Duke University

Nash Equilibrium

Benjamin Michael Nativi

Math 89S - Math of the Universe

Professor Hubert Bray

December 6, 2016

Introduction

The field of game theory has its roots in 1921 when Emile Borel first defined the idea of a finite and symmetric zero-sum two player game as a game of pure conflict. This extremely specific and abstract concept of a game is designed to study two people who are purely at conflict, where one person gains at the loss of another person, and the game is only played a set number of times. The rules of the game are also the same for each player. Better known than Borel, John Von Neumann then formalized the idea of a finite zero-sum game and developed a proof to show that pure strategies for such games existed. For 30 years, Borel, Von Neumann, and a partner of Von Neumann, Oskar Morgenstern dominated the field of game theory, but were scolded for their abstract ideas. It wasn't until John Nash proved his idea of Nash equilibrium that game theory was seen as a serious field with real world applications. Designed for games with any number of players and common knowledge between the players, the Nash equilibrium, and the later refined subgame perfect equilibrium have boosted game theory to a study involved in fields ranging from economics to evolutionary biology (Sethi et al., 2016).

Game Theory

Game theory, at its most fundamental level, is an understanding of how rational humans interact. Using a quantitative view of these interactions, game theorists are able to predict not why people interact the way they do, but simply what people will do when they interact.

In game theory, games are often simplified representations of more complex large scale interactions. Games that are used to represent or explain phenomena often involve just two players; Alice and Bob. Game theory attempts to predict how Alice and Bob will interact when each player is given the same rules for a game. One of the most significant requirements for the success of game theory is that both Alice and Bob have to be rational. Rationality is defined as consistency in choice, where a player would make the same decision if the exact same game were played twice. For Alice to make a choice in a game, Alice must know that Bob is rational and Bob must know that Alice is rational. If Alice and Bob are not rational players, then their choices cannot be predicted (Binmore, 2007).

Another component of a game, as important as the players, is the choices of a game. Each player is usually given a finite number of choices, and in the games we will consider, each player knows all of their choices and all of the choices for the other player. In games without perfect information, certain game theory solutions have been created but are more complex than this paper gets into.

The final component of a game that is important for both Alice and Bob to understand is called utility. Every choice for Alice and Bob has an associated payoff, defined as the benefit that Alice and Bob receive from that choice, but when Von Neumann and Morgenstern were developing their ideas on game theory, they realized that they needed some way to quantitatively compare payoffs. As a result, utility was invented. Utility is a ranking of how prefered a situation is. Utility can be thought of as a scale from 0 to 100, where 0 represents the worst possible situation for a player and

100 represents the best possible situation for a player. The actual unit for utility is then called a util. So if Alice is rational, then Alice will always pick a situation worth 75 utils over a situation worth 60 utils. Utility isn't a measurement of Alice's happiness, or the amount of money Alice receives, or any other tangible concept, because some people value money more than others, and ideas like happiness are difficult to define. Another way to think of a util is that a util is a gauge of the amount of risk Alice is willing to take to get a situation. Say Alice is given a choice of either getting a free car or getting a free lottery ticket. The lottery ticket is either worth 100 utils or 0 utils with some probability. If Alice chooses the lottery ticket over the car only when the lottery ticket has at least 75% chance of winning the 100 utils, then the free car is worth 75 utils (Binmore, 2007).

Our first example of a game is called the Coordination Game (Figure 1). In the coordination game, Alice and Bob both drive up to an intersection from opposite directions. Each player has the choice of turning either left or right. If one turns left and one turns right, their cars collide, costing them -1 util. If they both turn left or right, they continue on their way at the benefit of 1 util. The negative util represents that the situations are worse off than if the players did not play the game at all. In the table, Alice's payoff is the left value in each cell and Bob's payoff is the right value.

	Left	Right
Left	1, 1	-1, -1
Right	-1, -1	1, 1

Figure 1. The Coordination Game.

Utility doesn't have to be, but is often related to money. Most rational people consider a situation where the person receives more money to be worth more utils than a situation where the person receives less money, although the exact number of utils varies by person. When dealing with money, we have the concept of risk neutral and risk averse people. Someone who is risk neutral values all additional money at the same utility. For a risk neutral player, if gaining 1$ was worth 1 util, then gaining 3$ is worth 3 utils. Someone who is risk averse values each additional dollar less than the previous dollar. While gaining 1$ may be worth 1 util, gaining 3$ is only worth 2 utils (Binmore, 2007).

