Game Theory

1 Introduction

Game theory is the branch of mathematics that tries to model conflict or competition between two (or more) people. With the exception of Statistics, this is the branch of mathematics that is most used in the social sciences like Sociology, Psychology, Political Science, and also Economics. Game Theory had its roots in 1917, but began to pick up steam with John von Neumann in 1928 and John Nash in the 1950s.
The basic setup is that there are some number of players, each player has some number of strategies to choose from, and associated payoffs for each player for every combination of strategies. In general, this can get very complicated, but we will be focusing on the simple case where there are two players, and each player has two strategies. This would mean that there are four possible outcomes, with payoffs associated to each player. This can be shown in a matrix like the following:

		Player B	
		B_1	B_2
Player A	A_1	$2,2$	$2,0$
	A_2	$3,0$	$0,9$

This means that if Player A chooses strategy A_1 and Player B chooses strategy B_2 (denoted (A_1, B_2) for short), then Player A would win 2 and Player B would win 0. The goal of each player is to win the most (or lose the least). The only control each player has is in which strategy they themselves choose.

One basic situation and example matrix would be a situation many of us have been in, catching a bus. More specifically, suppose you want to cross a street to get to a bus stop, but the light is red. You see the bus coming down the road, and you know if you do not cross at this moment, you will miss the bus. Your choices (strategies): Cross the street or Don't cross the street. What can happen is in the next few seconds, either a car goes down the road, or there is nor car. This situation can be viewed by the following matrix:

		Nature	
		Car	No Car
You	Cross	-1000	10
	Don't Cross	-5	-5

The numbers are somewhat arbitrary, but the idea is that if you don't cross, you will be slightly annoyed that you missed the bus and have to wait for the next one; if you cross and there is no car, you would be quite happy you made it onto the bus and don't have to wait; if you cross and there is a car, you get hit, which is obviously very bad. This sounds all rather silly, but many of us have played this very game. We look at how busy the road is to tell how likely is it a car will be driving by. That is, we try to determine which strategy Nature is going to play, and we react accordingly. This highlights one of the main issues with game theory: If we knew what the opponent would do, we know what we would then do. This can quickly degenerate to the classic "If I know that she knows that I know that she knows..." kind of thinking.

2 Zero Sum Games

A zero sum game is one where any winnings for one player counts as a loss for the other player. We will generally simplify the diagram so it will be easier to read, where each number represents winnings for Player A and losses for Player B.

	B_1	B_2
A_1	$5, -5$	$-8, 8$
A_2	$3, -3$	$1, -1$

	B_1	B_2
A_1	5	-8
A_2	3	1

Figure 1: A zero-sum game matrix and its simplification.

Suppose you are Player A is the above situation. How do you choose which strategy to use? We want to have the best thing to happen to us, that is, win the most money. So, A_1 seems tempting. However, we would also have the potential to lose 8. In game theory, we generally take the pessimistic view of things. That is, for each of the strategies, what is the worst thing that can happen? Out of those, which is the least bad? This is the same as finding the minimum of each strategy, and then the maximum of the minimums. This is called the maximin.

Let us look at a specific example:

	B_1	B_2
A_1	6	5
A_2	5	4

Notice that the minimums are 5 for A_1 and 4 for A_2. So, to maximize these minimums, Player A would choose A_1 with worst case scenario of 5 if Player B chooses B_2. Now, these values represent *losses* for Player B, and so, would big numbers are bad. That is, Player B would mind the maximums of each strategy, and the minimums of those maximums (called the minimax). The maximums are 6 for B_1 and 5 for B_2. So, Player B would choose B_2 with worst case scenario of 5 if Player A chooses A_1. Notice how both players ended up at the outcome of (A_1, B_2). This is called the saddle point for this game.

We could have come at this another way, somewhat through trial and error. Again, consider yourself as Player A. Ask yourself, "If Player B chose B_1, what would I do?" "What if B_2 was chosen?" No matter what Player B does, you would want to choose A_1. Similarly, Player B would always want to choose B_2. If Player B were to deviate from this, it could only help you, and vice versa. So, (A_1, B_2) is the only "logical" choice.

