Game Theory and its Applications

Kristal Jameson
St. Catherine University

St. Catherine University

Department of Mathematics and Physics

Game Theory and its Applications

Author:
Kristal Jameson

Advisor:
Dr. Jessie Lenarz

December 17, 2013

Contents

Abstract

Game theory is the mathematical study of strategic decision making in situations of conflict. In game theory, a single interaction is defined as a *game*, and those involved in the decision-making are called the *players*, who are assumed to act rationally. We will explore the basic ideas of game theory, including the ideas of payoffs and outcomes; classification of games as *zero-sum* or *nonzero-sum*; types of strategies and counterstrategies and the idea of dominance.

Game theory has many applications in subjects such as economics, international relations and politics, and psychology as it can be used to analyze and predict the behavior and decisions of the players. We will apply game theory to the decisions made by the players during the Triwizard Tournament in *Harry Potter and the Goblet of Fire*.

1 Introduction

Game theory is the mathematical study of strategic decision making. It can be used to analyze the options, motivations, and rewards involved in a decision. This paper will discuss some of the basic concepts and modeling tools of game theory. We will also give examples of famous and common applications of game theory. Game theory is widely applied in the real world. Major areas of application include economics, diplomacy, and military strategy. Game theory can also be applied in fields such as psychology, biology, political science, computer science, sociology, and more.

2 The Basics of Game Theory

2.1 Definitions and Zero-Sum Games

Game theory is a branch of mathematics that analyzes interactions involving strategic decision making. Any such interaction is called a *game*. The parties involved in the game are the *players*, who can be individuals or groups. The players are assumed to act rationally. A *rational* player is one who makes decisions based on what will give themselves the greatest benefit. In game theoretical terms, a rational player maximizes their own payoff, a concept we will discuss in more depth later. One type of game is classified as *zero-sum*. A zero-sum game is one in which one player's gain causes the other player's loss of an equal magnitude.

Games in which two players have m and n options, respectively, can be represented by an $m \times n$ matrix. For example, a simple game consisting of two players where each has two options can be represented by a 2×2 matrix, as shown in Figure 1. We can let player one's options be represented by the rows, and player two's options be represented by the columns. The entries of the matrix contain possible *outcomes* of the game. The outcome of a single game is the entry in which the row of player one's decision and the column of player two's decision meet. For example, if each player chooses option 1, the outcome of the game is the entry in position $(1, 1)$ of the matrix representing the game.

<table>
<tr><td></td><td></td><td colspan="2">Player Two</td></tr>
<tr><td></td><td></td><td>Option 1</td><td>Option 2</td></tr>
<tr><td rowspan="2">Player One</td><td>Option 1</td><td>Outcome if each chooses option 1</td><td>P1 chooses 1, P2 chooses 2</td></tr>
<tr><td>Option 2</td><td>P1 chooses 1, P2 chooses 2</td><td>Outcome if each chooses option 2</td></tr>
</table>

Figure 1: The general form of a payoff matrix for a two person zero-sum game.

One simple example of a game that can be modeled with a payoff matrix is *Penny Matching* [1]. Two players each have a penny. They display their pennies simultaneously, with either heads or tails facing up. If both players' pennies match, i.e., both are heads-up or both are tails-up, player one gets both pennies. If they display opposite faces, player two gets both pennies.

		Player Two	
		Heads	Tails
Player One	Heads	1	-1
	Tails	-1	1

Figure 2: The payoff matrix for *Penny Matching* from the perspective of player one.

As shown in Figure 2, we can construct the matrix for this zero-sum game from the perspective of player one. That is, an entry of 1 is a gain of one penny for player one and -1 is a loss of one penny for player one. This gain or loss of one penny represents player one's *payoff*. While it may be intuitive to think of players' payoffs as monetary, this is not necessarily the case. Payoff is formally defined as the player's utility gained in an outcome. It could be money, but it could also be cups of coffee or pairs of socks or anything else the player might gain from a given outcome. Utility that is difficult to quantify can also be numerically represented. For example, if a player valued one outcome twice as much as another, the outcomes could be expressed as 2 and 1, respectively.

