dilemma (years in prison).

		Prisoner 2	
		Confess	Don’t Confess
Prisoner 1	Confess	(8, 8)	(0.25, 10)
	Don’t Confess	(10, 0.25)	(1, 1)

□

Since this game, as presented, is probably only played once, we can begin by looking for dominated strategies and equilibrium points.

Exercise 4.2.9 Dominated strategies. Does the matrix in [Table 4.2.8](#) have any dominated strategies for Player 1? Does it have any dominated strategies for Player 2? Keep in mind that a prisoner prefers smaller numbers since prison time is bad.

If you were to be one of the prisoners, what would you do? Do you think everyone would do that, too? What would our perfectly rational player do? Would your strategy change if you are allowed to communicate? We examine some of these questions in the next few exercises.

Exercise 4.2.10 A prisoner’s strategy. Suppose you are Prisoner 1. What should you do? Why? Suppose you are Prisoner 2. What should you do? Why? Does your choice of strategies result in an equilibrium pair?

Exercise 4.2.11 The best outcome. Looking at the game as an outsider, what strategy pair is the best option for both prisoners.

Exercise 4.2.12 Two rational prisoners. Now suppose both prisoners are perfectly rational, so that any decision Prisoner 1 makes would also be the decision Prisoner 2 makes. Further, suppose both prisoners know that their opponent is perfectly rational. What should each prisoner do?

Exercise 4.2.13 A crazy prisoner. Suppose Prisoner 2 is crazy and is likely to confess with 50/50 chance. What should Prisoner 1 do? Does it change if he confesses with a 75% chance? What if he confesses with a 25% chance.

Exercise 4.2.14 Communication between prisoners. Suppose the prisoners are able to communicate about their strategy. How might this affect what they choose to do?

Exercise 4.2.15 The dilemma. Explain why this is a “dilemma” for the prisoners. Is it likely they will chose a strategy which leads to the best outcome for both? You might want to consider whether there are dominated strategies.

You should now have some idea about why we call this game a dilemma, since the players may in fact have difficulty deciding on whether to confess or not. Even two perfectly rational players may not be able to get the best payoff.

We now turn to another classic example. We can ask similar questions about Chicken that we ask about Prisoner’s Dilemma.

Example 4.2.16 Chicken. Two drivers drive towards each other. If one driver swerves, he is considered a “chicken.” If a driver doesn’t swerve (drives straight), he is considered the winner. Of course if neither swerves, they crash

and neither wins. A possible payoff matrix for this game is given in [Table 4.2.17](#).

Table 4.2.17 The game of Chicken.

		Driver 2	
		Swerve	Straight
Driver 1	Swerve	(0, 0)	(-1, 10)
	Straight	(10, -1)	(-100, -100)

□

Again, since this game as presented is probably only played once, we can begin by looking for dominated strategies and equilibrium points.

Exercise 4.2.18 Dominated strategies. Does the Chicken game in [Table 4.2.17](#) have any dominated strategies?

Exercise 4.2.19 The best outcome. What strategy results in the best outcome for Driver 1? What strategy results in the best outcome for Driver 2? What happens if they both choose that strategy?

Exercise 4.2.20 Equilibrium points. Consider the strategy pair with outcome $(-1, 10)$. Does Driver 1 have any regrets about his choice? What about Driver 2? Is this an equilibrium point? Are there any other points in which neither driver would regret his or her choice?

If you were to be one of the drivers, what would you do? Do you think everyone would do that, too? What would our perfectly rational player do? Would your strategy change if you are allowed to communicate? We examine some of these questions in the next few exercises.

Exercise 4.2.21 A driver's strategy. Can you determine what each driver should do in Chicken? If so, does this result in an equilibrium pair?

Exercise 4.2.22 Two rational drivers. Now suppose both drivers in the game of Chicken are perfectly rational, so that any decision Driver 1 makes would also be the decision Driver 2 makes. Further, suppose both drivers know that their opponent is perfectly rational. What should each driver do?

Exercise 4.2.23 A remote control driver. Suppose Driver 2 is a remote control dummy and will swerve or drive straight with a 50/50 chance. What should Driver 1 do? Does it change if the remote control dummy swerves with 75% chance?

Exercise 4.2.24 Communication between drivers. Can it benefit drivers in the game of Chicken to communicate about their strategy? Explain.

Exercise 4.2.25 Compare games. Compare Prisoner's Dilemma and Chicken. Are there dominated strategies in both games? Are there equilibrium pairs? Are players likely to reach the optimal payoff for one player, both players, or neither player? Does a player's choice depend on what he knows about his opponent (perfectly rational or perfectly random)?

