Harvard CS 1360, Spring 2026
Instructor: Ariel Procaccia
Scribe: Taig Singh, Jakob de Raaij

Nash Equilibrium

—

Lecture 1

Definition 1 (Normal-Form Game)**.** A *game in normal form* consists of a set of players $N = \{1, \ldots, n\}$, a set of pure 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 $s_j \in S$.[1]

When there are only $n = 2$ players and the set of pure strategies $S = \{s^1, ..., s^m\}$ is finite, we'll often represent that game with a *payoff matrix*.

	s^1	s^2	$\ldots$	s^m
s^1	$u_1(s^1, s^1), u_2(s^1, s^1)$	$u_1(s^1, s^2), u_2(s^1, s^2)$	$\ldots$	$u_1(s^1, s^m), u_2(s^1, s^m)$
s^2	$u_1(s^2, s^1), u_2(s^2, s^1)$	$u_1(s^2, s^2), u_2(s^2, s^2)$	$\ldots$	$u_1(s^2, s^m), u_2(s^2, s^m)$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
s^m	$u_1(s^m, s^1), u_2(s^m, s^1)$	$u_1(s^m, s^2), u_2(s^m, s^2)$	$\ldots$	$u_1(s^m, s^m), u_2(s^m, s^m)$

Player 1 is the *row player* and player 2 is the *column player*. The cell in the row corresponding to s^r and column corresponding to s^c contains the payoffs (utilities) of the row player and the column player, in that order, if the row player chooses strategy s^r and the column player chooses strategy s^c.

Example 1 (The Ice Cream Wars)**.** Ed and Ted are selling identical ice cream bars on a beach, which we model as the interval $[0, 1]$, with customers uniformly distributed along that interval. The ice cream bars are identical, so customers always go to the vendor closest to them. Initially, Ed sets up his cart at the $1/4$ mark on the beach and Ted sets up his cart at the $3/4$ mark, so both vendors get $1/2$ of the customers:

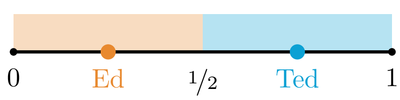

After some time, Ted realizes that he can employ a useful deviation from his original strategy and moves to the $1/2$ mark on the beach. Ted now gets $5/8$ of the customers while Ed only gets $3/8$ of the customers:

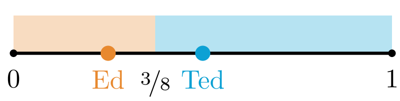

This game can be modeled as a normal-form game. We have $N = 1, 2$ and $S = [0, 1]$, where a strategy $s \in S$ represents setting up a cart at s on the beach. Further, we have that

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

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. If the two vendors are at the same position, we assume that customers choose randomly between them, giving both utility $1/2$.

[1] In our definition, all players have the same set of strategies S available to them. In a scenario where the players have different strategies, we can add all individual player's strategies to S while making the strategies in S that a given player cannot play be undesirable for this player, for example by giving this player utility $-\infty$ for such a strategy.

Example 2 (The Sucker's Dilemma)**.** There are two players, with each choosing between getting \$100 (from the "bank") and giving \$300 (from the "bank") to the other player. Here, $N = \{1, 2\}$ and $S = \{C, D\}$ where C represents "cooperating" (giving money) and D represents "defecting" (taking money). This can be written as a game in normal-form with the following payoff matrix:

	C	D
C	$300, 300$	$0, 400$
D	$400, 0$	$100, 100$

Remember, the first entry in each cell is the utility of the row player and the second entry is the utility of the column player.

In this game, defection is a *dominant strategy*: No matter what the other player does, you are always better off defecting. However, instead of both players defecting, they can do much better by both cooperating (and thus both not playing the dominant strategy).

