1.0 What is Game Theory?

Game theory can best be described as the careful use of mathematics and logic in order to predict what decision players will make anytime two or more (two for our examples) are engaged in a scenario where each of the players has something to gain or lose depending on the decision the other player makes. The hypothetical scenario we are talking about here is known as a game. However, in the context of game theory, we usually are not concerned with what we typically know as games, e.g. poker. Although one could certainly use the techniques we will develop here in order to better understand the decisions one could witness poker players make in a given game.

So what are some examples of games? A game is any situation where one or more players are set upon a conflict course. With this in mind, the interaction between OPEC pricing their oil and our government deciding on how much oil it should buy is certainly a game. The participants in this game will more than likely make rational decisions since the decisions both parties would make will drastically affect the other. In an effort to summarize here, the reader should know that game theory is the use of logic and mathematical tools to understand and even predict what decision the players (individual parties on the course of conflict) would make to maximize their respective payoffs, provided all players behave logically.

2.0 Elements of a Game, Some Practical Examples, and Equilibria

In order to analyze a game, three different types of information must be known. First, the analyst must clearly know the participants in the game. Secondly, one must know the strategies that both parties can choose from. Finally, one must know enough about the game to ascertain the payoffs each player will have for each available combination of strategies. These abstract notions will become clearer in the examples that follow.

2.1 Example 1: The Classical Prisoner's Dilemma and Dominant Strategy

		Prisoner 2	
		Deny	Confess
Prisoner 1	Deny	-1,-1	-10,0
	Confess	0,-10	-8,-8

(Prisoner 1, Prisoner 2)

The above colored 2x2 grid along with the accompanying labels lists all the information we said we needed in section 2.0—we know the players, that is prisoner 1 and prisoner 2, we know the payoffs of each player (the ordered pairs in the colored boxes), and we know the available strategies in this game, to deny or to confess. Before one can really make much sense of the colored grid shown above, one must be familiar with a description of the classical prisoner's dilemma, which will now be discussed.

The police have apprehended a pair of suspects of whom they think are responsible for a recent bank robbery. The police do not have enough evidence mounted against the two prisoners in order to have them incarcerated for a long amount of time, however the police decide to interrogate them separately to see if they can learn more about the robbery and the roles each of these suspects may have played.

In order to make the following rationale less confusing and ambiguous, we will pick a reference point of one of the players in question—in this case we will tell the story from prisoner 1's perspective. The police detective tells prisoner 1 that another friend of his is currently chatting things up with his buddy, prisoner 2. The detective goes on to tell prisoner 1 that he can gain full immunity if he will just come clean and tell them some key things about the robbery that only the perpetrators could know. In order to do this, he would have to rat out his friend, that is prisoner 2, but prisoner 1 will walk away without any jail time. Alternatively, he could deny the charges altogether, and suffer one year of jail time. However, the detective goes on to tell prisoner 1 that they don't normally offer such a generous deal and that he really thinks his buddy, prisoner 2 is going to rat out on him. If this occurs, the detective makes it clear to prisoner 1 that he will face 10 years of jail while his buddy makes out without any jail time. Alternatively, prisoner 1 could confess and rat himself and prisoner 2 out. This would enable prisoner 1 to not face any jail time while his buddy, prisoner 2, does ten years. Finally, there is the case that both of the prisoners could confess. If this occurs, the police have no reason to give either of them a bargain since they both independently admitted guilt, and they both earn 8 years in jail. Now that the reader has a basic idea of the premise behind the problem, we can go on to decode the mysterious grid we are presented with above.

The first thing the reader should notice in this example is that most of the payoffs come with a negative sign attached. Not surprisingly, a negative sign denotes an undesirable payoff, in this case time in jail. Also, one should notice that the best payoff the players can hope for in this game is for nothing to happen to them. This is kind of unusual as typically participants in a game will have a strategy that will result in them gaining something ; however, in this scenario walking away with no time in jail is the best thing the players, the prisoners, can possibly hope for.

Since similar nxn grids will be used in our subsequent examples, we will talk about every colored box in detail. The first box one comes to is the yellow box. As one can hopefully tell, the yellow box depicts the payoffs the prisoners would obtain if each of them denied taking part in the crime. The first number in the ordered pair (as illustrated below the grid) pertains to Prisoner 1's payoff, whereas the second number pertains to prisoner 2's payoff. Thus, in the case of both prisoners denying the crime, both prisoner 1 and prisoner 2 will have to spend one year in jail.