The matrix could also be constructed from the perspective of player two; in this case, the positive and negative 1s would simply be switched to display player two's payoffs in each possible scenario. Throughout this paper, matrices for zero-sum games will be constructed in terms of payoff for player one, that is, the player whose choices are represented by the rows, unless otherwise stated.

2.2 Nonzero-sum Games

A *nonzero-sum* game is where one player's gain does not strictly imply another player's loss. This means that outcomes cannot be expressed by only one number; rather, the outcomes of two person nonzero-sum games can be expressed as an ordered pair. The first number of the ordered pair represents the payoff to player one for a given outcome, and the second number represents the payoff to player two. These games can still be represented with a matrix, as shown in Figure 3. In the interest of brevity, we can denote player one's outcomes by a, and player two's outcomes by b.

		Player Two	
		Option 1	Option 2
Player One	Option 1	(a_1, b_1)	(a_1, b_2)
	Option 2	(a_2, b_1)	(a_2, b_2)

Figure 3: The general form of a payoff matrix for a two person, nonzero-sum game.

2.3 Strategies and Expected Payoff

A player's *strategy* is defined as the frequency with which he or she chooses each of the options available to him or her. For a two person, zero sum game, a player's strategy is denoted $[1-p, p]$. For the example of *Penny Matching*, if a player chooses heads or tails with equal probability, this strategy is expressed $[.5, .5]$.

The *expected payoff* for a player can be calculated if we know each player's strategy. The expected payoff is similar to the concept of expected value in probability. To find expected payoff, we multiply each payoff by the probability of that payoff occurring, and sum all these together. For *Penny Matching*, we can find player one's expected payoff as follows:

$$.5 \times .5 \times 1 + .5 \times .5 \times (-1) + .5 \times .5 \times 1 + .5 \times .5 \times (-1) = 0 \tag{1}$$

This makes sense because if each player chooses heads or tails with equal frequency, we would expect player one to win exactly half the games (and therefore lose half the games as well).

2.3.1 Optimizing Payoff

If one player's strategy is known by the other, the other player can use this information to decide what strategy to use to produce the greatest payoff. For a simple example, consider the game *Rock, Paper, Scissors*. In this game, two players simultaneously show rock, paper, or scissors. Rock beats scissors, scissors beats paper, and paper beats rock. If you were playing *Rock, Paper, Scissors* and you knew your competitor favored paper over the other two options, you would choose scissors more often. In fact, your optimal strategy would be to choose scissors 100% of the time. Then, your strategy for *Rock, Paper, Scissors* would be expressed as $[0, 0, 1]$. A strategy where a player chooses one option 100% of the time is called a *pure* strategy. Conversely, a *mixed* strategy is one where options have differing probabilities of being chosen. According to Theorem 1 in *A Gentle Introduction to Game Theory*, if one player of a game employs a fixed strategy, then the opponent has an optimal counterstrategy that is pure [1].

2.4 Maximin and Minimax Strategies

The *maximin* and *minimax* strategies involve essentially choosing the option with the "best worst case scenario." These ideas are among the first ever considered in game theory, as British ambassador James Waldegrave found a minimax solution for two-person games in 1713 [2].

2.4.1 Maximin: determining a strategy for player one

As discussed above, there is a pure strategy that is best for player two if player one's strategy is fixed. However, player one can incorporate this into their choice of strategy, anticipating what player two will choose. For a given strategy of player one, player two will choose the option that yields player one the least possible benefit. With this in mind, player one can choose the strategy that has the greatest minimum payoff. Player one would consider the payoffs for every possible pure strategy of player two. Figure 4 shows the two matrices representing the possible outcomes from player two's two pure strategies.

		Player Two	
		1	0
Player	$1-p$	a	b
One	p	c	d

		Player Two	
		0	1
Player	$1-p$	a	b
One	p	c	d

Figure 4: Payoff matrices for player two's two possible pure strategies.

