CS295 Topics in Algorithmic Game Theory September 29, 2021

Lecture 1-2: Notation, Nash Equilibrium, and Zero-sum Games

Lecturer: Ioannis Panageas *Scribes: Sam Showalter and Antony Alexos*

1 Introduction

In our first lecture, we were first introduced to the concept of game theory, defined below.

Definition 1. ***Game Theory*** *is defined as the analysis of experiments theoretically and empirically to* ***help us understand rational behavior in situations of conflict.***

Game theory applies to various real-world situations, including auctions, routing, resource allocation, and voting. The components of game theory are defined as follows.

1. **Conflict:** Everyone's actions affect other people
2. **Rational Behavior:** Every player wishes to individually maximize their benefit, or *utility*.
3. **Prediction:** We want to know what will happen and what, if any, equilibrium points will be reached in behavior. This is typically accomplished through the use of *solution concepts*. Moreover, if we can effectively discover this, we wish to develop structural mechanisms to facilitate optimal group outcomes.

Some common examples of games we can apply this type of analysis to include rock-paper-scissors and *The Prisoner's Dilemma*. In this thought experiment, two prisoners can cooperate to receive a lower punishment but inevitably do not due to their sole interest in maximizing their utility. More details of these examples can be found on the lecture slides, but in general, these games give examples of **Nash Equilibrium**, defined later on in these notes. Nash Equilibrium, named after the scientist John Nash, who discovered that an equilibrium point exists for strategies in any finite game. Colloquially, we can think of this as a strategy that each game player or *agent* keeps even *if they know the strategies of the other players*. In other words, at a Nash Equilibrium *all agents* simultaneously play the best responses to each other's strategies.

Theorem 1. *For any game with a finite number of players and action, a Nash Equilibrium exists (John Nash, 1951) [1].*

Building on the concept of Nash Equilibrium, we were also introduced to **The Price of Anarchy**. This concept mathematically defines the ratio between the worst-case Nash-Equilibrium that can be achieved in a game over the optimal outcome if each agent was completely controlled. Ideally, we want this ratio to be as close to 1 as possible, implying that the actual equilibrium achieved will be near-optimal. Such considerations apply to building roadways or other routing applications where individuals, driven by their self-interest, jointly slow each other down due to their lack of coordination. The remainder of the lecture explored a different situation that can be explored with Game Theory: **Auctions**, defined below.

Definition 2. Auctions describe a situation where:

- An auctioneer has *one item* to sell.
- There are n bidders interested in the item.
- Each bidder i has a valuation v_i for the item that they determine ahead of the auction itself

- At the time of the auction, each bidder i places a bid b_i to the auctioneer. The auctioneer collects all of these bids $\{b_1, ..., b_n\}$and privately decides which bidder i should receive the item, as well as the price p he or she should pay.
- If bidder i gets the item and pays price p, his or her utility is $v_i - p$. All non-winning bidders have utility zero.

What is fascinating about auctions defined in this manner is that the optimal choice for the auctioneer, who wishes to maximize the price p paid for the item, is to give the item to the highest bidder but make them pay the *second highest bid*. This is because if the auctioneer took the highest bid, the bidders have no incentive to bid anything and bid zero in a rational world. The second-highest bid encourages bidders to bid as high as they can.

2 Notation and Context

Before formalizing the notion of Nash Equilibrium, we need to define the games for which it applies. In general, Nash Equilibrium is explored for what are called **Normal Form Games**.

Definition 3. Normal Form Games are defined by:

- A set of n players $[n] = \{p_1, ..., p_n\}$
- Each player p_i is given a set of strategies (actions) $S_i = \{s_1, ..., s_{m_i}\}$. As shown, the number of strategies m_i can vary per player p_i.
- Set $\mathfrak{S} = \{S_1 \times S_2 \times ... \times S_n\}$ represents the set of all possible strategy profiles for *all players*. Let us consider $\mathcal{S}$ to be a sample set of chosen strategies drawn from $\mathfrak{S}$ according to some distribution (generally the mixed strategy profile of the game $\mathcal{X}$, defined below).
- Each player also is given a utility function $u_i = f(\mathcal{S}) \to \mathbb{R}$ which determines the payout given the strategies of all players.

