Game Theory: Normal Form Games

Michael Levet

June 23, 2016

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. These notes introduce the basic strategic form game, also known as the normal form game.

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 $u_i : \prod_{i \in N} S_i \to \mathbb{R}$ is player i's payoff or utility function. The sequence of strategies $(s_1, ..., s_n) \in \prod_{i \in N} S_i$ 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 year.

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 $u_i : S_i \times S_i \to \mathbb{R}$. 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	Player 2: 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 $a, b > 0$:

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

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 $u_i : \mathbb{R}^2_+ \to \mathbb{R}$ 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 $\max_{s_i \in S_i} u_i(s_i, s_{-i})$.

Remark: Note that the Best Response Correspondence is only defined for pure strategies here. If we are interested in mixed strategies (which will be introduced later), we consider the *mixed-extension* of the normal form game Γ'. The mixed extension of a normal form game considers the same set of players and utility functions. However, each player i's strategy set in Γ' is $\Delta(S_i)$, where $\Delta(S_i)$ is the set of all probability distributions over i's strategy set S_i in Γ. We then consider the Best-Response Correspondence over $\Delta(S_i)$ rather than 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 $\overline{s_i} \in S_i$ is *weakly dominated* if there exists a second strategy $\hat{s}_i \in S_i$ such that: $u_i(\overline{s_i}, s_{-i}) \leq u_i(\hat{s}_i, s_{-i})$ for every $s_{-i} \in S_{-i}$, with strict inequality for at least one s_{-i}. We say that $\overline{s_i}$ 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. For finite normal form games, 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 $q_{-i} \geq \dfrac{a-c}{b}$, 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. $\square$

Theorem 3.1. *Let Γ_0 be a finite game. Let $\Gamma_1, ..., \Gamma_k$ be the sequence of games produced by the Iterated Elimination of Strictly Dominated Strategies. For every $i \in \{0, ..., k-1\}$, 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. $\square$

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. $\square$

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

4 Mixed Strategies

In this section, we introduce the notion of mixed-strategies. One important motivator for mixed-strategies is that not every game has a pure strategies Nash equilibrium. Consider the matching pennies game:

		Player 2: H	Player 2: T
Player 1	H	$1, -1$	$-1, 1$
	T	$-1, 1$	$1, -1$

In any pure strategy profile, one player incurs utility 1 and the other player incurs utility -1. The player incurring utility -1 can unilaterally deviate by switching its choice to improve its utility. This inverts the payoffs- the first player incurs utility -1 while the second player incurs utility 1. Iterating on the above argument, we see that no Nash equilibrium exists in pure strategies.

Definition 6 (Mixed Strategies)**.** Let Γ be a normal form game. Let $i \in N$. A mixed strategy is a sequence $(s_j)_{j=1}^{k} \in S_i$ and a probability distribution $\sigma = (\sigma_j)_{j=1}^{k}$ where player i selects strategy s_j with probability σ_j. Note that $\sum_{j=1}^{k} \sigma_j = 1$. The set of mixed strategies for player i is denoted $\Sigma_i := \Delta(S_i)$, where $\Delta(S_i)$ is the simplex in $\mathbb{R}^{|S_i|}$. That is, $\Delta(S_i) = \{x \in \mathbb{R}^{|S_i|} : x_i \geq 0 \ \forall i \in \{1, ..., |S_i|\}, \sum_{i=1}^{|S_i|} x_i = 1\}$.

Note that pure strategies are a special case of mixed strategies. The mixed extension will now be defined, to formalize the notion of games with mixed strategies. A mixed strategies Nash equilibrium in a normal form game is equivalent to a pure strategies Nash equilibrium in a mixed extension.

Definition 7 (Mixed Extension)**.** Let $\Gamma = [N, (S_i)_{i \in N}, (u_i)_{i \in N}]$ be a normal form game. The *mixed extension* of Γ is the three-tuple $[N, (\Sigma_i)_{i \in N}, (u_i)_{i \in N}]$, where $\Sigma_i := \Delta(S_i)$.

