CCICADA

Command, Control, and Interoperability Center for Advanced Data Analysis

A Department of Homeland Security Center of Excellence

Competition or Collusion?
Game Theory in Security, Sports, and Business

A CCICADA Homeland Security Module

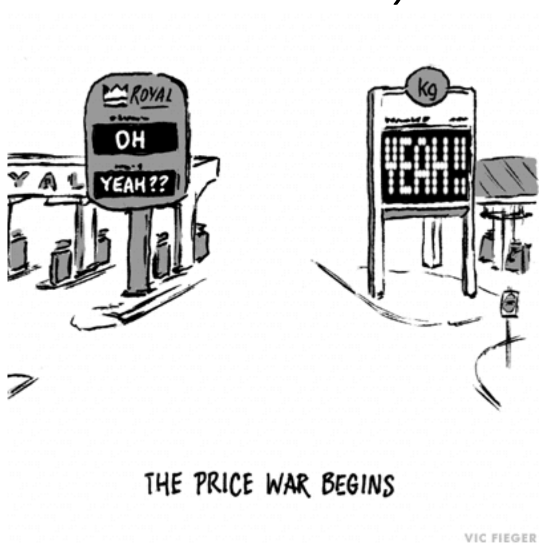

James Kupetz, Luzern County Community College kupetzj@comcast.net

Choong-Soo Lee, St. Lawrence University clee@stlawu.edu

Steve Leonhardi, Winona State University sleonhardi@winona.edu

Competition or Collusion?
Game Theory in Security, Sports, and Business

<u>Note to teachers</u>: Teacher notes appear in dark red in the module, allowing faculty to pull these notes off the teacher version to create a student version of the module.

Module Summary

This module introduces students to game theory concepts and methods, starting with zero-sum games and then moving on to non-zero-sum games. Students learn techniques for classifying games, for computing optimal solutions where known, and for analyzing various strategies for games in which no optimal solution exists. Finally, students have the opportunity to transfer what they've learned to new game-theoretic situations.

Prerequisites

Students should be able to use the skills learned in High School Algebra 1, including the ability to graph linear equations, find points of intersection, and algebraically solve systems of two linear equations in two unknowns. Knowledge of basic probability (such as should be learned by the end of 9th grade) is also required; experience with computing expected value would be helpful, but can be taught as part of the module. No computer programming experience is required or involved, although students with some programming knowledge may be able to adapt their knowledge to optional projects.

Suggested Uses

This module can be used with students in grades 10-14 in almost any class, but is best suited to students in mathematics, economics, political science, or computer science courses. For post-secondary students, the module can be used in either Mathematics Survey or Math for Liberal Arts courses.

Module Structure: Units, Suggested Timeline, Materials, and Objectives

The suggested times refer to the number of 45-minute class periods. These times are intended as rough guides and will need to be adapted to the length and format of your classes and to the levels of your students. The time needed to teach particular units will also depend on how you choose to use the activities in any given unit as some can be used equally effectively as homework or as in-class work.

Unit/Time	*Topic*
Unit 1 1 day	Zero-Sum Games: What's In It For Me? • Rock, Paper, Scissors • The Matching Game • Fair, Balanced, and Unbalanced Games • Computing Expected Value
Unit 2 1 day	Zero-Sum Games: Saddle Points for Determined Games • A Modified RPS Game • *Dominated Strategies (optional)* • Maximin, Minimax, and Saddle Points • Soccer Penalty Kicks • *Modified PK's (optional)*
Unit 3 1 day	Zero-Sum Games: What's Your Best Strategy? • Calculating the Optimal Mixed Strategy • Oddments (optional)
Unit 4 1 day	Non-Zero Games: Pranks and Prisoners • Prank on the Dean (Prisoner's Dilemma) • Split or Steal (Golden Ball) • Price Fixing and Price Wars • Weird Haircuts
Unit 5 1 day	Non-Zero Games, Continued: Halo or Just Dance? • Halo or Just Dance? (Battle of the Sexes) • Chicken • Finding Dates
Unit 6	Applications of Game Theory to Homeland Security and Other Topics

out of class Homework plus 1 day	• Air Strike or Diplomacy? • Hunting Down Terrorists (Stag Hunt) • *The Hunger Games:* Swallow Berries or Bluff? • Tragedy of the Commons • Leader or Follower? • Applying for Jobs • Union Negotiations

*Unit 1 Objective:*Students will be able to interpret and identify the appropriate strategies and payoff matrix used to model a zero-sum game, and compute the expected value of the game for given combinations of mixed strategies.

Unit 2Objective: Students will be able to identify the maximin and minimax outcomes for a zero-sum game, and identify any saddle point(s) for a strictly determined game.

Unit 3 Objective: Students will be able to calculate and interpret the optimal mix of strategies (Nash equilibrium) and expected value for a zero-sum game, using one or more of the following methods: technological assistance (phone app or web site), or algebraic formula, or graphical derivation, or “oddments.”

Unit 4 Objective: Students will be able to identify, interpret, and analyze the appropriate strategies and payoff matrix used to model a non-zero-sum game—especially in the Prisoner’s Dilemma and other games in which players can achieve their best possible mutual outcome by cooperating with the other player.

Unit 5 Objective: Students will be able to identify, interpret, and analyze the appropriate strategies and payoff matrix used to model a non-zero-sum game—especially in the Battle of the Sexes and other repeatedgames in which players can achieve their best possible mutual outcome by alternating between their first preference and the other player’s first preference.

Unit 6 Objective: Students will be able to adapt the skills they have developed in earlier units tomodel, classify, and analyze new game-theoretic situations.

Before beginning the module:

The day before you begin the module, ask your students to find an article or web page discussing the applicability of game theory to some topic that they find interesting, such as:

- Sports (especially baseball, football, soccer, tennis, or basketball)
- Economics and business
- Fashion Trends
- Politics
- Biology (especially evolutionary biology)
- Warfare
- Homeland Security and Law Enforcement
- Popular Culture (e.g., The Big Bang Theory, Arrested Development, etc.)
- Social Media Marketing (e.g., trying to increase one's number of followers)

Introduction

"Every day's a competition…"—Storyhill, Ballad of Joe Snowboard

We take part in some form of competition every day. Students compete for higher grades in their classes. Businesses compete for higher profits and larger market share. Politicians compete to win elections, and then compete to have their policies enacted. Most of us compete in sports and games just for fun in our free time, or enjoy watching professional teams compete.

We usually think of defeating our opponent as the goal of competition. Yet, there are many situations in which the "players" may be better off if they cooperate rather than simply trying to beat each other. For example, suppose a small town has only two gas stations located across the street from each other. Whichever station has the lower price will usually attract more customers, and will therefore gain higher sales and total profit than the station with the higher price. This may lead to a "price war" (both stations repeatedly lowering prices as they attempt to offer a lower price than the other station) in which profits for both stations are lower than they would have been, had they cooperated.

http://24.media.tumblr.com/tumblr_lroqx2xlRT1qg19pgo1_400.gif

However, if the two gas station managers can agree to stick to an equal higher price, they should receive roughly the same sales at the higher price (assuming their customers do not go to a third gas station), and therefore both increase their profits.

Of course, such "price-fixing" hurts customers, and for that reason it is often prohibited by law. Cooperation between businesses (or two players in any "game") that is harmful to their customers (or any third party) is called "collusion."

http://espn.go.com/nfl/story/_/id/9291834/nfl-mailbag-pro-bowl-shaky-ground

Such collusion really does occur, and often leads to multi-million dollar lawsuits.
For example, the National Football League player's union recently accused the 32 NFL team owners of secretly agreeing to an unofficial salary cap during the uncapped 2010 NFL season. Google has filed a complaint with the European Commission alleging that Microsoft and Nokia colluded to drive up handset prices together with Canadian IP licensing firm MOSAID. As a third example, gasoline prices on the US West Coast are often higher than the rest of the country, possibly because there is very little gasoline brought in from other parts of the country.

In this module, we will develop general-purpose concepts and computational methods to answer questions such as:

- How can we describe a particular competition or conflict between two or more players using a game matrix as a mathematical model? (The competition could be athlete vs. athlete, business vs. business, nation vs. nation, management vs. employees, politician vs. politician, etc.)

- How can we predict the expected average outcome when the competition consists of several repeated rounds, if we know the strategies or mixture of strategies that each player will use?

- How can we compute optimal strategies for each player, and the resulting expected outcome of the conflict?

- If we cannot compute optimal strategies, how can we identify possible strategies for each player and analyze the advantages and disadvantages of each such strategy?

- When are the players better off cooperating than competing, and under what circumstances will they be able to maintain their cooperation?

- How can we apply the concepts and methods of game theory to a wide variety of situations in sports, economics, homeland security, politics, and biology?

Unit 1: Zero-Sum Games: What's In It For Me?

Rock, Paper, Scissors: A Familiar Example

We can classify games into three different types: probability games, combinatorial games, and strategic games. Probability games might involve rolling dice or spinning a wheel; the winner or payoff relies entirely upon chance. Combinatorial games are challenging to win becauseeven though we know our opponent's prior moves and possible follow-up moves, the number of possible sequences of moves and countermovesis too enormous to evaluate every possible sequence. Examples of combinatorial games include checkers, chess, and Connect Four. Strategic games are situations in which each player secretly chooses a strategy from a number of two or more options and simultaneously reveals their choice. The options might be whether to run or pass for a football team, whether to throw a fastball or changeup for a pitcher, or whether to reduce prices or keep prices higher for a business. Strategic games are challenging because neither player knows which option the opponent is going to choose. (Of course, some games involve two or even three of these aspects; for example, poker involves both probabilities—what cards are dealt—and strategy—bluffing or folding.)

This module will focus solely on strategic games. We will assume throughout that the players choose their strategies secretly and reveal them simultaneously each round.

(Note to instructor: Similar computational thinking modules and curriculum resources focusing instead on probabilistic or combinatorial games are available from VCTAL at http://dmac.rutgers.edu/VCTALand elsewhere. Other modules focusing on applications of game theory to Homeland Security may be found at http://www.ccicada.org/education.)

One simple example of a strategic game is Rock, Paper, Scissors. In each round, the two players each secretly choose one of the three options (we will call each of these options a "strategy"). If the two players have chosen identical strategies, that round is a draw. Otherwise, the winner is determined as follows: Rock breaks scissors, scissors cuts paper, and paper covers rock. We may represent these outcomes in the form of the table below, which we will call a "payoff matrix." We will use positive values to represent wins for the row player Robin, and use negative values to represent losses for the row player Robin (which will be wins for the column player Cal). Thus, the row player tries to build up a positive cumulative score as large as possible, and the column player tries for a cumulative score as far negative as possible.

(Robin, Cal)	Rock	Paper	Scissors
Rock	0	-1	+1
Paper	+1	0	-1
Scissors	-1	+1	0

Discussion Questions, Unit 1, Set 1.

1.1. A. Suppose that you and an opponent were to play repeated rounds of this game. (Don't actually do it; we will soon consider a more interesting variation of this game.) If you played 100 rounds, roughly what percentage of the rounds do you think you and your opponent would choose each of the three options (Rock, Paper, or Scissors)? Why?

Note on Terminology: The term "strategy" is sometimes used to mean the option chosen in a single round of play (such as "Rock"), and sometimes used to mean the overall fraction or percentage with which each option is (randomly) chosen over several repeated rounds. For example, if we choose 20% Rock, 30% Paper, and 50% Scissors, we could say we are using "mixed strategy" [0.2, 0.3, 0.5]. This ambiguous usage of the term "strategy" is unfortunate but standard in game theory books and articles. The intended meaning will usually be clear from the context. We will sometimes use the non-standard term "metastrategy" to indicate the overall mixture or percentages of strategies that is used over several repeated rounds.

1.1.A. Response: choose each strategy 1/3 of the time. Use symmetry of the payoffs. Try to randomize choices so that the option chosen in any particular round is unpredictable. We could say that our "metastrategy" is [⅓, ⅓, ⅓].

1.1.B. If you played 100 rounds, which player do you think would come out ahead in total points (or would you expect to come out approximately even)?Does either player have an advantage over the other?

1.1.B. Response: We would expect a draw about one-third of the rounds, and each player to win approximately one-third of the rounds, for a total score near zero. Depending upon the previous knowledge of the class, if it comes up in discussion, you could discuss the concept of "expected value" and discuss the concept of a "fair" game (expected value of zero); otherwise, you may want to wait until this concept is presented later in the module. If each player uses metastrategy[⅓, ⅓, ⅓], then the expected value of the game is

(1/3)*(1/3)*(0) + (1/3)*(1/3)*(-1) + (1/3)*(1/3)*(+1)
+ (1/3)*(1/3)*(+1) + (1/3)*(1/3)*(0) + (1/3)*(1/3)*(-1)
+ (1/3)*(1/3)*(-1) + (1/3)*(1/3)*(+1) + (1/3)*(1/3)*(0) = 0.

(These concepts and this calculation will be presented formally within the module in a few pages. In the meantime, students can access http://myslu.stlawu.edu/~clee/rps/ to simulate their strategies.)

1.1.C. Suppose you noticed that your opponent played RPSRPS or some other detectable pattern during the first 6 rounds of the game. What mistake might your opponent be making? How could you take advantage of this mistake?

1.1.C. Response: If we can predict which option our opponent will play in any round, then we can choose an option that gives us the best possible outcome in response to our opponent's play. For example, if we knew our opponent was going to play RPS the next three rounds, then we could play PSR and win each of these rounds. This demonstrates that it is not enough to use the optimal percentages for each option in a mixed strategy; we must also randomize our play, even though we are playing each strategy with a fixed percentage or probability (such as ⅓ or 33.3% each). We could use a spinner, a random number generator, a die, or some other randomizing method to help us avoid any detectable pattern in our play.

The Matching Game (Two-Finger Morra). This game is somewhat similar to Rock-Paper-Scissors, except that each player has just two rather than threepossible options. Each round, Player A (Alice) and Player B (Bob) each simultaneously show either one finger or two fingers. (Alternately, they could each show "Heads" or "Tails" on a coin that they place down on the table, or choose "Left" or "Right," "Red" or "Black," etc.) This game, known as "Morra," dates back to ancient Roman times. Different versions are still played around the world, especially in Italy; the game is sometimes used as a betting game, or to settle disputes.

The standard ("balanced" or "symmetric")two-finger version of the game follows these rules: If both players show the same number of fingers, Alice wins 1 point (and Bob loses 1 point; assume that any gain for Alice is a loss for Bob, and vice versa). If one player shows one finger and the other player shows two fingers, then Bob wins 1point (so Alice loses 1 point).

Discussion Questions, Unit 1, Set 2.

1. 2.A. Write down the payoff matrix for the "balanced" matching game.

1.2.A. Response:

(Alice, Bob)	One finger	Two fingers
One finger	+1	-1
Two fingers	-1	+1

1.2.B. Why should each player use a mixed strategy rather than a pure strategy?

1.2.B. Response: If one player makes their option known ahead of time (or if the other player can guess correctly), the other player can use that knowledge to always "win" each round (the row player will match, or the column player will avoid a match).

1.2.C. With what percentage would you expect each player to choose each option?

1.2.C. Response: By the symmetry of the payoffs, we would expect each player to choose each option with equal probability, that is, ½ of the time.

1.2.D. Does either player have an advantage over the other in this game, or would you predict that they would come out about even after some number of repeated rounds? Explain.

1.2.D. Response: This appears to be a fair game. We can easily calculate that if each player chooses each outcome with probability ½, then the expected value of the game is zero. Using a process of analysis that we will develop in Unit 3, it can be shown that our intuition is correct in this game—it is a fair game.

Fair Games, Balanced Games, and Unbalanced Games

The standard version of the matching game given above is “balanced” or “symmetric” in the sense that the amount Alice wins with a match (1 point) is equal to the amount Bob wins if they don’t match. Informally, we think of this kind of game as “fair.”

Later on, we will develop a more formal and more technical definition of the notion of a “fair” game, but for now we will use the following informal, intuitive definition: we call a game “fair” if neither player has an advantage over the other, that is, if we would predict the two players to come out about even or with about the same number of points as each other after a number of repeated rounds. Note that in a zero-sum game, this would mean that each player would have a cumulative score of about zero points, since they would have won about the same amount as they would have lost.

Quick exercise: write down a payoff matrix for a different fair game, and a payoff matrix for an unfair game.

Observe that our notion of “fairness” assumes that each player is using his or her “best” or “optimal” strategy. For example, if Alice decided to use the pure strategy of choosing One finger every single round and Bob guessed or became aware of this, then Bob would play Two fingers and win a point from Alice every round, so Bob would come out ahead. But this doesn’t mean the game isn’t fair; it just means that Alice was not using her best possible mixed strategy. By “best” or “optimal” strategy, we mean a combination of options chosen (possibly either pure or mixed) that makes that player’s average payoff per round as large as possible, assuming that the opponent is also employing their optimal strategy. (Later in this module, we will show how to calculate the “optimal” strategy for each player using the concepts of “maximin” and “minimax.”)

To make this game more interesting, let’s consider an “unbalanced” (“asymmetric”) version. Each player still shows one or two fingers each round, but the payoffs are changed as follows: If both players show one finger, Alice wins 20 points from Bob (we still assume that any gain for Alice is a loss for Bob, and vice versa). If both players show two fingers, Alice wins 1 point from Bob. If Alice shows one finger and Bob shows two fingers, Bob wins 4 points (so Alice loses 4 points); and if Alice shows two fingers and Bob shows one finger, then Bob wins 5 points from Alice. We summarize the “payoffs” in the “game matrix” below. Remember that a positive entry means a gain for Alice (and a loss for Bob), and a negative entry means a loss for Alice (and a gain for Bob).

“Unbalanced” Matching Game		Bob	
		One finger	Two fingers
Alice	One finger	+20	-4
	Two fingers	-5	+1

Activity 1: Students should pair up and quickly play 12 rounds of this game. Use Handout 1 to record each player’s choice in each round and the total points gained or lost over the 12 rounds.

Discussion Questions, Unit 1, Set 3.
1.3. A. What percentage of the rounds did each player choose one or two fingers? If you played 100 rounds, roughly how many times do you think each player would choose each of the two strategies? Why?

1.3.A. Responses will vary widely. A common response may be to choose each option 1/2 of the time, randomizing choices so that the option chosen in any particular round is unpredictable. We could say that our “mixed strategy” or “metastrategy” is [½ , ½]. Or, Alice might choose one finger more than half the time hoping for a payoff of +20, and Bob might choose one finger less than half the time hoping to avoid the payoff of +20 for Alice.

1.3.B. How did the total points you gained or lost over the 12 rounds compare to what you thought would happen? If you played 100 rounds, which player do you think would come out ahead (or would you expect to come out approximately even)? Is this a “fair” game, or does one player have an advantage over the other? Is it possible for this to be a fair game even though the payoff matrix is not symmetric and the payoffs do not add to zero? Explain.

1.3.B. Response: This question could be discussed at a very informal level or at a more advanced level, depending upon the previous knowledge of the class regarding the concept of “expected value” and a “fair” game (expected value of zero).

The following calculations could be carried out right away if the class is more knowledgeable, or could wait until after the class has gone over the presentation of these concepts in the next section of this Unit 1 if the class has less knowledge.
If each player employs the mixed strategy [½ , ½], the expected value per round would be as shown below:

$$\left(\frac{1}{2}\right)\left(\frac{1}{2}\right)(+20)+\left(\frac{1}{2}\right)\left(\frac{1}{2}\right)(-4)+\left(\frac{1}{2}\right)\left(\frac{1}{2}\right)(-5)+\left(\frac{1}{2}\right)\left(\frac{1}{2}\right)(+1)$$
$$=\left(\frac{1}{4}\right)(+20-4-5+1)=+3.$$

That is, if each player chooses each option ½ of the time, Alice should come out ahead by about 3 points per round.
In general, if Alice uses mixed strategy [p , 1-p] and Bob uses mixed strategy [q , 1-q], with $0 \leq p \leq 1$ and $0 \leq q \leq 1$, then the expected value per round would be:

$$(p)(q)(+20) + (p)(1-q)(-4) + (1-p)(q)(-5) + (1-p)(1-q)(+1)$$
$$= 30pq - 5p - 6q + 1$$

1.3.C. Next, Bob will receive some computational assistance from a game theory app, but Alice will not. The instructor should make assistance available to Bob, using either the supplementary phone app, or a web site, or simply a written note. Alice and Bob then play 12 more rounds of the game, recording their choices each round and the outcomes as before, on page 2 of Handout 1.

1.3.C. Response: Bob should choose one finger q = 1/6 of the time, and choose two fingers 1-q = 5/6 of the time, randomizing his or her choices. This simplifies the expected value of the game to 30pq-5p-6q+1 = 30p(1/6)-5p-6(1/6)+1 = 0. That is, if Bob uses mixed strategy [q , 1-q]=[1/6, 5/6], then the expected value of the game is 0 regardless of what mixed strategy Alice uses.