With this very general definition, there are many games and situations that we can formulate as normal-form games. Now, at this point it is important to explain the difference between *pure* and *mixed* strategies. In the most general case, we can consider all player strategies from the mixed strategy perspective.

Definition 4. A **mixed strategy** represents a valid probability mass function (pmf) over the strategy set S_i of player p_i. Mixed strategies encode the randomness of player decision-making while also offering a clear method of defining the **expected utility** of the agent. Let $x_i \in \mathbb{R}^{m_i}$ represent the mixed strategy of player p_i and $\mathcal{X} = \{x_1, ..., x_n\}$ be the set of mixed strategies from all players. By representing the probability of a player p_i taking strategy S_{ij} (the j^{th} strategy of the i^{th} player, simplified to just s_i to denote the strategy taken by the i^{th} player) as $x_i(s_i)$,we can represent **expected utility** as:

$$u_i(\mathcal{X}) = \mathbb{E}_{\mathcal{S} \sim \mathcal{X}}[u_i(\mathcal{S})] = \sum_{\mathcal{S}_j \in \mathfrak{S}} u_i(\mathcal{S}_j) x(\mathcal{S}_j) = \sum_{(s1, ..., s_n) \in \mathfrak{S}} u_i(s_1, ..., s_n) \prod_{j=1}^{n} x_j(s_j) \tag{1}$$

NOTE: In future formulations, we refer to s_i as the strategy taken by player p_i instead of the i^{th} strategy option in arbitrary strategy profile $S = \{s_1, ..., s_m\}$. With this information, the concept of a *pure strategy* is simple: it is simply a mixed strategy in which player p_i has all of its probability density on a single strategy $s_j \in S_i$. That is, the player always takes the same strategy regardless of what other players do. Better understanding the possible strategies (including mixed strategies) a player can take is essential to grasp the proof of the existence of Nash Equilibrium in all finite games so that we will explore that next:

3 Strategy (action) Spaces and Simplexes

Discovering a Nash Equilibrium can be interpreted as an optimization over all the possible strategies that all players can take. An arbitrary player p_i can choose from m_i pure strategies. However, this is far smaller than their total number of possible strategies once mixed strategies are included. There is an infinite number of mixed strategies x_i player p_i can take. This infinite set is represented by the following m_i-dimensional simplex Δ_i:

$$\Delta_i = \{x_i : \sum_{s_i \in S_i} x_i(s_i) = 1 \text{ and } x_i \geq 0\} \tag{2}$$

Visualized in fig. 1 below, we can see what a 3-D simplex would look like for rock-paper-scissors. Each point on this simplex represents a valid mixed strategy x_i for player p_i. Moreover, we can extend this line of thinking to define a simplex of good strategies for all players and, in turn, the entire game. We represent this simplex as Δ:

$$\Delta = \{\Delta_1 \times \Delta_2 \times ... \times \Delta_n\} \tag{3}$$

Example (Rock-Paper-Scissors).

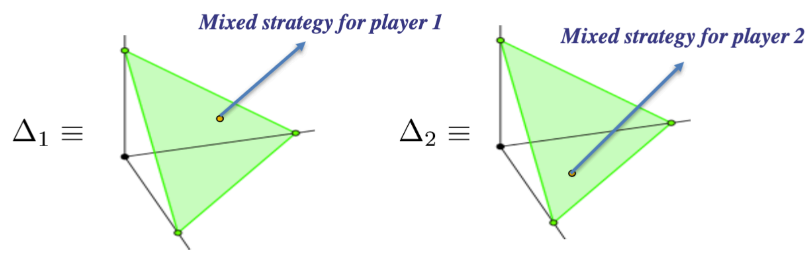

Figure 1: Visualization of the one-player strategy simplex set for Rock-Paper-Scissors

Another useful tool is to refer to the set of all valid mixed strategies for all players *except player* p_i as Δ_{-i}. We can sample from this set to attain valid strategies for all players except p_i by choosing $x_{-i} \in \Delta_{-i}$. Now that our notation is complete, we can formally define Nash Equilibrium and validate its existence with John Nash's proof.