The notion of mixed strategies is rather unintuitive from a behavioral perspective, as a normal form game is played simultaneously. So how is a mixed strategies Nash equilibrium formulated? Recall that each player is a rational, utility maximizing agent that is aware of the structure of the game. Each player still seeks to mix its strategies in such a way to maximize its utility. In mixing strategies, a player runs the risk that another player can take advantage of a given mixing. Thus, in a Nash equilibrium, each player i mixes strategies such that $-i$ is indifferent to whichever pure strategy ends up being played. That is, $-i$'s expected utility for each of i's pure strategies in the mixing is the same. This is formalized as follows.

Theorem 4.1. *Let Γ be a normal form game. A mixed strategy profile σ^* is a mixed strategy Nash equilibrium if and only if, for each player i, the following two conditions are satisfied:*

1. *Every pure strategy $s_i \in S_i$ which is given positive probability by σ_i^* yields the same expected payoff against σ_{-i}^*; that is, $u_i(s_i, \sigma_{-i}^*) = u_i(\sigma^*)$.*
2. *Every pure strategy $s_i \in S_i$ which is given probability zero by σ_i^* yields no more than the pure strategies that are assigned positive probability: $u_i(s_i, \sigma_{-i}^*) \leq u_i(\sigma^*)$.*

Proof. Suppose first that the mixed-strategy profile σ^* satisfies conditions (1) and (2). Let $i \in N$. If i unilaterally deviates by shifting positive probability to a strategy $s_i \in S_i$ given zero probability in σ^*, then i's utility does not increase by condition (2). Let $S_i' = \{s_i \in S_i : \sigma_i^*(s_i) > 0\}$ be the set of strategies given positive probability by σ_i^*. By condition (1), each pure strategy in S_i' results in the same expected utility $\overline{u}$. Thus, any mixing γ of strategies in S_i' results in expected utility:

$$\sum_{s_i \in S_i'} \gamma(s_i) u_i(s_i, \sigma_{-i}^*) = \overline{u} \sum_{s_i \in S_i'} \gamma(s_i) = \overline{u}$$

Thus, player i cannot unilaterally deviate and improve its outcome, so σ^* is a mixed strategies Nash equilibrium.

Conversely, suppose σ^* is a mixed-strategies Nash equilibrium. As no player can unilaterally deviate and improve its outcome, condition (2) follows immediately. Suppose to the contrary that condition 1 does not hold. Let $i \in N$ and $s_i \in S_i$ such that the $u_i(s_i, \sigma^*_{-i}) \neq u_i(\sigma^*)$. If $u_i(s_i, \sigma^*_{-i}) > u_i(\sigma^*)$, then player i could assign more weight to s_i in σ_i^* and improve its outcome. Similarly, if $u_i(s_i, \sigma^*_{-i}) < u_i(\sigma^*)$, then player i could assign less weight to s_i in σ_i^* and improve its outcome. Either occurrence contradicts the assumption that σ^* is a mixed-strategies Nash equilibrium. □

Example 1: Let's now use Theorem 2.1 to find a mixed-strategies Nash equilibrium for the Matching Pennies game. Player 1 mixes strategies such that Player 2 is indifferent to H and T. Suppose Player 1 plays H with probability p and T with probability $1-p$. Player 2's payoff from playing H is $-p + (1-p) = 1 - 2p$. Player 2's payoff from playing T is $p - (1-p) = 1 - 2p$. In equilibrium, Player 2 is indifferent between playing H and T. Setting $1 - 2p = 2p - 1 \implies p^* = \frac{1}{2}$. By symmetry, we have Player 2 mixing between H and T with frequencies $(\frac{1}{2}, \frac{1}{2})$ as well.

In addition to guaranteeing the existence of a Nash equilibrium, mixed strategies are also useful in selecting realistic Nash equilibria. Consider the following example.

Example 2: Consider a traffic routing game on the following network. The weight of each edge denotes the latency cost of traversing that edge. The variable x denotes the number of players traversing the edge (A, B), and the variable y denotes the number of players using the edge (C, D). So for example, if $x = 50$, then the latency cost of (A, B) is 1.5 for every each of the 50 players. Each player starts at A and ends at D, seeking to minimize latency.