Both Prisoner's Dilemma and Chicken are models of games where we describe the choice of strategy as "Cooperate" and "Defect". In Prisoner's Dilemma, we think of "cooperating" as cooperating with the other player, and "defecting" as turning against the other player. So if both players cooperate (with each other, not the police), they will get the higher payoff of only one year in prison. They defect by ratting on each other. In Chicken, players cooperate by swerving and defect by driving straight. Using the examples of Prisoner's Dilemma and Chicken, think about how these games can model other everyday interactions where you could describe your choices as cooperating and defecting. The

remainder of the chapter looks more closely at situations where players can cooperate or defect.

4.3 A Class-Wide Experiment

We are going to look at a class-wide game.

Each member of the class secretly chooses a single letter: “C” or “D,” standing for “cooperate” or “defect.” This will be used as your strategy choice in the following game with each of the other players in the class. Here is how it works for each pair of players: if they both cooperate, they each get 3 points. If they both defect, they each get 1 point. If one cooperates and one defects, the cooperator gets nothing, but the defector gets 5 points. Your one choice of “C” or “D” will be used to play the game with all the other players in the class.

Thus, if everyone chooses “C,” everyone will get 3 points per person (not counting yourself). If everyone chooses “D,” everyone will get 1 point per person (not counting yourself). You can’t lose! And of course, anyone chooses “D” will get at least as much as everyone else will. If, for example in a class of 20 people, 11 people choose “C” and 9 choose “D,” then the 11 C-ers will get 3 points apiece from the other C-ers (making 30 points), and zero from the D-ers. So C-ers will get 30 points each. The D-ers, by contrast, will pick up 5 points apiece from each of the C-ers, making 55 points, and 1 point from each of the other D-ers, making 8 points, for a grand total of 63 points. No matter what the distribution is, D-ers always do better than C-ers. Of course, the more C-ers there are, the better everyone will do!

By the way, I should make it clear that in making your choice, you should not aim to be the winner, but simply to get as many points for yourself as possible. Thus you should be happier to get 30 points (as a result of saying “C” along with 10 others, even though the 9 D-sayers get more than you) than to get 19 points (by saying “D” along with everybody else, so nobody “beats” you).

Of course, your hope is to be the only defector, thus really cleaning up: with 19 C-ers, you’ll get 95 points, and they’ll each get 18 times 3, namely 54 points! But why am I doing the multiplication or any of this figuring for you? You’ve been studying game theory. So have all of you! You are all equally versed in game theory and understand about making rational choices. Therefore, I hardly need to tell you that you are to make what you consider to be your maximally rational choice. In particular, feelings of morality, guilt, apathy, and so on, are to be disregarded. Reasoning alone (of course including reasoning about others’ reasoning) should be the basis of your decision.

So all you need to do is make your choice. Write it down.

It is to be understood (it almost goes without saying, but not quite) that you are not to discuss your answer with anyone else from the class. The purpose is to see what people do on their own, in isolation. Along with your answer you should include a short explanation for why you made your particular choice.

Adapted from Douglas Hofstadter, *Metamagical Themas*, p. 740.

Once everyone in class has made his or her choice, share your answers with the class. Then answer the following questions about the class’s responses.

Exercise 4.3.1 Summary of responses.

- a How many C’s were there?
- b How many D’s were there?

c What was the payoff to each C?

d What was the payoff to each D?

Exercise 4.3.2 Payoff matrix. Determine the payoff matrix for class-wide Prisoner's Dilemma.

Hint. Although you played this game with each other person in the class, this is still a 2 person game!

Exercise 4.3.3 Reasons for choice. What are some reasons people chose C? What are some reasons people chose D?

Although we can now see what everyone chose, we might not agree that everyone made the most rational choice. How might perfectly rational players play the game?

Exercise 4.3.4 The rational choice. Thinking about the idea of *rationality*, what appears to be the most rational choice, C or D? If everyone is *equally* rational, then what would everyone do? If everyone is equally rational, should everyone choose the same thing?

Exercise 4.3.5 Everyone is rational. Now suppose everyone is equally (and perfectly) rational. AND everyone knows that everyone else is equally (and perfectly) rational. What should everyone choose?

Hint. If everyone knows that everyone will choose the same answer, what should everyone choose to do?

The next two exercises look at two more examples of games where players can "Cooperate" or "Defect". How does changing the payoffs change the players' incentive to cooperate or defect?