1.3.D. What mix of strategies gave the best results for Bob? What mix of strategies gave the best results for Alice? Survey the class and compare results.

1.3.D. Responses will vary widely. Observe that Bob's outcome is much better if he chooses one finger q = 1/6 of the time (for an expected value of 0 regardless of Alice's mixed strategy [p, 1-p]), and much worse if Bob chooses one finger q = 1/2 of the time (for an expected value of 30p(1/2)-5p-6(1/2)+1 = 10p - 2 depending on Alice's mixed strategy [p, 1-p]).

1.3.E. Collect the results from the entire class and jointly fill in the table on the third page of Handout 1. Are the class results closer to what you expected than your individual group results? Explain.

1.3.E. Response: We would expect most total scores near zero but a few significantly positive and a few significantly negative. (Probability distribution with mean zero.) Discuss the concept of a "fair" game (expected value of zero).

1.3.F. On a blackboard/whiteboard and/or on the fourth page of Handout 1, create a frequency graph of the total overall average points per round from the different pairs of players in the class. What do you observe about the mean and the distribution of the results in the two graphs?

1.3.F. Response: We would expect most total scores near zero but a few significantly positive and a few significantly negative. (Probability distribution with mean zero.) Discuss informally the concept of a "fair" game (expected value of zero).

Computing the Expected Value of a Game

In the Rock, Paper, Scissors game, you probably guessed that each player should play each strategy ⅓ of the time, and this is a "fair" game in the sense that we would expect the overall net value of the game to be approximately zero points to each player, if we play a large enough number of rounds, since the number of wins for the Row Player should be roughly equal to and therefore cancelled out by the wins for the Column player.

We can formalize the notation of a “fair” game by first defining the **expected value** of a game. In the RPS game, each player has three possible options to choose from, and the choice of the Row Player is “independent” of the choice of the Column Player (since they choose their options secretly and reveal simultaneously round by round), for a total of nine possible (independent and mutually exclusive) combinations of options chosen. Each of these possible combinations (or “events”) $\mathrm{E_i}$ has a probability $\mathrm{p_i} = \mathrm{prob(E_i)}$ of occurring, and an associated payoff or outcome $\mathrm{y_i}$.

We may calculate the expected value of a game relative to any particular combination of mixed strategies used by each player. Usually no reference to a particular mixed strategy is given, and in this case we will assume that each player is using the “optimal” mixed strategy (in the sense of maximizing their personal payoff, ignoring any gain or loss for other players) when calculating the expected value of a game.

More generally, suppose that based on the mixture of strategies (or “metastrategy”) that each player uses, we have determined that there are n possible (independent and mutually exclusive) events or combinations of options, with respective probabilities $\mathrm{p_1}$, $\mathrm{p_2}$, …,$\mathrm{p_n}$ and corresponding payoffs $\mathrm{y_1}$, $\mathrm{y_2}$, …, $\mathrm{y_n}$. Then the **expected value** of the game (or **average points per round**) is defined as the sum $\mathrm{E} = \mathrm{p_1y_1 + p_2y_2 + \ldots + p_ny_n} = \sum_{\mathrm{i=1}}^{\mathrm{n}} \mathrm{p_iy_i}$.

If we assume that Robin plays each strategy with probability ⅓ and Cal also plays each strategy with probability ⅓, then each of the nine possible combinations of strategies will occur with probability (1/3)*(1/3) = 1/9. The expected value of the game under this combination of metastrategies will be

$$\begin{gathered}\left(\frac{1}{3}\right)\left(\frac{1}{3}\right)(0)+\left(\frac{1}{3}\right)\left(\frac{1}{3}\right)(-1)+\left(\frac{1}{3}\right)\left(\frac{1}{3}\right)(+1)\\+\left(\frac{1}{3}\right)\left(\frac{1}{3}\right)(+1)+\left(\frac{1}{3}\right)\left(\frac{1}{3}\right)(0)+\left(\frac{1}{3}\right)\left(\frac{1}{3}\right)(-1)\\+\left(\frac{1}{3}\right)\left(\frac{1}{3}\right)(-1)+\left(\frac{1}{3}\right)\left(\frac{1}{3}\right)(+1)+\left(\frac{1}{3}\right)\left(\frac{1}{3}\right)(0)=0\end{gathered}$$

Note how the rectangular table structure of this computation corresponds to the entries in the game matrix, especially after we fill in the probabilities that each player use to randomly choose the various possible options, as we have done below.

(Robin, Cal)	Cal:	Rock	Paper	Scissors
Robin:		⅓	⅓	⅓
Rock	⅓	0	-1	+1
Paper	⅓	+1	0	-1
Scissors	⅓	-1	+1	0

Suppose instead that Robin uses mixed strategy [0.2, 0.3, 0.5] and Cal uses mixed strategy [0.4, 0.5, 0.1], as shown in the game matrix below.

(Robin, Cal)	Cal:	Rock	Paper	Scissors
Robin:		0.4	0.5	0.1
Rock	0.2	0	-1	+1
Paper	0.3	+1	0	-1
Scissors	0.5	-1	+1	0

Then the expected value would be

(0.2)*(0.4)*(0) + (0.2)*(0.5)*(-1) + (0.2)*(0.1)*(+1)
+ (0.3)*(0.4)*(+1) + (0.3)*(0.5)*(0) + (0.3)*(0.1)*(-1)
+ (0.5)*(0.4)*(-1) + (0.5)*(0.5)*(+1) + (0.5)*(0.1)*(0) = +0.06.

This is only a slight advantage to the Row Player, but if the game is repeated a large number of times, say 100 rounds, then we would expect the Row Player to come out ahead with a net score of approximately (+0.06 points per round)(100 rounds) = 6 points advantage over the Column Player.

Side note: Casinos enjoy huge profits based on exactly this scenario. For example, in roulette, the house advantage is just 5.26%. Another way of saying this is that the expected value of 100 bets of $1 each (or one bet of $100) is -5.26, that is, the bettor should expect to lose about $5.26 per $100wagered, on average, over a large number of spins. It doesn’t hurt one bettor on vacation much to lose about 5% of the amount they wager, but with millions of bettors, the casino makes billions of dollars!

Discussion Questions, Unit 1, Set 4.
These questions refer to the game represented by payoff matrix:

(Robin, Cal)	Cal:	Option V	Option W
Robin:		q	1-q
Option A	0.1	+1	-1
Option B	0.2	-2	+2
Option C	0.7	-3	+3

1.4.A. Compute the expected value of the Rock, Paper, Scissors game if Robin uses mixed strategy $[0.1, 0.2, 0.7]$ and Cal uses mixed strategy $[0.6, 0.4]$.

1.4.A. Response: The expected value would be

(0.1)*(0.6)*(+1) + (0.1)*(0.4)*(-1)
+ (0.2)*(0.6)*(-2) + (0.2)*(0.4)*(+2)
+ (0.7)*(0.6)*(-3) + (0.7)*(0.4)*(+3) = -0.48 per round.

1.4.B. Suppose that Cal somehow learns that Robin is using mixed strategy $[0.1, 0.2, 0.7]$. What mix of strategies should Cal use in order to optimize the overall payoff to Cal? Hint: use unknown variable q to represent the probability that Cal chooses option V. Explain why the probability the Cal chooses option W must be 1-q. Then write out the expected value in terms of q, simplify this expression and graph expected value as a function of q, and then determine the allowed value of q that results in the least (i.e., most negative) possible expected value for Cal.

1.4.B. Response: The probabilities with which Cal chooses options V and W must add to 1 or 100%, so if the probability of choosing V is q, then the probability of choosing W must be 1-q.

Sophistication of student responses may vary widely with their background. Students might intuitively guess that since Robin is playing Option C most frequently, therefore Cal should choose Option V most frequently in order to make the expected payoff as far negative as possible.

The expected value in terms of q is:
(0.1)*(q)*(+1) + (0.1)*(1-q)*(-1)
+ (0.2)*(q)*(-2) + (0.2)*(1-q)*(+2)
+ (0.7)*(q)*(-3) + (0.7)*(1-q)*(+3) = -2.4(q)+2.4(1-q) = -4.8q+2.4 per round.

Since $0 \leq q \leq 1$ and the expected value is decreasing (as Cal desires) as q increases, therefore Cal should use q =1. That is, Cal should choose Option V 100% of the time and never choose

Option W (assuming that Robin does not waiver from the fixed strategy [0.1, 0.2, 0.7]). This leads to an expected payoff of -2.4 per round.

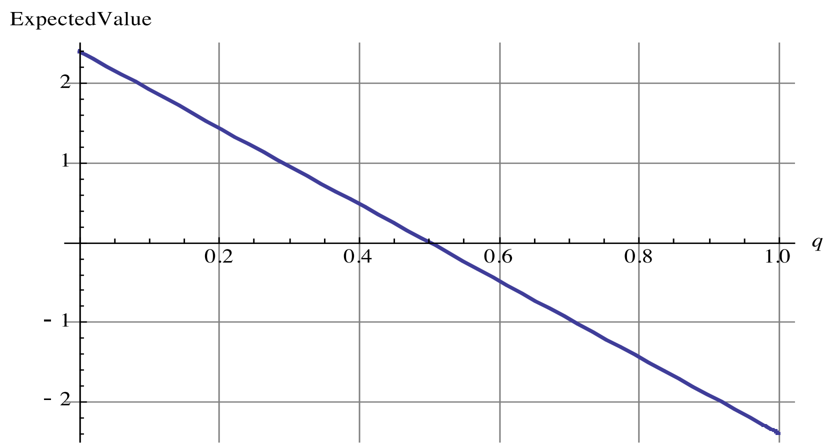

Unit 2: Zero-Sum Games: Saddle Points for Determined Games

A Modified Version of Rock, Paper, Scissors

Activity 2: Now play 12 more rounds of the Rock, Paper, Scissors game, with the modification that the column player is allowed to choose either Rock or Paper but cannot choose Scissors. (The row player is still allowed to choose any of the original three options.) Use Handout 2 to record the results for later reference.

Discussion Questions, Unit 2, Set 1.

2.1.A. Write down the payoff matrix in this modified game.

2.1.A. Response:

(Robin, Cal)	Rock	Paper
Rock	0	-1
Paper	+1	0
Scissors	-1	+1

2.1.B. How often should the row player choose the "Rock" strategy? Explain. How can we "edit" the payoff matrix to reflect this?

2.1.B. Response: The row player should never choose Rock, because it is "dominated" by the Paper strategy. Discuss the concept of a "dominated strategy." We can delete the "Rock" row for the row player to obtain edited payoff matrix below.

(Robin, Cal)	Rock	Paper
~~Rock~~	~~0~~	~~-1~~
Paper	+1	0
Scissors	-1	+1

(Robin, Cal)	Rock	Paper
Paper	+1	0
Scissors	-1	+1

2.1.C. What percentage of the rounds did each player choose each strategy?

2.1.C. Response: Most likely, each player will choose each strategy about ½ the time, more or less randomly. If anyone chose a different percentage of mixed strategies, discuss their reasons and their overall outcome.

2.1.D. How did the total points you gained or lost over the 12 rounds compare to your predictions for this game? If you played 100 rounds, which player do you think would come out ahead (or would you expect to come out approximately even)? Explain.

2.1.D. Response: The sophistication of responses will depend upon prior knowledge about probabilities (if any) and ability of the students. The most basic response would be that we would guess that the row player would come out ahead, since there are more positive than negative entries, and because the sum of the entries is positive rather than zero.

If students understand the probability concept of "expected value" introduced earlier in this unit, they should be able to carry out the following analysis: If each player chooses each of the two remaining strategies ½ of the time, then the expected value of the game will be E = (1/2)*(1/2)*(+1+0-1+1)=1/4 per round. Over 12 rounds, we would expect a cumulative result of about (1/4)*(12)=3 points for the row player. Over 100 rounds, we would expect a cumulative result of about (1/4)*(100)=25 points for the row player.

If students do not have prior knowledge of "expected value," the teacher may need to present additional examples and provide students with more practice problems.

2.1.E. Which game is more "fair"—the original version of Rock, Paper, Scissors, or the modified version in which the column player has only two options? Explain.

2.1.E. Response: The original game is "fair" (in the sense of having expected value of zero for each player). The modified game does not appear to be fair, since the expected value is +1/4 (advantage to the row player) per round if each player chooses each of his or her two viable options ½ of the time.

2.1.F. Write out the payoff matrix for the game "Rock Paper Scissors Lizard Spock" played by Sheldon and Raj on "The Big Bang Theory" as demonstrated in the video and graphic below. (This game was invented by Sam Kass with Karen Bryla, according to http://www.samkass.com/theories/RPSSL.html .) What percentage of the time would you expect each player to select each of the five strategies? Does this game appear to be fair? Explain your reasoning.

http://www.youtube.com/watch?v=iapcKVn7DdY

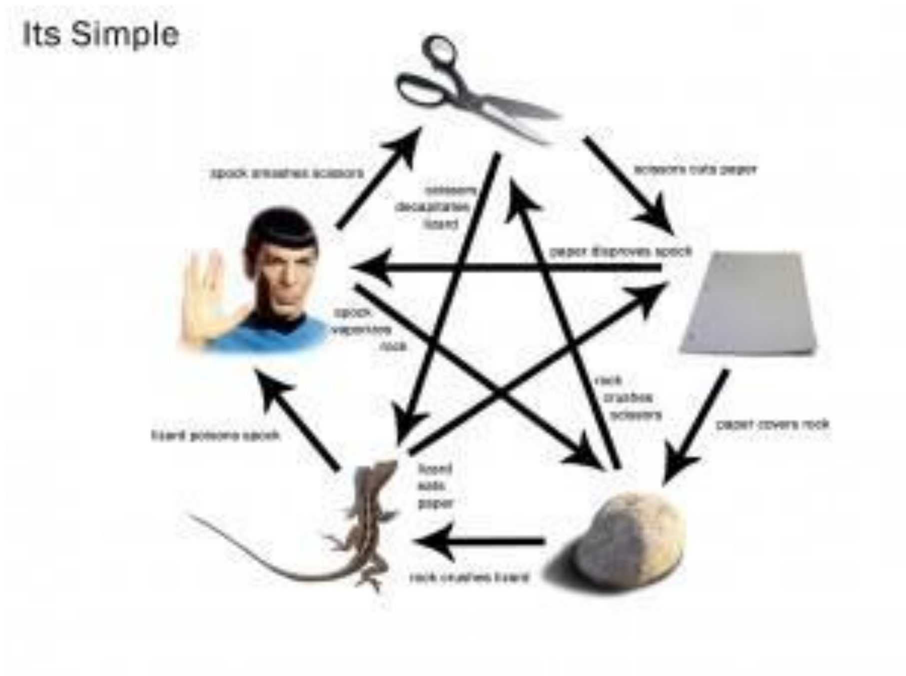

http://wordpress.morningside.edu/cdl001/rpsls/

2.1.F. Response:

(Robin, Cal)	Rock	Paper	Scissors	Lizard	Spock
Rock	0	-1	+1	+1	-1
Paper	+1	0	-1	-1	+1
Scissors	-1	+1	0	+1	-1
Lizard	-1	+1	-1	0	+1
Spock	+1	-1	+1	-1	0

Dominated Strategies (optional but recommended section)

In the modified Rock-Paper-Scissors game of Activity 2 (in which the Column Player is not allowed to choose Scissors), we find that there is no reason for the Row Player to choose Rock, since choosing Paper leads to an outcome for the Row Player that is sometimes better and never worse than the outcome when choosing Rock, that is, the entry in the Paper row is greater than or equal to the payoff in the Rock row (from the point of view of the Row Player) in each column. In this case, we say that the Paper strategy **dominates** the Rock strategy for the Row Player. We can simplify our analysis of a game if we first remove any “dominated” strategies from the game matrix.

For example, consider the abstract game represented by the payoff matrix below.

(Robin, Cal)	V	W	X	Y	Z
A	0	-3	-2	-1	+3
B	0	0	-1	-1	+3
C	-1	+1	0	+1	-1
D	-1	+1	0	0	-1
E	0	-1	+1	-1	0

Note that from Robin's point of view (choosing a row and trying to maximize the payoff), strategy A is dominated by strategy B, and strategy D is dominated by strategy C. For this reason, we would guess that Robin would never choose these strategies, and we can essentially delete rows A and D from the game matrix, leading to the reduced matrix shown below.

(Robin, Cal)	V	W	X	Y	Z
B	0	0	-1	-1	+3
C	-1	+1	0	+1	-1
E	0	-1	+1	-1	0

Now consider Cal's point of view (choosing a column and trying to minimize the payoff) after Robin's reduction. Strategy W is now dominated by Y, and strategy Z is now dominated by V so we would expect Cal never to choose these columns from the game matrix, leaving us with the newly reduced matrix below.

(Robin, Cal)	V	X	Y
B	0	-1	-1
C	-1	0	+1
E	0	+1	-1

Reconsidering from Robin's viewpoint, row B is dominated by row E in the remaining reduced matrix, so we can delete row B.

(Robin, Cal)	V	X	Y
C	-1	0	+1
E	0	+1	-1

Reconsidering from Cal's viewpoint, column X is dominated by column V, so we can delete column X, leaving the reduced matrix below.

(Robin, Cal)	V	Y
C	-1	+1
E	0	-1

We now see that no more reductions due to domination will occur, so the matrix has been reduced as far as possible.

Homework Exercises, Unit 2, Set 2.

2.2.A. Reduce the game matrix below by deleting dominated rows.

(Robin, Cal)	V	W	X	Y	Z
A	4	0	-2	-1	3
B	3	-2	-5	-1	2
C	-3	2	0	-2	-1
D	1	-4	-2	0	1

2.2.A. Response: First note that for Robin, Row A dominates Row B, so we can delete Row B:

(Robin, Cal)	V	W	X	Y	Z
A	4	0	-2	-1	3
C	-3	2	0	-2	-1
D	1	-4	-2	0	1

Then note that for Cal, Column Y dominates Column Z, so we can delete Column Z:

(Robin, Cal)	V	W	X	Y
A	4	0	-2	-1
C	-3	2	0	-2
D	1	-4	-2	0

There is no remaining domination, so the game matrix is reduced as much as possible.

2.2.B. (for computer science students) Describe an algorithm for reducing a matrix by iteratively deleting dominated rows, then dominated columns, and so on.

<u>Maximin, Minimax, Saddle Points, and Pure-Strategy Nash Equilibrium</u>

Most games are played out as a series of repeated challenges over time: a pitcher throwing to a batter, a business offering a price that will hopefully result in greater weekly market share than its competitors, and so on. You have also seen that in most competitions, the best approach is to "keep our opponent guessing" by using a mix of strategies rather than just selecting the same option each round.

In this section, we introduce a general guiding principle in choosing strategies, and we describe the special circumstances in which the best strategy is to stick with the same pure strategy each round rather than using a mixed strategy.

What do you think will happen when Robin and Cal play the game with the payoff matrix shown below?

(Robin, Cal)	V	W	X	Y
A	3	0	-2	-1
B	6	-2	-4	-1
C	-3	2	-2	-2
D	1	-8	-2	0

(If you've read the previous section, you will notice that Column Y is dominated by Column X and therefore can be deleted, but our discussion here does not require this reduction, so we will continue with the full, original matrix here.)

Robin could play aggressively by choosing row B in hopes of gaining the +6 payoff, but this also runs the risk of the -4 payoff and a big loss. Likewise, Cal could play aggressively by choosing column W in hopes of the -8 payoff (a big win for Cal), but since row D does not contain very good payoffs for Robin, it is unlikely that Robin would make this choice very often.

A different, more cautious approach would be for each player to choose the option (or mixture of options) that results in the best possible outcome from their point of view, regardless of which option their opponent chooses. For Robin, this means choosing the "maximin" (the maximum of the minimum possible outcomes in each row), and for Cal, this means choosing the "minimax" (the minimum of the maximum possible outcomes in each column). We illustrate how to calculate these in the example matrix given above.

(Robin, Cal)	V	W	X	Y	Row min	
A	3	0	-2	-1	-2	maximin
B	6	-2	-4	-1	-4	
C	-3	2	-2	-2	-3	
D	1	-8	-2	0	-8	
Col max	6	2	-2	0		
			minimax			

In this game, each row min represents the "worst that could happen" (from Robin's viewpoint) when choosing that particular row, and each column max represents the worst that could happen (from Cal's viewpoint) when choosing that particular column. Robin can guarantee that his or her payoff will never be less than -2 by always choosing the maximin row, that is, choosing row A. Cal can guarantee that the payoff will never be any greater than -2 by always choosing the minimax column, that is, choosing column X.

