What is Nash Equilibrium

Mario Gómez Flores

July 25, 2019

Problems in biological, social or economic environments deal with several agents taking decisions that affect one another. The main issue is that each agent has a particular way of seeing the world, and wants to achieve a specific goal. In doing so, it might make it easier or harder for the other agents to achieve their own objectives. Game theory developed as a formal analysis of these situations and Nash equilibrium is its central concept. It soon became the gold standard against which all other proposed solutions are compared, and it provided mathematical footing to social sciences and biology – which would otherwise have to rely on intuition.

We will see the definition of Nash equilibrium and a couple of examples in simple two player games. I will prove the existence of equilibria using fixed point theorems. Lastly, I will (partly) present the computational complexity of finding Nash equilibria and provide references for the interested reader.

1 Normal form games

Definition 1 (Normal form game)**.**
A finite n-person normal-form game is a triple (N, A, u) where:

- $N = \{x_1, \ldots, x_n\}$ is a set of n players.
- $A = A_1 \times \cdots \times A_n$, and A_i is a finite set of actions that player i can take. Each element $a \in A$ is called a strategy profile.
- $u = (u_1, \ldots, u_n)$ where $u_i : A \to \mathbb{R}$ is called the utility of x_i.

This definition encodes the interaction between the players. Each point $a \in A$ corresponds to every player choosing an action to perform. The function u gives the payoff that each player receives in this scenario. Each player's objective is then to maximize their own utility u_i. Let's see some examples.

1.1 Examples

Example 2 (Rock-paper-scissors)**.** Games are usually represented in a matrix:

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

The rows encode the actions player 1 can take. In other words, $A_1 = \{\text{Rock, Paper, Scissors}\}$. In this game, $A_2 = A_1$. Each entry of the matrix is the payoff which represents the usual hierarchy of moves: rock beats scissors, which beats paper, which beats rock, and everything else is a tie. For example, if player 1 chooses rock and player 2 plays scissors, $u(\text{Rock, Scissors}) = (1, -1)$.

Example 3 (Penalty kick)**.** Another simple game is a penalty kick. Here, a kicker tries to hit the ball into the goal by shooting either to the right or to the left. The goalie tries to prevent this by jumping to either side. If the goalie predicts the shot correctly, he stops it. Otherwise, the kicker scores. We can encode this as a normal form game:

	Jump left	Jump right
Kick left	$(-1, 1)$	$(1, -1)$
Kick right	$(1, -1)$	$(-1, 1)$

Example 4 (Prisoner's dilemma)**.** In this game, we have two people that are accused of a crime. They are given the option to remain silent or to betray and testify against the other player. The payoffs are as follows:

	Stay silent	Betray
Stay silent	$(-1, -1)$	$(-3, 0)$
Betray	$(0, -3)$	$(-\mathbf{2}, -\mathbf{2})$

Here, $-n$ means n years of jail. Notice that the best option is for both players to remain silent. That way they collectively get the smallest sentence of 1 year.

2 Nash equilibrium

With these examples at hand, we define the central concept:

Definition 5 (Nash equilibrium)**.**
A (pure) Nash equilibrium of the normal form game (N, A, u) is a strategy profile $a^* \in A$ such that for all $i \in N$,

$$u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*),$$

where $a_{-i}^* = (a_1^*, \dots, a_{i-1}^*, a_{i+1}^*, \dots, a_n^*)$. In other words, it is a choice of moves a^* where no player can improve their position by choosing a different move while the others pick the same strategy.

In the previous examples, only the Prisoner's Dilemma (Example 4) has a Nash equilibrium: both players have to betray one another. This is a little paradoxical because the optimal strategy is cooperation (*i.e.* both remain silent). To understand why, let's say that the players somehow managed to talk before the interviews and agreed to cooperate (this is not part of the original game, but will serve our example). In other words, they set $a^* = (\text{Stay silent, Stay silent})$. Of course, player 1 might be lying and thinking of betraying player 2, that is, player 1 will choose $a_1 = \text{Betray}$. Then

$$\begin{aligned} u_1(a_1^*, a_{-1}^*) &= u_1(\text{Stay silent, Stay silent}) = -1, \text{ while} \\ u_1(a_1, a_{-1}^*) &= u_1(\text{Betray, Stay silent}) = 0. \end{aligned}$$

Player 1 can obtain a better result by playing something different. This is why (Stay silent, Stay silent) is not a Nash equilibrium. In contrast, (Betray, Betray) is an equilibrium because the result of a single player changing strategy is getting more jail time.

However, our intuition suggests that this is not the complete picture. Looking back at Example 3, we can intuitively say there is an optimal strategy: the kicker has to shoot to the left (or to the right) half of the time. Otherwise, if he shoots left with probability $p \neq \frac{1}{2}$, the goalie will always jump to the side that has higher probability and outperform the kicker in the long run. We incorporate this idea by defining:

Definition 6 (Mixed strategy)**.**
Let S_i be the space of probability distributions over the set A_i. The product $S = S_1 \times \cdots \times S_n$ is called the mixed strategy profile of the game (N, A, u) and each point $s \in S$ is called a mixed strategy. The utility function is replaced by the expected value of the mixed strategy.

We extend the definition of Nash equilibrium to that of mixed equilibrium by replacing strategy profiles in Definition 5 with mixed strategies and utilities with expected values.