Now that the reader understands what the yellow scenario is saying, we can take a look at the green box. The green box represents the payoffs the prisoners would get if Prisoner 1 denied the crime and Prisoner 2 rats Prisoner 1 out and confesses to the crime. In this case, Prisoner 1 will get ten years in jail, whereas Prisoner 2 would be allowed to walk away a free man for his cooperation.

Alternatively, Prisoner 1 could confess to the crime and Prisoner 2 could deny all the accusations. If this scenario occurred, Prisoner 1 would be allowed to walk a free man and Prisoner 2 would have to spend ten years in jail.

Finally, there is the red scenario. The red scenario actually denotes a very special case and a description as to why will soon follow. Back to discussing the proper interpretation of our 2x2 grid drawn above, the red box represents the payoffs the prisoners will win if Prisoner 1 confesses and Prisoner 2 confesses. As you can see from the ordered pair written in the box, both prisoners will win 8 years in jail. But why do they get 8 years you ask? After all, they had both cooperated with the authorities. Well, the truth of the matter is the authorities no longer have a reason to cooperate with them because they got what they ultimately wanted from the prisoners in the first place. As you will soon see, it is this logic that makes this scenario a very effective technique for interrogators to gain confessions in the real world. Furthermore, you now know why suspects are always interrogated alone!

As mentioned above, the red box represents a special case for several reasons. In order to understand why, we will ask a series of questions. Our first question is what strategy should Prisoner 1 select if he thought Prisoner 2 was going to confess? So, given that Prisoner 2 is confessing in this hypothetical scenario, Prisoner 1 has the power to place us within the green box if Prisoner 1 denies or the red box if he chooses to confess. Given these two situations the red box results in only eight years for Prisoner 1 as compared to a possible ten years if he denies. Thus, logically Prisoner 1 should elect to confess if he thought Prisoner 2 was going to confess.

Our next question is what strategy should Prisoner 1 employ if he thought Prisoner 2 was going to deny the charges? Given that Prisoner 2 denies, this gives Prisoner 1 the option of placing us within the yellow box or the blue box. In this case, Prisoner 1 should logically choose to confess because confessing would result in no time in jail as opposed to one year in jail.

Now we will flip the questions around and look from the viewpoint of Prisoner 2. What strategy should prisoner 2 employ given that Prisoner 1 decided to deny the charges? Well, this gives Prisoner 2 the option of selecting the yellow or green box. Logically, Prisoner 2 would want to confess in this case because confessing would yield him zero years in jail whereas denying would earn him one year in jail. Finally, let's ask the question of what Prisoner 2 should do given that Prisoner 1 decided to confess to the charges? This scenario gives Prisoner 2 the choice of placing us in the blue or red box, depending on whether he decides to deny or confess, respectively. Logically, Prisoner 2 would want to confess to the charges in this case as well, since the eight years in jail Prisoner 2 would receive from confessing is less than the ten years of jail he would receive from denying.

At this point, the reader has probably noticed a very interesting thing. Regardless of what the other player does, it is actually in each prisoner's best interest to confess. Thus (Confess, Confess) is said to be a **dominant strategy**. Now here is the tricky thing; the best decision each of the prisoner's can make given the uncertainty of the decision the other player will make actually results in the maximum combined jail time for the two prisoners. It is paradoxes such as these that make Mathematics and Game Theory seem very peculiar at times.

2.2 The Prisoner's Dilemma and Nash Equilibrium

At this point, it is important to realize that the majority of games do not have dominant strategies in the real world. It is for this reason that life seems so complex at times. However, even in games that do not have dominant strategies, economist and mathematician John Nash proved something of such profound importance that his discovery, the Nash Equilibrium, resulted in his winning a Nobel Prize. What John Nash proved is that every finite game has at least one Nash equilibrium. In a nutshell, this means that even though a game may not have a dominant strategy, there is still plenty we can say about it anyway, assuming we cannot predict the outcome outright, which is often the case. In simple terms, the **Nash Equilibrium** represents the intersection of strategies such that neither player will want to abandon his or her strategy. Like all things in mathematics, this concept can best be understood by an example. In order to get a firsthand account of what Nash equilibrium is, we will find that the Prisoner's dilemma presented above does indeed possess Nash equilibrium. We will find the Nash equilibrium of the Prisoner's dilemma by "walking around" the previously given grid. So, without further adieu, let us begin.