In this particular game, it happens that the maximin is the same value as the minimax, and this occurs when Robin chooses row A and Cal chooses column X. If both players choose cautiously, we would expect both players to select these pure strategies. That is, we would expect Robin to always choose row A; we would expect Cal to always choose column X; and the expected value of the game would be -2.

For a game in which the maximin is the same as the minimax, we call this entry in the matrix a "saddle point," and the expected value of the game is the payoff at that point. The name comes from the fact that at a saddle point, the payoff values go up as we move left or right in the matrix (just as a saddle curves upward to the front or rear), but the payoff values go down as we move up or down in the matrix (just as a saddle curves downward in the direction of either leg of the rider). We also call this entry a "pure equilibrium," in the sense that we expect players to stick with their corresponding pure strategy round after round, with a predictable payoff.

Another way of defining this notion is as follows: a "saddle point" in a game matrix is an entry that is both a minimum in its row and a maximum in its column. Such an entry may or may not exist in a game matrix. We can always calculate the maximin and the minimax for a game matrix, but it is the exception rather than the rule for a given game to have a saddle point. If a saddle point exists, it is also called a pure strategy "Nash equilibrium" in honor of John Nash, who won a Nobel Prize for Economics based on his work in Game Theory, and who was portrayed in the best-selling book and Academy Award-winning movie "A Beautiful Mind."

<u>Discussion Questions, Unit 2, Set 3.</u>

2.3.A. Find the maximin and minimax of the game. Determine whether or not the game has a saddle point. If it does, predict the players' play and the expected value of the game.

(Robin, Cal)	X	Y	Z	Row min	
A	-2	-3	5		
B	-1	0	0		
C	-2	6	-4		
Col max					

2.3.A. Response:

(Robin, Cal)	X	Y	Z	Row min	
A	-2	-3	5	-3	
B	-1	0	0	-1	maximin
C	-2	6	-4	-4	
Col max	-1	6	5		
	minimax				

The game has a saddle point at (B, X) with expected payoff -1.

2.3.B. Find the maximin and minimax of the game. Determine whether the game has a saddle point. If it does, predict the players' play and the expected value.

(Robin, Cal)	X	Y	Z	Row min	
A	4	0	-3		
B	-1	2	0		
C	1	-4	4		
Col max					

2.3.B. Response:

(Robin, Cal)	X	Y	Z	Row min	
A	4	0	-3	-3	
B	-1	2	0	-1	maximin
C	1	-4	4	-4	
Col max	4	2	4		
		minimax			

Since the maximin does not equal the minimax, therefore the game does not have a saddle point. We expect both players to use a mixture of strategies rather than a pure strategy.

Penalty Kicks in Soccer

Many sport competitions involve two opponents trying to outguess each other: pitcher versus batter in baseball or softball, where the options might be fastball or curveball; offense versus defense in football, where the options might be run or pass; and so on. These competitions are difficult to model realistically because usually a pitcher has three or four different pitches he or she might throw, and a football team has several formations and dozens of plays to run and defend against. One of the simplest situations to model is the taking of a penalty kick. Usually, the kicker will aim at either the left or right side of the net, and the goalie will either dive right or dive left at the moment the ball is kicked.

Discussion Questions, Unit 2, Set 4.

2.4.A. Write down a payoff matrix that models a penalty kick in soccer. Suppose that the striker (shot-taker) is the row player, and the goalie is the column player.

2.4.A. Response:

(Striker, Goalie)	Dive Left	Dive Right
Kick Left	0	+1
Kick Right	+1	0

2.4.B. In what ways is this game similar to the Matching Game? In what ways is this game different from the Matching Game?

2.4.B. Response: Similarities: One player (the goalie) prefers the outcome when the players' chosen option are a "match," and the other player (the shooter) prefers theoutcome when the players' chosen options do not match. Also, the optimal mixture of strategies for each player seems to be choose Left and Right each with probability 50%. Difference: The Matching Game appears to have expected value 0, which would make it a fair game, whereas the Striker-Goalie game has expected value + ½ , which is an advantage to the Striker.

2.4.C. For what reasons is this matrix an over-simplification of the actual situation? What assumptions are we making when we use this payoff matrix?

2.4.C. Response (other responses are possible):

- We are assuming that both striker and goalie have only these two options.
- We are assuming that the striker can always hit the net, given the incorrect dive by the goalie.
- We are assuming that the goalie can always block the shot, if he or she dives in the correct direction.

Refining the Model: Strong-Side versus Weak-Side Penalty Kicks (optional section)

The above model is an oversimplification of the actual situation. We could "refine" our model by making modifications that are somewhat less oversimplified (and hence more realistic), but these modifications also make the model somewhat more complicated to analyze and understand. For example, suppose that the kicker has a very high probability of scoring a goal (essentially 1 or 100%) if he or she kicks to the right and the goalie dives left, but only has probability s (where "s" is some fixed constant parameter value between 0.00 and 1.00) of scoring the goal if the kicker aims left and the goalie dives right.

Discussion Questions, Unit 2, Set 5.

2.5.A. Write down the modified payoff matrix for this situation.

2.5.A. Response:

(Striker, Goalie)	Dive Left	Dive Right
Kick Left	0	+s
Kick Right	+1	0

2.5.B. Suppose that s = 2/3. Do you think the striker should always kick one way, or use a mixed strategy? Doe you think the goalie should always dive to one side, or use a mixed strategy? What percentage of the time do you think the striker should kick left or kick right? What percentage of the time do you think the goalie should dive left or dive right?

2.5.B. Response: Since the striker has a higher probability of scoring when kicking right (if the goalie dives the wrong way), it would seem that the striker should kick right more often than kick left. We would also guess that the goalie should probably dive right more often than dive left.

We can show mathematically that the optimal strategy for the goalie is, as we expected, to dive left less often, with probability $q = 0.4$ (or 40%), and dive right more often, with probability $1 - q = 0.6$ (or 60%). However, we can also show (much to our surprise!) that the optimal strategy for the striker is to kick left with probability $p = 0.6$ (60%) and kick right with probability $1 - p = 0.4$ (40%). In general, for any value of s strictly between 0 (no chance of scoring when kicking left) and 1 (original problem), the optimal strategy for the striker is to kick left with probability $p = 1/(1+s)$, and the optimal strategy for the goalie is to dive left with probability $q = s/(1+s)$. We will later go through a method for determining these optimal probabilities. It's interesting that as the probability s of scoring when kicking left decreases, the kicker should kick left more often!

Unit 3: Zero-Sum Games: What's Your Best Strategy?

Calculating the Optimal Mixed Strategy

Remember the "unbalanced" matching game that you played in Activity 1?
Each round, each player secretly chooses either One or Two fingers. (Players could also choose Heads or Tails, Left or Right, etc.) The payoff matrix is reprinted below.

(Alice, Bob)	One/Heads/Left	Two/Tails/Right
One/Heads/Left	+20	-4
Two/Tails/Right	-5	+1

Discussion Questions, Unit 3, Set 1.

3.1.A. With what percentage would you expect player to choose each option?

3.1.A. Response: We might expect each player to choose each option with roughly equal probability, that is, ½ of the time. But it seems unlikely this is the best mix of strategies for either player as it is in the "balanced" matching game.

3.1.B. Do you think that this is a fair game? Explain.

3.1.B. Response: This does not appear to be a fair game, since the expected value is not zero if both players choose each strategy with probability 1/2. However, we will show in this section that if the column player Bob uses an optimal mixed strategy, it is a fair game—that is, the expected value is zero under optimal play by both players.

Suppose that each player chooses each strategy with probability 1/2. We can calculate that under this combination of mixed strategies, the expected value of the game is positive:

$$\left(\frac{1}{2}\right)\left(\frac{1}{2}\right)(+20)+\left(\frac{1}{2}\right)\left(\frac{1}{2}\right)(-4)+\left(\frac{1}{2}\right)\left(\frac{1}{2}\right)(-5)+\left(\frac{1}{2}\right)\left(\frac{1}{2}\right)(+1)=+3.$$

This means that we would expect the row player Alice to eventually come out ahead in repeated play, at a rate of approximately 3 points per round, if each player used the mixed strategy [0.5, 0.5]. However, we should ask, is this really the best that each player can do?

Suppose that Bob finds out (or just guesses) that Alice intends to play One and Two fingers equally often, that is, each with probability 1/2. Can Bob choose a mix of strategies such that the game will have an expected value of less than +3 points per round?

If Bob chooses One finger, the expected value will be (1/2)*(+20)+(1/2)*(-5) = +7.5. If Bob chooses Two fingers, the expected value will be (1/2)*(-4)+(1/2)*(+1) = -1.5. Remember that a positive outcome is good for Alice, and a negative outcome is good for Bob. So Bob should choose Two fingers every round, as long as Bob thinks that Alice will continue to use the mix p =

1/2, 1-p = 1/2. This improves (for Bob) the expected value of the game from E = +3.0 (advantage Alice) to expected value
E = -1.5 (an advantage for Bob), assuming that Alice continues her mixed strategy.

Of course, if Alice knows or guesses that Bob is going to choose the Two fingers option every round, then Alice should switch from choosing One and Two each with probability 1/2, to always choosing Two, and the game would have a value of +1 (still an advantage for Alice) each time.

But then if Bob knows that Alice is always going to choose Two fingers, then Bob should always choose One finger, which will give payoff -5 points per round, a big advantage for Bob.

Continuing in this manner, it can be shown that a "pure" strategy (always choosing one option over repeated rounds) will not give the best results for either player in this game. How can we calculate the best mix of strategies for each player?

Suppose we use the unknown variable "p" to represent the percentage of the time that Alice chooses One, and use the unknown variable "q" to represent the percentage of the time that Bob chooses One. Then we have bounds
$0 <= p <= 1 \text{and} 0 <= q <= 1.$ Since the probabilities for One and Two must add up to 1 (or 100%) for each player, it follows that Alice will choose Two with probability 1 - p, and Bob will choose Two with probability 1 - q.

We can reformulate the payoff matrix with this additional notation included:

(Alice, Bob)	(Bob)	One	Two
(Alice)	probability	q	1-q
One	p	+20	-4
Two	1-p	-5	+1

We next calculate the expected values of the game as two different functions.

If Bob chooses One, the expected value of the game will be:
E[Bob,One] = E[B1] = (+20)(p)+(-5)(1-p) = 25p - 5.

If Bob chooses Two, the expected value of the game will be:
E[Bob,Two] =E[B2] = (-4)(p)+(+1)(1-p) =1 – 5p.
Alice wants to find a value for p that makes the expected value of the game as large as possible, regardless of whether Bob chooses One or Two. In other words, Alice should find a value for p such that Bob would be indifferent between choosing One or Two.

This guiding principle is sometimes called the "Principle of Indifference:" if one player uses the optimal (pure or mixed) strategy, then the expected payoffs will not depend on which strategy is chosen by the other player, and in that sense both players are "indifferent" as to what strategy is used against an optimal strategy.

Since Alice is trying for a maximum game value and Bob is trying for a minimum game value (remember that -5 is less than -4, for example), therefore what Alice is trying to achieve is called a “maximin”: Alice knows that Bob is trying to minimize the outcome, so Alice is trying to maximize the possible minimums. We can picture this in a graph as below:

Graph of expected values E[B1] = 25p - 5 and E[B2] = 1 - 5p as functions of p:

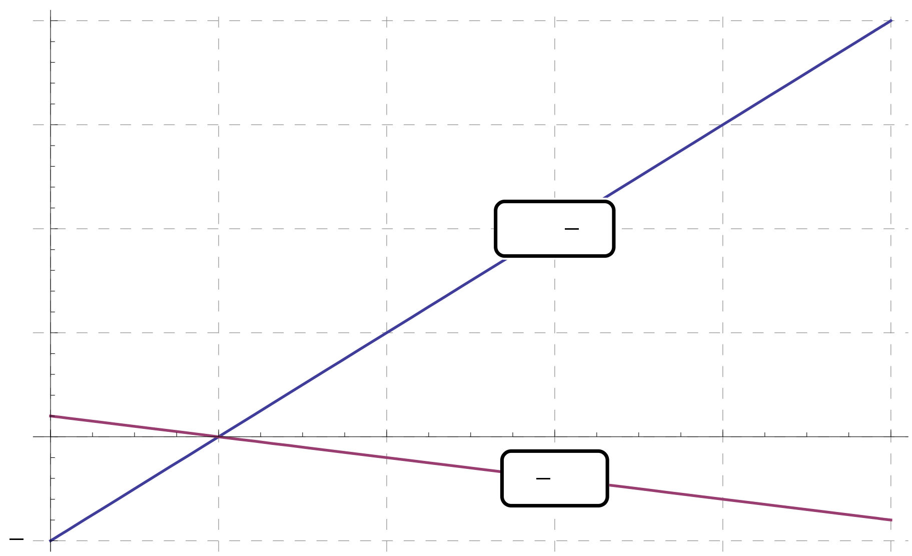

To find the point of intersection of these two lines, we set the expected value functions equal to each other and solve for p. This gives us 25p - 5 = 1 - 5p, which has solution p = 1/5 = 0.2. That is, the optimal strategy for Alice (assuming that Bob plays optimally) is to play One finger 20% and play Two fingers 80% of the time, in some randomized, unpredictable order. Note that this is different from the natural guess that Alice should play left and right each 50% of the time. In fact, you might think that if anything, Alice should choose One more often, since it has a potential payoff of +20rather than just +1, but the mathematical analysis implies that Alice should choose One fingeronly 20% of the time.

To determine an optimal mixed strategy for Bob, we go through the corresponding analysis from Bob’s point of view, as you will do by answering the questions below.

Note: Now that we have shown how to calculate the optimal strategy for each player, we can make the formal definition: A game is called “fair” if it has expected value zero when each player uses their optimal (pure or mixed) strategy. Later in this Unit, we will give an algebraic formula based only on the payoff matrix values that tells us quickly whether a two by two zero sum game is fair or not.

<u>Discussion Questions, Unit 3, Set 2.</u>

3.2.A. To carry out the analysis from the column player Bob’s point of view, find an expression for E[Alice,One]=E[A1], the expected value of the game when Alice chooses One, as a function of q, the probability that Bob chooses One. Then simplify this expression algebraically as much as possible.

3.2.A. Response: $E[A1] = (+20)(q)+(-4)(1-q) = 24q - 4$.

3.2.B. Find an expression for E[Alice,Two], the expected value of the game when Alice chooses Two, as a function of q, the probability that Bob chooses One. Then simplify this expression as much as possible.

3.2.B. Response: $E[A2] = (-5)(q)+(+1)(1-q) = 1- 6q$.

3.2.C. Graph the two functions E[A1] and E[A2] on the same graph, as a function of q. Then find the value of q at the point of intersection of these two graphs.

3.2.C. Response: $24q - 4 = 1-6q$ implies that $q = 1/6 \approx 16.7\%$.

Graph of expected values $E[A1] = 24q - 6$ and $E[A2] = 1 – 6q$ as functions of q:

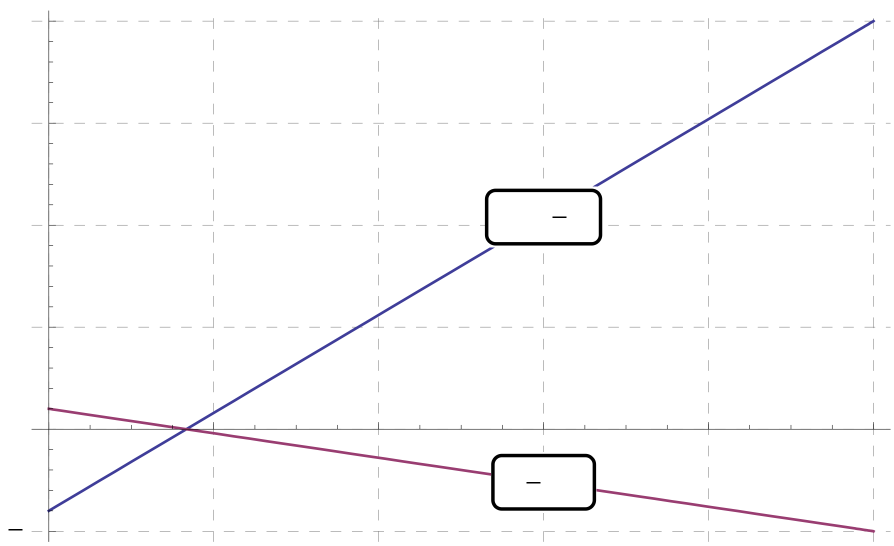

3.2.D. Explain why the point you just found may be called a “minimax” for Bob.

3.2.D. Response: As seen on the graph above, the point of intersection gives the minimum expected value when Alice is trying to choose between One and Two in a way that maximizes the game value. That is, Bob chooses a value for q that minimizes the possible maximums.

3.2.E. What is the expected value of the game if each player uses his or her optimal mix of strategies (assuming that their opponent is also playing their optimal mix of strategies)? If they play 12 rounds, who do you think would probably be ahead, and by approximately how many points?

3.2.E. Response: The expected value of the game in terms of p and q is

E = (+20)(p)(q) + (-4)(p)(1-q) + (-5)(1-p)(q) + (+1)(1-p)(1-q).

The optimal strategies found above are: Alice plays [One, Two]with mixed strategy [p, 1-p] $=\left[\frac{1}{5},\frac{4}{5}\right]$ and Bob plays [One, Two]with mixed strategy [q, 1-q] = $\left[\frac{1}{6},\frac{5}{6}\right]$.

Therefore, if each player uses his or her optimal mix, the value of the game is
E = $(+20)\left(\frac{1}{5}\right)\left(\frac{1}{6}\right) + (-4)\left(\frac{1}{5}\right)\left(\frac{5}{6}\right) + (-5)\left(\frac{4}{5}\right)\left(\frac{1}{5}\right) + (+1)\left(\frac{4}{5}\right)\left(\frac{5}{6}\right) = 0.$

That is, if both players use an optimal mix of strategies, then the expected value of the game is exactly zero, meaning that there is no advantage to either player and this is a fair game! Playing 12 or 100 rounds does not change the expected outcome, other than the observation that the more rounds we play, the greater the probability that the cumulative outcome would be approximately zero.

Activity 3: Pair off into groups of two players in each group and play 12 rounds of the game using the payoff matrix shown below. Each round, each player secretly chooses either Left or Right. (Players could also choose One or Two fingers, Heads or Tails, etc.) Follow the instructions given on Handout 3 and record the results.

(Robin, Cal)	Left	Right
Left	+8	-6
Right	-4	+2

Discussion Questions, Unit 3, Set 3.

3.3.A. Does this appear to be a fair game? If not, which player do you think has an advantage?

3.3.A. Response: The game might appear to be fair, since the sum of the matrix entries is zero. However, we don't really know for sure whether or not the game is fair until we calculate each player's optimal mix of strategies and then calculate the expected value assuming that each player uses their optimal mix of strategies. As shown later, the general expected value of a 2 by 2 zero sum game is E = (ad-bc)/e, which for this payoff matrix works out to E = [(+8)(+2)-(-6)(-4)]/20 = -0.4 points per round. That is, the column player Cal has an advantage when playing the optimal mixed strategy, so the game is not fair.

3.3.B If each player chose each strategy ½ of the time, what would be the expected value of the game?

3.3.B. Response: E = (1/2)*(1/2)*(+8-6-4+2) = 0 points per round. Under this combination of mixed strategies, we would expect the players to come out about even.

3.3.C. What percentage of the rounds did each player actually choose each strategy? What were each player's reasons for the mixed strategy they used? Did anyone's choice of percentages surprise you?

3.3.C. Responses will probably vary widely. Some students might still choose each strategy roughly ½ of the time. Some row players might choose Left more often (to try for the +8 outcome), and some column players might choose Right more often (to try for the -6 outcome). Some might attempt to calculate the optimal strategies (as requested in a later question) and use those.

3.3.D. How did the total points you gained or lost over the 12 rounds compare to your expectations for this game? If you played 100 rounds, which player do you think would come out ahead (or would you expect to come out approximately even)? Explain.

3.3.D. Response: The sophistication of responses will depend upon the ability of the students. The most naive response would be to guess that playerswould come out about even, because the sum of the entries is zero. However, as shown by our calculation below, the expected value is -0.4 points per round, so the column player has an advantage when using the optimal mixed strategy. This would translate into -4.8 points over 12 rounds, or -40 points over 100 rounds.

3.3.E. Use the algebraic maximin/minimax method outlined in this "Matching Game Revisited" section to calculate the optimal strategy and the expected value of the game (assuming that each player uses their optimal strategy). Then discuss how you might use this information to revise your strategy if you played another 12 rounds of this game, and how this would affect your answers to questions 3.3.C and 3.3.D.