Exercise 4.3.6 A game of cooperation and defection. Consider the game in [Table 4.3.7](#)

Table 4.3.7 Matrix for Exercise 4.3.6

	<i>C</i>	<i>D</i>
<i>C</i>	(3, 3)	(0, 50)
<i>D</i>	(50, 0)	(.01, .01)

What would you do? Why? What seems to be the most rational thing to do? Why?

Exercise 4.3.8 Another game of cooperation and defection. Consider the game in [Table 4.3.9](#)

Table 4.3.9 Matrix for Exercise 4.3.8

	<i>C</i>	<i>D</i>
<i>C</i>	(1000, 1000)	(0, 100)
<i>D</i>	(100, 0)	(100, 100)

What would you do? Why? What seems to be the most rational thing to do? Why?

Exercise 4.3.10 Motivation to cooperate or defect. Looking at all three of the above games, can you think of what sort of payoffs you would need in order to cooperate (C)? What about to defect (D)?

Not every game where player's cooperate or defect is a Prisoner's Dilemma, or even a dilemma. You can certainly change the payoffs in the above matrices so that it is very clear what each player should do. But as you've seen with this section's experiment, there is something special about the Prisoner's Dilemma.

Everyone does better if they all cooperate, but any one player does better to defect. The next section will look more specifically at what makes a game a Prisoner's Dilemma.

4.4 What Makes a Prisoner's Dilemma?

In this section we give a mathematical description of Prisoner's Dilemma and compare it to some similar games.

The Class-wide Prisoner's Dilemma game we played in [Section 4.3](#) has the payoff matrix given in [Table 4.4.1](#) for each pair of players.

Table 4.4.1 A Class-wide Prisoner's Dilemma.

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	(3, 3)	(0, 5)
	Defect	(5, 0)	(1, 1)

We can classify each of the values for the payoffs as follows:

- Reward for Mutual Cooperation: $R = 3$.
- Punishment for Defecting: $P = 1$.
- Temptation to Defect: $T = 5$.
- Sucker's Payoff: $S = 0$.

In order for a game to be a variation of Prisoner's Dilemma it must satisfy two conditions:

1. $T > R > P > S$
2. $(T + S)/2 < R$

Let's apply this description of Prisoner's Dilemma to a few games we've seen. We can use the conditions to check if a game is really a Prisoner's Dilemma.

Exercise 4.4.2 Description of conditions. Describe each of the conditions (A) and (B) in words.

Hint. $(T + S)/2$ is the average of T and S .

Exercise 4.4.3 The conditions for Classwide Prisoner's Dilemma. Show that the two conditions hold for the Class-wide Prisoner's Dilemma.

Exercise 4.4.4 The conditions for Prisoner's Dilemma. Recall the matrix for Prisoner's Dilemma from [Example 4.2.7](#).

Table 4.4.5 Prisoner's Dilemma (again).

		Prisoner 2	
		Confess	Don't Confess
Prisoner 1	Confess	(8, 8)	(0.25, 10)
	Don't Confess	(10, 0.25)	(1, 1)

Determine R, P, T , and S for this game. Be careful: think about what cooperating versus defecting should mean. Show the conditions for Prisoner's Dilemma are satisfied.

Hint. Time in jail is bad, so the bigger the number, the worse you do; thus, it might be helpful to think of the payoffs as negatives.

Exercise 4.4.6 The conditions for Chicken. Recall the matrix for Chicken from [Example 4.2.16](#).

Table 4.4.7 Chicken (again).

		Driver 2	
		Swerve	Straight
Driver 1	Swerve	(0, 0)	(-1, 10)
	Straight	(10, -1)	(-100, -100)

Determine R, P, T , and S for this game. Again, think about what cooperating and defecting mean in this game. Determine if the conditions for Prisoner's Dilemma are satisfied. If not, which condition(s) fail?

Exercise 4.4.8 The conditions on another game. Consider the game:

	C	D
C	(3, 3)	(0, 50)
D	(50, 0)	(.01, .01)

Determine R, P, T , and S for this game. Determine if the conditions for Prisoner's Dilemma are satisfied. If not, which condition(s) fail?

Exercise 4.4.9 A little more practice. Consider the game:

	C	D
C	(1000, 1000)	(0, 100)
D	(100, 0)	(100, 100)

Determine R, P, T , and S for this game. Determine if the conditions for Prisoner's Dilemma are satisfied. If not, which condition(s) fail?