Minimax Theorem

Von Neumann made the first major strides in game theory when he refined Borel's idea of a finite and zero-sum two player game. A finite and zero-sum two player game is a game between two players where the game is played a set number of times, usually once, and one player only gains utility at the cost of utility for the other player.

An example of a finite zero-sum game is called Matching Pennies (Figure 2).

	Heads	Tails
Heads	1, -1	-1, 1
Tails	-1, 1	1, -1

Figure 2. Matching Pennies.

In this game, Alice and Bob each have a coin with heads on one side and tails on the other. Each player chooses a side of the coin and reveals the coin at the same time. Alice wins if both coins are heads or tails, and Bob wins if the coins are different. Here, Alice's two strategies are either heads or tails and are listed on the left, and Bob's are the same strategies listed on the top. In this game, Alice and Bob are pitted in pure conflict. To be zero-sum, the payoffs for Alice and Bob have to sum to zero in each cell. This is easily seen in Matching Pennies because the sum of each cell is 1-1=0.

A solution to this game can easily be found using the minimax theorem, first proved by Von Neumann in 1928. The minimax theorem shows that all zero sum finite two player games have a solution called the minimax strategy that is equal to the maximin strategy. The goal of a minimax strategy is to look at the best payoffs of each available strategy and to then choose the worst of those payoffs. The goal of a maximin strategy is to take the worst payoffs of each available strategy and choose the best of those payoffs. The worst best payoff is always equal to the best worst payoff. Slightly convoluted, but what the minimax theorem really suggests is that you first assume that the other player already knows what move you will make. In Matching Pennies, Bob already knows that Alice plans to play heads. Then Bob will play tails to get +1 util while Alice gets -1 util. The same happens if Alice plans to play tails. Thus the minimax strategy for Alice is to play tails and heads each with a 50% probability each turn. Bob, knowing this strategy, can play all heads, all tails, or any combination of heads and tails and still only wins 50% of the time. But if Bob does anything other than 50% heads and tails, Alice would change her strategy to increase her utility. So both Bob and Alice should play the minimax strategy of 50% heads and 50% tails with an average payoff of 0 utils for each of them (Binmore, 2007).

The minimax theorem has some serious drawbacks. The theorem only applies in the very seldom number of games that are pure conflict between two players, and the vast majority of human interactions are not pure conflict. Also, the theorem assumes that Bob already knows Alice's move ahead of time, in which case Bob isn't really a player making a choice, so Alice isn't really playing a game but simply making a choice between a set of payoffs. In comparison, the Nash equilibrium is much broader.

Nash Equilibrium

John Nash proved in 1952 that finite games of any number of players, that can be non zero-sum and nonsymmetric, always have a Nash equilibrium solution. Nash proved this idea while he was a graduate student and later won the 1994 Nobel Prize in economics for his work.

A Nash equilibrium exists when both Alice and Bob, or any number of players, knowing the other player's move would also make the same move. As an example, let's again look at the Coordination Game (Figure 3).

	Left	Right
Left	1 , 1	-1, -1
Right	-1, -1	1 , 1

Figure 3. The Coordination Game with Nash equilibrium.

In this game, for each decision that Alice makes, Bob has circled his best response in blue. The same has been done for Alice. When these two circles line up, this is a pure Nash equilibrium. If Alice plays left, Bob should also play left, and vice versa. Thus (left, left), where Alice and Bob both play left is an equilibrium strategy with payoff (1,1). Neither Alice or Bob want to deviate from playing left or else both players will receive (-1,-1).

Another game that easily exemplifies Nash equilibrium is the game Chicken (Figure 4). In a classic version of the game Chicken, Alice and Bob are driving towards each other in a narrow alley way. Alice and Bob can only pass each other safely if at least one of them drives slowly. Otherwise, both cars crash. This game can be represented where if both play the strategy slow, both get 3 utils. For one speeding and one driving slowly, the speedy driver gets 4 utils and the slow driver gets 2. If both speed, both drivers get -1 util.

	Slow	Speed
Slow	3, 3	2 , 4
Speed	4 , 2	-1, -1

Figure 4. The game Chicken with Nash equilibrium.

