

Game Theory: Normal Form Games

Michael Levet

December 22, 2015

1 Introduction

Game Theory is a mathematical field that studies how rational agents make decisions in both competitive and cooperative situations. It has widespread applications in economics, political science, psychology, biology, computer science, and data science. Some of the applications include radio spectrum auctions, voting, and organ donations. This tutorial introduces the basic strategic form game, also known as the normal form game. Attention will be restricted to pure strategies.

2 Model

In this section, we formally define the normal form game. Let's begin with some intuition. A normal form game has a set of players. Each player has a set of strategies. These players each select a strategy and play their selections simultaneously. In this manner, no player is responding to another's selection. Furthermore, we think of the players' strategies as setting the rules of the game. Finally, the selection of strategies results in payoff or utility for each player. Each player's goal in a game is to maximize utility. Each player is aware of the structure of the game; that is, the other players' strategy sets and payoffs. The normal form game will now be formally defined.

Definition 1 (Normal Form Game). A normal form game Γ is a three-tuple $[N, (S_i)_{i \in N}, (u_i)_{i \in N}]$ where N is the set of players, S_i is player i 's strategy set, and $u_i : \prod_{i \in N} S_i \rightarrow \mathbb{R}$ is player i 's payoff or utility function. The sequence of strategies $(s_1, \dots, s_n) \in \prod_{i \in N} S_i$ is referred to as a strategy profile.

In addition to the assumptions that the players are economically rational and play at the same time, it is also assumed that the structure of the game is perfectly known. In other words, each player knows every player's strategy set and utility function.

Let's examine an example of a normal form game, the standard Prisoner's Dilemma.

Example 1 (Prisoner's Dilemma). In this game, the police have two accomplices of a crime in separate rooms. They are each offered a deal: implicate the other prisoner and earn a reduced sentence if the other player remains silent. If both players remain silent, they each end up in jail for two years. If both players implicate each other, they each go to jail for five years.

Formally, we have two players $N = \{1, 2\}$. Each player has the strategy set $S_i = \{\text{Quiet}, \text{Fink}\}$, and the utility function of the form $u_i : S_i \times S_i \rightarrow \mathbb{R}$. If both players play Quiet, they each earn utility of -2 ; and if both play Fink, they each earn utility of -5 . If one player plays Quiet and the other Fink, they earn utilities of -10 and -1 respectively.

We represent the normal form game using the following matrix known as a **payoff matrix**. Player 1's strategies are on the left-side while Player 2's strategies are on the top of the matrix. Each cell represents the payoffs of the form $(u_1(s_1, s_2), u_2(s_1, s_2))$ for the selected strategies $s_1 \in S_1$ and $s_2 \in S_2$.

		Player 2	
		Quiet	Fink
Player 1	Quiet	$-2, -2$	$-10, -1$
	Fink	$-1, -10$	$-5, -5$

Not every normal form game can be represented as a matrix. When we have more than two players or continuous strategies, tables are not very helpful. Let's consider a second example of a normal form game: the Cournot duopoly.

Example 2 (Cournot Duopoly). In this game, we have two players again: $N = \{1, 2\}$. Each player is a firm producing the same, identical good. The market sets the price for the good based on the total amount produced by the two firms. The two firms compete in the quantities of the good they each produce, incurring a fixed cost $c > 0$ for each unit of good produced. Each firm seeks to maximize its own profit. Suppose the inverse demand function (the price function) is given as follows where $a, b > 0$:

$$P(q_1, q_2) = \begin{cases} a - b(q_1 + q_2) : & \text{if } q_1 + q_2 \leq a/b \\ 0 : & \text{if } q_1 + q_2 > a/b \end{cases} \quad (1)$$

Each firm has the strategy set $S_i = \mathbb{R}_+$, indicating the quantity of the good it can produce. Each firm also has the utility function of the form $u_i : \mathbb{R}_+^2 \rightarrow \mathbb{R}$ given by $u_i(q_1, q_2) = q_i P(q_1, q_2) - cq_i$. Since the strategies are uncountable, it is pointless to construct a payoff matrix to help in analyzing the game.