The first point we come to is the yellow colored box. From this point, Prisoner 1 can decide to stay in the yellow box or move to the blue box. Since $0 > -1$, Prisoner 1 will want to change to the blue box strategy. Thus, the yellow scenario cannot be a Nash equilibrium. Now we come to the green box. Prisoner 1 can either decide to stay in the green box or move to the red box. Since $-8 > -10$, Prisoner 1 will want to move to the red box. Since even one player (Prisoner 1) does not like this strategy, we know the green box does not represent Nash equilibrium. Now, let's look at the blue box. From the blue box, Prisoner 1 can decide to stay in the blue box or move to the yellow box. Since $0 > -1$, Prisoner 1 is actually happy with staying in the blue box. Now, we have to see if Prisoner 2 is content with the blue box strategy. From the blue box, Prisoner 2 can either decide to stay in the blue box or migrate on over to the red box. Since $-8 > -10$, Prisoner 2 would like to migrate over to the red box. Thus, blue cannot represent a Nash equilibrium.

Finally, we come to the red box. From the red box, Prisoner 1 can either stay in the red box or jump ship to the green box. However, since $-8 > -10$, Prisoner 1 is right at home with the red box strategy. Now, let's see what Prisoner 2 thinks. From the red box, Prisoner 2 can elect to stay there or migrate over to the blue box. However, -8 is once again greater than -10. Thus Prisoner 2 is also happy with the red box strategy. Thus, since each prisoner is happy with the red box strategy of (Confess, Confess) that represents our **Nash Equilibrium** in addition to a **dominant strategy**. It is for this reason that independent interrogation is extremely effective at eliciting the maximum number of confessions.

At this point, the reader is probably thinking the Prisoner's Dilemma has been explained to death, and it has. However, if the reader can understand the Prisoner's Dilemma, he can say he is familiar with all of the basic high level ideas of Game Theory.

2.3 A Tale of Two ~~Cities~~ Pigs

The game that follows tells the story of two pigs—one little pig and one big pig (the dominant pig)—which are locked in a box. At one end of the box, there is a button one of the pigs can push that will cause food to come down a chute at the complete other end of the box. The payoffs and strategies are given as shown below.

		Little Pig	
		Press	Wait
Big Pig	Press	5,1	4,4
	Wait	9,-1	0,0

(Big Pig, Little Pig)

Now, let us explain the payoffs. The yellow box represents both the big pig and little pig rushing to push the button. This results in the big pig netting 5 food units and the little big pig netting a single food unit. This is due to the fact that they both expend a food unit while they rush to push the button. Moreover, the big pig is capable of pushing the little pig out of the way on the way back to the chute, so the little pig only gets a single food unit, i.e. what the big pig does not want. The green box represents the scenario where the big pig rushes to press the button while the little pig waits at the chute. This enables the little pig to eat a fair share before the big pig gets back, as the big pig is fat and slow. It is in this way a roughly equal share of the food is had by both the big pig and the little pig. The red box represents both the big and the little pig sitting at the chute and waiting. Obviously, no food is dispensed until one of them presses the button at the other end of the box. This results in neither the little pig nor the big pig getting any food. Finally, the blue box represents the big pig waiting at the chute while the little pig rushes off to push the button. This scenario results in the big pig getting the lion's share of the food and the little pig, not surprisingly, starving. So, the question that this problem raises is what can we expect the little pig and the big pig to do assuming they both act rationally? In other words, what is the Nash equilibrium in this problem?

In order to find the Nash equilibrium we will again "walk around" the grid as we did in solving the Prisoners' Dilemma. So, let us begin with the yellow scenario, which is with the big pig rushing to press the button and the little pig likewise rushing to press the button. From this location in the game, the big pig has the power to move to the blue box if he so chooses. Since 9 > 5, the big pig would undoubtedly rather wait as opposed to pushing the button. Since the big pig would rather be in the blue box outcome, the yellow box cannot possibly be a Nash equilibrium.