In Chicken, (speed, slow) and (slow, speed) are both pure Nash equilibriums. For Alice, if she knows that Bob will play slow, then Alice will speed to get 4 utils instead of 3. If Alice knows Bob will play speed, then Alice will play slow to get 2 utils instead of -1.

So far, our games have only had pure Nash equilibriums, where each player plays one strategy all the time. We can also see that some games have mixed Nash equilibriums, where each strategy is played with a certain probability. In Mixed Pennies, the minimax strategy of 50% heads and 50% tails is an example of a mixed Nash equilibrium. Another example can be seen in the game Battle of the Sexes (Figure 5). Imagine Alice and Bob argue over dinner whether to go a boxing match or a ballet. Alice would rather go to the ballet and Bob would rather see boxing. Alice and Bob are somehow split up after dinner and must decide each whether to go to ballet or boxing.

	Ballet	Boxing
Ballet	2, 1	0, 0
Boxing	0, 0	1, 2

Figure 5. The Battle of the Sexes with Nash equilibrium.

Battle of the Sexes has two pure Nash equilibriums (ballet, ballet) and (boxing, boxing). Both of them would rather go to the same place then go to different locations, but each wants to go to their location more than the other location. This leaves the possibility of a third Nash equilibrium. Imagine you are Alice. To get Bob to go to ballet, you must make Bob's payoff from going to ballet equal to his payoff from going to boxing. Thus, Alice must play ballet twice as much as boxing, so that she plays ballet $\frac{2}{3}$ of the time and boxing $\frac{1}{3}$ of the time. For Bob, if he plays ballet, he has $\frac{2}{3}$ of a chance at 1 util, or if he plays boxing, he has $\frac{1}{3}$ chance at 2 utils, so Bob always has an average payoff of $\frac{2}{3}$

util. But if Bob plays ballet more than ⅓ of the time, Alice will play ballet more often and the same holds for boxing. Thus Bob must play boxing ⅔ of the time and ballet ⅓ of the time. This (⅔,⅔) payoff is not as high as (1,2) or (2,1) but is the best that can be done without knowing the other player's move. This is also true for the coordination game, where a mixed Nash equilibrium exists at playing left and right each 50% for a payoff of (½, ½).

Our final Nash equilibrium example is called the Prisoner's Dilemma (Figure 6). Imagine that Alice and Bob are gangsters who have committed a crime and were caught. The police don't have enough evidence to convict, so the police make up a weaker charge to accuse both Alice and Bob with. Alice and Bob are separated into different rooms and offered a deal. If Alice rats on Bob, Alice goes free, and Bob gets a full sentence. If Bob doesn't stay silent and also rats on Alice, both Bob and Alice receive almost the full sentence for cooperating. If both stay silent, they get the minor charge.

	Silent	Rat
Silent	-1, -1	-3, 0
Rat	0, -3	-2, -2

Figure 6. The Prisoner's Dilemma with Nash equilibrium.

Each negative util represents a year in prison, so 1 year is given for the minor charge, 3 for the major crime, and 2 for the major crime if both confessed. The only Nash equilibrium for the Prisoner's Dilemma is (rat, rat) resulting in a payoff of (-2, -2). This concept has sparked controversy since Nash first proved his theorem in the 50's. The reason for this Nash equilibrium is to imagine that Alice plays silent. Bob would then play rat to get a util of 0 instead of -1, so one less year in jail. If Alice played rat, then Bob should also play rat to get -2 utils instead of -3. Alice knows that Bob will always play rat, so Alice also always plays rat. This seems counterintuitive as clearly the strategy (silent, silent) would give (-1,-1) which is better for Alice and Bob, but if Alice and Bob are rational players attempting to maximize utility, then Alice and Bob will always rat.

Finding Nash equilibriums are beneficial for multiple reasons. First, Nash equilibriums are solutions for what rational players would play, so game theorists can accurately predict the behavior of rational players. Second, even if the players of a game don't know the Nash equilibriums of a game, when a game is played many times, rational players eventually adopt the Nash equilibrium strategies. If we look at the game Chicken, for example, we can imagine that Alice and Bob play Chicken many times. Say Alice notices that Bob is playing speed most of the time, so Alice would adjust her strategy to play slow more often. Bob would see Alice playing slow more often and would play speed more often. Eventually, Alice would always play slow and Bob would always play speed (slow, speed) and either player would be punished for deviating from this strategy. Finally, Nash equilibriums can be used in a variety of applications. Because of Nash's theorem and further work, game theory is now being used in diverse applications of economics, evolutionary biology, and political science (Binmore, 2007).