We can compute player one's expected payoff for each scenario. Player one's strategy can be denoted $[1-p, p]$. Equation 2 gives the expected payoff to player one when player two's strategy is $[1, 0]$, and Equation 3 gives the expected payoff when player two's strategy is $[0, 1]$.

$$(1-p) \times 1 \times a + p \times 1 \times c = (c-a)p + a \tag{2}$$

$$(1-p) \times 1 \times b + p \times 1 \times d = (d-b)p + b \tag{3}$$

Because a, b, c, and d are constants, the best strategy for player one can be determined from Equations 2 and 3. The actual outcome will be the smaller of these two results, because player two will choose the option that minimizes player one's payoff. Effectively, player one is able to secure a guaranteed minimum payoff.

We can find a maximin strategy for player one in *Penny Matching* [1]. Recall the payoff matrix from the perspective of player one, as shown in Figure 2. We can substitute the values of this matrix into Equations 2 and 3. From the payoff matrix, $a = d = 1$ and $b = c = -1$. Equations 4 and 5 give the two possible payoffs to player one when player two's strategy is $[1, 0]$ and $[0, 1]$, respectively. Player one's strategy is denoted $[1 - p, p]$, $0 \leq p \leq 1$. In words, Equation 4 gives the payoff to player one when player two only displays heads, and Equation 5 gives the payoff when player two only displays tails. Player one displays tails with probability p.

$$(1 - p) \times 1 \times 1 + p \times 1 \times -1 = (-1 - 1)p + 1 = -2p + 1 \tag{4}$$

$$(1 - p) \times 1 \times -1 + p \times 1 \times 1 = (1 - (-1))p - 1 = 2p - 1 \tag{5}$$

We can then graph Equations 4 and 5 to help determine the optimal strategy for player one, as shown in Figure 5 [1].

Player two will choose the strategy that gives the lower payoff to player one, so player one's expected payoff will be on the bolded line, with the exact value dependent on their choice of p. From this, we can see that player one's best expected payoff is 0, which is obtained by the strategy $[.5, .5]$. Intuitively, this makes sense. If player one were to display heads more than tails, for example, player two would respond by always displaying heads, which gives the advantage to player two. If player one employs strategy $[.5, .5]$, their expected payoff is 0 regardless of which pure strategy player two chooses.