3.3.E. Response:
We can reformulate the payoff matrix with this additional notation included:

(Robin, Cal)	(Cal)	Left	Right
(Robin)	probability	q	1-q
Left	p	+8	-6
Right	1-p	-4	+2

The expected value of the game is
E = (p)*(q)(+8) + (p)*(1-q)*(-6) + (1-p)*(q)*(-4) + (1-p)*(1-q)*(+2).
From Robin's point of view,
When q =1, E[CL] = (p)*(+8) + (1-p)*(-4) = 12p – 4.
When q =0, E[CR] = (p)*(-6) + (1-p)*(+2) = 2 – 8p.
We can picture this in a graph of expected values E[CL] = 12p - 4 and E[CR] = 2 - 8p as functions of p:

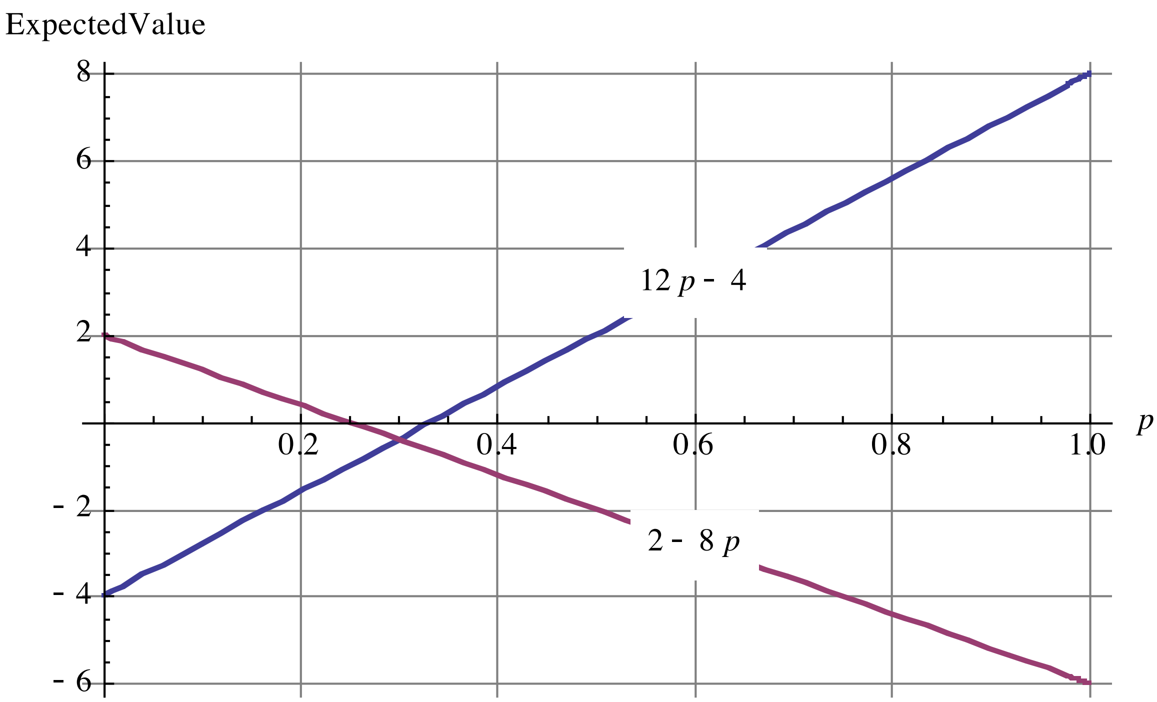

To find the point of intersection of these two lines, we set the expected value functions equal to each other and solve for p. This gives us 12p - 4 = -8p +2, which has solution p = 3/10 = 0.3. That is, the optimal strategy for Robin (assuming that Bob plays optimally) is to play Left 30% and play Right 70% of the time, in some randomized, unpredictable order. Note that this is different from the natural guess that Robin should play left and right each 50% of the time. In fact, you might think that if anything, Robin should choose Left more often, since it has a potential payoff of +8 rather than just +2, but the mathematical analysis implies that Robin should choose Left a little less than one-third of the time.

To determine an optimal mixed strategy for Cal, we go through the corresponding analysis from Cal's point of view. The expected value of the game when Robin chooses Left (so p=1), as a function of q, the probability that Cal chooses Left, is given by E[RL] = (+8)(q)+(-6)(1-q) = 14q - 6.

When Robin chooses Right (so p =0), the expected value of the game is
E[RR] = (-4)(q)+(+2)(1-q) = 2 - 6q.

We can picture this in a graph of expected values E[RL] = 14q - 6 and E[RR] = 2 – 6q as functions of q:

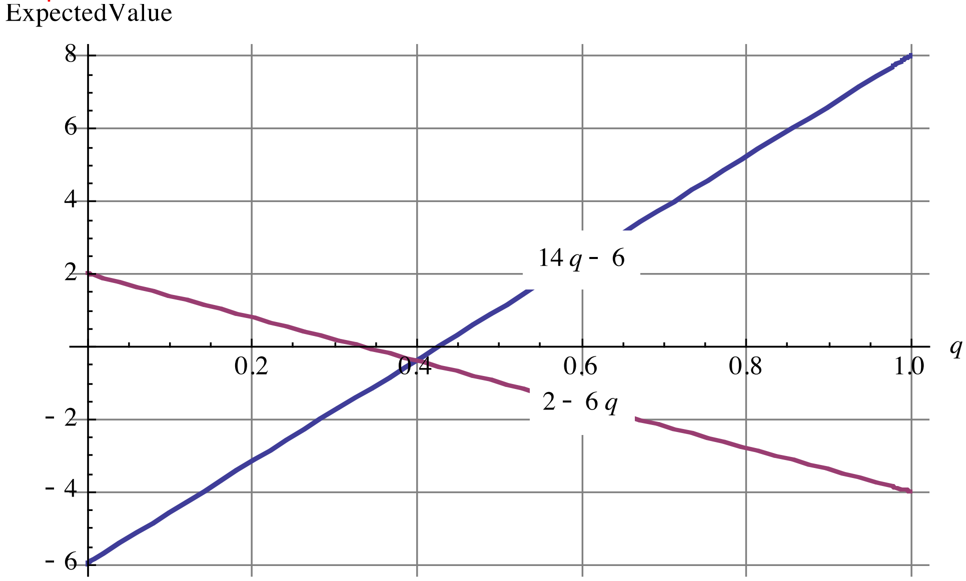

To find the point of intersection of these two lines, we set the expected value functions equal to each other and solve for q. This gives us 14q - 6 = 2 -6q , which has solution q = 4/10 = 0.4.

The optimal strategies we found above are: Robin plays Left p = 0.3, Right 1-p = 0.7;
Cal plays Left q = 0.4, Right 1-q = 0.6.

Therefore, if each player uses his or her optimal mix, the value of the game is
E = (+8)(0.3) (0.4)+ (-6)(0.3)(0.6) + (-4)(0.7)(0.4) + (+2)(0.7)(0.6)= -0.4

That is, if both players use an optimal mix of strategies, then Cal should win by an average outcome of -0.4 points per round. So if they play 12 rounds, the expected value would be (12)(-0.4) = -4.8 points (in Cal's favor).

3.3.F. Suppose that we play a different revised matching game, which has the payoff matrix shown below. How, if at all, would this change the optimal mixed strategy for each player? How, if at all, would this change the expected value of the game? Explain.

(Robin, Cal)	(Cal)	Left	Right
(Robin)	probability	q	1-q
Left	p	+4	-3
Right	1-p	-2	+1

3.3.F. Response: Because the new payoff matrix is a constant scalar multiple in each entry compared to the old payoff matrix (each old entry is multiplied by ½ to obtain the new entry), therefore the optimal strategies will not change; they are still p = 0.3 for Robin and q = 0.4 for Cal. The new expected value is just ½ times the old expected value, that is, the new expected value is (1/2)*(-0.4) = -0.2 points per round. Note that multiplying a payoff matrix by the same constant scalar in each entry changes the expected value by that multiple, but does not change the optimal mixed strategy for either player.

Homework Questions, Unit 3, Set 4.

Homework 3.4.A. Apply the algebraic maximin/minimax method outlined in this section to the Modified Rock-Paper-Scissors game (which reduced to a 2 by 2 payoff matrix shown below) to find the optimal mixed strategy for the row player Robin to choose [Paper, Scissors] with probabilities [p, 1-p], and the optimal mixed strategy for the column player Cal is to choose [Rock, Paper] with probabilities [q, 1-q].

(Robin, Cal)	Rock	Paper
Paper	+1	0
Scissors	-1	+1

Homework 3.4.ASolution. The expected value of the game is
E = (p)*(q)(+1) + (p)*(1-q)*(0) + (1-p)*(q)*(-1) + (1-p)*(1-q)*(+1).
From Robin's point of view,
When q =1, E[CR] = (p)*(+1) + (1-p)*(-1) = 2p – 1.
When q =0, E[CP] = (p)*(0) + (1-p)*(+1) = 1 – p.
These lines intersect (giving the maximin) when p = 2/3, so Robin's optimal strategy is to choose Paper $\frac{2}{3} \approx 66.7\%$ of the time and Scissors $\frac{1}{3} \approx 33.3\%$ of the time.
From Cal's point of view,
When p =1, E[RP] = (q)*(+1) + (1-q)*(0) = q.
When p =0, E[RS] = (q)*(-1) + (1-q)*(+1) = 1 – 2q.
These lines intersect (giving the minimax) when q = 1/3, so Cal's optimal strategy is to choose Rock $\frac{1}{3} \approx 33.3\%$ of the time and choose Paper $\frac{2}{3} \approx 66.7\%$ of the time.
The expected value of the game (assuming that each player uses their optimal mixed strategy) is
E = $\left(\frac{2}{3}\right)*\left(\frac{1}{3}\right)$ (+1) + $\left(\frac{2}{3}\right)*\left(\frac{2}{3}\right)$*(0) + $\left(\frac{1}{3}\right)*\left(\frac{1}{3}\right)$*(-1) + $\left(\frac{1}{3}\right)*\left(\frac{2}{3}\right)$*(+1) = $\frac{+2+0-1+2}{9} = +\frac{1}{3} \approx +0.33$ points per round.

Homework 3.4.B. Use the algebraic maximin/minimax method to find the optimal mixed strategy for both players, and the resulting expected value, in the game with the payoff matrix given below. (If preferred, this problem can be more easily solved using the "oddments" method described in the next section.)

(Robin, Cal)	Heads	Tails
Heads	+1	0
Tails	-10	+1

Homework 3.4.B Solution.The expected value of the game is
E = (p)*(q)(+1) + (p)*(1-q)*(0) + (1-p)*(q)*(-10) + (1-p)*(1-q)*(+1).
From Robin's point of view,
When q =1, E[CH] = (p)*(+1) + (1-p)*(-10) = 11p – 10.
When q =0, E[CT] = (p)*(0) + (1-p)*(+1) = 1 – p.
These lines intersect (giving the maximin) when p = 11/12, so Robin's optimal strategy is to play Heads $\frac{11}{12} \approx 91.7\%$ of the time and Tails $\frac{1}{12} \approx 8.3\%$ of the time.
This corresponds to our intuition that Robin may play Heads more often.
From Cal's point of view,
When p =1, E[RH] = (q)*(+1) + (1-q)*(0) = q.
When p =0, E[RT] = (q)*(-10) + (1-q)*(+1) = 1 – 11q.
These lines intersect (giving the minimax) when q = 1/12, so Cal's optimal strategy is to play Heads $\frac{1}{12} \approx 8.3\%$ of the time and Tails $\frac{11}{12} \approx 91.7\%$ of the time. This runs somewhat counter to our intuition that Cal will try most often for the -10 outcome. However, it makes intuitive sense that Cal would respond in this way to his or her guess that Robin will usually be playing Heads.
The expected value of the game (assuming that each player uses their optimal mixed strategy) is
E = $\left(\frac{11}{12}\right)*\left(\frac{1}{12}\right)$ (+1) + $\left(\frac{11}{12}\right)*\left(\frac{11}{12}\right)$*(0) + $\left(\frac{1}{12}\right)*\left(\frac{1}{12}\right)$*(-10) + $\left(\frac{1}{12}\right)*\left(\frac{11}{12}\right)$*(+1) = $\frac{+11+0-10+11}{144} = +\frac{1}{12} \approx +0.083$ points per round.
That is, the game is not fair, but it is close to fair. There is actually a slight advantage to the row player rather than the column player, if the row player uses the optimal strategy! Over 12 rounds, assuming that each player uses their optimal mixed strategy, the expected outcome is (+1/12)*(12) = +1 points, and over 100 rounds, the expected result is (+1/12)*(100) = +8.3 points.

Homework 3.4.C. Use the algebraic maximin/minimax method to find the optimal mixed strategy for both the striker and the goalie, and the resulting expected value, in the Strong-Side versus Weak-Side Penalty Kick example of Unit 1, if we assume that s = 2/3. Then find the optimal mix of strategies in terms of the constant unknown parameter s (that is, without assuming any particular number value for s).

(Striker, Goalie)	Dive Left (prob.= q)	Dive Right (prob. = 1-q)
Kick Left (prob. = p)	0	+s
Kick Right (prob. = 1-p)	+1	0

Homework 3.4.C Solution. (Note: students who are comfortable working with unknown constant parameters may find the solution in terms of s first, and substitute s = 2/3 only at the end. But most students will be more comfortable doing the calculation with s = 2/3 plugged in initially.)

The expected value of the game is
E = (p)*(q)(0) + (p)*(1-q)*(s) + (1-p)*(q)*(+1) + (1-p)*(1-q)*(0).
From the Striker's point of view,
When q =1, E[GL] = (p)*(0) + (1-p)*(+1) = 1 - p.
When q =0, E[GR] = (p)*(+s) + (1-p)*(0) = (+s)*p.
These lines intersect (giving the maximin) when p = 1/(s+1) and 1-p = s/(s+1).
With s = 2/3, this gives p = 3/5 and 1-p = 2/5, so the Striker's optimal strategy is to kick Left $\frac{3}{5} = 60\%$ of the time and kick Right $\frac{2}{5} = 40\%$ of the time.
From the Goalie's point of view,
When p =1, E[SL] = (q)*(0) + (1-q)*(s) = s - s*q.
When p =0, E[SR] = (q)*(+1) + (1-q)*(0) = q.
These lines intersect (giving the minimax) when q = s/(s+1) and 1-q = 1/(s+1).
With s = 2/3, this gives q = 2/5 and 1-q = 3/5, so the Goalie's optimal strategy is to dive Left $\frac{2}{5} = 40\%$ of the time and dive Right $\frac{3}{5} = 60\%$ of the time.
The expected value of the game when s = 2/3 (assuming that each player uses their optimal mixed strategy) is
E = $\left(\frac{3}{5}\right)*\left(\frac{2}{5}\right)$ (0) + $\left(\frac{3}{5}\right)*\left(\frac{3}{5}\right)$*(+2/3) + $\left(\frac{2}{5}\right)*\left(\frac{2}{5}\right)$*(+1) + $\left(\frac{2}{5}\right)*\left(\frac{3}{5}\right)$*(0)
= $\frac{0+6+4+0}{25} = \frac{2}{5} = 0.4$ goals per kick.
In terms of general parameter s, the expected value of the game is E = s/(s+1) goals per kick.

"Oddments": A Faster Method for a Special Case (optional section)

The algebraic maximin/minimax method we have outlined above is "robust" in the sense that it can be generalized to games in which one or both players have three or more options to choose from. (See homework problem below.) However, this algebraic method has the disadvantage of requiring a bit of equation solving, and sometimes graphing, to carry out.

In the special case of a game with two players each having two options, there is a faster method of computing the optimal strategies, sometimes called the "oddments" method.We will demonstrate this method with an example. Suppose Robin and Cal are playing the 2 by 2 game with the payoff matrix shown below, and they each want to quickly calculate their optimal mixed strategies.

(Robin, Cal)	V	W
X	+5	-2
Y	-4	+1

As long as the matrix does not contain a saddle point, they can each calculate their optimal mixed strategy by using the following steps.

For Robin's calculation, find the difference of the payoffs in each row to form a new column. Then take the absolute value of each number (that is, make it positive if it is negative), switch or swap the differences in columns 1 and 2 with each other, and switch or swap the differences in rows 1 and 2 with each other. Finally, compute the rational fraction with that difference as the numerator and the sum of the differences as the denominator.

This leads to the calculation that Robin's optimal strategy is to play [X, Y] with ratios [5/12, 7/12], as shown in the expanded matrix below.

A similar, corresponding calculation shows that Cal's optimal mix is to play [V, W] with ratios [1/4, ¾].

The expected value of this game is

$$E = \left(\frac{5}{12}\right)*\left(\frac{1}{4}\right)(+5) + \left(\frac{5}{12}\right)*\left(\frac{3}{4}\right)*(-2) + \left(\frac{7}{12}\right)*\left(\frac{1}{4}\right)*(-4) + \left(\frac{7}{12}\right)*\left(\frac{3}{4}\right)*(+1) = \frac{25-30-28+21}{48} = -\frac{1}{4}$$

$= -0.25$ points per round.

(Robin, Cal)	V	W	difference	switch	ratio
X	+5	-2	5-(-2)=7	5	$\frac{5}{(5+7)} = \frac{5}{12}$
Y	-4	+1	-4-1=-5	7	$\frac{7}{(5+7)} = \frac{7}{12}$
difference	5-(-4)=9	-2-1=-3			
switch	3	9			
ratio	$\frac{3}{(3+9)} = \frac{1}{4}$	$\frac{9}{(3+9)} = \frac{3}{4}$			

Homework. Use the "oddments" method to find the optimal mixed strategy for each player and the expected value in the Modified Rock-Paper-Scissors game by filling in the table.

(Robin, Cal)	Rock	Paper	difference	switch	ratio
Paper	+1	0			
Scissors	-1	+1			
difference					
switch					
ratio					

(Robin, Cal)	Rock	Paper	difference	switch	ratio
Paper	+1	0	1-0=1	2	$\frac{2}{(2+1)} = \frac{2}{3}$
Scissors	-1	+1	-1-1=-2	1	$\frac{1}{(2+1)} = \frac{1}{3}$
difference	1-(-1)=2	0-1=-1			
switch	1	2			
ratio	$\frac{1}{(2+1)} = \frac{1}{3}$	$\frac{2}{(2+1)} = \frac{2}{3}$			

The expected value of the game is

E = $\left(\frac{2}{3}\right)*\left(\frac{1}{3}\right)(+1) + \left(\frac{2}{3}\right)*\left(\frac{2}{3}\right)*(0) + \left(\frac{1}{3}\right)*\left(\frac{1}{3}\right)*(-1) + \left(\frac{1}{3}\right)*\left(\frac{2}{3}\right)*(+1) = \frac{+2+0-1+2}{9}$

$= +\frac{1}{3} \approx +0.33$points per round.

Homework. Use the "oddments" method to find the optimal mixed strategy for each player and the expected value in the Strong-Side versus Weak-Side Penalty Kick game if we assume that the probability of a goal in the Kick Left, Dive Right outcome is s = 2/3.

(Striker, Goalie)	Dive Left (prob.= q)	Dive Right (prob. = 1-q)
Kick Left (prob. = p)	0	+ 2/3
Kick Right (prob. = 1-p)	+1	0

Strike/Goalie	Dive Left (prob.= q)	Dive Right (prob. = 1-q)	difference	switch	ratio
Kick Left (prob. = p)	0	+ 2/3	$0 - \frac{2}{3} = -\frac{2}{3}$	1	$\frac{1}{\left(\frac{2}{3}+1\right)} = \frac{3}{5}$
Kick Right (prob. = 1-p)	+1	0	1-0=1	$\frac{2}{3}$	$\frac{\frac{2}{3}}{\left(\frac{2}{3}+1\right)} = \frac{2}{5}$
difference	0-(1)=-1	$\frac{2}{3} - 0 = \frac{2}{3}$			
switch	$\frac{2}{3}$	1			

ratio	$\frac{\frac{2}{3}}{\left(\frac{2}{3}+1\right)} = \frac{2}{5}$	$\frac{1}{\left(\frac{2}{3}+1\right)} = \frac{3}{5}$			

The expected value of the game is
E = $\left(\frac{3}{5}\right)$*$\left(\frac{2}{5}\right)$*(0) + $\left(\frac{3}{5}\right)$*$\left(\frac{3}{5}\right)$ *(+$\frac{2}{3}$) + $\left(\frac{2}{5}\right)$*$\left(\frac{2}{5}\right)$*(+1) + $\left(\frac{2}{5}\right)$*$\left(\frac{3}{5}\right)$*(0) = $\frac{0+6+4+0}{25}$
$= +\frac{2}{5} = +0.4$points per round.

<u>Homework.</u> Calculate a general solution to the Strong-Side versus Weak-Side Penalty Kick example. That is, show that for general unknown s with $0 < s < 1$, the optimal strategy for the striker is to kick left with probability
p =1/(s+1), and the optimal strategy for the goalie is to dive left with probability
q = s/(s+1), leading to an expected value of s/(s+1) goals per kick.