Definition 2 (Tragedy of the Commons)**.** The Tragedy of the Commons is a social dilemma where individuals have an incentive to over-consume a common resource or act in their own self-interest at the expense of society. Scottish economists first observed this in the 19th century, when farmers would not reign in their cows, which lead to overgrazing and subsequently the grass not growing back in subsequent years. The farmers would have collectively been better off if they all reigned in their cows, guaranteeing a lasting grass supply; however, any farmer individually was always better of not reigning in their cows, since reigning the cows in limits their access to grass. For each farmer, it was individually optimal to deviate from the social optimum, leading to all farmers being worse off eventually.

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.[2]

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 an 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.

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

One can quickly verify that in this game, there are no dominant strategies. The best strategy for a player depends on what the other player is doing: If the professor knows that the students are going to listen, they are better off making an effort, but if they know the students are going to sleep, it is in their interest to slack off. Conversely, the students prefer to listen if the professor makes an effort and prefer to sleep if the professor slacks off.

Even though no dominant strategies exist in the Professor's Dilemma, there does exist a situation where no player would want to unilaterally change their strategy; for example, when the professor makes an effort and the students listen. This concept was formalized by John Forbes Nash (1928–2015), a mathematician and Nobel laureate in economics:

Definition 3 (Pure Nash Equilibrium)**.** A *(pure) Nash equilibrium* is a pure strategy profile $\mathbf{s} = (s_1, \ldots, s_n) \in S^n$ such that for all $i \in N$ and $s'_i \in S$,

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

$\mathbf{s}_{-i}$ denotes the vector $\mathbf{s}$ without s_i and $u_i(s'_i, \mathbf{s}_{-i})$ is as shorthand notation for the utility of agent i in the event that agent i deviates from strategy s_i to s'_i.

[2] You can read more about this here.

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 the other players — even if they knew which strategy everyone else was selecting, they would not change their strategy since the would not benefit from it.

Nash equilibria in the Professor's Dilemma The strategy profile (Make effort, Listen) is a Nash equilibrium: If the class were to deviate from Listen to Sleep, their payoff would decrease from 10^6 to 0. If the professor was to deviate to Slack off, their payoff would also decrease from 10^6 to 0. Similarly, one can check that the strategy profile (Slack off, Sleep) is also a Nash equilibrium.

Nash equilibrium in the Ice Cream Wars Assume that Ed starts at $1/4$ and Ted went back to his starting spot of $3/4$. Then, Ed has a useful deviation to move just left of Ted, say $3/4 - \epsilon$, so that Ed captures almost $3/4$ of the customers instead of his original $1/2$.

Then, Ted will have a useful deviation to move just left of Ed. This process of moving a tiny bit to the left of each other will continue until Ted and Ed both sit at the $1/2$ mark when no one has a useful deviation and both vendors collect exactly $1/2$ of the customers.

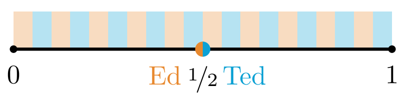

Note that if either vendor moves left or right from $1/2$ here, they will collect less than $1/2$ of the customers. Thus, the strategy profile $(1/2, 1/2)$ is a Nash equilibrium.

We now turn to a game where no pure Nash equilibrium exists:

Example 4 (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}. Rock beats Scissors, Scissors beats Paper, and Paper beats Rock; if both players choose the same strategy the game is a tie. The payoff matrix of this game is

	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 Nash equilibria in which a player deterministically picks one of the pure strategies (rock, paper, or scissors): No matter which pure strategy the player chooses, the other player can always deviate to the pure strategy that beats them, thus being better off.

Luckily, this is not the end of the story about Nash equilibria. As you may know from your own experience playing Rock-Paper-Scissors, a decent strategy is to pick an option at random. We'll now formalize this.

Definition 4 (Mixed Strategies)**.** A *mixed* strategy is a probability distribution over the pure strategies S. The mixed strategy of player $i \in N$ is $x_i \in \Delta(S)$[3] where

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

[3] $\Delta(S)$ denotes the set of all probability distributions over a set S. Thus, $x_i \in \Delta(S)$ is any function $x_i : S \to [0, 1]$ for which $\sum_{s \in S} x_i(s) = 1$.