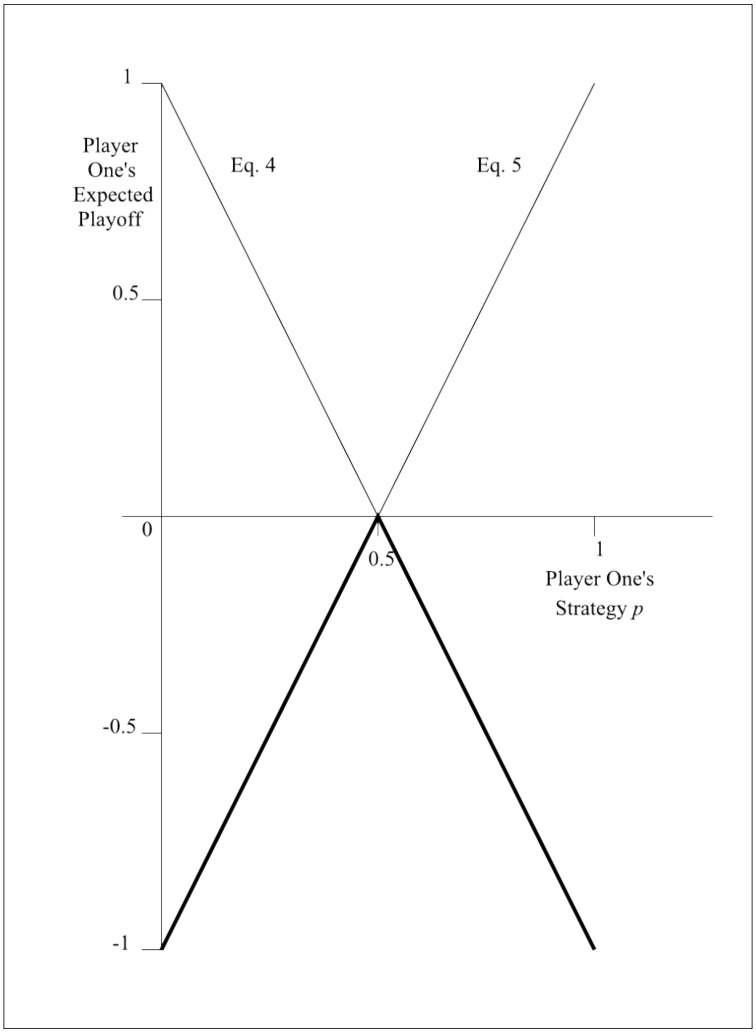

Figure 5: A graph of player one's expected payoff.

2.4.2 Minimax: determining a strategy for player two

The minimax strategy is the best strategy for player two, as opposed to player one. While both the minimax and maximin strategies have the same general idea, the minimax strategy determines player two's optimal strategy $[1-q, q]$ considering the outcomes of player one's two pure strategies. Because the outcomes in the matrix describe the payoff to player one, player two knows player one will choose the strategy that gives the highest of the two possible payoffs. Effectively, player two wants to choose a strategy $[1-q, q]$ that gives the "smallest maximum" because player one will choose the strategy with maximum payoff. The two possible matrices are shown in Figure 6.

		Player Two	
		$1-q$	q
Player	1	a	b
One	0	c	d

		Player Two	
		$1-q$	q
Player	0	a	b
One	1	c	d

Figure 6: Payoff matrices for player one's two possible pure strategies.

We can also compute player one's two expected payoffs with this matrix, as shown in Equations 6 and 7.

$$1 \times (1-q) \times a + 1 \times q \times b = (b-a)q + a \tag{6}$$

$$1 \times (1-q) \times c + 1 \times q \times d = (d-c)q + c \tag{7}$$

Payoff is dependent on player two's strategy, denoted $[1-q, q]$. For a given strategy of player two, the actual payoff to player one will be the higher result of those given by Equations 6 and 7. This is because, of the two choices available, player one will choose the option with the higher payoff. Player two is able to choose a strategy that places a guaranteed upper limit on player one's payoff.

2.5 Dominance

While we have so far focused on games represented by 2×2 matrices, matrices for two person games can have any dimension depending on how many options each player has. In such cases, it may seem more complicated to determine the best move for either player because there are more options, but in some of these cases we can eliminate one or more options. We say option i *dominates* option j if, for all possible outcomes of choosing j, the possible outcomes of choosing i are higher. In words, for every potential outcome in a row or column, the potential outcomes in another row or column are higher.

Consider the game represented by the matrix in Figure 7, where each of the two players have three options. Of player one's three options, option 3 has a better outcome than option 2 regardless of the choice player two makes. Therefore, we can remove the second row from the matrix because player one would never choose that option when acting rationally. This step is shown in Figure 8.

Player Two

Player One	-3	3	4
	2	-4	1
	3	-2	2

Figure 7: A two person game where each player has three options.

Player Two

Player One	-3	3	4
	2	-4	1
	3	-2	2

Player Two

Player One	-3	3	4
	3	-2	2

Figure 8: The matrix after removing player one's dominated option.

Now we have a 2×3 matrix representing the game. We can also eliminate one of player two's options. Here it is important to recall that the entries of this matrix are the payoffs to *player one* for every given outcome. This means that when player two is deciding the best option, they would choose the option that is consistently lower because less gain for player one means more gain for player two. Then option 2 dominates option 3, and we can remove the rightmost column from the matrix. Figure 9 shows this step and the final result, which is much easier to work with than the original matrix.

Player Two

Player One	-3	3	4
	3	-2	2

Player Two

Player One	-3	3
	3	-2

Figure 9: The matrix after removing player two's dominated option.