Strikr/Goalie	Dive Left (prob.= q)	Dive Right (prob. = 1-q)	difference	switch	ratio
Kick Left (prob. = p)	0	+ s	$0 - s = -s$	1	$\frac{1}{(s+1)}$
Kick Right (prob. = 1-p)	+1	0	1-0=1	s	$\frac{s}{(s+1)}$
difference	0-(1)=-1	$s - 0 = s$			
switch	s	1			
ratio	$\frac{s}{(s+1)}$	$\frac{1}{(s+1)}$			

<u>Homework.</u> (requires algebraic solution of simultaneous equations with unknown constants):Show that for the general 2 by 2 payoff matrix of the form below (assuming no saddle point), we obtain optimal mix of strategies
p = (d-c)/e,   q = (d-b)/e, where e = (a-c) + (d-b) = (a-b) + (d-c),
and the expected value of the game is thus E = (ad-bc)/e.

(Robin, Cal)		Left	Right
	probability	q	1-q
Left	p	a	B
Right	1-p	c	D

Algebraic maximin approach: The expected value of the game is
E = (p)*(q)(a) + (p)*(1-q)*(b) + (1-p)*(q)*(c) + (1-p)*(1-q)*(d).

From the Row player Robin's point of view,
When q =1 (Cal plays Left), E[CL] = (p)*(a) + (1-p)*(c) = (a-c)*p+c.
When q =0 (Cal plays Right), E[CR] = (p)*(b) + (1-p)*(d) = (b-d)*p+d.
These lines intersect (giving the maximin) when (a-c)*p+c=(b-d)*p+d, that is,
$[(a-c)+(d-b)]*p = d-c$,which has solution $p = \frac{(d-c)}{e}$, $1-p = \frac{(a-b)}{e}$,
where e = (a-c) + (d-b) = (a-b) + (d-c).

From the Column playerCal's point of view,
When p =1 (Robin plays Left), E[RL] = (q)*(a) + (1-q)*(b) = (a-b)*q+b.
When p =0 (Robin plays Right), E[RR] = (q)*(c) + (1-q)*(d) = (c-d)*q+d.
These lines intersect (giving the minimax) when (a-b)*q+b=(c-d)*q+d, that is,
$[(a-b)+(d-c)]*q = d-b$,which has solution $q = \frac{(d-b)}{e}$, $1-q = \frac{(a-c)}{e}$,
where e = (a-c) + (d-b) = (a-b) + (d-c).

The expected value of the game (assuming that each player uses their optimal mixed strategy) is
E = $\left(\frac{d-c}{e}\right)$*$\left(\frac{d-b}{e}\right)$ (a) + $\left(\frac{d-c}{e}\right)$*$\left(\frac{a-c}{e}\right)$ (b) + $\left(\frac{a-b}{e}\right)$*$\left(\frac{d-b}{e}\right)$ (c) + $\left(\frac{a-b}{e}\right)$*$\left(\frac{a-c}{e}\right)$ (d)

= $\frac{1}{e^2} * \begin{bmatrix} (ad^2 - abd - acd + abc) + (abd - abc - bcd + c^2b) \\ +(acd - abc - bcd + b^2c) + (a^2d - abd - acd + bcd) \end{bmatrix}$
= $\frac{1}{e^2} * [(ad^2 - acd) + (-bcd + c^2b) + (-abc + b^2c) + (a^2d - abd)]$
= $\frac{1}{e^2} * [ad(d-c) - bc(d-c) - bc(a-b) + ad(a-b)]$
= $\frac{1}{e^2} * (ad-bc)[(a-b)+(d-c)] = \frac{1}{e^2} * (ad-bc)[e] = \frac{(ad-bc)}{e}$.
Oddments approach, assuming that the desired quantities are all positive
(use absolute value if necessary):

(Robin, Cal)	Cal Left (prob.= q)	Cal Right (prob. = 1-q)	difference	switch	Ratio
Robin Left (prob. = p)	a	b	$a-b$	$c-d$	$\frac{c-d}{e}$
Robin Right (prob. = 1-p)	c	d	$c-d$	$a-b$	$\frac{a-b}{e}$
difference	$a-c$	$b-d$			
switch	$b-d$	$a-c$			
ratio	$\frac{b-d}{e}$	$\frac{a-c}{e}$			

Homework. Apply the algebraic maximin/minimax method to calculate the optimal mixed strategy for the Column player Cal in the game with the payoff matrix shown below. To calculate the optimal value of q, you will need to graphically find the value of q that gives the minimax for Cal, regardless of which of the three possible strategies (Left, Center, or Right) Robin chooses. Note: It is also possible to calculate the optimal mix of strategies [p, r, 1-p-r] for Robin, but that involves methods somewhat more advanced than those discussed in this module.

(Robin, Cal)		Left	Right
	probability	q	1-q
Left	p	+2	-2
Center	r	-2	+2
Right	1-p-r	+6	-4

From the Column player Cal's point of view,
When Robin plays Left, E[RL] = (q)*(+2) + (1-q)*(-2) = 4q - 2.
When Robin plays Center, E[RC] = (q)*(-2) + (1-q)*(+2) = 2 – 4q.
When Robin plays Right, E[RR] = (q)*(+6) + (1-q)*(-4) = 10q – 4.

The region on which we are trying to find a minimax is best visualized by graphing each of these three lines as a function of q, with $0 \leq q \leq 1$.

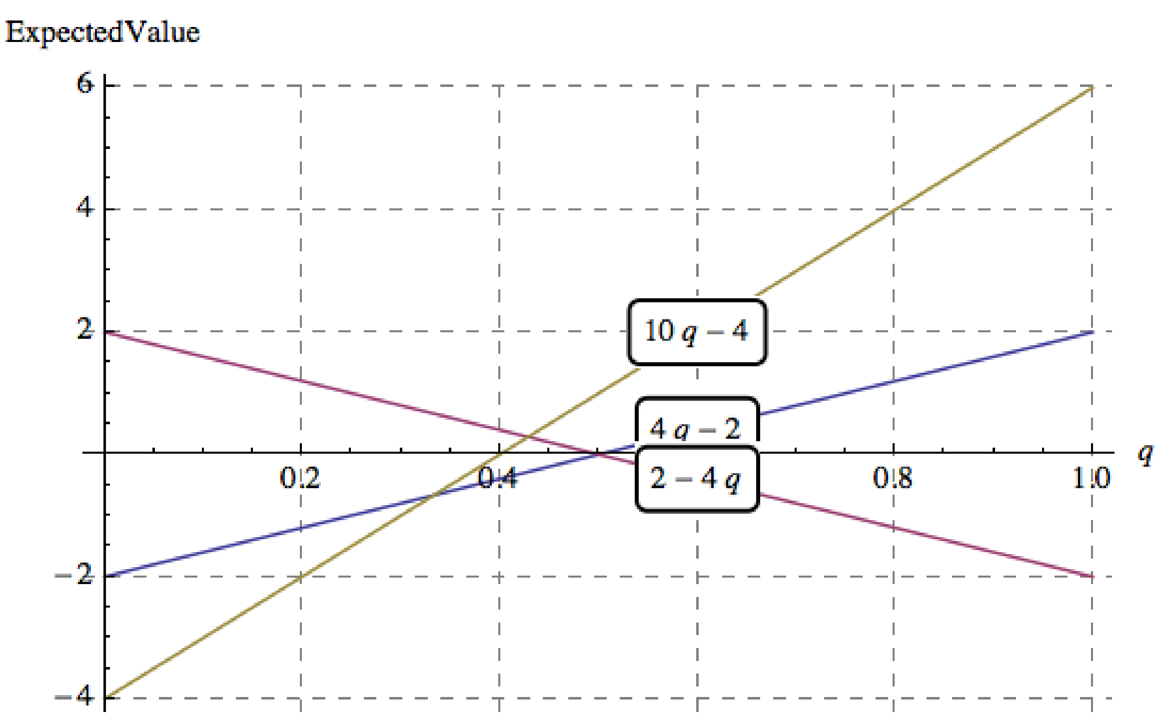

We are trying to find the value of q between 0 and 1 that gives the least possible upper bound of all three lines, which represent the payoffs of the three possible strategies of the Row player. We see graphically that this occurs at the intersection of E[RC] = 2 – 4q and E[RR] = 10q – 4.

Setting $2 - 4q = 10q - 4$, we solve for q to find $q = \frac{3}{7} \approx 0.429 \approx 43\%$, and $1 - q = \frac{4}{7} \approx 57\%$.

Note that from the Row player's viewpoint, playing Left (with E[RL] = 4q – 2) is dominated by Center (with E[RC] = 2 – 4q) if $q \leq \frac{3}{7}$, and playing Left is dominated by Right (with E[RR] = 10q – 4) if $q \geq \frac{3}{7}$. Therefore, Robin should never play Left.

With Robin's options reduced to just Center and Right, we can carry out the same analysis to calculate the optimal mixed strategy [p=0, r, 1-r] for Robin, which turns out to be to play Center with probability $r = \frac{5}{7}$ and to choose Right with probability $1 - r = \frac{2}{7}$.

The expected value of the game is therefore

$$E = \left(\frac{5}{7}\right)\left(\frac{3}{7}\right)(-2) + \left(\frac{5}{7}\right)\left(\frac{4}{7}\right)(+2) + \left(\frac{2}{7}\right)\left(\frac{3}{7}\right)(+6) + \left(\frac{2}{7}\right)\left(\frac{4}{7}\right)(-4) = \frac{2}{7} \approx 0.29.$$

Unit 4: Non-Zero-Sum Games: Pranks and Prisoners

Non-Zero-Sum Games

In the previous Unit, the only type of game we considered was the type in which any gain for one player corresponds to an equal and opposite loss for the other player, which we call a "zero-sum game." We showed how to use the concepts of maximin and minimaxto calculate an optimal solution (called the "Nash equilibrium") in any zero-sum game, which may consist of either a pure strategy or a mixed strategy for either player.

However, there are many situations in which the respective outcomes for the two players are not necessarily equal and opposite. In these situations, each player must decide whether to "cooperate" (work together for their common good, or yield to the other player's wishes) or to "defect" (i.e., betray the confidence of the other player for the possibility of achieving greater personal gain, even though this may produce a worse outcome for the other player). We can still use a payoff matrix to describe non-zero sum games, but now each cell will have an ordered pair of outcomes, where the first number of the pair gives the outcome to the Row Player, and the second number of the pair gives the outcome to the Column Player. For example, in the payoff matrix below, if both players choose Heads, then row player Robin gains 1 point, and column player Cal loses 2 points for that round.

(Robin, Cal)	Heads	Tails
Heads	(+1, -2)	(-2, +2)
Tails	(-3, +1)	(+6, -1)

Calculating an optimal strategy in these non-zero sum games is much less clear, and generally requires us to consider the practical level of communication and psychological trust between the players in addition to the usual mathematical calculations. We can still calculate the Nash equilibrium from each player's point of view, but it is no longer reasonable to expect that equilibrium to be the best outcome of the game. We will often find it more helpful in predicting and explaining likely outcomes to draw a "movement diagram" for non-zero sum games. To do this, from each outcome or cell in the matrix table, we draw a vertical arrow up or down if changing row strategy would benefit the Row player, and we draw a horizontal arrow left or right if changing column strategy would benefit the Column player. In the above game matrix, the movement diagram would look like this:

		Column player Cal		
		Heads		Tails
Row player Robin	Heads	(+1, -2)		(-2, +2)
	Tails	(-3, +1)		(+6, -1)

This gives us a nice visualization of the observations that if Cal plays Heads, Robin would prefer Heads over Tails, but if Cal plays Tails, Robin would prefer Tails over Heads. Likewise, if Robin plays Heads, Cal would prefer Tails over Heads, but if Robin plays Tails, Cal would prefer Heads over Tails. The movement arrows make it obvious that in this game there is no pure strategy Nash equilibrium; that is, there is no outcome in this game such that neither player would benefit by changing their strategy option unilaterally, with no change of strategy by the other player. In repeated play of this game, we would expect a “cycle” of outcomes rather than tending toward one stable, unchanging outcome in the long run.

Here is another non-zero sum game. Consider the following scenario.

Prank on the Dean

Bill and Ted play a prank on their university’s dean of students, and they get caught. The school security put them in separate rooms for interrogation. Bill and Ted are told the consequences of keeping silent versus confessing. If they both keep silent, (cooperating with their friend by refusing to admit their guilt), they will both get 4 days of community service. If they both confess, they will each get 8 days of community service. However, if one of them confesses (defecting by betraying their friend) while the other one keeps silent, the one who confesses will be rewarded with only 2 days of community service while the friend who kept silent will receive 12 days of community service. There is no way for them to communicate with each other, and they each get an hour to decide. What choice would Bill/Ted make? How do they reach their decision?

(Note: This example corresponds to a famous game theory problem known as the Prisoner’s Dilemma.)

Activity: Construct the payoff matrix and find the (mixed-strategy) Nash equilibrium.
(Note: We normally would use negative number entries for the days of community service, but since all possible payoffs are negative, we can simply enter the number of days of community service, and remember that in this particular example, both players are striving for the lowest possible payoff rather than the highest possible payoff.)

Use the number of days of community service as the coefficients, and thus, the lower the coefficient, the better off it is.

		Ted	
		Keep silent (cooperate)	Confess (defect)
Bill	Keep silent (cooperate)	(4, 4)	(12, 2)
	Confess (defect)	(2, 12)	(8, 8)

Consider the payoff matrix from Bill's perspective. Because Bill has no idea which choice Ted is going to make, Bill can only rely on what the consequences are based on his choice. If Bill assumes that Ted confesses, confession results in 8 days of community service while keeping silent punishes him with 12 days. Therefore, if Ted confesses, Bill's best option is to confess as well. Doing the same analysis based on the assumption that Ted keeps silent, it is still Bill's best interest to confess because confession costs him 2 days of community service while keeping silent punishes him with 4 days of community service. Thus, it makes sense for Bill to choose to confess. Ted goes through the same train of thoughts, and he chooses to confess as well. Therefore, the Nash equilibrium is the strategy of both of them confessing with the coefficients of (8, 8).

		Ted		
		Keep silent (cooperate)		Confess (defect)
Bill	Keep silent (cooperate)	(4, 4)	[illegible]	(12, 2)
		[illegible]		[illegible]
	Confess (defect)	(2, 12)	[illegible]	(8, 8)

Discussion question: Is this the best outcome for Bill and Ted? Why or why not?

This is clearly not the optimal outcome. If Bill and Ted both decided to keep silent, they would only have 4 days of community service each. The better outcome is unreached because any communication between Bill and Ted is prohibited. If they were put in the same room and allowed to discuss the situation before they made up their mind, they would agree to keep silent.

The Nash equilibrium is when both Bill and Ted confess (defect) but as we just discussed, it is not the best outcome. The best outcome is when both Bill and Ted keep silent (cooperate). This

set of strategies is called Pareto Optimality. “An outcome of a game is Pareto optimal if there is no other outcome that makes every player at least as well off and at least one player strictly better off.” http://www.gametheory.net/dictionary/ParetoOptimal.html
From the movement diagram, we can see that the arrows move away from the (4, 4) strategy. This means that there are better outcomes for each player individually. However, following the arrows does not result in an outcome where both players benefit from the movement.

Discussion question: Can Bill trust or betray Ted? What about Ted?

What happens if the trust is broken? Bill and Ted will get 4 days of community service if and only if both of them keep silent. Is there any benefit for one to betray the other? Assume Bill decides to confess instead after agreeing to keep silent. Given that Ted stays silent, the length of community service changes based on Bill’s decision. If he cooperates with Ted and keeps silent, he will still get 4 days of community service. If he changes his mind, then he will get only 2 days of community service! There is clearly a benefit for Bill to lie to Ted about keeping silent.

A mirror argument applies to Ted as well. This means that Ted also sees the benefit of betraying Bill to get a reduced penalty. But what happens if both Bill and Ted decide that betraying the other is a better play for themselves? We end up where we started: both of them confessing and receiving 8 days of community service.

Activity: Would extra information help? What if Bill has some idea of whether Ted would keep his promise or not? Assume that Bill knows the chance of Ted keeping his promise. Work on Bill’s strategy based on 10% and 90% as the chances of Ted keeping his promise.

If there is more information about Bill and/or Ted, you can compare your expected outcome. Bill knows that Ted is trusting and there is only 10% chance that Ted will betray Bill. This means that there is 90% chance that Ted will keep silent while 10% chance that Ted will confess. Bill can compute the expected number of days of community service based on these probabilities.

Expected # of days of community service for Bill if he

- keeps silent: 0.1 * 12 + 0.9 * 4 = 4.8
- confesses: 0.1 * 8 + 0.9 * 2 = 2.8

If Bill believes in these chances, then Bill is better off betraying Ted. Let’s assume the opposite case where Bill does not trust Ted at all by flipping the probabilities. This means that there is only 10% chance that Ted will keep silent and 90% chance that Ted will confess.

Expected # of days of community service for Bill if he

- keeps silent: 0.9 * 12 + 0.1 * 4 = 11.2
- confesses: 0.9 * 8 + 0.1 * 2 = 7.4

This case also shows that Bill is better off by confessing.

Activity: Is there any level of trust where Bill is better off by keeping his promise? Use an arbitrary chance q in the previous computation to see if there is a value for q that convinces Bill to keep his promise instead of confessing.

Assume that q is the probability that Ted would keep his promise (keeps silent). For Bill to choose to keep his promise, the expected number of days of community servicefor keeping silent must be less than that for confessing.

Expected # of days of community service for Bill if he

- keeps silent: $(1 - q) * 12 + q * 4 = 12 - 8q$
- confesses: $(1 - q) * 8 + q * 2 = 8 - 6q$

Therefore, we want

$12 - 8q < 8 - 6q$

Remember that less is better in this case.

$2q > 4$

$q > 2$ (This is an impossible probability. Probabilities range from 0 to 1 inclusive.)

There is no reason for Bill to keep his promise regardless of how much he trusts Ted. Again, this is because cooperation is impossible. Even from the payoff matrix viewpoint, the confession row and column are dominant strategies for Bill and Ted respectively.

Activity: Is there any set of coefficients that would make sense for Bill to trust Ted? Assume the following payoff matrix.

		Ted	
		Keep silent (cooperate)	Confess (defect)
Bill	Keep silent (cooperate)	(w, w)	(y, x)
	Confess (defect)	(x, y)	(z, z)

Prisoner's Dilemma requires that $x < w < z < y$. In our example, $x = 2$, $w = 4$, $z = 8$, and $y = 12$.

Ted's chance of keeping his promise (keeping silent) is q while that of betraying (confessing) is $(1 - q)$.

Expected # of days of community service for Bill if he

- keeps silent: $(1 - q) * y + q * w$
- confesses: $(1 - q) * z + q * x$

For Bill to keep his promise, $(1 - q) * y + q * w < (1 - q) * z + q * x$

$y + (w - y)\, q < z + (x - z)\, q$

$y - z < (x - z - w + y)\, q$

$q > (y - z) / (x - z - w + y)$

$q > (y - z) / ((y - z) - (w - x))$

We know that $(y - z)$ and $(w - x)$ are positive numbers because of $x < w < z < y$. What is the significance of $(y – z)$ and $(w – x)$? $(y – z)$ indicates how many more days of community service you get for keeping silent if your friend betrays you, while $(w – x)$ indicates how many more days of community service you sacrifice for keeping silent together with your friend. The answer depends on the relationship between these two quantities.

1. Case 1: $(y – z) >= (w – x)$
   This is the case where the risk of keeping silent is greater than or equal to the sacrifice you make for keeping silent. The fraction is positive and also greater than 1. Since p cannot be greater than 1, this suggests that betraying the friend is a better option regardless of how trustworthy the friend is.

2. Case 2: $(y – z) < (w – x)$

This is the case where the risk of keeping silent is less than the sacrifice you make for keeping silent. The fraction becomes negative! This means that it is always better to keep silent regardless of how trustworthy your friend is.

Split or Steal? (Golden Ball)

Watch the following video with the class.
http://www.youtube.com/watch?v=p3Uos2fzIJ0

The Golden Ball presents a situation where two people must decide if they want to split or steal the jackpot. If the contestants both choose to split (cooperate), each of them receives half of the jackpot. If the contestants both choose to steal (defect), neither of them receives any money. If one contestant chooses to steal (defect) and the other contestant chooses to split (cooperate), the contestant who chose to steal receives the full jackpot.

Activity: In the YouTube video, the jackpot was 100,150 British Pounds. Construct a payoff matrix for Steven and Sarah and find the Nash Equilibrium.

Use the amount of money each contestant receives as the coefficients.