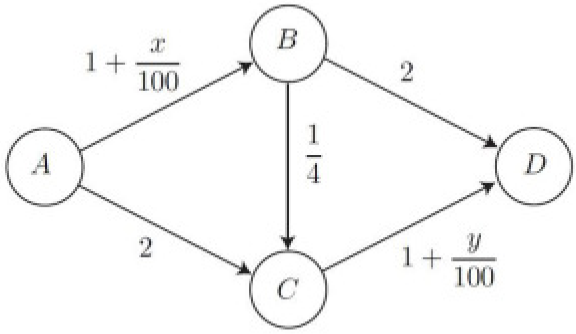

Suppose there are 100 players in the game. Denote n_1 as the number of players choosing the path ABD, n_2 as the number of players choosing the path ACD, and n_3 as the number of players choosing ABCD. Consider first the pure strategies Nash equilibrium of $n_1 = 25, n_2 = 25, n_3 = 50$. Both the edges (A, B) and (C, D) have 75 players traversing them, and so have latency costs 1.75. Players of each type incur latency cost 3.75. If a player of type n_1 unilaterally deviates, he increases the latency cost of the edge (C, D) to 1.76, resulting in a total latency cost of 3.76. By similar argument, players of type n_2 and n_3 cannot unilaterally deviate and decrease their costs as well.

While $n_1 = n_2 = 25$ and $n_3 = 50$ is a pure strategies Nash equilibrium, it is unlikely the 100 players will end up playing this strategy profile. However, this pure strategies equilibrium does provide the probabilities for a mixed strategies equilibrium. As the game is symmetric, there exists a Nash equilibrium in which each player selects the same strategy. Suppose each player selects the mixed strategy (ABD, ACD, ABCD) with probabilities $(\frac{1}{4}, \frac{1}{4}, \frac{1}{2})$. We apply Theorem 2.1 to verify this mixed strategy profile, denoted σ^*, is a mixed-strategies Nash equilibrium.

First, observe that $\mathbb{E}[u_i(\sigma^*)] = 3.75$. Consider each of the pure strategies ABD, ACD, ABCD.

- Suppose player i plays the pure strategy ABD. Under σ^*_{-i}, $n_1 = 24$, $n_2 = 25$, and $n_3 = 50$. So $\mathbb{E}[u_i(\text{ABD}, \sigma^*_{-i})] = 1.75 + 2 = 3.75 = \mathbb{E}[u_i(\sigma^*)]$.

- Suppose player i plays the pure strategy ACD. Under σ^*_{-i}, $n_1 = 25, n_2 = 24$, and $n_3 = 50$. So $\mathbb{E}[u_i(\text{ACD}, \sigma^*_{-i})] = 1.75 + 2 = 3.75 = \mathbb{E}[u_i(\sigma^*)]$.
- Suppose player i plays the pure strategy ABCD. Under σ^*_{-i}, $n_1 = n_2 = 25$ and $n_3 = 49$. So $\mathbb{E}[u_i(\text{ABCD}, \sigma^*_{-i})] = 1.75 + 0.25 + 1.75 = 3.75 = \mathbb{E}[u_i(\sigma^*)]$.

Thus, σ^* is a mixed-strategies Nash equilibrium.

5 Analysis and Topology Primer

So far, we have presented game theory from a practical perspective. The theorems above provide approaches for finding Nash equilibria in normal form games. However, the above results do not provide formal proofs for the existence of Nash equilibria. The goal of this section is to introduce enough analysis and topology to prove the existence of Nash equilibria in finite, normal form games. We begin wtih the definition of a metric space.