2.6 Nash Equilibria

In nonzero-sum games, a *Nash Equilibrium* occurs when each player individually cannot increase their expected payoff by switching to a different strategy. These strategies can be either pure or mixed. The concept of Nash Equilibrium is so named because John Nash proved that every two-person game contains a minimum of one equilibrium, whether it is a combination of mixed or pure strategies. Some games, especially those where each player has more choices, have multiple Nash Equilibria. Figure 10 shows a matrix of a nonzero-sum game with one Nash Equilibrium.

		Player Two	
		Option 1	Option 2
Player	Option 1	(-1, -1)	(-3,0)
One	Option 2	(0, -3)	(-2,-2)

Figure 10: A game matrix with a Nash Equilibrium in position $(2, 2)$.

Nash Equilibria can be determined by finding the entries (a, b) of the payoff matrix such that a is the maximum value for its column and b is the maximum value for its row. This method works because each player is guaranteed not to regret their decision, given the decision of the other player. In the game in Figure 10, the outcome $(-2, -2)$ is a Nash Equilibrium because neither player could get a better outcome by individually changing their choice; if either player one or player two switched, their payoffs would decrease to -3. This method works if a Nash Equilibrium is pure, meaning that each player employs a pure strategy to obtain the outcome of the Nash Equilibrium. For games without a pure strategy Nash Equilibrium, determining the mixed strategy solution is less straightforward.

Interestingly, it is clear that both players would have a better outcome if each chose option 1. We say an outcome is *Pareto optimal* if there is no outcome with better payoffs for all players. Sometimes the equilibrium outcome is not Pareto optimal, as in the case of the game in Figure 10. When this occurs, we call the outcome *Pareto inefficient* or *Pareto inferior*. Figure 11 shows a matrix of a game with an equilibrium that is Pareto optimal [3]. In this matrix, there is another Nash Equilibrium with outcomes $(2, 1)$ in position $(2, 2)$, but it is not Pareto optimal because player two has a better outcome in the other equilibrium point.

		Player Two		
		Option 1	Option 2	Option 3
	Option 1	(0, -1)	(0, 2)	(2, 3)
Player	Option 2	(0, 0)	(2, 1)	(1, -1)
One	Option 3	(2, 2)	(1, 4)	(1, -1)

Figure 11: A game matrix with a Pareto optimal Nash Equilibrium in position $(1, 3)$.

2.7 Cooperation

Up to this point, we have assumed that all players make decisions simultaneously and independently. We have also assumed their sole motivation is to maximize their own payoff (or minimize the opponent's payoff), and that play occurs exactly once. However, when games are played repeatedly, the phenomenon of *cooperation* emerges.

Cooperation can take many forms and arise from different motivations. One such form is essentially the Golden Rule: do unto others as you would have them do unto you. If I cooperate today, you might be more likely to cooperate with me in the future. Another form of cooperation is less optimistic, and based on retaliation. This means that a party would make their decision of whether to cooperate based on whether their opponent had cooperated in the past.

In order for cooperation to be a viable solution, there must be trust between the players. It is often the case that if only one player were to cooperate, they would suffer a great loss whereas the *defecting* player, that is, the one who betrays the other, would benefit. Thus, if there is not enough trust between the two parties, mutual cooperation is all but impossible. This idea can be extended to many real-world applications, such as in international relations. In these situations, if the two parties could only cooperate both would have a better outcome. This is also seen in the game in the previous section, in Figure 10, as well as the famous *Prisoner's Dilemma*, which will be discussed in Section 4.

3 History

Game theoretical concepts have been utilized to analyze problems for millennia, long before game theory was a formally-defined field. One interesting example is that the Talmud, the Jewish holy book that provides the basis for Jewish law, prescribes solutions for allocation of disputed resources that confounded scholars until the 1980s when mathematicians Robert Aumann and Michael Maschler solved the problem using the tools of modern game theory. As it turns out, the solution given by the Talmud is to split the disputed amount equally [4]. Another example is when James Madison considered the effects of different taxation systems with game theoretical concepts [2]. The list goes on, as conflict resolution and strategic decision-making have been important issues throughout all of human history.