Now, a more interesting example:

	B_1	B_2
A_1	3	6
A_2	5	4

Let's follow the same reasoning we just used in the last example. If you are Player A, what would you do if Player B chose B_1? Since $(A_1, B_1) = 3 < 5(A_2, B_1)$, you would want to choose A_2. However, if you were to choose A_2, Player B would want to choose B_2. But then, you would rather choose A_1. This would, in turn, make Player B want to choose B_1. So you would like to.... As you can see, this degenerates quickly to a case of "If she knows that I know that she knows that...".

We can also try to find the minimax and maximin. Notice that if Player A were to play conservatively, strategy A_2 would be chosen, with the worst case scenario of 4. Similarly, Player B's conservative move would be B_1 with worst case scenario of 5. Since these two values do not agree, there should be some sort of strategy that would balance things out.

This problem is not so cut-and-dry as the last one. The main difference is issue of *prior knowledge*. If you could spy on the other person, or if they gave some cue as to which strategy they will choose, it would have made no difference in the prior example since you would choose the same strategy regardless. Here, you would tailor your strategy based on what your opponent is planning. The only solution to this is to effectively get rid of the planning, in some sense. That is, we want to institute some chaos by randomly picking between the strategies. The question is now, how often should you pick A_1 versus A_2? To figure this out, we will follow a strange method, which hides a lot of math in the background. First, subtract one column from another (say, the second minus the first). We will also do similarly for Player B by subtracting the rows.

A_1	3
A_2	-1

B_1	B_2
2	-2

To average each of these to be zero, we would want to choose A_2 3 times as often as A_1. That is, 75% of the time A_2 should be chosen, and 25% of the time, A_1 should be chosen. If we played with this mixed strategy, what do we expect would happen? First, suppose Player B chooses B_1. Then, 75% of the time we would win 5, and 25% of the time we would win 3. That is, on the average, we would win $.75(5) + .25(3) = 4.5$. Similarly, If Player B chooses B_2, then we would expect to win $.75(4) + .25(6) = 4.5$. This shows that regardless how Player B plays, we will win the same amount on the average. So, Player B cannot outsmart us or trick us into getting a smaller amount. However, this also says if we play this way, we can't do any better than 4.5 no matter how "bad" Player B plays. Also, you can check that the mixed strategy for Player B should be to randomly choose between strategies, but pick each one equally often. Also, this would cause the expected value of the game to be, again, 4.5.

3 Non Zero Sum Games

Now, we shall (quickly) look at some games that are not zero sum. For example, consider the following situation:

	Stag	Hare
Stag	10, 10	0, 4
Hare	4, 0	4, 4

Figure 2: "Stag Hunt"

In this situation, consider two hunters trying to decide if they are going after big game or small game ("Stag" or "Hare"). They can get the Hare by themselves, but would need to work together in order to get a Stag, but each would see a large benefit. The minimax/maximin strategy would say they should go after Hares, but it would be in both their best interests to work together to get a Stag.

Another example (Warning: Gender stereotypes ahead!):

		Woman	
		Ballet	Football
Man	Ballet	6, 10	0, 0
	Football	0, 0	10, 6

Suppose a man and a woman decided to go out, the man wanting to go to a football game, and the women to the ballet. Further, suppose they know when they are going to meet, but forgot what they decided to do. If on the day of the event they couldn't contact each other, how should they decide where to go? If they end up in different places, no one would be happy. If they end up in the same place, they both would be happy, but whoever is doing the activity they prefer is happier.

	Silent	Talk
Silent	$-1, -1$	$-10, 0$
Talk	$0, -10$	$-7, -7$

Finally, let's finish with the Prisoner's Dilemma:
The situation here is that two people are arrested by the police for a crime, and each are offered a deal to rat out the other person. If they both stay silent, they are both convicted for a minor crime. If they both rat each other out, the both go to jail for a significant amount of time. However, if one talk and the other does not, the silent one goes to jail for a long time, while the other walks free. What is interesting here is that if both players follow the minimzx principle, they both get -7, when they could have agreed to both get -1. However, there is always the temptation to go for the 0 and the fear to stay away from the -10. This gets even more interesting when playing the game multiple times, as players can react to prior betrayals. A good resource to look at simulations of this is `http://www.iterated-prisoners-dilemma.net/`.