Subgame Perfect Equilibrium

Subgame perfect equilibrium is a specific case of Nash equilibrium that can only be applied to extensive form games; games where actions are taken at different times. Then each action forms its own subgame with its own Nash equilibrium. A Nash equilibrium is then only a subgame perfect equilibrium if the Nash equilibrium is also a Nash equilibrium for all of the subgames (Yildiz, 2012).

For example, imagine that Alice and Bob are playing the Ultimatum game (Figure 7) and (Figure 8). In this game, a rich benefactor is willing to give Alice and Bob a large sum of money only if they can agree on how to split the money. Alice must first decide to split the money in a fair way 50/50, or an unfair way 75/25. Bob then decides whether to accept or reject the offer. If Bob rejects, Alice and Bob both don't receive any money. Bob automatically accepts a fair deal, so (fair, reject) is not considered in Figure 7. For this game, suppose that the benefactor is offering 4 utils. Alice can split these utils (2,2) or (3,1) in her favor.

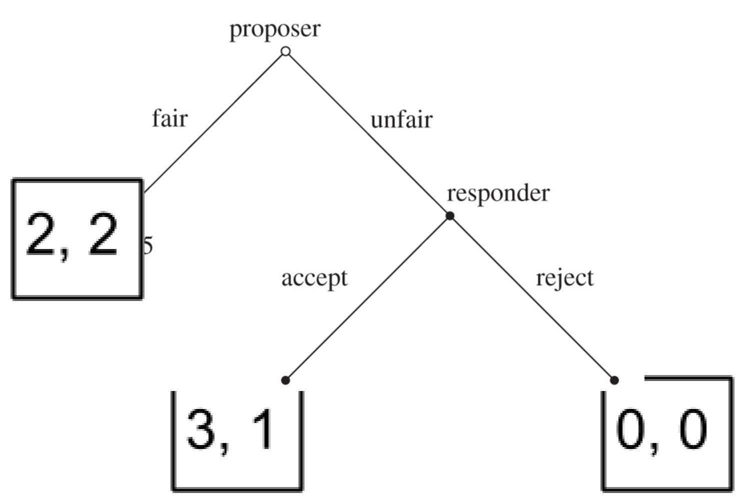

Figure 7. A game tree diagram for the Ultimatum Game.

	Accept	Reject
Fair	2, (2)	(2) , (2)
Unfair	(3) , (1)	0, 0

Figure 8. The Ultimatum Game with Nash equilibrium.

Using the concept of subgame perfect equilibrium, we can look just at the subgame where Alice chooses an unfair offer. If Alice chooses an unfair offer, then Bob will make the choice of accepting for 1 util instead of rejecting for 0 utils, so accept is the Nash equilibrium of this subgame. Now, knowing that Bob will accept an unfair offer, Alice is left with the choice of fair for (2,2) or unfair for (3,1), so Alice will choose unfair. Thus (unfair, accept) is the subgame perfect equilibrium for the Ultimatum Game. If we look at Figure 8, we can see that this game has two pure Nash equilibriums (unfair, accept) and (fair, reject). (fair, reject) is a Nash equilibrium because if Alice knew Bob will reject, Alice would play fair, and if Bob knew Alice played fair, then Bob would get (2,2) either with accept or reject. But reject is not a Nash equilibrium of all subgames so only (unfair, accept) is subgame perfect (Binmore, 2007).

Experimental Results

A main critic of game theory is that the theories don't accurately predict behavior in studies. One study by Ochs (1995) tested three different 2 player games where each player had two options and the games had a unique mixed equilibrium strategy. The players in the study did not often play according to the mixed Nash equilibrium. A study by Rankin et al. (2000) showed that players did not play as predicted in a simple coordination game similar to the coordination game shown above. In a study by Rydval and Ortmann (2005), a simple two player game was tested on players using all positive utils. This was compared to the results from a game structured in the exact same way except all utils were negative. Players showed different strategies in these two variants of the same game, which should not happen for rational players.