The first work that brought about game theory as a formal field of mathematics was Hungarian mathematician John von Neumann's paper *The Theory of Games* in 1928 [5]. This paper had three major results. The first was reducing a game to the cases where each player knows either everything or nothing about the other player's previous moves. He also proved the minimax theorem for two person zero-sum games, and he analyzed three person zero-sum games [5].

Economist Oskar Morgenstern connected with von Neumann in 1938, and the two then worked together on *Theory of Games and Economic Behavior*, published in 1944. This work was huge in the development of game theory. They expanded on von Neumann's previous work with an in-depth analysis of situations where players have only partial knowledge of other players' previous decisions, whereas *The Theory of Games* made the assumption that players knew either everything or nothing about previous decisions. They also expanded the definition of payoffs; previously payoffs were generally considered to be only monetary, but von Neumann and Morgenstern developed the theory of utility, which is still used today in many fields such as economics [5].

Since von Neumann and Morgenstern laid the foundation for game theory, it has been added to by many mathematicians, such as John Nash in the 1950s. However, the main development over the following decades was increasingly widespread application to many fields. While certainly important in the field of economics, the use of game theory has expanded to extensive use in biology, and it is also very important to the development of military strategy. Interestingly, the five game theorists who have won the Nobel Prize for economics also worked as advisors to the Pentagon over the courses of their careers [6]. Game theory has also been applied in fields such as computer science and moral philosophy [2].

4 More Examples of Games

As mentioned previously, game theory is applicable in many fields. This section will discuss some specific examples of game theory at work. For an easy and fun read about applications of game theory, see Fisher's *Rock, Paper, Scissors: Game Theory in Everyday Life* [6].

4.1 The Prisoner's Dilemma

One of the most famous problems of game theory is the *Prisoner's Dilemma*, a nonzero-sum game. The setup is as follows: two people are arrested for a crime, but there is not enough evidence to press charges. The police separate them for interrogation, and they have not worked out a story to tell beforehand. The decision before each of them is to either tell the police everything or to keep quiet. In game theoretical terms, they can either *defect* or *cooperate*. Cooperating means keeping quiet for this example, termed so because they are cooperating with each other. Defecting means to talk to the police. If both cooperate, then both will serve short sentences. If only one defects, the defector will be released and the other person will have a longer sentence. If both defect, then each is penalized, but to a lesser extent than if only the other person had defected. A payoff matrix can be constructed for the situation, as shown in Figure 12.

	Cooperate	Defect
Cooperate	(-5, -5)	(-15, 0)
Defect	(0, -15)	(-10, -10)

Figure 12: A payoff matrix for the *Prisoner's Dilemma.*

Even though it is clear that the best outcome is achieved if both cooperate, if the game is played exactly once the rational outcome is for both to defect. This is because there is a Nash Equilibrium at $(-10, -10)$ in the payoff matrix. In each player's case, regardless of whether the other player cooperates or defects, their better option is to defect. This is an example where the Nash equilibrium is not Pareto optimal.

4.2 Chicken

Another game that can be analyzed is *Chicken*, where two players are moving toward each other until they either crash or one moves out of the way. The payoffs are difficult to quantify if one or both players move because all that is gained by not crashing is glory or embarrassment for the winner or loser, respectively. However, if neither player moves, the payoff is certainly negative for both. This makes *Chicken* a nonzero-sum game because one player's loss is not necessarily the other's gain.

One example of real-life *Chicken* was the Cuban Missile Crisis. In 1962, the USSR attempted to build a missile base in Cuba, which is only 90 miles away from the US. The US set up a naval blockade to prevent supply ships coming in, so there was essentially a standoff for a few weeks before the USSR backed down. In game theoretical terms, the US and USSR (the players) each had the two possible strategies of proceed or back down. This example is very illustrative of how the payoffs when one or both players back down are not necessarily quantifiable. These outcomes are nonetheless important, however, especially in the context of the Cold War. On the other hand, if neither party had backed down, the result could very well have been the strongly negative outcome of nuclear annihilation.