5.1 Elementary Notions From Analysis and Topology

Definition 8 (Metric Space)**.** Let X be a set of elements, and let $d : X \times X \to \mathbb{R}^+$ be a function. The pair (X, d) is a metric space if d satisfies the following:

$$d(x, y) \geq 0 \ \forall x, y \in X \text{ with equality precisely when } y = x \tag{2}$$

$$d(x, y) = d(y, x) \ \forall x, y \in X \tag{3}$$

$$d(x, z) \leq d(x, y) + d(y, z) \ \forall x, y, z \in X \tag{4}$$

Example 4. Euclidean space $\mathbb{R}^n$ with the standard Euclidean metric $d(x, y) = \sqrt{\sum_{i=1}^{n} (x_i - y_i)^2}$ forms a metric space.

Example 5. Let $G(V, E, W)$ be a simple, weighted, undirected graph. Consider the shortest path function $d : V \times V \to \mathbb{R}^+$, which returns the length of the shortest path between any two vertices. (G, d) forms a metric space.

We now need some basic definitions from topology.

Definition 9 (Neighborhood)**.** Let (X, d) be a metric space. Let $p \in X$ and $\epsilon > 0$. The ϵ-neighborhood $N_\epsilon(p) = \{q \in X : d(p, q) < \epsilon\}$.

Definition 10 (Limit Point)**.** Let (X, d) be a metric space, and let $E \subset X$. A point $p \in X$ is a *limit point* of E if every neighborhood of p contains a point $q \neq p$ such that $q \in E$.

Example 6. Let $X = \mathbb{R}$ and $E = [0, 1)$. The point $1 \in X$ is a limit point of E.

Definition 11. Let (X, d) be a metric space. A set $E \subset X$ is *closed* if it contains all its limit points.

Example 7. The set $[0, 1] \subset \mathbb{R}$ is closed, and the unit circle is closed in $\mathbb{R}^2$.

Definition 12 (Interior Point)**.** Let (X, d) be a metric space, and let $E \subset X$. A point $p \in E$ is called an *interior point* if there exists an $\epsilon > 0$ such that $N_\epsilon(p) \subset E$.

Example 8. Observe that $(0, 0)$ is an interior point of the unit circle. Take $\epsilon \in (0, 1)$ and $N_\epsilon((0, 0))$ is contained in the unit circle.

Definition 13 (Open Set)**.** Let (X, d) be a metric space. A set $E \subset X$ is an *open set* if every point in E is an interior point.

Example 9. Every neighborhood of a point is open.

Some Important Facts:

- Let (X, d) be a metric space, and let $E \subset X$. The set E is said to be open if and only if $X \setminus E$ is closed.
- The only sets that are both open and closed in X are $\emptyset$ and X.
- Let (X, d) be a metric space. A set $E \subset X$ is closed if for every convergent sequence of points $(p_n)_{n \in \mathbb{N}} \in E$, $p_n \to p \in E$. (We need this when applying Kakutani's Fixed Point Theorem).

With some notions from topology in mind, we introduce the notion of a complete metric space. We need the notion of a complete metric space when discussing certain fixed-point theorems. However, we restrict attention to $\mathbb{R}^n$, which is the canonical example of a complete metric space. We begin with the definition of a Cauchy sequence:

Definition 14 (Cauchy Sequence)**.** Let (X, d) be a metric space, and let $(p_n)_{n \in \mathbb{N}}$ be a sequence of points in X. The sequence is said to be *Cauchy* if for every $\epsilon > 0$, there exists a $N \in \mathbb{N}$ such that for every $m, n \geq N$, $d(p_m, p_n) < \epsilon$.

Example 10. The sequences $(\frac{1}{n})_{n \in \mathbb{N}}$ and $(\frac{1}{n^2})_{n \in \mathbb{N}}$ are both Cauchy in $\mathbb{R}$. Consider $(\frac{1}{n})_{n \in \mathbb{N}}$. We pick $N > \frac{2}{\epsilon}$. So for $m, n \geq N$, we have:

$$|\frac{1}{n} - \frac{1}{m}| \leq |\frac{1}{n}| + |\frac{1}{m}| < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon \tag{5}$$

The argument is analogous with $(\frac{1}{n^2})_{n \in \mathbb{N}}$.

Example 11. The sequence $((-1)^n)_{n \in \mathbb{N}}$ is bounded but not Cauchy.