Further, the utility of player $i \in N$ for a mixed strategy profile $(x_1, ..., x_n)$ is the expected utility of the player for the vector of pure strategies coming from the distribution given by the mixed strategies of all players,

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

We'll often write x_i as $(x_i(s))_{s\in S}$, a vector containing $\Pr[i \text{ plays } s]$ for all $s \in S$, to simplify notation.

Definition 5 (Nash Equilibrium)**.** A (mixed) Nash equilibrium is a mixed strategy profile $\mathbf{x} = (x_1, \ldots, x_n) \in \Delta(S)^n$ such that for all $i \in N$ and mixed strategies $x'_i \in \Delta(S)$,

$$u_i(\mathbf{x}) \geq u_i(x'_i, \mathbf{x}_{-i}).$$

$u_i(x'_i, \mathbf{x}_{-i})$ is as shorthand notation for the utility of agent i in the event that agent i deviates from strategy x_i to x'_i.

Mixed strategies in Rock-Paper-Scissors Let's first consider the scenario where player 1 plays the mixed strategy $x_1 = (\frac{1}{2}, \frac{1}{2}, 0)$ and player 2 plays the mixed strategy $x_2 = (0, \frac{1}{2}, \frac{1}{2})$. We can calculate u_1 as

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

We can do similar calculations for the scenario in which both players play $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ to get that $u_1 = 0$.[4] In fact, this is a mixed Nash equilibrium for the game! To verify this, consider player 1 deviating to an arbitrary other mixed strategy $x'_1 = (r, p, s)$, while player 2 still plays $x_2 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. We have that

$$u_1(x'_1, x_2) = r(\frac{1}{3}\cdot 0+\frac{1}{3}\cdot(-1)+\frac{1}{3}\cdot 1)+p(\frac{1}{3}\cdot 1+\frac{1}{3}\cdot 0+\frac{1}{3}\cdot(-1))+s(\frac{1}{3}\cdot(-1)+\frac{1}{3}\cdot 1+\frac{1}{3}\cdot 0)=0,$$

so no more than $u_1(x_1, x_2)$. 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 mixed Nash equilibrium.

Importantly, mixed Nash equilibria are always guaranteed to exist:

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

We end by pointing out a notable caveat of Nash equilibria. While they are often helpful in predicting long-run outcomes of games, they do not always represent reality. Consider the following game and what you would do in real life for real money.

Example 5 (Undercutting game)**.** Two players play a game, and the pure strategy set is $S = \{2, \ldots, 100\}$. If both players choose the same number s, this is their payoff. If one player chooses s and the other chooses t where $s < t$, the former player gets $s + 2$ while the latter gets $s - 2$. Thus, the payoff matrix is

	2	3	4	…	99	100
2	2, 2	4, 0	4, 0	…	4, 0	4, 0
3	0, 4	3, 3	5, 1	…	5, 1	5, 1
4	0, 4	1, 5	4, 4	…	6, 2	6, 2
⋮	⋮	⋮	⋮	⋱	⋮	⋮
99	0, 4	1, 5	2, 6	…	99, 99	101, 97
100	0, 4	1, 5	2, 6	…	97, 101	100, 100

[4] Alternatively, this also follows from the fact that the game is completely symmetric in the two players and the utilities for any choice of pure strategies always add up to 0 (a game with this second property is often called a *zero-sum* game). Then, if the strategies are symmetric as well, both players will have a utility of 0.

The only Nash equilibrium in this game is $(2, 2)$: If one player chooses any number $s > 2$, the other player's best response will be to choose $s - 1$. 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 (s, t) where $s, t \geq 4$, both players are weakly better off than in this Nash equilibrium! Notably, this is yet another example of the Tragedy of the Commons.