Exercise 4.4.10 Compare the games. The games in [Exercise 4.4.6](#), [Exercise 4.4.8](#), and [Exercise 4.4.9](#) are not true Prisoner's Dilemmas. For each game, how do the changes in payoffs affect how you play? In particular, in Prisoner's Dilemma, a player will generally choose to defect. This results in a non-optimal payoff for each player. Is this still true in [Exercise 4.4.6](#), [Exercise 4.4.8](#), and [Exercise 4.4.9](#)? If possible, use the changes in the conditions (A) and (B) to help explain any differences in how one should play.

We can now define *defection* as the idea that if everyone did it, things would be worse for everyone. Yet, if only one (or a small) number did it, life would be sweeter for that individual. We can define *cooperation* as the act of resisting temptation for the betterment of all players.

Exercise 4.4.11 Example from real life. Give an example of defection and cooperation from real life.

4.5 Another Multiplayer Experiment

This activity needs to be played as a class. All players need to be able to respond without being able to see the responses of others. Answers may be revealed before moving on to the next question.

Without sharing your answers with others, select your answer to the following questions. Try to be as honest about your answer as possible. Make sure you have a reason for each answer.

1. The lights go out in the neighborhood. Someone needs to call the power company. If someone calls, everyone's lights go on.

- (a) Call
 - (b) Don't call
2. The same as in (1), but now you have to wait on hold for 5 minutes.
 - (a) Call
 - (b) Don't call
 3. The same as in (1), but now you have to wait on hold for 30 minutes.
 - (a) Call
 - (b) Don't call
 4. The same as in (1), but now you have to pay a \$.50 service fee.
 - (a) Call
 - (b) Don't call
 5. The same as in (1), but now you have to pay a \$2.50 service fee.
 - (a) Call
 - (b) Don't call
 6. The phone lines go down in your small mountain community. You have to hike 3 miles in the snow to notify the power company.
 - (a) Hike to notify the phone company
 - (b) Stay home and let someone else do the hiking
 7. Everyone in your class cheats on an exam. The professor says if someone confesses they receive an F. If no one confesses, everyone fails.
 - (a) Confess
 - (b) Don't confess
 8. Evil Dr. No captures the class and puts you in all in a shark tank separated so you can't communicate. If one person volunteers to be eaten then the rest go free. If no one volunteers after 10 minutes all get eaten by sharks.
 - (a) Volunteer
 - (b) Don't volunteer
 9. Evil Dr. No captures your family and puts you in all in a shark tank separated so you can't communicate. If one person volunteers to be eaten then the rest go free. If no one volunteers after 10 minutes all get eaten by sharks.
 - (a) Volunteer
 - (b) Don't volunteer
 10. For any "Big Brother" fans. Choose to eat all your favorite foods for a week or nasty "slop" for a week. If at least three people say slop, everyone gets what they asked for. Otherwise everyone is on slop.
 - (a) Favorite foods
 - (b) Slop

11. OK, now let's get serious about this. Answer 5 points or 1 point. If at least one person says 1 point, everyone gets the number of points they chose. Otherwise, everyone gets 0 points.
 - (a) 5 points
 - (b) 1 point
12. Answer 20 points or 1 point. If at least one person says 1 point, everyone gets the number of points they chose. Otherwise, everyone gets 0 points.
 - (a) 20 points
 - (b) 1 point
13. Answer 6 points or 5 points. If at least one person says 5 points, everyone gets the number of points they chose. Otherwise, everyone gets 0 points.
 - (a) 6 points
 - (b) 5 points
14. Answer 20 points or 1 point. If at least *five* people say 1 point, everyone gets the number of points they chose. Otherwise, everyone gets 0 points.
 - (a) 20 points
 - (b) 1 point

After answering the above questions and seeing the responses from your classmates, think about how you responded. Did this differ from how your classmates responded? Try to give reasons for how you chose your responses to the above questions. Ask classmates for their reasons for responding as they did. It can be particularly useful to share your answers with someone who chose a different response from you. You can summarize the various reasons for volunteering and not volunteering in the exercises below.

Exercise 4.5.1 Volunteer or not. After answering the questions, were you likely to volunteer or unlikely to volunteer? For example, were you likely to be the one to call the power company or get eaten by sharks, or were you generally hoping someone else would do it? If it depended on the situation, explain under what circumstances you were likely to volunteer.

Exercise 4.5.2 Always a volunteer. After sharing your answers as a class, did each situation have a volunteer? In other words, was there always someone willing to call the power company or take fewer points? If there was a question with no volunteer, can you suggest why?

Exercise 4.5.3 Unlikely to volunteer. For which questions was it unlikely that there would be very many volunteers? Did you take that into consideration when deciding if you were going to volunteer?