		Sarah	
		Split (cooperate)	Steal (defect)
Steven	Split (cooperate)	(50075, 50075)	(0, 100150)
	Steal (defect)	(100150, 0)	(0, 0)

Consider the payoff matrix from Steven's perspective. Because Steven has no idea which choice Sarah is going to make, Steven can only rely on what the consequences are based on his choice. If Steven assumes that Sarahsplits, splitting results in 50075 pounds while stealing earns him 100150 pounds. Therefore, if Sarahsplits, Steven's best option is to steal. Doing the same analysis based on the assumption that Sarahsteals, the choices do not make a difference to Steven. Regardless of his choice, Steven is left with 0 pounds if Sarah decides to steal. Sarah goes through the same train of thoughts, and she chooses to steal as well. Therefore, there are three Nash equilibria: both of them steal OR one of them steals and the other splits. In the payoff matrix, these correspond to the pairs of coefficients of (0, 0), (0, 100150) and (100150, 0).

Discussion question: Is this the best outcome for Steven and Sarah? Why or why not?

(0, 0) is clearly not the optimal outcome.(0, 100150) and (100150, 0) can be the optimal result for one of them. However, (50075, 50075) (where both cooperates) is the only strategy that does not leave either Steven or Sarah unhappy.

Discussion question: In the video, Steven and Sarah agreed to split. But Sarah betrayed Steven and walked away with 100,150 pounds.

Activity: Assume that Sarah knows the chance of Steven keeping his promise. Work on Sarah's strategy based on 10% and 90% as the chances of Steven keeping his promise.

If there is more information about Steven and/or Sarah, you can compare your expected outcome. Sarah knows that Steven is trusting and there is only 10% chance that Steven will betray her. This means that there is 90% chance that Steven will split while 10% chance that Steven will steal. Sarah can compute the expected amount of money earned based on these probabilities.

Expected amount of money for Sarah if she

- splits: 0.9 * 50075 + 0.1 * 0 = 45067.5
- steals: 0.9 * 100150 + 0.1 * 0 = 90135

If Sarah believes in these chances, then Sarah is better off betraying Steven. Let's assume the opposite case where Sarah does not trust Steven at all by flipping the probabilities. This means that there is only 10% chance that Steven will split and 90% chance that Steven will steal.

Expected amount of money for Sarah if she

- splits: 0.1 * 50075 + 0.9 * 0 = 5007.5
- steals: 0.1 * 100150 + 0.9 * 0 = 10015

This case also shows that Sarah is better off by stealing. Clearly, this is what Sarah ended up doing.

However, if Steven followed the same analysis and stole as well, they would have ended up with no money at all.

Coefficients in the Payoff Matrices

The non-zero sum games we have covered so far presents straightforward coefficients to use the payoff matrices. The coefficients in the Prank the Dean example are the number of days of community service, while the coefficients in the Split or Steal example are the amounts of money to walk away with. Sometimes, situations do not present clear coefficients as shown in these examples. And sometimes, you do not want to use the straightforward coefficients as presented.

Here is a modified version of the Prank the Dean example. Assume that Bill pranks the dean alone. Although the dean suspects Bill, he does not have any incriminating evidence. Bill is called into the dean's office to be interrogated. What should Bill do? Let's take a look at the following payoff matrix. (Note: in this payoff matrix we revert to the convention that a positive payoff is a win for the row player Bill, and a loss for the column player Dean.)

		Dean	
		No evidence	Incriminating Evidence
Bill	Deny Charges	10	0
	Admit Charges	7	5

If Bill denies the charges and the dean does not have any evidence, this is clearly the best case (10) for Bill. If Bill denies the charges and the dean has any incriminating evidence, this is the worst case (0) for Bill. If Bill admits the charges and the dean does not have any evidence, the dean probably will be a little more lenient (7) because he does not have any compelling evidence. If Bill admits the charges and the dean has any incriminating evidence, the dean will be harsher (5) because he is sure that Bill played the prank.

Unfortunately, Bill does not know if the dean has any incriminating evidence. There is a probability p that the dean has no evidence. Then, the probability that the dean has any incriminating evidence is $(1 - p)$. Bill's expected payoff for denying the charges is

$$10\,p + 0\,(1 - p) = 10p$$

whileBill's expected payoff for admitting the charges is

$$7\,p + 5\,(1 - p) = 5 - 2p$$

Compute the threshold probability p that helps Bill makes his decision.

$$10\,p \geq 5 - 2p$$
$$p \geq 5 / 12 \approx 0.42$$

If Bill believes that there is at least 42% chance that the dean does not have any evidence, he should deny all the charges. Otherwise, he should admit all the charges.

Now, let's alter the circumstances a little bit. Bill already has two strikes, and if he is found guilty of this prank, he is going to be expelled. This would have an impact on the payoff matrix.

		Dean	
		No evidence	Incriminating Evidence
Bill	Deny Charges	10	0
	Admit Charges	0	0

As long as the dean has confidence that Bill played the prank, he will expel Bill from the school. Therefore, admitting the charges is not an optionfor Bill because he does not want to get expelled. That is why we see 0s in the bottom row. If Bill denies the charges, he will not get expelled (10) if the dean does not have any evidence. However, if the dean has any incriminating evidence, he will still get expelled (0).

It is clear that Bill should deny any charges because the bottom row clearly is worse than the top row. Let's follow a similar analysis as before to see if it presents the same conclusion. Bill's expected payoff for denying the charges is

$$10\,p + 0\,(1 - p) = 10p$$

while Bill's expected payoff for admitting the charges is

$$0\,p + 0\,(1 - p) = 0$$

Compute the threshold probability p that helps Bill makes his decision.

$$10\,p \geq 0$$
$$p \geq 0$$

If Bill believes that there is at least0% chance that the dean does not have any evidence, he should deny all the charges. This means that he should deny all the charges regardless because there is no negative probability.

As demonstrated in these two examples, coefficients can be derived from situations and they may not be the same depending on a variety of factors regarding the situations.

Competition and Collusion Between Businesses: Price Wars and Price-Fixing

Consider two businesses that are vying for the same fixed market share of customers. They could be two gas stations on the same street corner, or two cable companies serving the same city. A business that sets a higher price will bring in more revenue per sale, but if their price is higher than their competitor's price, they will suffer reduced sales. To simplify the complexities of quantifying the payoffs, we will use an ordinal ranking of the outcomes for each business (with 4 being the best possible outcome and 1 being the worst possible outcome). We could represent this situation using the payoff matrix below.

		Gas Station B	
		High Price (cooperate)	Low Price (defect)
Gas Station A	High Price (cooperate)	(3, 3)	(1, 4)
	Low Price (defect)	(4, 1)	(2, 2)

Discussion Questions:

- In what sense can we characterize setting a high price as "cooperating" (with the other business) and in what sense can we characterize setting a low price as "defecting"?

  If both businesses agree to set a high price and stick to this agreement, it will benefit both businesses financially over the outcome of both setting a lower price.

- Do these rankings of payoffs seem to reflect the situation accurately? Which outcomes would correspond to "price-fixing" or "price war"? What is the difference between "cooperation" and "collusion"?

  We are assuming that the increase in sales when a station offers a lower price will more than make up for the decrease in revenue per sale. "Price-fixing" corresponds to both stations setting a high price. "Price war" corresponds to both stations setting a low price. "Collusion" is really just "cooperation" that is deemed to be illegal or unfair.

- Draw the "movement diagram" for the matrix, that is, from each cell, draw a vertical arrow in the direction that Gas Station A would prefer if starting from that cell, and draw a horizontal arrow in the direction that the Gas Station B would prefer to move from that cell. How does the movement diagram reflect the series of actions that the two businesses might choose?

		Station B		
		High Price (cooperate)		Low Price (defect)
Station A	High Price (cooperate)	(3, 3)		(1, 4)
	Low Price (defect)	(4, 1)		(2, 2)

The Pareto Optimal solution is for both gas stations to maintain a high price. However, this solution is unstable, that is, it is not a Nash equilibrium. The only Nash equilibrium is the outcome in which both stations set a low price.

- How does this situation resemble the Prank on the Dean game described earlier?
  This game has the same movement diagram, and in fact the same ordinal ranking of outcomes, as the Prank on the Dean (Prisoner's Dilemma).

- We have greatly oversimplified the situation by portraying each business as having only two possible pricing options. Suppose that we add a third option for each business—to set a medium-level price. The modified game matrix is shown below, but without ordinal payoffs. Draw the movement diagram for this game matrix. (You do not need to come up with any ordinal payoffs in order to draw the movement diagram.) Be prepared to give reasons justifying the direction in which you draw each movement arrow.

			Station B			
		High Price		Medium Price		Low Price
	High Price		?		?	
		?		?		?
Station A	Medium Price		?		?	
		?		?		?
	Low Price		?		?	

Assume that any time the two businesses have the same price level, either business will gain (in the short run) by lowering its price. The only time a business gains by raising its price is if the resulting price will still be lower than that of their competitor. These assumptions would give the movement diagram shown below.

			Station B			
		High Price		Medium Price		Low Price
	High Price		?		?	
		?		?		?
Station A	Medium Price		?		?	
		?		?		?
	Low Price		?		?	

- Suppose that instead of two businesses selling a product and choosing the price, we consider two sports teams (Team A and Team B) trying to sign a free agent and choosing whether to offer a low salary or a high salary. Show how the game matrix given originally for the gas stations (with two options for each station) can be relabeled to represent this “bidding war” game. Which type of salary offer corresponds to setting a high sales price, and which salary offer corresponds to setting a low sales price? Describe these options in terms of “cooperating” with the other sports team or “defecting”.

		Team B	
		Low Salary (cooperate)	High Salary (defect)
Team A	Low Salary (cooperate)	(3, 3)	(1, 4)
	High Salary (defect)	(4, 1)	(2, 2)

In this example, cooperation or collusion between the sports teams would correspond to each team offering low salaries. Defecting corresponds to offering higher salaries. It is in the teams' mutual best interest to agree on offering low salaries, but the only Nash equilibrium is for both teams to offer high salaries.

Optional or Homework: Weird Haircuts(Application of Prisoner's Dilemma)

http://www.freakingnews.com/funny-pictures/haircut-pictures.asp

You and a friend decide it would be cool if you both come to school tomorrow with weird haircuts. But you begin to worry. It would be embarrassing to come to school with a wild haircut if your friend does not. Prepare a payoff matrix for this situation, showing payoffs for you and your friend under the different possible outcomes, and then analyze the likely behavior by you and your friend.

Response:(Note: This is Activity 2.1 Question 3 on page 83 of Mathematics: Modeling Our World, published by COMAP. See the Resources list for publication details.)

The possible outcomes are that both of you have wild haircuts, or that only one of you has a wild haircut or that neither of you has a wild haircut. In the first case, both of you have high utility for pulling a trick off together. In the second case, the one who decided not to have a wild haircut has a higher utility because the other made a fool of himself/herself. In the third case, both of them have a lower utility compared to the first case.

Assigning reasonable utilities to each scenario is crucial as it changes the outcome of the analysis. Here's an example analysis.

		Your friend	
		Wild haircut	Normal haircut
You	Wild haircut	8, 8	2, 12
	Normal haircut	12, 2	4, 4

Assume that the probability that your friend comes to school with a wild haircut is p. Let's compute the expected outcome for you for each strategy.

1. Wild haircut = $8 * p + 2 * (1 - p) = 2 + 6p$
2. Normal haircut = $12 * p + 4 * (1 - p) = 4 + 8p$

If $2 + 6p$ is greater than $4 + 8p$, then it is beftter to do a wild haircut. This leads to the following computation.

$2 + 6p > 4 + 8p$

$2p < -2$

$p < -1$

Since probabilities cannot be negative, there is no reason to do a wild haircut with the given coefficients in the payoff matrix.

Unit 5: Non-Zero-Sum Games, Continued: Halo or Just Dance?

Halo or Just Dance?

Alice and Bob are siblings, and they would like to spend an afternoon playing Xbox. Unfortunately, they want to play different games. Alice's favorite game is Just Dance while Bob's is Halo. They start arguing about which game to play. Then their parents intervene and give them an ultimatum. At the count of three, Alice and Bob are to choose a game to play. If both of them choose the same game, they get to play the game. If not, neither can play Xbox. (You can substitute other games in place of Halo and Just Dance by surveying the class quickly to identify their favorite games.)

Activity: Have half of the class play the role of Bob and the other half to play the role of Alice. At the count of three, raise your hand if you choose Halo.

Discussion question: What did you choose to do? What influenced your decision?

Expected responses: Students may argue that it is 50/50. There is no way to know what the other sibling would choose.

Constructing a payoff matrix is trivial up to labeling the strategies for each sibling. The coefficients for strategies where the siblings choose different games are zeroes because neither of them will get to play Xbox. However, what are the coefficients when they do choose the same game? Let's make some assumptions. Both Alice and Bob get the same satisfaction of 10 when they get to play his/her favorite game. Otherwise, their satisfaction level is 5 when they get to play the sibling's favorite game. As before, we can characterize each option in terms of either "cooperating" (agreeing to play the sibling's favorite game) or "defecting" (putting one's own preference ahead of the sibling's preference). (The order in which we list the options of each player does not affect the final analysis, but sometimes makes it easier to compare this game matrix with the payoff matrix for other similar games.)

<table>
<tr><td colspan="2" rowspan="2"></td><td colspan="2">Alice</td></tr>
<tr><td>Halo (cooperate)</td><td>Just Dance (defect)</td></tr>
<tr><td rowspan="2">Bob</td><td>Just Dance (cooperate)</td><td>(0, 0)</td><td>(5, 10)</td></tr>
<tr><td>Halo (defect)</td><td>(10, 5)</td><td>(0, 0)</td></tr>
</table>

Let's analyze the game from Alice's perspective, using the "maximin" approach to calculating a Nash equilibrium that we presented in Unit 2, and assuming that Alice has no idea which game Bob is going to choose. If Bob chooses Just Dance, then Alice gets satisfaction of 10 for choosing Just Dance and 0 otherwise. If Bob chooses Halo, then Alice gets the satisfaction of 5 for choosing Halo and 0 otherwise. If we use variable q forAlice's probability of choosing Halo, then we can compute the expected payoff to Alice of each of Bob's game choices as a function of q.

Since Alice's goal is to make Bob's pay the same no matter which Xbox game he chooses, we will use Bob's payoff coefficients rather than Alice's payoffs.

If Bob chooses Just Dance, his payoff is: $0q + 5(1-q) = 5 - 5q$.
If Bob chooses Halo, his payoff is: $10q + 0(1 - q) = 10q$.

We equate the two expressions to each other to find the value of q that makes Bob's payoff the same regardless of which option he chooses.

$5 - 5q = 10q$
$15q = 5$
$q = 1/3$ *or* 33.3%, *so* $1 - q = 2/3$ *or* 66.7%.

Instead of 50/50, the Nash equilibrium has Alice choosing Halo 33.3% of the time, and choosingJust Dance 66.7% of the time! The analysis for Bob is exactly symmetrical, and it would show that the Nash equilibrium has Bob choosing Just Dance 33% of the time andHalo 66.7% of the time—again, using the "maximin" approach.

Let's check the compound probability to compare 50/50 strategy vs. 67%/33% strategy.

Chance of playing Xbox: Alice and Bob both choose Halo OR Alice and Bob both choose Just Dance.

The 50/50 strategy gives probability of playing = (1/2) * (1/2) + (1/2) * (1/2) = ½

The expected value of the 50/50 mixed strategy for Alice is
(1/2)*((1/2)*(5) + (1/2)* (1/2)*(10) = 15/4 = 3.75.

For comparison, the 67%/33%mixed strategy gives probability of playing
= (1/3) * (2/3) + (2/3) * (1/3) = 4/9

The expected value of the 67%/33% mixed strategy for Alice is
(1/3) * (2/3) *(10) + (2/3) * (1/3) *(5) = 30/9 = 3.33.

Following the 67%/33% mixed strategy is clearly worse than flipping coins (the 50/50 mixed strategy), whether from the point of view of probability of playing Xbox or expected value to each player! This is due to the fact that the goal of the maximin approach is to minimize the losses rather than maximize the satisfaction.

If eachplayer chooses their own favorite game, or if each player switches to their sibling's favorite game, then the two siblings would choose different games, and would not get to play either game! Therefore, the maximin approach does not give us a clear optimal strategy for each player. This example demonstrates a fundamental difference between zero sum games and non-zero sum games as regards the usefulness of the Nash equilibrium! That is, we can still calculate the Nash equilibrium for a non-zero sum game, but we should not expect it to give the optimal outcome for each player in the same way that the Nash equilibrium does in a zero sum game. This example also illustrates the potential value of communication and trust in a non-zero sum game.

Discussion question: What if Alice knows that Bob is selfish and unlikely to yield? What if Alice knows that Bob is more likely to yield? How would that change Alice’s strategy?

Assume Bob is selfish and unlikely to yield. 90% chance that Bob would stick to his favorite game Halo. The computation for Alice’s expected satisfaction changes.

Alice’s expected satisfaction for choosing Halo is: $0.9 * 5 + 0.1 * 0 = 4.5$

Alice’s expected satisfaction for choosing Just Dance is: $0.9 * 0 + 0.1 * 10 = 1.0$

Therefore, Alice is better off yielding to Bob’s choice.

Discussion question: Assume Bob is selfless and likely to yield. Only 10% chance that Bob would stick to his favorite game Halo. The computation for Alice’s expected satisfaction changes.

Alice’s expected satisfaction for choosing Halo is: $0.1 * 5 + 0.9 * 0 = 0.5$

Alice’s expected satisfaction for choosing Just Dance is: $0.1 * 0 + 0.9 * 10 = 9.0$

Therefore, Alice is better off sticking to her favorite game.

Activity: What if Alice and Bob were allowed to communicate before making the final choice? Pair up with your neighbor and play the roles of Alice and Bob.

Discussion question: What was the outcome? Do you get to play a game? Did communication improve your ability to reach an outcome satisfying both players, or did you reach a stalemate?

Many things can happen. For example, one of them can decide to limit the possible strategies that the other player can choose. For example, if Bob hides Just Dance where Alice cannot find it, this limits the choice for Alice. This is an extreme case of the previous example where we discussed having more information about the other player. Instead of having known probabilities of the other player’s action, we know what their action is going to be.

		Alice	
		Halo	Just Dance
Bob	Halo	10, 5	0, 0
	~~Just Dance~~	~~0, 0~~	~~5, 10~~

Looking at the row of Bob’s choice of Halo, Alice’s satisfaction scores of choosing Halo and Just Dance are 5 and 0 respectively. It is in Alice’s interest to choose Halo.

Activity: We all know that Alice and Bob will get to play Xbox on days other than today. This means that Alice and Bob will face these choices/strategies over and over again. Does repeating change their strategy? Let us assume that they are allowed to communicate.

Discussion question: Did your strategies change? If so, how so? If not, why not?

Repetition of this game facilitates players' cooperation. The players can come to an agreement to alternate the games they play. For example, today, Alice and Bob agree to play Halo while tomorrow, they will agree to play Just Dance.

Discussion question: How long can this arrangement last?

This agreement can potentially last forever. But, once Alice or Bob breaks the agreement, it reverts back to the analysis of being able to play just one day.

<u>Homework:</u> Suppose that Alice and Bob have a second Xbox in the house, so that if they choose different games, they still get to play, but not with each other, which is less enjoyable for them than playing their second preference game with a sibling. Suppose that this changes the payoff matrix to the one shown below. Find the Nash equilibrium (i.e., the optimal mixed strategy according to the maximin approach) for each player, and then discuss whether this is really the best mixed strategy for each player, with or without communication between the siblings allowed before they make their Xbox game choice.

		Alice	
		Halo (cooperate)	Just Dance (defect)
Bob	Just Dance (cooperate)	(0, 0)	(5, 10)
	Halo (defect)	(10, 5)	(3, 3)

Again, use q to represent the probability that Alice chooses Halo.

If Bob chooses Just Dance, his payoff is: $0q + 5\ (1 - q) = 5 - 5\ q.$
If Bob chooses Halo, his payoff is: $10q + 3\ (1 - q) = 7q + 3.$

We equate the two expressions to each other to find the value of q that makes Bob's payoff the same regardless of which option he chooses.

$5 - 5q = 7q + 3$
$12\ q = 2$
$q = 1/6$ *or* 16.7%, *so* $1 – q = 5/6$ *or* 83.3%.

The expected value of the 16.7%/83.3% mixed strategy for Alice and Bob is
(1/6) * (5/6) *(10) + (5/6) * (1/6) *(5) +(5/6) * (5/6) *(3) = 150/36 = 4.17.

The expected value of the 50%/50% mixed strategy for Alice and Bob is
(1/2) * (1/2) *(10) + (1/2) * (1/2) *(5) +(1/2) * (1/2) *(3) = 18/4 = 4.50.

Again we see that the Nash equilibrium is not really the best mixed strategy for this non-zero sum game.