Definition 15 (Complete Metric Space)**.** The metric space (X, d) is said to be *complete* if every Cauchy sequence of points in X converges in X.

Definition 16 (Compact Set)**.** A set $E \subset \mathbb{R}^k$ is said to be *compact* if it is closed and bounded. Equivocally, E is compact if every infinite subset of E has a limit point in E. (Note that this is the Heine-Borel characterization, and there is a more general definition of compact sets).

Example 12. The unit circle is compact, as is $[0, 1]$. The set $\mathbb{R}$ is closed but not bounded, so it is not compact.

Definition 17 (Continuous Function)**.** Let (X, d_X) and (Y, d_Y) be metric spaces. Let $E \subset X$ and $p \in E$. A function $f : E \to Y$ is said to be *continuous* at p if for every $\epsilon > 0$, there exists a $\delta > 0$ such that $d_Y(f(x), f(p)) < \epsilon$ for all points $x \in E$ such that $d_X(x, p) < \delta$.

Theorem 5.1 (Weierstrass Extreme Value Theorem)**.** *Let (X, d) be a metric space, and let $E \subset X$ be compact. A continuous function $f : E \to X$ acheives both a maximum and minimum value.*

Definition 18 (Convex Set)**.** Let $E \subset \mathbb{R}^n$. The set E is said to be convex if for any $x, y \in E$ and any $\lambda \in (0, 1)$, $\lambda x + (1 - \lambda)y \in E$.

Example 13. Let $p \in \mathbb{R}^n$ and fix $\epsilon > 0$. We have $N_\epsilon(p)$ is convex.

Example 14. The simplex $\Delta \subset \mathbb{R}^{n+1}$, which is defined below, is convex. The simplex is the set of all probability distributions over an $n + 1$ element set.

$$\Delta = \{(\lambda_0, \ldots, \lambda_n) \in \mathbb{R}^{n+1} : \sum_{i=0}^{n} \lambda_i = 1\} \tag{6}$$

More Important Facts:

- The simplex is compact.
- The Cartesian product of two compact sets is compact.
- The Cartesian product of two convex sets is convex.

5.2 Fixed Points

Definition 19 (Contraction)**.** Let (X, d) be a metric space. If $\phi : X \to X$ and if there is a constant $c \in [0, 1)$ such that for all $x, y \in X$:

$$d(\phi(x), \phi(y)) < c \cdot d(x, y) \tag{7}$$

Example 15. Let $f : [0, 1] \to [0, 1]$ by $f(x) = \frac{x}{2}$. This is a contraction, with $c = \frac{1}{2}$.

We first show that a contraction is continuous.

Lemma 5.1. *Let (X, d) be a metric space and let $\phi : X \to X$ be a contraction with constant $c \in [0, 1)$. We have that ϕ is continuous.*

Proof. Let $\epsilon > 0$ and let $\delta = \epsilon$. We have $d(x, y) < \delta \implies d(\phi(x), \phi(y)) < c \cdot d(x, y) < c \cdot \delta < \epsilon$. □

Theorem 5.2 (Contraction Mapping Principle)**.** *Let X be a complete metric space, and if $\phi : X \to X$ is a contraction, then there exists a unique $x \in X$ such that $f(x) = x$.*

Proof. Let $x_0 \in X$ be arbitrary, and let $c < 1$ such that $d(\phi(x), \phi(y)) \leq c \cdot d(x, y)$. Define $(x_n)_{n \in \mathbb{N}}$ recursively by setting $x_{n+1} = \phi(x_n)$. We now show this sequence is Cauchy. By induction, it follows that $d(x_{n+1}, x_n) \leq c^n d(x_1, x_0)$ for all $n \in \mathbb{N}$. For $n < m$, it follows that:

$$d(x_n, x_m) \leq \sum_{i=n+1}^{m} d(x_i, x_{i-1}) \tag{8}$$

This follows from the fact that $x_i = \phi(x_{i-1})$ and that ϕ is a contraction. We bound:

$$\sum_{i=n+1}^{m} d(x_i, x_{i-1}) \leq \left(\sum_{i=n}^{m-1} c^i \right) d(x_i, x_0) \tag{9}$$