Exercise 4.5.4 Reasons to volunteer. What are some reasons for volunteering? What are some reasons for not to volunteer?

4.6 Volunteer's Dilemma

In [Section 4.5](#) we played a game called Volunteer's Dilemma.

Example 4.6.1 A Volunteer's Dilemma. One example of a Volunteer's Dilemma is the game where everyone chooses "1 point" or "5 points." If at least

one person writes down 1 point, then everyone gets the number of points they wrote down. If no one chooses 1 point, then everyone gets 0 points. Choosing “1 point” is considered volunteering or cooperating. Choosing to not volunteer and take “5 points” is defecting.

You should note that it is difficult to put this game into a matrix form since payoffs depend on whether there is at least one volunteer or cooperator. \square

In this section we will compare Class-wide Prisoner’s Dilemma with Volunteer’s Dilemma. In particular, we want to think about the effect cooperating and defecting have on the group of players. How does one player’s choice affect everyone else? What happens to the group if there is a single cooperator or a single defector? What happens if everyone cooperates or everyone defects? We will use the payoffs for Prisoner’s Dilemma from [Section 4.3](#), given again in [Table 4.6.2](#).

Table 4.6.2 Class-wide Prisoner’s Dilemma (again).

		Player 2	
		Cooperate	Defect
Driver 1	Cooperate	(3, 3)	(0, 5)
	Defect	(5, 0)	(1, 1)

Exercise 4.6.3 Effect of a single defector in Class-wide Prisoner’s Dilemma. In *Class-wide Prisoner’s Dilemma* what effect does one defector have on the group? In other words, if a single player defects, how many points does he cost each of the other players?

Exercise 4.6.4 Effect of everyone’s defection in Class-wide Prisoner’s Dilemma. In *Class-wide Prisoner’s Dilemma* what effect does everyone’s defection have on the group? In other words, what is the most points lost by the group if everyone defects?

Exercise 4.6.5 Effect of your defection in Class-wide Prisoner’s Dilemma. In *Class-wide Prisoner’s Dilemma* what effect could your own defection have on the group? In other words, how many points does the group lose if you defect instead of cooperate? You may need to consider different cases depending on how many cooperators there are. For example what if there are no cooperators? What if there are no defectors? What if there are some of each?

Exercise 4.6.6 Effect of a single defector in Volunteer’s Dilemma. In *Volunteer’s Dilemma*, with the payoffs given in [Example 4.6.1](#), what effect does one defector have on the group? In other words, if there is a single defector, how many points do each of the other players lose?

Exercise 4.6.7 Effect of everyone’s defection in Volunteer’s Dilemma. In *Volunteer’s Dilemma*, with the payoffs given in [Example 4.6.1](#), what effect does everyone’s defection have on the group? In other words, if everyone defects, how many points does the group lose?

Exercise 4.6.8 Effect of your defection in Volunteer’s Dilemma. In *Volunteer’s Dilemma*, with the payoffs given in [Example 4.6.1](#), what effect could your own defection have on the group? In other words, how many points does the group lose if you defect instead of cooperate? You may need to consider different cases depending on how many cooperators there are. For example what if there are no cooperators? What if there are no defectors? What if there are some of each?

Now that we’ve considered how an individual decision can affect the group, we can think about what the most rational strategy is in a multiplayer Pris-

oner's Dilemma or a Volunteer's Dilemma.

Exercise 4.6.9 Rationality in Class-wide Prisoner's Dilemma. Considering your answers above and to previous work with Prisoner's Dilemma, in *Class-wide Prisoner's Dilemma*, which is more *rational* to be a cooperator or a defector? Why?

Exercise 4.6.10 Everyone is rational in Class-wide Prisoner's Dilemma. Whichever strategy you chose in [Exercise 4.6.9](#), explain what would happen if everyone was the most rational. Is it now more rational to do the opposite?

Exercise 4.6.11 Rationality in Volunteer's Dilemma. Considering your answers above, in *Volunteer's Dilemma*, which is more *rational* to be a cooperator (volunteer) or a defector? Why?

Exercise 4.6.12 Everyone is rational in Volunteer's Dilemma. Whichever strategy you chose in [Exercise 4.6.11](#), explain what would happen if everyone was the most rational. Is it now more rational to do the opposite?

Exercise 4.6.13 Class-wide Chicken. Volunteer's Dilemma can also be called "Class-wide Chicken." Try to describe this class-wide game in terms of "swerving" and "going straight." How do the payoffs for Volunteer's Dilemma relate to the payoffs for Chicken?