4 Nash Equilibrium

There are multiple formulations of Nash Equilibrium, some of which offer stricter bounds than others.

Definition 5. Nash Equilibrium represents a set of strategies where no player would change their existing strategy regardless of the actions of other players. They each are responding optimally to the other players' strategies. Mathematically, we can say that:

A mixed game strategy $\mathcal{X} = \{x_1, ..., x_n\} = \{x_i; x_{-i}\} \in \Delta$ is a Nash Equilibrium iff for all players p_i and other mixed strategies $x_i' \in \Delta_i$:

$$u_i(x_i; x_{-i}) \geq u_i(x_i'; x_{-i}) \tag{4}$$

we can extend upon this formulation slightly to also define a looser **ϵ-approximate Nash Equilibrium** as:

$$u_i(x_i; x_{-i}) \geq u_i(x_i'; x_{-i}) - \epsilon \tag{5}$$

Before starting the proof of Nash Equilibrium, we need to understand **Brouwer's Fixed Point Theorem**, a concept from topology that states:

Theorem 2. *Let $\mathcal{D}$ be a convex, compact subset of $\mathbb{R}^{\mathcal{D}}$ and $f : \mathcal{D} \to \mathcal{D}$ a continuous function. There always exists $x \in \mathcal{D}$ such that*

$$f(x) = x$$

Visualized in fig. 2 below, this theorem states that for any function operating on a compact, convex space, there exists a point x that is invariant to the function f and therefore will remain unchanged. This theorem will prove helpful in conducting the proof of the existence of Nash Equilibrium.

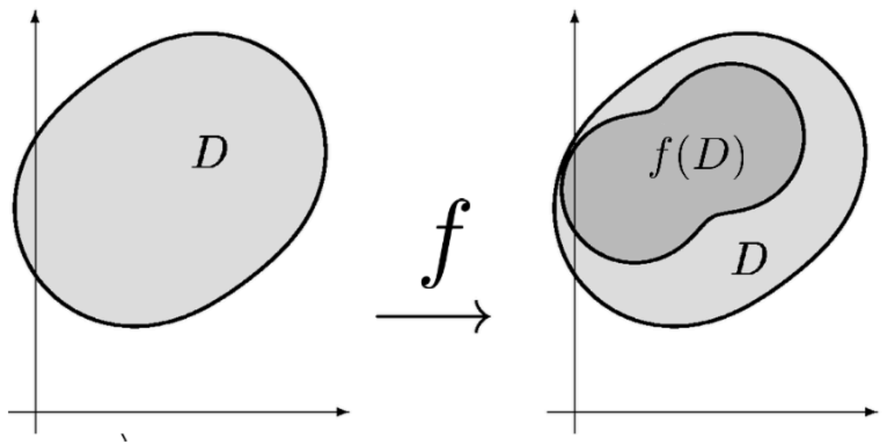

Figure 2: Visualization of Brouwer's fixed point theorem, where there always exists a point $x \in \mathcal{D}$ where $x = f(x)$

Proof:

Claim: Any game with a finite number of players and actions possesses a Nash Equilibrium.

