Harvard CS 1360, Spring 2025
Instructor: Ariel Procaccia
Scribe: Taig Singh

Nash Equilibrium

—

Lecture 1

Definition 1 (Normal-Form Game)**.** A game in normal form consists of a set of player $N = \{1, \ldots, n\}$, a set of strategies S, and a utility function $u_i : S^n \to \mathbb{R}$ for each player $i \in N$, where $u_i(s_1, \ldots, s_n)$ gives the utility of player i when each player $j \in N$ plays the strategy $s_j \in S$.

Example 1 (The Ice Cream Wars)**.** Ed and Ted are selling identical ice cream bars on a beach, which we model as the interval $[0, 1]$. Since the ice cream bars are identical, customers always go to the vendor closest to them. Initially, Ed sets up his cart at the $\frac{1}{4}$ mark on the beach and Ted sets up his cart at the $\frac{3}{4}$ mark, so both vendors got $\frac{1}{2}$ of the customers. After some time, Ted employs a useful deviation from his original strategy and moves to the $\frac{1}{2}$ mark on the beach, and Ted now gets $\frac{5}{8}$ of the customers while Ed only gets $\frac{3}{8}$ of the customers. This game can be modeled as a normal-form game. We have $N = 1, 2$ and $S = [0, 1]$, where a strategy $s \in S$ represents setting up a cart at s on the beach. Further, we have that

$$u_i(s_i, s_j) = \begin{cases} \frac{s_i + s_j}{2} & \text{if } s_i < s_j, \\ 1 - \frac{s_i + s_j}{2} & \text{if } s_i > s_j, \\ \frac{1}{2} & \text{if } s_i = s_j. \end{cases}$$

This follows because the utility of a vendor is described by the fraction of customers that will go to that vendor, and $\frac{s_i + s_j}{2}$ is the point on the beach where customers are indifferent between walking to either vendor.

Example 2 (The Prisoner's Dilemma)**.** Two men are charged with a crime. They are told that

- if one rats the other out while the other does not, the rat will be freed and the other will be jailed for nine years
- if both rat each other out, they will both be jailed for six years
- if neither rats the other out, they will both be jailed for one year

This can be modeled as a game in normal-form. $N = 1, 2$ and $S = C, D$ where C represents cooperating (not ratting the other out) and D represents defecting (ratting the other out). The utility functions are symmetric and can be defined using the following matrix:

	C	D
C	$(-1, -1)$	$(-9, 0)$
D	$(0, -9)$	$(-6, -6)$

where the first entry in each cell is the utility of prisoner 1 and the second entry is the utility of prisoner 2. In this game, defection is a dominant strategy: no matter what the other prisoner does, you are always better off defecting. However, instead of both players defecting, they can do much better by both cooperating. This is related to the tragedy of the commons.

Definition 2 (Tragedy of the Commons)**.** A social dilemma where individuals have an incentive to overconsume a common resource and act in their own self-interest at the expense of society. Scottish economists first observed this in the 19th century: farmers would not reign in their cows and overgrazing would occur such that grass would never grow back, but farmers would never have individual incentives to reign in their cows.

The Tragedy of the Commons exists all around us. Another example of this is tech companies hiring AI professors and pulling them away from academia. If tech companies in aggregate over-hire AI professors, no one will be there to teach the next generation, however, tech companies are individually incentivized to continue hiring these professors.

Example 3 (The Professor's Dilemma)**.** This is a game played between a professor and the class. The professor has two strategies when preparing a lecture: to make effort or to slack off. The class has two strategies in the lecture: to listen or to sleep. The payoffs of this game are shown below, where the first entry of each cell is the utility of the professor and the second entry is the utility of the student.

	Listen	Sleep
Make effort	$(10^6, 10^6)$	$(-10, 0)$
Slack off	$(0, -10)$	$(0, 0)$

In this game, there are no dominant strategies: the best strategy for a player depends on what the other player is doing.

Definition 3 (Nash Equilibrium)**.** A Nash Equilibrium is a vector of strategies $\mathbf{s} = (s_1, \ldots, s_n) \in S^n$ such that for all $i \in N$ and $s_i' \in S$,

$$u_i(\mathbf{s}) \geq u_i(s_i', \mathbf{s}_{-i})$$

This concept was introduced by John Forbes Nash (1928 - 2015), a mathematician and Nobel laureate in economics.

In this definition $u_i(s_i', \mathbf{s}_{-i})$ is the utility of agent i in the event that agent i deviates from strategy s_i to s_i'. In a Nash equilibrium, no player has a useful unilateral deviation. In other words, each player's strategy is a best response to the strategies of others - if they knew what everyone else was doing they would not change their strategy.

Going back to the Professor's Dilemma, the strategy profile (Make effort, Listen) is a Nash equilibrium. This follows because if the class were to deviate to Sleep, their payoff would decrease from 10^6 to 0, and if the professor were to deviate to Slack off, their payoff would also decrease from 10^6 to 0. Similarly, the strategy profile (Slack off, Sleep) is also a Nash equilibrium.