4.3 Tragedy of the Commons: the St. Catherine University Cafeteria

You may have noticed that every May our Sodexo overlords beg in vain for the St. Catherine University student body to return the silverware and plates that were taken throughout the year. "How can this *really* matter?" you might wonder. Surely a giant company such as this isn't taking a hit just because you've pocketed a few forks over the semester. While Sodexo is in all likelihood doing fine, this is actually an example of a phenomenon called *Tragedy of the Commons* [6]. Individual students rationally decide that if they take a fork home with them, the benefit to them is far greater than the loss to the university because the cafeteria has so many forks to begin with. However, when we all do this, there may actually be a shortage of utensils by the end of the year.

This is certainly an example of *Tragedy of the Commons* that is relatively unimportant in the grand scheme of things. Perhaps the disgruntled student body derives a sense of vigilante justice from this large-scale utensil thievery, an intangible payoff that may be even greater than Sodexo's loss. However, this phenomenon plays out in many other, more serious real-world examples, especially in relation to resource consumption. Whether it is forestry, driving an inefficient vehicle or unnecessary driving, or choosing not to recycle just this one time, countless seemingly isolated decisions add up to collectively create a large impact.

4.4 Rock, Paper, Scissors

Rock, Paper, Scissors, described in Section 2.3.1, is a very well-known zero-sum game. As you may have learned from playing in your childhood, the best strategy is $[\frac{1}{3}, \frac{1}{3}, \frac{1}{3}]$, or choosing each option with equal probability. As mentioned previously, if a player knows the other player is likely to choose one option, say scissors, more often than the others, they would simply switch their strategy to choosing rock 100% of the time. We can even construct a matrix to simulate *Rock, Paper, Scissors*, as shown in Figure 13.

		Player Two		
		Rock	Paper	Scissors
Player One	Rock	0	-1	1
	Paper	1	0	-1
	Scissors	-1	1	0

Figure 13: A payoff matrix for *Rock, Paper, Scissors.*

While it is intuitively and mathematically true that one should display each option with equal probability, in practice true randomness is nearly impossible for humans to achieve. Interestingly, in tournament play, it has been observed that scissors is only displayed 29.6% of the time [6]. To clarify, this does not mean that if you enter a *Rock, Paper, Scissors* tournament you should only play rock because in that case your opponent would quickly catch on and cover you with paper. Rather, the pure strategy response is only a viable option if your opponent's strategy is fixed. In real life, repeated gameplay strategies change frequently in response to what the other player is perceived to be doing.

5 The Triwizard Tournament

The Triwizard Tournament is the main event of *Harry Potter and the Goblet of Fire* [7]. In a tradition that began seven hundred years ago, the three wizarding schools of Hogwarts, Durmstrang, and Beauxbatons come together to compete. One student is to be chosen to represent each school in the tournament, which consists of three difficult tasks. Naturally, the process does not go as planned. The Goblet of Fire, which chooses the competitors, is bewitched to choose Harry in addition to the three it chooses as it should, so Hogwarts is represented by both Harry Potter and Cedric Diggory in the tournament. The Goblet of Fire chooses Viktor Krum to represent Durmstrang, and Fleur Delacour is the Beauxbatons champion.

Throughout the course of the tournament, countless decisions are made. As game theory is the analysis of strategic decision-making, there are certainly numerous aspects of the tournament that could be considered through a game theoretical lens. However, this application will focus on cooperation between the players.

5.1 Cooperation in the Triwizard Tournament

We see much cooperation between Harry and the other players during the Triwizard Tournament. One specific example of this is the exchange of information between Harry and Cedric. Thanks to Hagrid, the Hogwarts gamekeeper, Harry learns that the first task will involve facing a dragon. Hagrid brings the Beauxbatons headmistress into the forest to see the dragons, and Harry runs into the Durmstrang headmaster as he is sneaking back to the castle. This leads Harry to assume that both Fleur and Krum will know about the task ahead of time. Then Cedric, it is assumed, is the only competitor who doesn't know what is coming.