$$\leq \frac{c^n}{1-c} d(x_1, x_0) \tag{10}$$

With (10) following from the sum of a convergent geometric series. Fix $\epsilon > 0$. We solve for N such that:

$$\frac{c^N}{1-c} d(x_1, x_0) < \epsilon \tag{11}$$

So we let N such that (which is possible since $c \in [0, 1)$):

$$c^N < \frac{\epsilon \cdot (1-c)}{d(x_1, x_0)} \tag{12}$$

So $(x_n)_{n \in \mathbb{N}}$ converges to some point $x \in X$ (since X is complete). Since ϕ is continuous, we have:

$$\phi(x) = \phi(\lim_{n \to \infty} x_n) = \lim_{n \to \infty} \phi(x_n) = \lim_{n \to \infty} x_{n+1} = x \tag{13}$$

So we have a fixed point. We conclude with uniqueness. Suppose x, y are two fixed points of ϕ. Then:

$$0 \leq d(x, y) = d(\phi(x), \phi(y)) \leq c \cdot d(x, y) \tag{14}$$

Since $c \in [0, 1)$, it follows that $x = y$. □

The Contraction Mapping Principle is useful in dynamical systems theory when we study the stability of our system. Brouwer's Fixed Point Theorem is more general, though, as it removes the contraction assumption. However, we also lose the uniqueness of the fixed point. Nash's original theorem in his dissertation used Brouwer's fixed point theorem.

Theorem 5.3 (Brouwer's Fixed Point Theorem). *Let D be a convex, compact subset of $\mathbb{R}^n$. Any continuous function $f : D \to D$ has a point $x \in D$ such that $f(x) = x$.*

There are numerous proofs of Brouwer's Fixed Point Theorem. One of my personal favorites uses Sperner's Lemma, which is a famous result in combinatorics and discrete geometry. Other proofs rely on analytical techniques, and there is even a proof from mathematical logic. I won't prove Brouwer's Fixed Point Theorem in its full generality, but instead offer a proof of the one-dimensional case to illustrate the concept. We first recall the Intermediate Value Theorem:

Theorem 5.4 (Intermediate Value Theorem). *Let f be a continuous real-valued function on the interval $[a, b]$. If $f(a) < f(b)$ and if c is a number such that $f(a) < c < b$, then there exists a point x such that $f(x) = c$. (Analogously, if $f(a) > f(b)$ and $f(a) > c > f(b)$, there exists a point $x \in [a, b]$ such that $f(x) = c$).*

We now prove the one-dimensional case of Brouwer's Fixed Point Theorem.

Theorem 5.5. *Let $a, b \in \mathbb{R}$ with $a < b$ and consider $[a, b] \subset \mathbb{R}$. Let $f : [a, b] \to [a, b]$ be continuous. There exists a fixed point $x \in [a, b]$ such that $f(x) = x$.*

Proof. Let $g : [a, b] \to \mathbb{R}$ be given by $g(x) = f(x) - x$. Since g is the sum of two continuous functions, g itself is continuous. Since $f(a) \in [a, b]$, $g(a) = f(a) - a \geq 0$. By similar argument, $g(b) = f(b) - b \leq 0$. As g is continuous, we have $g(a) \geq 0 \geq g(b)$. So by the intermediate value theorem, there exists a point $x \in [a, b]$ such that $g(x) = 0$. This point x satisfies $f(x) = x$, so f has a fixed point. □

We conclude by introducing Kakutani's Fixed Point Theorem, which we will use for an alternative proof of Nash's theorem.

Definition 20 (Correspondence). Let X, Y be sets. A *correspondence* from X to Y is a function $f : X \to 2^Y$.

Definition 21 (Fixed Point (Correspondence)). Let X be a set, and let $f : X \to 2^X$ be a correspondence. A fixed point is a point $x \in X$ such that $x \in f(x)$.