3 Solving Games and Nash Equilibrium

Recall that each player in a given game seeks to maximize its utility. How do agents select strategies? What is the solution concept for a game? The first notion of comparison is rather intuitive and straight-forward. We refer to it as strategy dominance. Basically, if strategy s_i yields at least as good a payoff as strategy s_j regardless of the other agents' strategy selections, then why would the player ever choose strategy s_j ? We first introduce the following notation. Let $i \in N$. We denote $-i := N \setminus \{i\}$. Now we formalize the notion of desirable strategies by defining a *Best Response*.

Definition 2 (Best Response). Let $i \in N$ and let $s_i \in S_i$. The strategy s_i is a best response of player i against $s_{-i} \in S_{-i}$ if $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$ for every $s'_i \in S_i$.

We denote $B_i(s_{-i}) = \{s_i \in S_i : u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}) \forall s'_i \in S_i\}$ as i 's *best response correspondence* against s_{-i} , or the set of all i 's strategies that are best responses to s_{-i} . Note that every element of $B_i(s_{-i})$ solves $\max_{s_i \in S_i} u_i(s_i, s_{-i})$.

Let's consider an example with a new game, a voting game.

Example 3 (Voting Game). Suppose we have three players, $N = \{1, 2, 3\}$ and two candidates A, B . Each player can vote for exactly one of A or B , so $S_i = \{A, B\}$, for all $i \in N$. The three players cast their votes simultaneously. Players 1 and 2 incur utility 1 if A wins and utility 0 if B wins. Player 3 incurs utility 0 if A wins and utility 1 if B wins.

Observe that for any $s_{-1} \in S_{-1}$, $B_1(s_{-1}) = \{A\}$. If $s_{-1} = (B, B)$, player 1 selecting $s_1 = A$ incurs the same utility of 0 as selecting $s_1 = B$. Otherwise, if player 1 selects $s_1 = A$ and at least one player of -1 selects A , then player 1 incurs utility 1 while selecting B may result in utility 0 if only one of the players in -1 selects A . In fact, we can say more strongly that for player 1, voting for A is a *weakly dominant strategy*. This is formally defined as follows.

Definition 3 (Dominated Strategies). The strategy $\bar{s}_i \in S_i$ is *weakly dominated* if there exists a second strategy $\hat{s}_i \in S_i$ such that: $u_i(\bar{s}_i, s_{-i}) \leq u_i(\hat{s}_i, s_{-i})$ for every $s_{-i} \in S_{-i}$, with strict inequality for at least one s_{-i} . We say that \bar{s}_i is *strictly dominated* if the strict inequality holds for all $s_{-i} \in S_{-i}$.

Now that we have some notion of how a player compares its strategies, the next step is to discern how the players select their strategies in anticipation of each other. The solution concept is the Nash equilibrium. Formally, the Nash equilibrium is defined as follows.

Definition 4 (Nash Equilibrium). The strategy profile s^* is said to be a Nash Equilibrium if $s_i^* \in B_i(s_{-i}^*)$ for every $i \in N$.

Intuitively, a strategy profile s^* is a Nash equilibrium if no player can unilaterally change its strategy and improve its outcome. Let's apply this reasoning to deduce the Nash equilibrium for the Prisoner's Dilemma game from Example 1, with the payoff matrix included below. If the two players both select Quiet, then they each incur utility of -2 . One of the players can unilaterally deviate, playing Fink instead, decreasing its utility to -1 while the other player's utility is decreased to -10 . So (Quiet, Quiet) is not a Nash equilibrium. Consider (Fink, Quiet). Player 2 can unilaterally deviate, playing Fink instead of Quiet to improve its payoff from -10 to -5 . By symmetry, (Quiet, Fink) is not a Nash equilibrium. Finally, consider (Fink, Fink). By previous analysis, a single player unilaterally changing from Fink to Quiet will decrease its payoff from -5 to -10 . So no player can unilaterally deviate from (Fink, Fink). Thus, (Fink, Fink) is the Nash equilibrium of the Prisoner's Dilemma.