<u>Game of Chicken</u>

Consider a situation where two people are driving towards each other in a collision course. If they keep driving straight, they will die in a crash. If both drivers decide to swerve, they will not collide either. If one driver decides to swerve and the other continues to drive straight, no one will be hurt. However, the driver who decides to swerve will be called "chicken." This is why the situation is referred to as the "Game of Chicken."

Activity: Construct the payoff matrix for the game of chicken and find the Nash Equilibrium.

		Driver #2	
		Swerve	Straight
Driver #1	Swerve	(0, 0)	(-1, +1)
	Straight	(+1, -1)	(-10, -10)

The coefficients in the matrix do not have to match the values above. However, it should reflect the following:

- (swerve, swerve) should be close to (0, 0) since neither "wins"
- (swerve, straight) and (straight, swerve) should clearly show that the driver who chose to drive straight has a higher coefficient.
- (straight, straight) should have the lowest coefficients since both drivers will get hurt or die.

The Nash Equilibria are (-1, +1) and (+1, -1) where one of the drivers chooses to swerve while the other driver keeps straight.

		Driver #2		
		Swerve		Straight
Driver #1	Swerve	(0, 0)	□	(-1, +1)
		□		□
	Straight	(+1, -1)	□	(-10, -10)

Activity: What should a driver do?

Let's analyze in the perspective of Driver #1. The analysis would be the same for Driver #2. Let's use p as the probability that Driver #2 will swerve. This means that $(1 - p)$ is the probability that Driver #2 will go straight. Since p is unknown, we can calculate the expected outcome for each strategy for Driver #1.

Swerve: $p * 0 + (1 - p) * -1 = p - 1$
Straight: $p * 1 + (1 - p) * -10 = 11p - 10$

What p values result in the swerve strategy to yield a higher expected outcome than the straight strategy?

$p - 1 > 11p - 10$
$9 > 10p$
$p < 0.9$

If Driver #1 believes that the probability that Driver #2 will swerve is less than 0.9, it is better for Driver #1 to swerve. If the probability that Driver #2 will swerve is greater than 0.9, Driver #1 is better off to drive straight.

Activity: In the previous example, we did not specify the direction in which a driver can swerve. If the drivers swerve in the same direction (left and right OR right and left), they may still collide and get hurt. Construct a payoff matrix by expanding the swerve strategy into left and right strategies.

		Driver #2		
		Straight	Left	Right
Driver #1	Straight	(-10, -10)	(+1, -1)	(+1, -1)
	Left	(-1, +1)	(0, 0)	(-5, -5)
	Right	(-1, +1)	(-5, -5)	(0, 0)

Again, the coefficients in the matrix do not have to match the values above. However, it should reflect the following:

- (left, left) and (right, right) should be close to (0, 0) since the drivers avoid collision.
- (straight, straight) should have the lowest coefficients since both drivers will get hurt or die.
- (left, right) and (right, left) should have low coefficients but higher than (straight, straight) since drivers may collide but probably will not get hurt as much as a head-on collision caused by (straight, straight)
- (left, straight), (right, straight), (straight, left) and (straight, right) should clearly show that the driver who chose to drive straight has a higher coefficient.

Finding Dates

Suppose that Bob, Charlie, Alicia, and Danielle are all friends. Bob and Charlie would both prefer to ask Alicia to a dance (maybe because Alicia is better at math); however, they would rather go with Danielle than not go at all. They are fairly confident that if the two boys ask two different girls, each of the girls will say "yes." However, if they both ask the same girl, each boy has only a 50% of receiving a "yes" answer, and the other girl will feel slighted and say "no" if the other boy asks her later.

- Quick discussion question: How can we model this situation using a game matrix? What is the order of preference of outcomes for each boy?

To create the matrix model, we will use ordinal ranks rather than cardinal values. For each boy, the best outcome would be if he gets to go out with Alicia and the other asks Danielle. The second best outcome would be if he gets to go with Danielle and the other asks Alicia. The third best outcome would be if each asks Alicia; although they would like to go with her, they each only have a 50% chance of going with anybody in this scenario. And the worst scenario for each is if both ask Danielle; they still have only a 50% chance of going with anybody, and even if they do go, it will be with their less-desired date. The matrix below summarizes these preferences.

		Charlie	
		Ask Alicia	Ask Danielle
Bob	Ask Alicia	(2, 2)	(4, 3)
	Ask Danielle	(3, 4)	(1, 1)

- Quick discussion question: Based on this game matrix, what should Bob and Charlie do?

- Possible response: Bob and Charlie could agree to ask two different girls, and flip a coin to see which one gets to ask Alicia. Or, if they think Alicia is really more likely to say "yes" to one of them, that person could ask Alicia.

Unit 6: Applications of Game Theory to Homeland Security and Other Topics

Air Strike or Diplomacy? In September 2013, U.S. President Obama accused Syrian President Bashar al-Assad of using chemical weapons, and threatened to punish Syria with a U.S. air strike if Syria did not reveal and surrender their chemical weapons. In this example, it is more difficult to "quantify" the payoffs under each combination of strategies chosen—but we can still describe outcomes as positive or negative for each country, and assign number values in relative order of preference. Also, a "win" for one country is not necessarily a "loss" of the same amount for the other country. One way to represent this situation is given in the payoff matrix below. For example, the entry (+2, -2) means that if Syria surrenders weapons while the U.S. launches an air strike, this outcome has values of +2 for the U.S. and -2 for Syria.

One way to model this situation might be called the "Aggression" model:

Aggression model		Syria	
		Surrender weapons	Keep weapons
United States	Diplomacy	(+1, +1)	(-2, +2)
	Air strike	(+2, -2)	(-1, -1)

Another way to model this situation might be called the "Negotiation" model:

Negotiation model		Syria	
		Surrender weapons	Keep weapons
United States	Diplomacy	(+2, +1)	(-2, +2)
	Air strike	(-1, -1)	(+1, -2)

Discussion Questions:

- Interpret in words (good, bad, better, worse, etc.) the payoffs for each country under the four possible combinations of strategies chosen, in each of the two models. Do you agree with these values and rankings? Which of the two models makes more sense to you, if either? Which of the two models might be taking world opinion and reputation into account?

  According to the "Aggression" model, the ranking of outcomes from the U.S. viewpoint, from best to worst, would be: (Air Strike, Surrender) or (A, S),
  (D, S), (A, K), and then (D, K).The ranking of outcomes from the Syrian viewpoint according to the "Aggression" model would be: (Diplomacy, Keep) or (D, K), (D, S), (A, K), and then (A,

S). The Aggression model rankings imply that military power is the primary value, and ignore world opinion and reputation.

According to the "Negotiation" model, the ranking of outcomes from the U.S. viewpoint, from best to worst, would be: (D, S), (A, K), (A, S), and then (D, K). The ranking of outcomes from the Syrian viewpoint according to the "Negotiation" model would be: (D, K), (D, S), (A, S), and then (A, K). The Negotiation model rankings take world reputation into account.

- What motivation does each country have to choose each option? Draw the movement diagram for each matrix. How does the movement diagram correspond to the name given to each model? In particular, if Syria keeps its weapons and the U.S. is using diplomacy only, what motivation does either country have to unilaterally change its choice of "strategy" in each model?

Aggression model		Syria	
		Surrender weapons (cooperate)	Keep weapons (defect)
United States	Diplomacy (cooperate)	(+1, +1)	(-2, +2)
	Air strike (defect)	(+2, -2)	(-1, -1)

In the Aggression model, Air Strike is a dominant strategy for the U.S., and Keep Weapons is a dominant strategy for Syria. This means that the outcome
(Air Strike, Keep Weapons) is a saddle point/Nash equilibrium, even though it is not Pareto optimal. The movement diagram and relative order of outcomes in the Aggression model is the same as in the Prisoner's Dilemma.

Negotiation model		Syria	
		Surrender weapons (cooperate)	Keep weapons (defect)
United States	Diplomacy (cooperate)	(+2, +1)	(-2, +2)
	Air strike (defect)	(-1, -1)	(+1, -2)

In the Negotiation model,neither side has a dominant strategy, and there is no pure strategy Nash equilibrium. We can calculate the mixed strategy Nash equilibrium, which turns out to be the U.S. using Diplomacy 50% of the time, and Syria surrendering weapons 50% of the time. The movement diagram and relative order of outcomes in the Negotiation model is a cycle, very similar to our second model of the Management versus Employee's Union example later in Unit 6. (In fact the movement diagram is exactly the same, relative to cooperating versus defecting, if we recast Management as Syria and the Union as the U.S.). Each side's choice of strategy in this situation may be highly influenced by their belief regarding what strategy the other side will choose.

In particular, if Syria keeps its weapons and the U.S. is using diplomacy only, the U.S. has a strong motive (according to either model) to launch an Air Strike, whereas Syria does not have a motive to change its strategy unilaterally unless it believes that the U.S. is about to change strategies.

- Suppose we view each day of the conflict as a "round" of play over several weeks. How is a "repeated" game different from playing just one round?
  This game is probably best modeled as a "sequential" game, in the sense that either side can maintain a certain strategy (Diplomacy, or Keep Weapons in particular) for a week at a time, but will still have the option to change strategies in a later week, possibly in response to the opponent's previous move.

- Investigate the Cuban missile crisis of 1962 as described in the Stephen Brams web article on a "Theory of Moves" found at:
  http://plus.maths.org/content/game-theory-and-cuban-missile-crisis
  How does the 2013 Syrian standoff compare to the 1962 Cuban missile crisis? How does either of these situations compare this game to the game of "Chicken" discussed earlier in Unit 5? See the article for discussion points.

<u>Hunting Down Terrorists (Stag Hunt)</u>

Two countries are trying to remove a terrorist threat that exists within the borders of one of the countries. If they work together, they have a high probability of successfully removing the threat which will increase the popularity of the leaders within the countries and the countries' influence on a global level. If one country gets distracted by other political concerns and pursues those instead, that country's leader will still have a small increase in political influence at home, but will not get the credit with the global community, while the other countries' leader, who is still trying to remove the terrorists, will not succeed. If both countries abandon the attempt to remove the terrorists and pursue their own agenda, they will each have some level of success with those agendas, but less than if they had worked together to remove the terrorists as a threat to the global community. *Historical note:* This game was originally described as "Stag Hunt" by social philosopher Jean-Jacques Rousseau in <u>Discourse on Inequality</u>, published in 1755. Philosopher David Hume discussed similar examples (two oarsmen in a boat, two neighbors wanting to drain a flooded meadow), also in the mid-18th century.

Discussion Questions:

- What are the options of each country? Which option would you characterize as cooperating, and which as defecting?

Each country can choose to pursue the terrorists with the other country, which is cooperating, or pursue a separate agenda, which is defecting.

- Use ordinal values (4 as the best outcome, 3 as next best, down to 1 as the worst possible outcome) as payoff entries and write the game matrix for this situation, listing the "cooperate" option for each country first and the "defect" option second. Then draw the "movement diagram" for this matrix. Identify all equilibria. What action do you expect each country to take, and what could go wrong with your prediction?

		Country B	
		Pursue Terrorists (cooperate)	Pursue Own Agenda (defect)
Country A	Pursue Terrorists (cooperate)	(4, 4)	(1, 3 or 2)
	Pursue Own Agenda(defect)	(3 or 2, 1)	(2, 2)

The outcome of both pursuing the terrorists is an equilibrium, and we would expect this to be a stable outcome, because it is in the best interest of both countries, both collectively and individually. However, if one of the countries gets distracted, then the other country is best off switching to internal matters as well. Both countries pursuing internal affairs is the other equilibrium outcome.

- Can you think of other situations that resemble the Stag Hunt in terms of the movement diagram, characterizing the first option as cooperation and the second option as defection?

- Military alliances: two neighboring countries can form an alliance, and thereby provide greater mutual defense than either country could individually. (A listing of the alliances before any large scale conflict, and the subsequent actions taken by the countries involved could lead to a large number of scenarios to model.)
- Business mergers: two businesses can often improve their profits by merging rather than competing individually.
- Biology: some species such as orcas and wolves hunt more effectively in packs. Others cooperate to watch for predators (prairie dogs and meerkats) or care for young (differences between mule deer and white tail deer). (See also “The Game of Life: Evolutionary Game Theory:” Fefferman, Gabric and Kupetz, published by COMAP for more examples of game theory in animal behavior.)

Game Theory in *The Hunger Games*: Swallow Berries or Bluff?

The fictional book and movie “The Hunger Games” portrays a repressive government that annually forces twelve new contestants to entertain the populace by fighting a bloody battle to the death until only one victor remains. Naturally, alliances arise between contestants as the battle proceeds, just like on the TV show “Survivor,” but in the end, some contestant eventually betrays the others to win the competition. In the year this story takes place, only two “tributes” remain: Katness and Peeta. The government expects them to duel until one is killed, either by the other tribute or by the natural hazards introduced by the “GameMaker.” Katness and Peeta have survived largely because of their alliance, and do not wish to fight. Katness decides to challenge the rules of the game: she and Peeta hold poisonous Nightlock berries up to their mouths and threaten to commit suicide rather than fight each other. The moment is climatic: Will they carry through and take their own lives? Are they only bluffing? How will the GameMaker respond to their threat?

No spoilers here, in case you haven’t read the book or seen the movie, but we will proceed to model this situation using Game Theory. One way to view the situation is to model it as Katness versus the government’s GameMaker. We present the options of each in table form:

		GameMaker	
		Change rules to allow two victors	Don’t allow two victors
Katness&Peeta	Swallow berries	Won’t happen if GM acts quickly	Both tributes die, govt has no victor
	Don’t swallow	Both tributes live, govt loses conflict	Tributes fight, govt wins conflict

Questions for Discussion:

- Which outcome would Katness prefer? Which outcome would the GameMaker prefer? In what order would each rank their preference of outcomes?
- How might each "player" (Katness and the GameMaker) be influenced by an estimate of the probabilities that the other player might choose either strategy?
- This conflict is fictional, but resembles some real-life conflicts: for example, negotiations between management and an employees' union, or between two countries on the brink of war over territorial disputes. Discuss these and other similar real-life conflicts. Who are the "players"? What strategy options does each player have? What is the payoff of each outcome, and the likelihood that each player will choose one or the other strategic option?
- Suppose that a terrorist has been surrounded but not yet apprehended because the terrorist is holding a hostage and threatening to kill the hostage if his demands are not met. How does this situation resemble the *Hunger Games* scenario?
- This conflict only occurs once, but many real-life conflicts occur repeatedly: negotiations as described above, two animals fighting over food or mating privileges, etc. What combinations of strategies might each side choose if there were a way for this game to be repeated (say, if the tributes suffered some significant loss without dying or if the tributes were forced to take part in next year's games)?

The Tragedy of the Commons

Suppose that a resource is freely available to all who wish to make use of it, at no charge. The resource may be a pasture in medieval England, or the Boston Commons in colonial times, or the Brazilian rainforest in modern times. Any farmer with sheep or goats needing pasture for grazing will profit by making use of this free resource rather than using his or her own land. In addition, the more animals the farmer brings to graze on this common pasture, the greater profit the farmer will realize from this free resource.

The "tragedy" in this arrangement is that while any individual's financial interest is best served by using the common resource, in the long run this common pasture will be overgrazed after some threshold is reached, and the pasture will lose its value, not just for that farmer but for all the farmers who were using that resource.

This "tragedy of the commons" was first mentioned in an 1833 pamphlet by William Forster Lloyd, and popularized in a 1968 *Science* journal article by biologist Garrett Hardin. This general scenario is often used to describe situations in which individual interest appears to be in conflict with the common, collective good. Some analysts would view this as an argument for government regulation of common resources; others would view this as an argument for privatization.

We will attempt to describe this scenario game-theoretically. The most accurate portrayal would be an "n-person game" where n is some large positive integer that represents the total number of farmers (or resource users). However, for simplicity, we will consider this as a 2-person game, with one farmer trying to decide what to do while the other (n-1) farmers decide simultaneously.

The options could be either to use or not use the resource. Instead, we will suppose that each farmer is going to use the resource, but has the option of either paying a voluntary, suggested user fee/donation, or not paying the fee. We denote the fee as "C" ("Cost"). Suppose that if some minimum number M of farmers pay the voluntary fee, the resource will retain its value for all, but if this minimum level of support M is not reached, the resource will completely lose its value. That is, assume that the common resource has value B ("Benefit") to each user (regardless of whether or not they pay the fee) if the minimum level of support is reached, but has value 0 if the minimum level of support is not reached.

The dilemma is most vividly portrayed from the point of view of Farmer M. That is, suppose that the number of other farmers who have chosen to voluntarily pay their share of the maintenance cost is right at the threshold of the minimum necessary support: if Farmer M pays the fee, the resource value will be maintained, and if Farmer M does not pay the fee, the resource will lose its value due to insufficient maintenance and support. We can represent this interpretation of the situation using the payoff matrix below.

Assuming Farmer M is on threshold of minimum support needed		Other farmers (as a group)	
		Enough pay fee (cooperate)	Not enough pay (defect)
Farmer M	Pay fee (cooperate)	(B - C, B - C)	(-C, 0)
	Don't pay (defect)	(0, -C)	(0, 0)

Discussion questions:

- Assume that the benefit B is greater than the cost C. Replace the number values in the matrix with ordinal rankings (4 for best outcome, 3 for next best, down to 1 for worst outcome). Then draw the movement diagram for this game matrix, and discuss the implications for what actions might be taken by Farmer M as an individual, or by the other farmers as a collective group. How does this compare to the Stag Hunt?

Assuming Farmer M is on threshold of minimum support needed		Other farmers	
		Enough pay fee (cooperate)	Not enough pay fee (defect)
Farmer M	Pay fee (cooperate)	(4, 4)	(1, 2.5)
	Don't pay (defect)	(2.5, 1)	(2.5, 2.5)

If Farmer M is really on the threshold of maintaining the value of the common resource (and he knows this), then it is in his best interest to pay the fee. Same for the other farmers (as a collective group). This situation is essentially the same as the Stag Hunt, as described by the movement diagram and ordinal payoffs.

- Rewrite the game matrix if Farmer M is not on the threshold of the minimum level of support. That is, rewrite the game matrix assuming that whether or not there are enough other farmers paying the fee to support the common resource will not be affected by Farmer M's choice as to whether to pay the fee or not. Use B for Benefit and C for Cost as before. Then draw the movement diagram for this new game matrix, and discuss the implications for Farmer M's choice of actions.

Assuming Farmer M is not on threshold of minimum support needed		Other farmers	
		Enough pay fee (cooperate)	Not enough pay fee (defect)
Farmer M	Pay fee (cooperate)	(B - C, B - C)	(-C, 0)
	Don't pay (defect)	(B, B - C)	(0, 0)

If Farmer M is really not on the threshold of maintaining the value of the common resource (and he knows this), then it is in his best interest to not pay the fee. He will still reap whatever the benefit of the resource is, but at no cost. In fact, the "Don't Pay" strategy dominates the "Pay Fee" strategy for Farmer M.

Note that considered as a group, the other farmers are always better off when the minimum necessary number pay the maintenance fee. However, if every farmer in the group makes the same analysis as Farmer M from their own individual viewpoint, then not enough farmers will pay the fee, and the common resource will lose its value. Another way to express this "tragedy" is that the group would always choose to "cooperate," but the individual would always choose to "defect."

<u>Leader or Follower</u>: Suppose two nominally allied countries(A and B) are interested in acquiring resources from a third country (C). (For example, these resources could be oil, fresh water, precious minerals, or permission to build a strategically located military base.) Each allied country A and B has the option to either interfere with the internal politics of C, or to stand aside and not interfere. If countries A and B both interfere, then C's stability will erode to the point that C will move into the influence of another country (D), which is antagonistic to both A and B. If only one interferes, that country will have greater influence in C, but if neither interfere, country D's influence will grow, although not to the point of the chaos caused by both interfering.

Exercise: Determine the order of preference of outcomes for each country (A and B), model this situation using a game matrix, and then predict the behavior of each country. How is this situation related to the “Finding Dates” example?

		Country B	
		don’t interfere	interfere
Country A	don’t interfere	(2, 2)	(3, 4)
	interfere	(4, 3)	(1, 1)

For each country, their preference order from best to worst would be: interfere while the other doesn’t, don’t interfere while the other does, don’t interfere while the other doesn’t as well, and both interfere. The optimal solution would be to determine in some fair way which country intervenes and then have the other country stand aside.

This is similar to the Finding Dates example in that the best overall outcome is for the countries to cooperate, but is different in the sense that the scenario in which both pursue their highest preference outcome (interfering) leads to the worst payoff rather than the second worst payoff.

Applying for Jobs: Alice and Bob are each deciding whether to apply for a job at Company C or Company D. Company C offers a higher salary, so is more desired by both Alice and Bob. If the two apply for two different jobs, they will each get the job they applied for at full salary. If the two apply for the same job, they will each get a half-time position in that capacity and will receive half the normal salary.