Other studies, though, have shown high correlation with equilibrium play. An experiment by Neugebauer et al. (2008) showed that players often played the harsher selfish strategies in games similar to the Prisoner's Dilemma. A study by Rey-Biel (2009) showed that, in 10 different games where players are given 3 strategies, there was a high proportion of players playing Nash equilibriums. In separate studies by Walker and Wooders (2001) on tennis players and by Chiappori et al. (2002) on soccer players and penalty kicks, it was seen that players randomize which side they shoot on as predicted with mixed Nash equilibriums.

There are three main reasons that players didn't play Nash equilibrium strategies. The first is that the players were not given enough trials to develop or recognize the Nash equilibriums. Many experiments likely only used participants each a single time. The second is that the incentives used in the studies likely did not accurately reflect utility. Money is usually used in these experiments, but some people may value their own morality, feeling like they help some other player, as a component of utility that needs to be accounted for. Third, Nash equilibriums only work if both players know the other player is rational. If one player assumed that the other player was irrational, the player may attempt a different strategy to get a higher utility.

Conclusion

While the maximin theorem was its precursor and the subgame perfect equilibrium its refinement, much research is still being done on Nash equilibriums. One focus of research since the 1990's has been work on the Generalized Nash Equilibrium Problem, essentially a model that allows for real world applications (Facchinei, 2007). Game theory is constantly used to describe new phenomenon and has become more of a mainstream science over the past couple decades. Research is also being focused on increasing the speed for finding solutions to Nash equilibriums. It has already been proven that all Nash equilibriums can be found in exponential time, but work is being done to find Nash equilibriums in polynomial time. This discovery would greatly increase the computer applications of Nash equilibriums as games could be solved much quicker for large numbers of players and parameters (Daskalakis, 2009). In conclusion, the discovery of Nash equilibrium greatly changed our understanding of human interaction and and has given us a new lens for viewing our environment.

References

Binmore, K. G., 1940-. (2007). Game theory : A very short introduction. Oxford ; New York: Oxford University Press.

Chiappori P-A, Levitt S, Groseclose T (2002) Testing mixed-strategy equilibria when players are heterogeneous: the case of penalty kicks in soccer. Am Econ Rev 92:1138–1151

Daskalakis, C., Goldberg, P. W., & Papadimitriou, C. H. (2009). The complexity of computing a nash equilibrium. SIAM Journal on Computing, 39(1), 195-259. doi:10.1137/070699652

De Souza, F. C., & Rêgo, L. C. (2014). Mixed equilibrium, collaborative dominance and burning money: An experimental study. Group Decision and Negotiation, 23(3), 1-24. doi:10.1007/s10726-013-9341-y

Facchinei, F., & Kanzow, C. (2007). Generalized nash equilibrium problems. 4or, 5(3), 173-210. doi:10.1007/s10288-007-0054-4

Muhamet Yildiz. 14.12 Economic Applications of Game Theory. Fall 2012. Massachusetts Institute of Technology: MIT OpenCourseWare, https://ocw.mit.edu. License: Creative Commons BY-NC-SA.

Neugebauer T, Poulsen A, Schram A (2008) Fairness and reciprocity in the hawk-dove game. J Econ Behav Organ 66:243–250

Ochs J (1995) Games with unique mixed strategy equilibria: an experimental study. Games Econ Behav
10:200–217

Rankin FW, Van Huyck JB, Battalio RC (2000) Strategic similarity and emergent conventions: evidence from similar stag hunt games. Games Econ Behav 32:315–337

Rydval O, Ortmann A (2005) Loss avoidance as selection principle: evidence from simple stag-hunt games. Econ Lett 88:101–107

Sethi, R., Weibull, J., Matematik (Inst.), Skolan för teknikvetenskap, (., KTH, & Matematisk statistik. (2016). WHAT IS… nash equilibrium? Notices of the American Mathematical Society, 63(5), 526-528. doi:10.1090/noti1375

Walker M, Wooders J (2001) Minimax play at Wimbledon. Am Econ Rev 91:1521–1538

Yildiz, M. (2012). Subgame-perfect nash equilibrium   MIT Open Courseware, Retrieved from https://ocw.mit.edu/courses/economics/14-12-economic-applications-of-game-theory-fall-2012/lecture-notes/MIT14_12F12_chapter11.pdf