Theorem 5.6 (Kakutani's Fixed Point Theorem). *Let $A \subset \mathbb{R}^n$ be non-empty, convex, and compact. Let $f : A \to 2^A$ be a correspondence which has a closed graph, and the property that $f(x)$ is non-empty and convex for all $x \in S$. Then f has a fixed-point.*

5.3 Existence of Nash Equilibrium

Theorem 5.7 (Nash). *Every finite, normal form game has a Nash equilibrium in mixed strategies.*

Proof. Let $\Gamma = [N, (S_i)_{i \in N}, (u_i)_{i \in N}]$ be a normal form game. Let $\Sigma = \prod_{i=1}^{n} \Delta(S_i)$. Define the correspondence $f : \Sigma \to 2^\Sigma$ where for $\sigma \in \Sigma$, $f_i(\sigma_{-i})$ is the best-response correspondence for player i against the strategy profile σ_{-i}. We apply Kakutani's Fixed Point Theorem to show that f has a fixed point; that is, each player's strategy is a best response to the other players' strategies. Recall that this is the definition of a Nash equilibrium.

Claim 1: Σ is convex, compact, and non-empty.

Proof. For each i, $\Delta(S_i)$ is convex and compact as it is a simplex. Since Σ is the Cartesian product of simplex polytopes, Σ itself is convex and compact. As each player's strategy set S_i is non-empty, Σ is non-empty. □

Claim 2: For each $\sigma \in \Sigma$, $f(\sigma) \neq \emptyset$.

Proof. Fix $\sigma \in \Sigma$. Recall that for each $i \in N$, $f_i(\sigma) = \arg\max_{s \in \Sigma} u_i(s, \sigma_{-i})$. As each player's utility function is continuous and $\Delta(S_i)$ is compact, u_i acheives a maximum by the Weierstrass Extreme Value Theorem. So $f_i(\sigma) \neq \emptyset$. □

Claim 3: For each $\sigma \in \Sigma$, $f(\sigma)$ is convex.

Proof. Fix $\sigma \in \Sigma$ and $i \in N$. Each strategy $s \in f_i(\sigma)$ maximizes u_i with respect to σ_{-i}. So taking convex combinations of strategies in $f_i(\sigma)$ results in the same payoff. Thus, $f_i(\sigma)$ is convex, and so $f(\sigma)$ is convex. □

Claim 4: The correspondence f has a closed graph.

Proof. The graph relation on a correspondence h is the set $G = \{(x, y) : y \in h(x)\}$.We show the graph is closed by showing that any sequence of points from G converges within G. We prove this by means of contradiction. Let $(\sigma^n, \tau^n)_{n\in\mathbb{N}}$ be sequences of strategy profiles such that $\tau^n \in f(\sigma^n)$ and $(\sigma^n, \tau^n) \to (\sigma, \tau)$. Suppose to the contrary that $\tau \notin f(\sigma)$. So for some player i, we have a strategy $\mu \in \Delta(S_i)$ and $\epsilon > 0$ such that:

$$u_i(\mu, \sigma_{-i}) \geq u_i(\tau_i, \sigma_{-i}) + 3\epsilon \tag{15}$$

As $\sigma^n_{-i} \to \sigma_{-i}$ and each player's utility function is continuous, there exists an $N \in \mathbb{N}$ such that for every $n \geq N$:

$$u_i(\mu, \sigma^n_{-i}) \geq u_i(\mu, \sigma_{-i}) - \epsilon \tag{16}$$

$$u_i(\tau^n_i, \sigma^n_{-i}) - u_i(\tau_i, \sigma_{-i}) < \epsilon \tag{17}$$

From (15), (16) and (17) we obtain:

$$u_i(\mu, \sigma^n_{-i}) > u_i(\tau_i, \sigma_i) + 2\epsilon \geq u_i(\tau^n_i, \sigma^n_{-i}) + \epsilon \tag{18}$$

This implies μ dominates τ^n_i with respect to the strategy profile σ^n_{-i}, contradicting the assumption that τ^n_i is a best response to σ^n_{-i}. So f must have a closed graph. □

Claims 1-4 satisfy the hypotheses of Kakutani's Fixed Point Theorem. So f contains a fixed point, which is a Nash equilibrium. □