Exercise: Determine the order of preference of outcomes for each applicant (Alice and Bob), model this situation using a game matrix, and then predict the behavior. How is this situation related to the “Finding Dates” example? How is this situation different? How may this model be unrealistic?

		Bob	
		Apply to C	Apply to D
Alice	Apply to C	(2, 2)	(3, 4)
	Apply to D	(4, 3)	(1, 1)

For each applicant, their preference order from best to worst would be: apply to C while the other applies to D, apply to D while the other applies to C, apply to C while the other applies to C (so they each work half-time for a higher salary), and both apply to D (so they each work half-time for a lower salary). The optimal solution would be to determine in some fair way which applicant applies to C (such as a coin flip, or more realistically, agreeing on which applicant is better qualified) and then having the other person apply to the other company.

This is similar to the Finding Dates example in that the best overall outcome is for the applicants to cooperate. However, it seems more likely and more realistic that Company C will hire the better-qualified applicant full-time, and then the second applicant can apply at and be hired by Company D later. Most applicants apply to several companies at once, and most companies will not feel “slighted” if applicants would rather work at a different company that offers a higher salary.

Game Theory Glossary

Dominated strategy: We say that a strategy A is "dominated" by another strategy B for a given player if the payoff for that player under strategy B is greater than or equal to the payoff under strategy A, regardless of which strategy the other player chooses.

Equilibrium (or "Nash equilibrium"): We call an outcome a "Nash equilibrium" if neither player can improve their payoff by changing their choice of strategy unilaterally. Such an equilibrium may be a combination of either pure strategies or mixed strategies. John Nash proved that such an equilibrium always exists in any two-person game, and showed how to calculate this equilibrium using the minimax approach. However, in practice this equilibrium is more useful in zero sum games than in non-zero sum games.

Expected value: The "expected value" (or just "value") of a game for a given player is the sum of the products of the probability that each given combination of strategies will be chosen by the players, multiplied by the payoff for that combination of strategies. If no reference to the players' chosen mixed strategies is given, it is assumed that each player is using their optimal mixture of strategies.

Fair: A zero-sum game is called "fair" if the expected value of the game is zero for each player, when each player uses their optimal mixture of strategies. A non-zero-sum game is called "fair" if each player has the same expected value.

Optimal strategy: A (pure or mixed) strategy is called "optimal" (or "best") if that mixture of strategies optimizes the payoff to that player (ignoring any gain or loss to the other player).

Pareto optimal outcome: We call an outcome "Pareto optimal" if there is no other outcome that improves one player's payoff without making the payoff worse for at least one other player.

Payoff matrix: a table or matrix listing the payoffs for each player under each possible combination of strategies chosen by the players.

Principle of Indifference:the principle that each player should choose a (pure or mixed) strategy that guarantees a certain minimum payoff regardless of which strategy the other player uses, thereby making both players indifferent to which strategy the other player uses. This implies that the maximin strategy is the safest strategy for the Row player, and the minimax strategy is the safest strategy for the Column player.

Zero-sum game: a two-person game is called "zero-sum" if under any combination of strategies chosen, the payoff for one player is equal and opposite to the payoff for the other player.

References

- The Mathematics of Games, John Beasley, Dover Publications, 2006.
- Luck, Logic, & White Lies: The Mathematics of Games,JorgBewersdorff, translated by David Kramer, AK Peters, 2005.
- Mathematics: Modeling Our World, 2nd Edition, Course 2, Chapter 2: Strategies, Consortium for Mathematics and Its Applications (COMAP), 2011.
- Discrete Mathematics: Modeling Our World (Formerly Discrete Mathematics through Applications), Nancy Crisler and Gary Froelich, 4th Edition, Chapter 7: Matrix Models, Consortium for Mathematics and Its Applications (COMAP), 2014.
- Thinking Strategically: The Competitive Edge in Business, Politics, and Everyday Life, by Avinash K. Dixit and Barry J. Nalebuff, 1993.
- Games of Strategy, Avinash K. Dixit and Susan Skeath, W.W. Norton and Company, 1999.
- Taking Chances: Winning with Probability, John Haigh, Oxford Press, Second Edition, 2003.
- Ross, Don, "Game Theory", *The Stanford Encyclopedia of Philosophy* (Winter 2012 Edition), Edward N. Zalta (editor), http://plato.stanford.edu/archives/win2012/entries/game-theory, last accessed: January 12, 2014.
- Shor, Mikhael (editor), “Game Theory .net”, http://www.gametheory.net, last accessed: January 12, 2014
- A Gentle Introduction to Game Theory, Saul Stahl, American Mathematical Society, Mathematical World: Volume 13, 1999.
- Game Theory and Strategy, Philip D. Straffin, Mathematical Association of America, New Mathematical Library, 1993.

Internet Resources

Paper based on real MLB pitching and NFL play selection data:
http://www.nber.org/digest/oct09/w15347.html

When should animals share food? Game theory applied to kleptoparasitic populations with food sharing: **http://beheco.oxfordjournals.org/content/23/5/977.abstract**

Game theory research in baseball (at the SABR conference):
http://www.hardballtimes.com/main/article/game-theory-is-the-next-moneyball/

Game theory applied to base stealing in baseball:
http://theincidentaleconomist.com/wordpress/the-game-theory-within-the-game/

Weekly columns on game theory
(e.g., Arrested Development):http://mindyourdecisions.com/blog/index.php?cat=7 e-book price fixing: http://www.nytimes.com/2013/06/06/business/media/publishers-tell-of-disputes-with-apple-on-e-book-prices.html?_r=0

Resources on penalty kicks in soccer: http://gametheory101.com/Penalty_Kicks.html
http://gametheory101.com/Comparative_Statics.html
http://gametheory101.com/Comparative_Statics.html
http://www2.owen.vanderbilt.edu/mike.shor/courses/gametheory/docs/lecture05/MixedSoccer.html
http://www2.owen.vanderbilt.edu/mike.shor/courses/game-theory/docs/lecture05/MixedSoccer.html

Paper based on real soccer penalty kicks data:
http://www.palacios-huerta.com/docs/professionals.pdf

Game Theory Module: Competition or Collusion? Handouts

Handout #1: Unbalanced Matching Game: Version #1

Unbalanced Matching Game: Version 1 (with **no** computational assistance):

(Alice, Bob)	Bob:	One finger	Two fingers
Alice:			
One finger		+20	-4
Two fingers		-5	+1

Directions: For each of twelve rounds, the Row and Column Players each secretly select a strategy (One or Two fingers) and then reveal their choice simultaneously. Place a check in the box corresponding to each player's choice, and a check to indicate the outcome or payoff of each round. When the twelve rounds are completed, find the total number of times that each choice was made and that each payoff occurred. Then find the ratio and percentage for each.

	Alice/Row		Bob/Column		Outcome/Payoff			
Round	One	Two	One	Two	+20	-4	-5	+1
1								
2								
3								
4								
5								
6								
7								
8								
9								
10								
11								
12								
Totals								
Ratio	$/_{12}$	$/_{12}$	$/_{12}$	$/_{12}$	$/_{12}$	$/_{12}$	$/_{12}$	$/_{12}$
Pct.								

Finally, rounding to one decimal place, compute your group's average payoff per round = (sum of all payoffs)/(# of rounds) =

Handout #1: Unbalanced Matching Game: Version #1, page 2

Unbalanced Matching Game: Version 1 (**with** computational assistance):

(Alice, Bob)	Bob:	One finger	Two fingers
Alice:			
One finger		+20	-4
Two fingers		-5	+1

Directions: For each of twelve rounds, the Row and Column Players each secretly select a strategy (One or Two fingers) and then reveal their choice simultaneously. Place a check in the box corresponding to each player's choice, and a check to indicate the outcome or payoff of each round. When the twelve rounds are completed, find the total number of times that each choice was made and that each payoff occurred. Then find the ratio and percentage for each.

	Alice/Row		Bob/Column		Outcome/Payoff			
Round	One	Two	One	Two	+20	-4	-5	+1
1								
2								
3								
4								
5								
6								
7								
8								
9								
10								
11								
12								
Totals								
Ratio	/12	/12	/12	/12	/12	/12	/12	/12
Pct.								

Finally, rounding to one decimal place, compute your group's

average payoff per round = (sum of all payoffs)/(# of rounds) =

Handout #1: Unbalanced Matching Game: Version #1, page 3

Directions: Based on the 12 rounds with computational assistance (recorded on page 2), for each pair/group in the class, enter the frequency (total number of occurrences) of each Row Player's choice, each Column Player's choice, and the frequency of each outcome for that group. Then add the numbers in each column to obtain class total frequencies for each of these.

Group	Freq. of RowChoice		Freq. of Column Choice		Freq. of Outcome Payoff			
Num.	One	Two	One	Two	+20	-4	-5	+1
1								
2								
3								
4								
5								
6								
7								
8								
9								
10								
11								
12								
13								
14								
15								
16								
Class Total								
Ratio								
Pct.								

Finally, compute the Class Average Payoff Per Round using the two methods below, rounding to one decimal place. Why does either method give the same result?

(Sum of payoffs from all groups)/(product of # of groups times # of rounds) =

(Sum of average payoff per rounds from all groups)/(# of groups) =

Handout #1: Unbalanced Matching Game: Version #1, page 4

Frequency of Average Points Per Round (including all groups in the class)

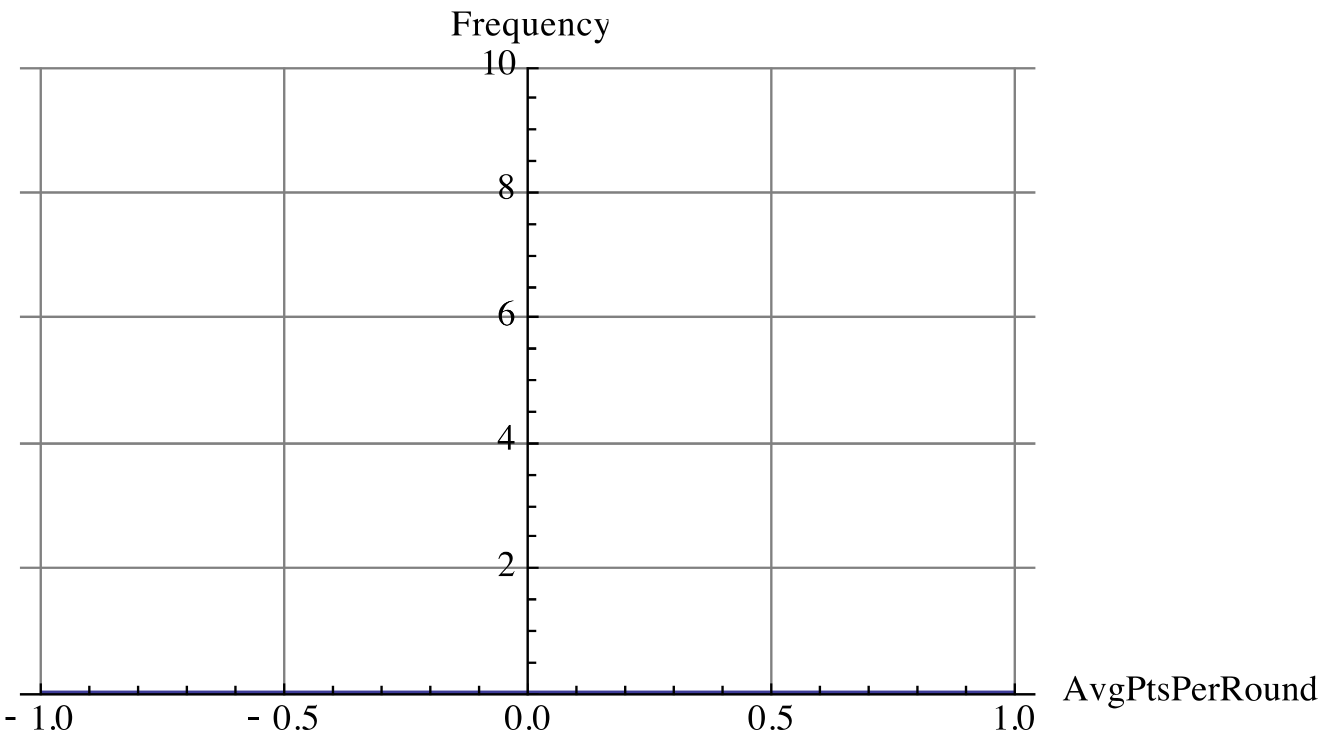

Handout #2: Modified Rock-Paper-Scissors

Modified Rock-Paper-Scissors (Cal **cannot** play Scissors):

(Robin, Cal)	Cal:	Rock	Paper
Robin:			
Rock		0	-1
Paper		+1	0
Scissors		-1	+1

Directions: For each of twelve rounds, the Row and Column Players each secretly select a strategy (Rock, Paper, or Scissors) and then reveal their choice simultaneously. Place a check in the box corresponding to each player's choice, and a check to indicate the outcome or payoff of each round. When the twelve rounds are completed, find the total number of times that each choice was made and that each payoff occurred. Then find the ratio and percentage for each.

	Row Player's choice			Column Player's choice			Outcome/Payoff		
Round	Rock	Paper	Scssrs	Rock	Paper	Scssrs	+1	0	-1
1						NA			
2						NA			
3						NA			
4						NA			
5						NA			
6						NA			
7						NA			
8						NA			
9						NA			
10						NA			
11						NA			
12						NA			
Totals						0			
Ratio	$/_{12}$	$/_{12}$	$/_{12}$	$/_{12}$	$/_{12}$	$^{0}/_{12}$	$/_{12}$	$/_{12}$	$/_{12}$
Pct.						0.00			

Handout #2: Modified Rock-Paper-Scissors, page 2

Directions: For each pair/group in the class, enter the frequency (total number of occurrences) of each Row Player's choice, each Column Player's choice, and the frequency of each outcome for that group. Then add the numbers in each column to obtain class total frequencies for each of these.

Group	Freq. of Row choice			Freq. of Column choice			Freq. of Outcome		
Num.	Rock	Paper	Scssrs	Rock	Paper	Scssrs	+1	0	-1
1						NA			
2						NA			
3						NA			
4						NA			
5						NA			
6						NA			
7						NA			
8						NA			
9						NA			
10						NA			
11						NA			
12						NA			
13						NA			
14						NA			
15						NA			
Class Total						0			
Ratio	/	/	/	/	/	0/	/	/	/
Pct.						0.00			

Finally, compute the Class Average Payoff Per Round using the two methods below, rounding to one decimal place. Why does either method give the same result?

(Sum of payoffs from all groups)/(product of # of groups times # of rounds) =

(Sum of average payoff per rounds from all groups)/(# of groups) =

Handout #2: Modified Rock-Paper-Scissors, page 3

Frequency of Average Points Per Round (including all groups in the class)

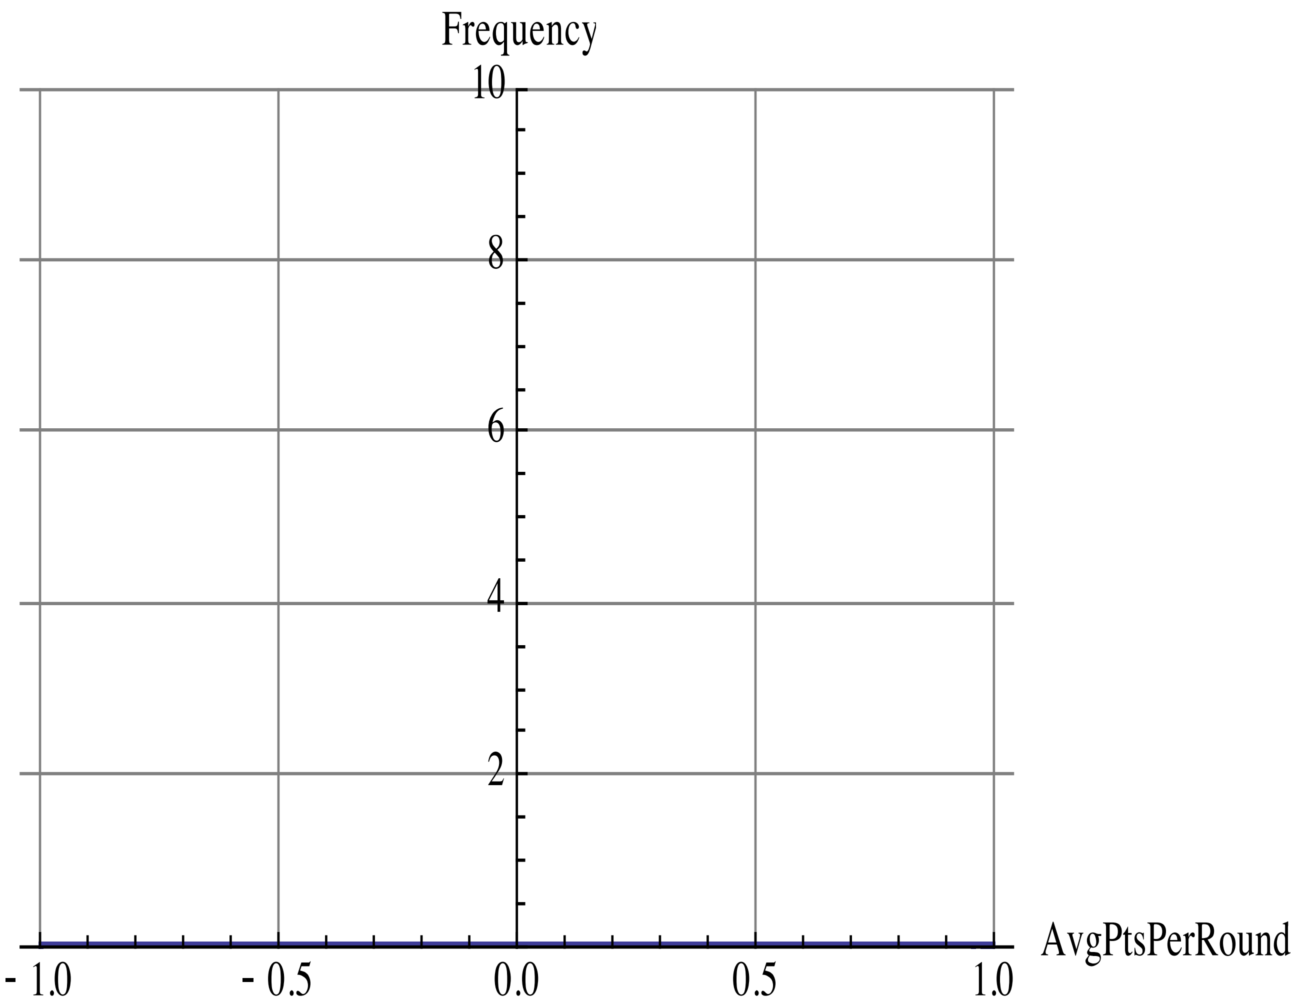

Handout #3: Unbalanced Matching Game: Version #2

Matching Game (unbalanced version):

(Robin, Cal)	Cal:	Left	Right
Robin:			
Left		+8	-6
Right		-4	+2

Directions: For each of twelve rounds, the Row and Column Players each secretly select a strategy (Heads or Tails) and then reveal their choice simultaneously. Place a check in the box corresponding to each player's choice, and a check to indicate the outcome or payoff of each round. When the twelve rounds are completed, find the total number of times that each choice was made and that each payoff occurred. Then find the ratio and percentage for each.

	Row Player		Column Player		Outcome/Payoff			
Round	Left	Right	Left	Right	+8	-6	-4	+2
1								
2								
3								
4								
5								
6								
7								
8								
9								
10								
11								
12								
Totals								
Ratio	$/_{12}$	$/_{12}$	$/_{12}$	$/_{12}$	$/_{12}$	$/_{12}$	$/_{12}$	$/_{12}$
Pct.								

Finally, rounding to one decimal place, compute your group's average payoff per round = (sum of all payoffs)/(# of rounds)

Handout #3: Unbalanced Matching: Version #2, page 2

Directions: For each pair/group in the class, enter the frequency (total number of occurrences) of each Row Player's choice, each Column Player's choice, and the frequency of each outcome for that group. Then add the numbers in each column to obtain class total frequencies for each of these.

Group	Frequency of Row choice		Frequency of Column choice		Outcome/Payoff			
Num.	Left	Right	Left	Right	+8	-6	-4	+2
1								
2								
3								
4								
5								
6								
7								
8								
9								
10								
11								
12								
13								
14								
15								
Class Total								
Ratio								
Pct.								

Finally, compute the Class Average Payoff Per Round using the two methods below, rounding to one decimal place. Why does either method give the same result?

(Sum of payoffs from all groups)/(product of # of groups times # of rounds) =
(Sum of average payoff per rounds from all groups)/(# of groups) =

Handout #3: Unbalanced Matching: Version #2, page 3

Frequency of Average Points Per Round (including all groups in the class)