		Player 2	
		Quiet	Fink
Player 1	Quiet	- 2, -2	-10, -1
	Fink	-1, -10	- 5, -5

The first concern is whether every game has a Nash equilibrium. The answer is yes, but not necessarily in pure strategies. A pure strategy is the selection of a single strategy from the set S_i which player i always uses. The Nash equilibrium of (Fink, Fink) is the pure strategy Nash equilibrium for the Prisoner's Dilemma. Nash equilibria are guaranteed to exist in mixed strategies, which will be introduced later. Additionally, we note that in a symmetric game (that is, a game where each player has the same strategy set and utility function), there exists a Nash equilibrium where each player selects the same strategy. The equilibrium of (Fink, Fink) in the Prisoner's Dilemma is actually a symmetric Nash equilibrium.

The problem of computing Nash equilibria is difficult. Formally, it is a complete problem for the complexity class PPAD. This means that given an arbitrary game, there is (believed to be) no efficient procedure to compute a Nash equilibrium for an arbitrary game. Additionally, a game may have multiple Nash equilibria. Enumerating all such equilibria as well as selecting the most realistic equilibria are both difficult problems of interest. We examine two strategies to help compute pure strategies Nash equilibria: leveraging games with continuous strategy sets and the elimination of dominated strategies.

3.1 Leveraging Continuity.

Recall Example 2, the Cournot Duopoly. Each player's strategy set is \mathbb{R}_+ , an amount to produce. Furthermore, each player has the same utility function, so this game is symmetric. Therefore, we solve for a symmetric Nash equilibrium. Each player seeks to solve $\max_{q_i} u_i(q_1, q_2) = (a - c)q_i - b(q_1 + q_2)q_i$. We consider the first partial derivative of u_i with respect to q_i , as player i cannot vary q_{-i} , and set it to 0 to identify potential maximizers: $a - c = 2bq_i + bq_{-i}$. Solving for q_i yields:

$$q_i = \frac{a - c}{2b} - \frac{q_{-i}}{2}$$

As this game is symmetric, there exists a Nash equilibrium where each player selects the same strategy. So we set: $q_i = q_{-i}$ and solve:

$$q_{-i} = \frac{a - c}{2b} - \frac{q_{-i}}{2} \implies q_{-i} = \frac{a - c}{3b}$$

Note that if $q_{-i} \geq \frac{a - c}{b}$, then $q_i = 0$ is player i 's best response.

3.2 Elimination of Dominated Strategies.

Recall the approach in reasoning the Nash equilibrium for the Prisoner’s Dilemma. Checking each strategy profile to determine if it is a Nash equilibrium is tedious. We prove a simple lemma, upon which the approach of eliminating dominated strategies is based.

Lemma 3.1. *Let $i \in N$. Suppose $s, t \in S_i$. If s strictly dominates t , then i will not play t in any Nash equilibrium.*

Proof. If s strictly dominates t , then $u_i(s, s_{-i}) > u_i(t, s_{-i})$ for every $s_{-i} \in S_{-i}$. If i plays t , then i can unilaterally deviate and choose s instead. So i will never play t in a Nash equilibrium. \square

The procedure below is the same regardless if attention is restricted to strictly dominated strategies or weakly dominated strategies. Note that Lemma 3.1 is not a necessary condition of weakly dominated strategies. However, a Nash Equilibrium found in a game produced from eliminating weakly dominated strategies is also a Nash Equilibrium in the original game. This will be proven later. Let’s begin by examining the procedures.

Definition 5 (Iterated Elimination of Strictly (Weakly) Dominated Strategies). : The procedure begins by accepting a game Γ where each player’s strategy set is finite. While there is a player $x \in N$ with strategies $x_1, x_2 \in S_x$ such that x_1 strictly (weakly) dominates x_2 , set $S_x := S_x \setminus \{x_2\}$. Consider Γ with the updated S_x at the next iteration of the procedure.

We apply the Iterated Elimination of Strictly Dominated Strategies to the Prisoner’s Dilemma, eliminating Quiet for the two respective players. This yields the sole strategy profile (Fink, Fink), which we recognize as the Nash Equilibrium for the game.

Now let’s apply the Iterated Elimination of Weakly Dominated Strategies to the game given by the following payoff matrix. Observe first this game has three Nash equilibria: (T, L) , (B, L) , and (B, R) .