Now, let's look at the blue box. The big pig has the power to stay in the blue box or to move to the yellow box. Since 9 > 5, if left up to the big pig, he would rather stay in the blue box. The question now becomes whether the little pig would prefer the blue box outcome or the red box outcome. Since 0 > -1, the little pig would prefer the set of strategies resulting in the red box outcome. Thus, the blue box cannot possibly be a Nash equilibrium.

It is at this point that we find ourselves analyzing whether the red box is a Nash equilibrium. The big pig has the deciding power to stay in the red box or to migrate to the green box. Since 4 >0, the big pig would prefer to move to the green box. Therefore, the set of strategies (wait, wait) cannot be a Nash equilibrium.

This brings us to the green box at long last. Since we have a theorem showing that every finite game has at least one Nash equilibrium the green box had better be it. But is it? From the green box location, the big pig has the option of staying there or moving to the red box. However, since 4 > 0, the big pig would rather remain at the green box. How about the little pig? Would he rather stay in the green box or migrate over to the yellow box? Well, since 4 > 1, the little pig is happiest at the green box. It is for this reason—the fact that neither pig would prefer another strategy once they arrive at the green box—that we can say (Press, Wait) is the Nash equilibrium for this game. Thus, one would expect that the little pig would wait for the big pig to press the button and start eating while he is away. The big pig would then be able to kick the little pig out of the way to get his share, and both pigs are filled.

2.4 The Government and the Pauper

We will start the analysis of this problem as we have done the previous problems by first displaying the players, strategies, and payoffs for this game.

		Pauper	
		Work (y)	Loaf (1-y)
Government	Aid (X)	3,2	-1,3
	No Aid (1-x)	-1,1	0,0

(Government, Pauper)

One can think of this as being a simplified representation of the payoffs and strategies associated with the welfare system present in our country today. So, what is the Nash equilibrium in the above problem? In an effort to answer this question we will once again “walk” our way through each of possible pairs of strategies and their resulting payoffs.

Let us begin with the yellow box. The government has the option of staying in the yellow box or moving to the blue box. Since 3 > -1 the government would like to stay in the yellow region. How about the Pauper? The Pauper has the option of staying in the yellow box or migrating to the green box. Since 3 > 2, the Pauper would like to move to the green box. Therefore (Aid, Work) cannot be a Nash equilibrium.

This brings us to the green box. The government has the option of staying in the green box or jumping ship to the red box. Since 0> -1, the government would prefer to be in the red region; thus (Aid, Loaf) cannot be a Nash equilibrium.

How about the red box? Is it a Nash equilibrium? Well, the government has the power to stay in the red box or to migrate to the green box. However, we know the government would prefer the green box from the scenario above. So, how about the Pauper? The Pauper has the option of staying in the red box strategy location or moving to the blue box. Since 1>0 the pauper would rather be in the blue box. So, (No Aid, Loaf) cannot be a Nash equilibrium.

It seems our last hope of having a Nash equilibrium resides in the blue region or (No Aid, Work). The government has the deciding power to stay within the blue box or to migrate to the yellow box. However, 3>-1 so the government would rather be in the yellow region. Thus, (No Aid, Work) cannot possibly be a Nash equilibrium. Now, wait just a minute. Did Nash not prove that every finite game has a Nash equilibrium? Every game does indeed have a Nash equilibrium; however, it is hidden a bit in our game between the government and the pauper. It turns out that the government and the Pauper can utilize a **mixed strategy** to arrive at a Nash equilibrium. Simply put, a **mixed strategy** simply means that the players do not use any one given strategy one hundred percent of the time. Rather, they use a combination of strategies available to them with varying probabilities. Working with our example here, the government could decide to aid the pauper 70% of the time and to not aid him 30% of the time. That would be a possible mixed strategy the government could employ. So, how are we going to find the **mixed strategy Nash equilibrium** for this problem?

We begin by assigning an arbitrary probability to each of the strategies that government and the pauper can engage in. So, let's say the government will aid the Pauper with P(Aid) = x. Thus, we know the government will choose to not Aid the Pauper with P(No Aid) = 1-x since $Aid \bigcup NoAid$ is equal to the sample space of the government's strategies and the axiom of probability that states that $P(\bigcup_{SampleSpace} x_i) = 1$. Similarly, we can assign P(work) = y and P(Loaf) = 1-y for the pauper. The next thing that we need to do is calculate the payoff for the government and the pauper in terms of these variable probabilities.