Even though the Class-wide Prisoner's Dilemma and the Volunteer's Dilemma games were played with multiple players, each game was only played once. In the next section we look at what might happen if we repeatedly play Prisoner's Dilemma with the same opponent.

4.7 Repeated Prisoner's Dilemma

In this section we look at two players playing Prisoner's Dilemma repeatedly. We call this game an *iterated* Prisoner's Dilemma. Recall the general Prisoner's Dilemma matrix from previous sections, given again in [Table 4.7.1](#).

Table 4.7.1 A Prisoner's Dilemma matrix.

		Player 2	
		Cooperate	Defect
Driver 1	Cooperate	(3, 3)	(0, 5)
	Defect	(5, 0)	(1, 1)

Before playing the iterated version, think about how you would play the above game if you only play it once with an opponent. But let's also give the game some context as in the following exercise.

Exercise 4.7.2 A single internet purchase. Suppose the above matrix represents the situation of purchasing an item (say, a used textbook) on the internet where both parties are untraceable. You agree to send the money at the same time that the seller agrees to send the book. Then we can think of Cooperating as each of you sending money/ book, and Defecting as not sending money/ book. Why might a player Cooperate? Why might a player Defect?

Exercise 4.7.3 Repeated internet purchases. Now suppose you wish to continue to do business with the other party. For example, instead of buying a used textbook, maybe you are buying music downloads. Why might a player cooperate? Why might a player defect? Do these reasons differ from your reasons in [Exercise 4.7.2](#)?

It is possible that your answers to the above questions depended on the context, so let's go back to just thinking about the game as a simple matrix

game. Think about how you played the Class-wide Prisoner's Dilemma, but now think about repeating the game several times with the same player. Do you think your strategy would change? Remember, if we repeat the game we are not restricted to playing a pure strategy.

Exercise 4.7.4 Strategy for repeated Prisoner's Dilemma. Suggest a strategy for playing the Prisoner's Dilemma in Table 4.7.1 repeatedly. DON'T SHARE YOUR STRATEGY WITH ANYONE YET!

Now let's see how your strategy works by actually playing the game several times.

Exercise 4.7.5 Play Prisoner's Dilemma repeatedly. Play 10 rounds of Prisoner's Dilemma with someone in class. Use your suggested strategy. Keep track of the payoffs. Did your strategy seem effective? What should it mean to have an "effective" strategy?

We are now going to play a Prisoner's Dilemma Tournament! Several strategies are suggested below. Choose one of the strategies below (or keep playing with your own strategy). You are to play your strategy against everyone else in the class. Keep a running total of your score. Don't tell your opponents your strategy.

Some possible strategies:

- Strategy: **Defection.** Your strategy is to *always* choose DEFECT (D).
- Strategy: **Cooperation.** Your strategy is to *always* choose COOPERATE (C).
- Strategy: **Tit for Tat.** Your strategy is to play whatever your opponent just played. Your first move is to COOPERATE (C), but then you need to repeat your opponent's last move.
- Strategy: **Tit for Two Tats.** Your strategy is to COOPERATE (C) unless your opponent DEFECTS twice in a row. After two D's you respond with D.
- Strategy: **Random (1/2, 1/2).** Your strategy is to COOPERATE (C) randomly 50% of the time and DEFECT (D) 50% of the time. [Note: it will be hard to be truly random, but try to play each option approximately the same amount.]
- Strategy: **Random (3/4, 1/4).** Your strategy is to COOPERATE (C) randomly 75% of the time and DEFECT (D) 25% of the time. [Note: it will be hard to be truly random, but try to play C more often than D.]
- Strategy: **Random (1/4, 3/4).** Your strategy is to COOPERATE (C) randomly 25% of the time and DEFECT (D) 75% of the time. [Note: it will be hard to be truly random, but try to play D more often than C.]
- Strategy: **Tit for Tat with Occasional Surprise D.** Your strategy is to play whatever your opponent just played. Your first move is to COOPERATE (C), but then you need to repeat your opponent's last move. Occasionally, you will deviate from this strategy by playing D.

Exercise 4.7.6 A Prisoner's Dilemma tournament. WITHOUT SHARING YOUR STRATEGY, play Prisoner's Dilemma 10 times with each of the other members of the class. Keep track of the payoffs for each game and your total score.

After playing Repeated Prisoner's Dilemma as a class, can you determine who used which strategy? At this point you may share your strategy with others.

Was your strategy more effective with some strategies than with others? If some of the above strategies were not used, can you guess how your strategy would have done against them?

Exercise 4.7.7 Effectiveness of your strategy. Describe which opponents' strategies seemed to give you more points, which seemed to give you fewer?