In this "game," the players are Harry and Cedric. We will exclude all other characters. Voldemort may be orchestrating everything from Harry's entrance in the tournament to his success in the tasks, but Harry does not yet know this. We can also exclude the other two champions for simplicity.

Harry could have kept the information of the first task to himself, which would have given him a huge advantage over Cedric. However, Harry's payoff in this situation is not the thousand-galleon prize for winning the tournament he didn't even want to enter. Rather, his decision to inform Cedric is motivated by his sense of fairness and tendency to do what he thinks is right regardless of the rules. His decision is a form of cooperation. While Harry likely was not thinking about the situation explicitly in terms of "If I help Cedric now, he'll feel obligated to help me later," Harry tends to make decisions based on what he perceives to be fair – whether this entails informing the fourth tournament participant or getting back at Malfoy for whatever he did. This is another example of a game with outcomes that are not easily quantified, as in the game *Chicken*. Although Harry's outcomes cannot be quantified easily, his desire to do the right thing is still a very real motivation.

The fact that Cedric is the other Hogwarts champion could also have played a role in this decision. Because he did not even want to enter, Harry does not care much about winning. It only makes sense that he would be rooting for the other Hogwarts champion to win if he doesn't. Additionally, the fact that they are both from Hogwarts and have generally been on good terms in the past could have increased the trust between them, which is essential for cooperation to take place.

Cedric is the first to learn how to decipher the clue of the second task. He explicitly states that he gives Harry a hint because he feels obligated to return the favor from the previous task, so it is clear that in this case Harry's previous cooperation paid off greatly for him.

This is only one of the many examples of the applications of game theory to the Triwizard Tournament. For further reading, Highfield's *The Science of Harry Potter* discusses a few other Harry Potter plot lines in terms of game theory, including Harry's decisions made during the second Triwizard task [8].

6 Conclusion

Game theory is all around us, from our favorite stories, to our everyday decisions, to the critical negotiations made by global leaders. Nearly anywhere a decision is being made, the situation can be examined with game theoretical concepts. This paper has only explored a basic overview of game theory; when it is applied in the real world the problems become far more complicated and challenging to define. For further reading on the basic definitions and concepts of game theory, see Stahl's *A Gentle Introduction to Game Threory*, and for more on everyday applications see Fisher's *Rock, Paper, Scissors:*

Game Theory in Everyday Life [1, 6].

References

[1] S. Stahl, *A gentle introduction to game theory.* Providence, R.I.: American Mathematical Society, 1999.

[2] "Game theory." `http://www.policonomics.com/game-theory/`, 2012.

[3] P. D. Straffin, *Game theory and strategy.* Washington: Mathematical Association of America, 1993.

[4] R. J. Aumann and M. Maschler, "Game theoretic analysis of a bankruptcy problem from the talmud," *YJETH</cja:jid> Journal of Economic Theory*, vol. 36, no. 2, pp. 195–213, 1985.

[5] J. Eatwell, M. Milgate, and P. Newman, *Game theory.* New York: W.W. Norton, 1989.

[6] L. Fisher, *Rock, paper, scissors : game theory in everyday life.* New York: Basic Books, 2008.

[7] J. K. Rowling, *Harry Potter and the goblet of fire.* Scholastic Inc., 2000.

[8] R. Highfield, *The science of Harry Potter : how magic really works.* New York: Viking, 2002.

[9] G. Owen, *Game theory.* Philadelphia: Saunders, 1968.

[10] D. Fudenberg and J. Tirole, *Game theory.* Cambridge, Mass.: MIT Press, 1991.

[11] R. McCain, *Game theory and public policy.* Cheltenham, UK; Northhampton, Mass.: Edward Elgar, 2009.

[12] G. P. Miller, ed., *Economics of Ancient Law.* Northampton, MA: Edward Elgar Publishing, Inc., 2010.