◇

Consider a finite game with strategy space Δ. Define a function $f : \Delta \to \Delta$ that defines an operation on each player p_i and their strategy choice s_i as:

$$f_{is_i}(\mathcal{X}) = f_{is_i}(x_i; x_{-i}) = \frac{x_i\left(s_i\right) + \max\left\{u_i\left(s_i; x_{-i}\right) - u_i(\mathcal{X}), 0\right\}}{1 + \sum_{s' \in S_i} \max\left\{u_i\left(s'; x_{-i}\right) - u_i(\mathcal{X}), 0\right\}} \tag{6}$$

Intuitively, the numerator is the sum of the probability of player p_i choosing *pure strategy* s_i and the maximum between 0 and the difference in utility between taking pure strategy s_i (with all other players taking a combined mixed strategy of x_{-i}) and taking the mixed strategy $\mathcal{X} = \{x_i; x_{-i}\}$. Regarding the last term, we can interpret this maximum as *how much better off* a player would be taking pure strategy s_i instead of mixed strategy x_i. The denominator normalizes this value by summing the possible probability density of 1 with the total potential utility gain of considering all pure strategies $s' \in S_i$ relative to mixed strategy x_i.

Remark 1. There are three observations we can make by considering the following game formulation with included function f.

- *Δ is both a convex and compact space*
- *Based on our definition, f is a continuous mapping from $\Delta \to \Delta$*
- *Since all of these terms are normalized by the same term, $\sum_{s' \in S_i} f_{is'}(\mathcal{X}) = 1$*

Therefore, we can conclude that *a fixed point always exists for some $\mathcal{X}$ passed through f*. We can now complete the proof by considering this fact in greater detail.

Claim: If $\mathcal{X}^*$ is a fixed point of f, then $\mathcal{X}^*$ is a Nash Equilibrium. $\diamond$

Let $\mathcal{X}^*$ represent the fixed point of f. This implies that:

$$f(\mathcal{X}^*) = \mathcal{X}^* \tag{7}$$

Moreover, this implies that $f_{is_i}(\mathcal{X}^*) = x_i(s) \qquad \forall s \in S_i$ for each player p_i. Further decomposing this equivalence based on the definition of f, we see this implies:

$$x_i^*(s) \sum_{s' \in S_i} \max\left\{u_i\left(s'; x_{-i}^*\right) - u_i\left(x^*\right), 0\right\} = \max\left\{u_i\left(s; x_{-i}^*\right) - u_i\left(x^*\right), 0\right\} \tag{8}$$

To satisfy this equation, we must consider two cases which would make it true:

- $\forall s \in S_i, \quad x_i(s) = 0 \to u_i(s; x_{-i}^*) \leq u_i(\mathcal{X}^*)$
- $\forall s \in S_i, \quad x_i(s) > 0 \to u_i(s; x_{-i}^*) \geq u_i(\mathcal{X}^*) \to u_i(s; x_{-i}^*) = u_i(\mathcal{X}^*)$

The first bullet states that whenever we place zero probability on a strategy s, this implies the utility from taking action must be *at most the expectation of utility* $u_i(\mathcal{X}^*)$. If the utility was better, we would have a non-zero probability of taking that strategy.

The implication of the second bullet is slightly more subtle. Since the expected utility is $u_i(\mathcal{X}^*)$ simply the weighted sum of individual pure strategies and their probabilities of occurring, having a positive probability $x_i(s) > 0$ for all actions s implies that the utility of that strategy *must be equal* to the overall expectation. If not, then we reach a contradiction where the expectation of the utility $u_i(\mathcal{X}^*) = \sum_{s' \in S_i} u_i(s'; x_{-i}) x_i(s')$ is actually lower than each of its component utilities $u_i(s; x_{-i}^*) \quad \forall s \in S_i$.

Recapping for this proof, we have shown that if we play something with probability 0, the expectation that we get is less than the expected utility. If we play something with positive probability, then we get our own expected utility. Both of these hold for all players i without needing to make any assumptions. Continuing our proof, we claim that these two cases are sufficient to prove Nash Equilibrium. To show that something is a Nash Equilibrium, we arbitrarily choose a player and assume that this player deviates from his strategy. We need to show that the utility of this player before the deviation is at least his utility after the deviation. In that way, he does not have an incentive to deviate. By showing that for every player, we have a Nash Equilibrium. Defining that formally, for any $\widetilde{x_i}$ we have $u_i\left(x_i^*; x_{-i}^*\right) \geq u_i\left(\tilde{x}_i; x_{-i}^*\right)$. We also note that all players keep the same strategy apart from player i who deviates from his strategy and gains nothing from deviating.