Exercise 4.7.8 Winning strategies. Describe the strategy or strategies that had the highest scores in the tournament. Does this seem surprising? Why or why not? How do the winning strategies compare to the strategy you suggested in [Exercise 4.7.4](#)?

What strategies seem the most rational? Are pure strategies the most rational? Does it depend on what sort of strategy your opponent is playing?

Exercise 4.7.9 Rationality in repeated Prisoner's Dilemma. How does Repeated Prisoner's Dilemma differ from the "one-time" Prisoner's Dilemma? Try to think in terms of rational strategies.

As [Exercise 4.7.3](#) suggests we can think of real-life interactions that can be modeled as a Prisoner's Dilemma.

Exercise 4.7.10 Example of Repeated Prisoner's Dilemma in real life. Describe a situation from real life that resembles a Repeated Prisoner's Dilemma.

Repeated or Iterated Prisoner's Dilemma has applications to biology and sociology. If you think of higher point totals as "success as a species" in biology or "success of a society" in sociology, we can try to determine which strategies seem the most effective or successful. Individuals do not need the highest point total to be successful, but low point totals will not succeed. Just like grades in a course, you don't need the highest score to pass a class, but very bad scores will fail. In order to model the situation of a society, think about what happens to defectors in a society of cooperators or cooperators in a society of defectors. Who will be able to succeed?

Exercise 4.7.11 Only a few defectors. How do a few defectors fare in a society of mostly cooperators? How do the cooperators fare? (In other words, who will be more successful?) Keep in mind that everyone is playing with lots of cooperators and only a few defectors. Who will have the most points, cooperators or defectors?

Exercise 4.7.12 Only a few cooperators. How do a few cooperators fare in a society of mostly defectors? How do the defectors fare? (In other words, who will be more successful?) Keep in mind that everyone is playing with lots of defectors and only a few cooperators. Who will have the most points, cooperators or defectors?

After thinking about individuals in a society playing pure strategies, what happens to individuals who are playing some of the mixed strategies listed above?

Exercise 4.7.13 A society of TIT-FOR-TATers. Now consider a society of mostly TIT-FOR-TATers. How do a few defectors fare in a society of mostly TIT-FOR-TATers? How do the TIT-FOR-TATers fare? How would a few cooperators fare with the TIT-FOR-TATers? Would the evolution of such a society favor cooperation or defection?

The TIT-FOR-TAT strategy is particularly interesting in an iterated Prisoner's Dilemma. It has a few special characteristics that lead to success. First it is *responsive* in that it responds to the strategy of the other player. If the other player is cooperating, the TIT-FOR-TAT strategy will be able to gain the 3 points on each round. If the other player is defecting, it will defect so as

to not keep getting the sucker's payoff of 0. The random strategies and pure strategies, for example, are not responsive. They do not respond to how the other player is playing. Chances are when you played in the tournament, you wanted to be able to adapt your strategy to respond to how your opponent was playing.

The TIT-FOR-TAT strategy is also *nice* in that it starts by cooperating. If it meets another cooperator it will continue to cooperate. If the opponent at some point begins cooperating, it will, too. However, it is also *unexploitable*. This means that a defector cannot take advantage of the “niceness.” It “punishes” any defection with an in-kind defection.

The TIT-FOR-TAT behavior has been observed in some animal populations, but you also might be able to think of situations in your own life where you or your friends have employed such a strategy with each other! Has it been effective for you?

4.8 Popular Culture: Prisoner's Dilemma and Chicken

In this section, we will look at applications of Prisoner's Dilemma and Chicken in popular culture.

The movie *Return to Paradise* (1998) explores a Prisoner's Dilemma throughout the film. The two main characters, Tony and Sheriff, must decide if they will cooperate by returning to Malaysia to serve time in prison, or defect by not returning to Malaysia. If both defect, their friend will die in prison. If both cooperate, their friend will be released and they will each serve short sentences.

Question 4.8.1 Give a payoff matrix to model the Prisoner's Dilemma in the film. By the end of the film have the payoffs changed? Is it still a Prisoner's Dilemma? Explain.

Question 4.8.2 In the classic Prisoner's Dilemma, communication is not allowed between the players. In the film, Tony and Sheriff can communicate all they want. How does this communication impact the Prisoner's Dilemma. Does it help or hinder their choice of strategy? Explain.

The movie *Rebel Without a Cause* (1955) contains an iconic Chicken scene, in which the two characters race towards a cliff. The last one to jump out of his car is declared the winner.