We begin by calculating the expected payoff of the government, denoted by π_g.

Keeping in mind the fact that one can calculate the expected payoff by multiplying the probability of each strategy times its associated payoff. We will now go through step-by-step for the government's payoff. We will start with (Aid, Work). The government's associated payoff from (Aid, Work) is by this notion 3xy. Similarly, the government's payoff from (Aid, Loaf) is given by -1(x)(1-y). Next, the government's payoff from (No Aid, Work) is given by -1(1-x)(y). Finally, the government's expected payoff from (No Aid, Loaf) is simply zero. Therefore the total expected payoff of the government is equal to: $\pi_g = 3xy + -1(x)(1-y) + (-1)(1-x)(y) + 0$. We can expand this result to obtain $\pi_g = 3xy - x + xy + xy - y \rightarrow \pi_g = 5xy - x - y$. Very similarly, the payoff for the pauper, π_{Pauper} can be given as follows: $\pi_{Pauper} = -2xy + y + 3x$. Now that the expected payoff for both the government is known we can use the definition of Nash equilibria and a bit of Calculus to find the **mixed strategy Nash equilibrium** for the government and the pauper. Recall that the Nash equilibrium for any given game is the point at which none of the players are willing to change strategies. Thus, it stands to reason that if we can find the point at which both the government and the pauper are receiving their maximum payoffs then we have found the mixed strategy Nash equilibrium.

With the definition of Nash equilibrium in mind we set about differentiating and setting equal to zero the government's payoff. However, there are two possible variables that one can use to differentiate π_g, so what variable should we pick? The answer to this question is surprisingly simple. The government can only dictate what its strategy is, so we want to find $\frac{\partial \pi_g}{\partial x}$. $\frac{\partial \pi_g}{\partial x} = 5y - 1$. Equating 5y-

1 to 0 results in y being equal to 1/5. Thus, the government is receiving its maximum payoff if the Pauper is working 20% of the time and Loafing 80% of the time. Using very similar logic we find that $\frac{\partial \pi_{pauper}}{\partial y} = -2x + 1$. Setting $\frac{\partial \pi_{pauper}}{\partial y} = 0$ yields x = 1/2. Thus, the pauper is receiving his maximum payoff as long as the government gives aid 50% of the time. To summarize, the Nash equilibrium of this seeming paradox is given as follows:

		Pauper	
Government		Work (1/5)	Loaf (4/5)
	Aid (1/2)	3,2	-1,3
	No Aid (1/2)	-1,1	0,0

(Government, Pauper)

So, with the mixed strategy given above, all the above regions—yellow, green, blue, and red—are Nash equilibria. So, this begs the final question of how one can determine what is most likely to come about from this symbiotic relationship between the government and the Pauper. Well, one can assume that each of these events: (Aid, Work), (Aid, Loaf), (No Aid, Work), (No Aid, Loaf) are independent from one another. With this assumption one finds the following to be true:

Strategy	Probability
(Aid, Work)	1/10
(Aid, Loaf)	4/10
(No Aid, Work)	1/10
(No Aid, Loaf)	4/10

Thus, we see that the **dominant mixed strategy Nash equilibria** are (Aid, Loaf) and (No Aid, Loaf).

3.0 Exercise To the Reader

Can you find the **mixed strategy Nash equilibrium** for Rock, Paper, and Scissors? In order to get you started I have provided the payoffs below. Hint: It is probably what you think it is, but you must prove your intuition is right using the method presented in 2.4. Additionally, why do the sums of the payoffs in each box equal zero?

		Player 2		
		Rock (x)	Paper (y)	Scissors (1-x-y)
Player 1	Rock (a)	0,0	-1,1	1,-1
	Paper (b)	1,-1	0,0	-1,1
	Scissors (1-a-b)	-1,1	1,-1	0,0
(Player 1, Player 2)				

Bibliography

Rasmusen, Eric. *Games and Information (Second Edition).* Malden: Blackwell Publishers Inc, 1994

An Introduction to Nonstochastic Game Theory

By Tyler Rainey

Math 4513: Senior Seminar