We multiply the first case with $\widetilde{x_i}$ and we have $\tilde{x}_i(s) u_i\left(s, x_{-i}^*\right) \leq \tilde{x}_i(s) u_i\left(x^*\right)$. By taking the summation of that we finally have:

$$u_i\left(\tilde{x}_i, x_{-i}^*\right) = \sum_{s'} \tilde{x}_i\left(s'\right) u_i\left(s'; x_{-i}^*\right) \leq \sum_{s'} \tilde{x}_i\left(s'\right) u_i\left(x^*\right) = u_i\left(x^*\right) \tag{9}$$

We observe that the expected utility is at least as much as the left-hand side of eq. (9), and because this is true in general, as we see from above, we sum over all possible s strategies. So basically, the utility of player i from the law of total probability is the utility of s' given x^* while the rest is fixed, times the probability of playing $\tilde{x}\left(s'\right)$. This is at most the right-hand side of eq. (9) since it is true for every s'. Here $u_i\left(x^*\right)$ does not depend on s' so we can put it outside of the sum and because $\tilde{x}\left(s'\right)$ are probabilities they sum to 1.

We have shown that the utility either increases or stays the same if we deviate, but it does not increase. We have shown that x^*, the fixed point of this function that we have defined in eq. (6), is a Nash Equilibrium. So finding Nash Equilibrium is easy cause we compute eq. (6) and find a fixed point for this function. Nevertheless, this is a computationally hard problem, so finding fixed points of functions is very hard, and finding Nash Equilibrium is of the same difficulty. Although it might be computationally hard in general, finding Nash Equilibrium is tractable in some classes of games, like Zero-Sum Games.

□

5 Zero-Sum Games

For Zero-Sum Classes of Games we do not have to solve a fixed point problem, and we can reduce to find Nash Equilibrium to something that is computationally easy. Zero-Sum Games are defined below:

Definition 6. Zero-Sum Games describe a situation where:

- Zero-sum games are a specific example of constant sum games where the sum of each outcome is always zero.
- They are tractable classes of games.
- They consist of 2 players, the Row player and the Column player.
- They have a payoff matrix R size of nxm.
- They have n and m strategies available.

We can see a toy example of Zero-Sum Games in fig. 3. This matrix shows the money that the column player pays for playing j when the row player plays i. It also shows the money that the row player gets for playing i when the column player plays j.

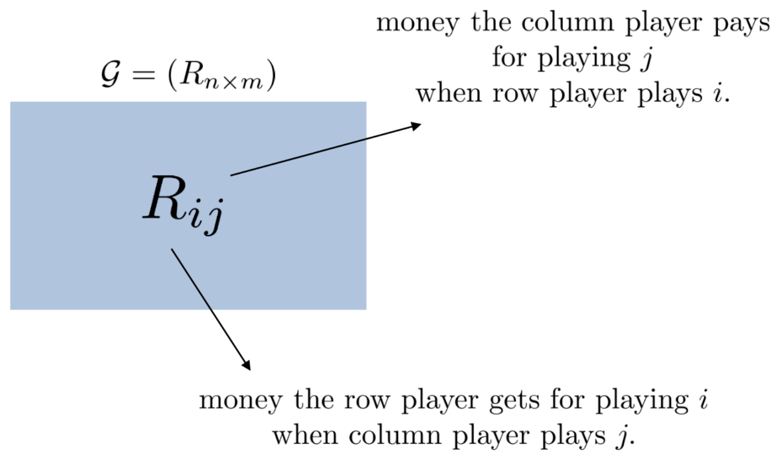

Figure 3: Visualization of a toy Payoff Matrix G, with the money that the Row and Column players pay and get in every action.