Question 4.8.3 Does Jim win or lose the game of Chicken? Explain your answer.

Question 4.8.4 The movie *Footloose* (1984) also has a Chicken scene (this time with tractors). Compare the Chicken scenes in *Rebel* and *Footloose*. Is the Chicken game used similarly in each? In both scenes, one player has no choice of strategy. Why might the writer have made this choice in each of these films?

Now try to apply the ideas of rationality and perfect information to your own popular culture examples.

Question 4.8.5 Suppose players are allowed to communicate in a Prisoner's Dilemma. Explain the relationship between trust and communication in a Prisoner's Dilemma. Give an example from a film demonstrating the relationship.

Question 4.8.6 Why might a writer include a Chicken scene in a film? What key attributes might the director be trying to display about the winner of Chicken and the loser? Use an example from popular culture to demonstrate your answer. \square

Question 4.8.7 One of the interesting comparisons between Prisoner's Dilemma and Chicken is that in a Prisoner's Dilemma is with their equilibrium points. Players in a Prisoner's Dilemma can reach an equilibrium point if they play the *same* way. Players in a Chicken game can reach an equilibrium if they choose *different* strategies. Find examples of how these games are used in popular culture to emphasize differences between characters. \square

References

- [1] Roy Gardner, *Games for Business and Economics*, Wiley, 2003.
- [2] Douglas Hofstadter, *Metamagical Themas: Questing for the Essence of Mind and Pattern*, Basic Books, 1985.
- [3] Jennifer Nordstrom, “Battles of Wits and Matters of Trust: Game Theory in Popular Culture”, *Mathematics and Popular Culture: Essays on Appearances in Film, Fiction, Games, Television and Other Media*, eds. E. Sklar and J. Sklar, McFarland, 2012, pp. 86-98.
- [4] William Poundstone, *Prisoner’s Dilemma: John von Neumann, Game Theory, and the Puzzle of the Bomb*, Anchor Books, 1993.

Index

- Dr. Strangelove*, 27
- Footloose*, 71
- Rebel Without a Cause*, 71
- Return to Paradise*, 71
- Sage*, 39
- The Princess Bride*, 27

- assumptions, 3
- augmented matrix, 48, 49

- Battle of the Sexes, 54

- Cake Division, 2
- Chicken, 59, 71
- Class-Wide Prisoner's Dilemma, 61
- Competitive Advantage, 13
- constant-sum game, 8
- cooperate, 61
- cooperation, 64
- cooperative game, 58

- defect, 61
- defection, 64
- dominated by, 15
- dominates, 15

- equally likely events, 19
- equilibrium
 - mixed strategy, 37
- equilibrium pair, 11, 24
- equilibrium point, 24
- expected payoff, 34
- expected value, 20, 23, 30, 43

- finite game, 4

- game matrix, 5
- Generalized One-Card Stud Poker, 24
- graphical solution, 32

- infinite game, 4
- iterated Prisoner's Dilemma, 68

- Liar's Poker, 46

- Matching Pennies, 5, 43
- maximin mixed strategy, 36, 41
- maximin strategy, 18
 - graphical solution, 32
 - repeated games, 35
- maximin strategy, repeated games, 41
- minimax, 18
- minimax strategy
 - repeated games, 37
- mixed strategies, expected value solution, 43
- mixed strategy, 30
 - graphical solution, 32
- mixed strategy equilibrium, 37

- net gain, 52
- non-cooperative game, 58
- non-cooperative games, 56
- non-zero-sum game, 54

- One-Card Stud Poker, 22

- payoff, 3, 5
- payoff vector, 5
- perfect information, 4, 27
- perfectly logical, 6
- player, 2
- popular culture, 27, 71
- Prisoner's Dilemma, 59, 71
- probability, 19
- proof by contradiction, 24
- pure strategy, 30

- rational, 3
- rationality, 27
- Rock-Paper-Scissors, 12, 30, 46

- self-interest, [3](#)
- self-interested, [6](#)
- solution, [4](#)
- Solution Theorem for Zero-Sum Games, [24](#)
- solution to two player game
 - graphical, [32](#)
- strategy pair, [9](#)
- strategy, mixed, [30](#)
- strategy, pure, [30](#)
- strategy, repeated game, [29](#)
- symmetric game, [52](#)
- systems of equations, [48](#)
- Tic Tac Toe, [4](#)
- Tit for Tat strategy, [69](#)
- Undercut, [51](#)
- Volunteer's Dilemma, [66](#)
- zero-sum game, [8](#)

Colophon

This text was authored in PreTeXt.