		Player 2		
		L	C	R
Player 1	T	0, 1	1, 0	0, 0
	B	0, 0	0, 0	1, 0

Observe that Player 1 does not have a dominant strategy in this game. However, L dominates both C and R for Player 2. We remove R from Player 2’s strategy set and consider the reduced game:

		Player 2	
		L	C
Player 1	T	0, 1	1, 0
	B	0, 0	0, 0

In this new game, T dominates B for Player 1. So we eliminate B to obtain the following game:

		Player 2	
		L	C
Player 1	T	0, 1	1, 0

Player 2 will play L in this new game, yielding the Nash equilibrium (T, L) .

The Elimination of Weakly Dominated Strategies suffers from the fact that the order in which strategies are eliminated may result in a different Nash Equilibrium. This is due to the fact that weak dominance only requires one strict inequality, while strict dominance requires all strict inequalities. Consider again the above game:

		Player 2		
		L	C	R
Player 1	T	0, 1	1, 0	0, 0
	B	0, 0	0, 0	1, 0

If we instead eliminated M first from Player 2's strategy set initially, we would eliminate T from Player 1's strategy set at the next iteration. This would result in finding the Nash equilibrium (B, L) . By exhaustion, it is possible to verify that the Nash equilibrium (B, R) cannot be found by eliminating weakly dominated strategies.

These algorithms do not always yield pure strategies equilibria, but they do reduce the search spaces considerably. Eliminating a strictly dominated strategy preserves all Nash equilibria in a game. We have already seen that this is not the case when eliminating weakly dominated strategies. However, in a game derived from eliminating a weakly dominated strategy has a Nash equilibrium, which is also a Nash equilibrium in the original game. Let's prove these results formally.

I begin with the following Lemma:

Lemma 3.2. *Let Γ be a finite game and let Γ' be the game produced by eliminating a strictly dominated strategy in Γ . Then Γ and Γ' have the same set of Nash equilibria.*

Proof. Let Γ be a game and suppose the strategy s_j is eliminated from player j 's strategy set by the algorithm. Let Γ' be the resulting game. Let s^* be a Nash equilibrium for Γ' . Player j cannot unilaterally deviate in Γ' and improve its outcome. As s_j is strictly dominated, player j will not deviate to s_j in Γ . As every player $k \in N \setminus \{j\}$ has the same strategy set in Γ and Γ' , no other player can unilaterally deviate and improve its outcome. It follows that s^* is a Nash equilibrium of Γ .

Conversely, suppose q^* is a Nash equilibrium of Γ . By Lemma 3.1, s_j does not appear in any Nash equilibrium of Γ . It follows immediately that q^* is a Nash equilibrium of Γ' .

It follows that Γ and Γ' have the same set of Nash equilibria. □

Theorem 3.1. *Let Γ_0 be a finite game. Let $\Gamma_1, \dots, \Gamma_k$ be the sequence of games produced by the Iterated Elimination of Strictly Dominated Strategies. For every $i \in \{0, \dots, k - 1\}$, Γ_i and Γ_{i+1} have the same set of Nash equilibria.*

Proof. Theorem 3.1 follows immediately by applying induction and Lemma 3.2. □

Lemma 3.3. *Let Γ be a game and let Γ' be the game produced by eliminating a weakly dominated strategy in Γ . Then every Nash equilibrium in Γ' is also a Nash equilibrium in Γ .*

Proof. As Γ' is a finite game, it has a Nash equilibrium. Suppose that s_i was eliminated from player i 's strategy set in the construction of Γ' . Suppose to the contrary that there exists a Nash equilibrium s^* of Γ' that is not a Nash equilibrium of Γ . As deviating from s^* in Γ and Γ' is equivalent for $-i$, only player i can unilaterally deviate and improve its outcome. The only such option is for i to deviate in Γ is s_i , contradicting the assumption that s_i was eliminated as a weakly dominant strategy. □

Applying induction and Lemma 3.3 immediately implies the correctness of the Iterated Elimination of Weakly Dominated Strategies.

4 Conclusion

In this tutorial, the normal form game was introduced. In addition, the Nash equilibrium was defined and we explored ways to compute pure strategies Nash equilibria. The next tutorial will explore the role of mixed strategies.