Going back to the Ice Cream Wars, assume that Ed starts at $\frac{1}{4}$ and Ted goes back to his starting spot of $\frac{3}{4}$. Then, Ed has a useful deviation to move just south of Ted, say $\frac{3}{4} - \epsilon$ so that Ed captures $\frac{3}{4}$ of the customers instead of his original $\frac{1}{2}$. Then, Ted will have a useful deviation to move just south of Ed. This process of moving just south of each other will continue until Ted and Ed both sit at the $\frac{1}{2}$ mark when no one has a useful deviation and both vendors collect exactly $\frac{1}{2}$ of the customers. Note that if either vendor moves north or south from $\frac{1}{2}$ here, they will collect less than $\frac{1}{2}$ of the customers, and thus the strategy profile $(\frac{1}{2}, \frac{1}{2})$ is a Nash equilibrium.

Nash equilibria are often helpful in predicting long-run outcomes for games but do not always represent reality. Consider the following game, and what you would do in real life for real money:

Example 4. Two players play a game, and the strategy set is $\{2, \ldots, 100\}$. If both players choose the same number, that is what they get. If one chooses s and the other chooses t and $s < t$, then the former player gets $s + 2$ while the latter gets $s - 2$.

Note that the only Nash Equilibrium in this game is $(2, 2)$. This is because if one player chooses any number $n > 2$, the other player's best response will be to choose $n - 1$. This progression of undercutting the other player continues until both players settle at $(2, 2)$ in which there is no useful deviation. However, in any strategy profile (a, b) where $a, b \geq 4$, both players are weakly better off than in this Nash equilibrium.

Example 5 (Rock-Paper-Scissors)**.** The game of Rock-Paper-Scissors is played between two players, where both players simultaneously choose from the strategy set of {Rock, Paper, Scissors}, and Rock beats Scissors, Scissors beats Paper, and Paper beats Rock. The payoff matrix of this game looks like the following:

	Rock	Paper	Scissors
Rock	$(0, 0)$	$(-1, 1)$	$(1, -1)$
Paper	$(1, -1)$	$(0, 0)$	$(-1, 1)$
Scissors	$(-1, 1)$	$(1, -1)$	$(0, 0)$

In this game, there are no pure strategy Nash equilibria because if you know what the other player is doing, you can always deviate to the strategy that beats them. However, there is a mixed strategy Nash equilibrium...

Definition 4 (Mixed Strategies)**.** A mixed strategy is a probability distribution over (pure) strategies. The mixed strategy of player $i \in N$ is $x_i : S \to [0,1]$ where

$$x_i(s_i) = \Pr[i \text{ plays } s_i]$$

Further, the utility of player $i \in N$ is just the expected utility according to the vectors of pure strategies and their associated probabilities:

$$u_i(x_1, \dots, x_n) = \sum_{(s_1,\dots,s_n)\in S^n} u_i(s_1, \dots, s_n) \cdot \prod_{j=1}^{n} x_j(s_j)$$

Continuing with the example of Rock-Paper-Scissors, if player 1 plays the mixed strategy $(\frac{1}{2}, \frac{1}{2}, 0)$ and player 2 plays the mixed strategy $(0, \frac{1}{2}, \frac{1}{2})$, we can calculate u_1 as the following:

$$\underbrace{\frac{1}{2}\cdot 0 \cdot 0 + \frac{1}{2}\cdot\frac{1}{2}\cdot(-1) + \frac{1}{2}\cdot\frac{1}{2}\cdot 1}_{\text{player 1 plays rock}} + \underbrace{\frac{1}{2}\cdot 0 \cdot 1 + \frac{1}{2}\cdot\frac{1}{2}\cdot 0 + \frac{1}{2}\cdot\frac{1}{2}\cdot(-1)}_{\text{player 1 plays paper}} + \underbrace{0\cdot 0\cdot(-1) + 0\cdot\frac{1}{2}\cdot 1 + 0\cdot\frac{1}{2}\cdot 0}_{\text{player 1 plays scissors}} = -\frac{1}{4}$$

Further, if we do similar calculations for the scenario in which both players play $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, we will get that $u_1 = 0$. This also follows from the fact that the game is a symmetric zero-sum game, so if the strategies are symmetric as well, both players will have a utility of 0. Note further that this is a Nash equilibrium for the game. We justify this below. Consider the situation in which player 1 plays (r, p, s) and player 2 plays $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Then we have that

$$\begin{aligned} u_1 &= r(\frac{1}{3}\cdot 0 + \frac{1}{3}\cdot(-1) + \frac{1}{3}\cdot 1) + p(\frac{1}{3}\cdot 1 + \frac{1}{3}\cdot 0 + \frac{1}{3}\cdot(-1)) + s(\frac{1}{3}\cdot(-1) + \frac{1}{3}\cdot 1 + \frac{1}{3}\cdot 0) \\ &= 0 \end{aligned}$$

and thus player 1 has no useful deviations, as their utility will always be 0. By symmetry, the same conclusion follows for player 2. Thus, both players playing $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ is a Nash equilibrium.

We end with an existence theorem for mixed strategy Nash equilibria:

Theorem 1 (Nash, 1950)**.** *In any (finite) game there exists at least one (possibly mixed) Nash equilibrium*