Let us investigate another Zero-Sum Games example with a Payoff Matrix. Hypothetically the row player chooses $x \in \Delta_n$ and the column player choose $y \in \Delta_m$. Every cell of the Payoff matrix is denoted as $R_{i,j}$. Row player gets $x^\top Ry$ and the column player pays $x^\top Ry$.

	Tax-Cuts	Society
Economy	3,-3	-1,1
Education	-2,2	1,-1

Another example of a Zero-Sum Game is the following in which we have two candidates aiming for the presidency, which is depicted in section 5. Each candidate focuses on two strategies. The first one focuses on Economy and Education, and the second one focuses on Tax-Cuts and Society. All the combinations are shown in section 5, and both candidates want to win the election, so they need to maximize their utility. Let's assume that the row player plays (x_{11}, x_{12}). The column player has two choices, to play either Tax-Cuts or Society. In the first scenario the expectation is $u_2\left(x_1,' \text{Tax} - \text{cuts}'\right) = -3x_{11} + 2x_{12}$, and in the second scenario the expectation is $u_2\left(x_1,' \text{ Society }'\right) = x_{11} - x_{12}$. The best option for the column player is to play the maximum of the two options, which formally is $\max\{3x_{11} - 2x_{12}, x_{11} - x_{12}\}$. The column player needs to put all probability to the one that gives the highest utility to maximize the gain. Since the column player will play the best option, the row paper should play the best option to minimize the gain of the column player. More formally this can be written as $\min\{-3x_{11} + 2x_{12}, 1x_{11} + x_{12}\}$. So if the row player wants to maximize her utility, she needs to play:

$$(x_{11}^*, x_{12}^*) = \arg\max_{x_{11}, x_{12}} \min\left\{3x_{11} - 2x_{12}, -x_{11} + x_{12}\right\} \tag{10}$$

Now let's assume that each of the two numbers in eq. (10) will be at least z. So we can write a linear program for this max min problem:

$$\begin{aligned} &\max z \\ &\text{s.t } 3x_{11} - 2x_{12} \geq z \\ &-x_{11} + x_{12} \geq z \\ &x_{11} + x_{12} = 1 \\ &x_{11}, x_{12} \geq 0 \end{aligned} \tag{11}$$

If we solve this problem we get the solution $x_1 = \left(\frac{3}{7}, \frac{4}{7}\right), z = \frac{1}{7}$, so $z = \frac{1}{7}$. Which means that no matter what the column player does, the row player will get at least $\frac{1}{7}$.

This is the scenario in which the row player plays first. Now let's do see the same problem when the column player plays first. Row's player best response is $\max\left\{3x_{21} - x_{22}, -2x_{21} + x_{22}\right\}$ so the column player gets $\min\left\{-3x_{21} + x_{22}, 2x_{21} - x_{22}\right\}$. The column player maximizes her utility by solving the following max min problem: $(x_{21}^*, x_{22}^*) = \arg\max_{x_{21}, x_{22}} \min\left\{-3x_{21} + x_{22}, 2x_{21} - x_{22}\right\}$. Then we can write the following linear program:

$$\begin{aligned} &\max w \\ &\text{s.t } -3x_{11} + 2x_{12} \geq z \\ &x_{11} - x_{12} \geq z \\ &x_{11} + x_{12} = 1 \\ &x_{11}, x_{12} \geq 0 \end{aligned} \tag{12}$$

If we solve this problem we get the solution $x_1 = \left(\frac{2}{7}, \frac{5}{7}\right), w = \frac{-1}{7}$, so $w = \frac{1}{7}$. Which means that no matter what the row player does, the column player will get at least $\frac{-1}{7}$.

Since the problem of the two solutions $\left(\frac{3}{7}, \frac{4}{7}\right), \left(\frac{2}{7}, \frac{5}{7}\right)$ is a zero sum game, it must also be a Nash Equilibrium.

References

[1] J. Nash. Noncooperative Games. *Annals of Mathematics*, 54:289–295, 1951.