Example 7. If $A_i = \{c_1, \dots, c_m\}$, then $S_i = \{(p_1, \dots, p_m) \in [0,1]^m : p_1 + \cdots + p_m = 1\}$. The number p_k represents the probability that player i chooses to play the move c_k. Notice that $S_i \simeq \Delta^m$, a compact and convex set.

Example 8. In Example 3, let $A_k = \{\text{Kick left, Kick right}\}$ and $A_g = \{\text{Jump left, Jump right}\}$ be the strategy sets of the kicker and the goalie respectively. A mixed strategy for the kicker will be a probability distribution over A_k, or in other words, a choice of a number $p_\ell \in [0,1]$ that represents the probability of kicking to the left. Notice that the probability of kicking to the right is $p_r = 1 - p_\ell$. Analogously, the goalie chooses probabilities $q_\ell, q_r \in [0,1]$ that represent jumping to the left and to the right respectively. Thus, we can represent a mixed strategy as a tuple $(p_\ell, p_r; q_\ell, q_r) \in S_k \times S_g \simeq \Delta^1 \times \Delta^1$. The utility is replaced by its expected value over the chosen distribution. For example, the utility of the goalie when the players chose the strategy $(p_\ell, p_r; q_\ell, q_r)$ is

$$\begin{aligned} E[u_g(p_\ell, p_r; q_\ell, q_r)] &= \sum_{x \in A_k} \sum_{y \in A_g} u_g(x,y) \cdot \Pr[\text{Kicker shoots } x] \cdot \Pr[\text{Goalie jumps to } y] \\ &= p_\ell q_\ell + p_r q_r - p_\ell q_r - p_r q_\ell \\ &= (p_\ell - p_r)(q_\ell - q_r). \end{aligned}$$

When choosing a mixed strategy, this is the value that the goalie wants to maximize. Analogously, the utility function of the kicker is the negative of the above value.

The mixed Nash equilibrium in this situation, as remarked before Definition 6, occurs when $p_\ell = p_r = \frac{1}{2}$ and $q_\ell = q_r = \frac{1}{2}$. To see why, let $s^* = (\frac{1}{2}, \frac{1}{2}; \frac{1}{2}, \frac{1}{2})$ be a mixed strategy profile and $s^*_{-k} = (\frac{1}{2}, \frac{1}{2})$ be the mixed strategy of the goalie. Even if the kicker chooses a different strategy $s_k = (p_\ell, p_r)$, $E[u_k(s_k, s^*_{-k})] = E[u_k(p_\ell, p_r; \frac{1}{2}, \frac{1}{2})] = 0$. The same holds for the goalie: $E[u_g(\frac{1}{2}, \frac{1}{2}; q_\ell, q_r)] = 0$. Since $E[u(s^*)] = E[u(\frac{1}{2}, \frac{1}{2}; \frac{1}{2}, \frac{1}{2})] = (0,0)$ (*i.e.* the expected utility is 0 for both players), this shows that s^* is a mixed Nash equilibrium. The expected utility of a player does not improve when deviating from the equilibrium strategy.

Example 9. The mixed Nash equilibrium in Rock-Paper-Scissors is playing each option with probability $\frac{1}{3}$.

The above example shows that games that have no pure Nash equilibrium can have mixed ones. This is no coincidence, but a consequence of the following result [3].

Theorem 10. *Every finite n-person normal form game has a mixed Nash equilibrium.*

Proof: First of all, notice that $S_i \simeq \Delta^{m_i}$ means that $S = S_1 \times \cdots \times S_n \simeq \Delta^{m_1} \times \cdots \times \Delta^{m_n} \simeq \Delta^M \simeq \mathbb{D}^M$ for some M. In other words, S is compact and convex.

For $x, y \in S$, say that x counters y if $\forall i \in N$ and $\forall z \in S$,

$$E[u_i(x_i, z_{-i})] \geq E[u_i(y_i, z_{-i})].$$

That is, the expected utility for player i when choosing x_i is always higher than that of y_i, even when considering any possible choice of strategy z_{-i} from the opponents, and this happens for all choices of i. Now, define $F : S \to \mathcal{P}(S)$ by $F(x) = \{y \in S : y \text{ counters } x\} \subset S$.

The Kakutani fixed point theorem states that if a function $F : \mathbb{D}^M \to \mathcal{P}(\mathbb{D}^M)$ has $F(x)$ convex for all $x \in \mathbb{D}^M$ and the graph $\Gamma = \{(x,y) : x \in \mathbb{D}^M, y \in F(x)\}$ is closed, then F has a fixed point. Namely, there exists $x \in \mathbb{D}^M$ with $x \in F(x)$. We check these two conditions. Since the expected value is linear, a convex combination of mixed strategies that counter x will also counter x. That is, $F(x)$ is convex. Also, if $(P_n, Q_n)_n \subset \Gamma$ is a sequence where Q_n counters P_n and coverges to (P, Q), then Q will counter P because the expected value of the utility is a continuous function of the mixed strategy. More explicitly, $E[u(Q_n)] \to E[u(Q)]$ and $E[u(P_n)] \to E[u(P)]$. This shows that Γ is closed, and thus, there exists a mixed strategy $x \in F(x)$. This is a self countering strategy by definition of F and a Nash equilibrium. $\square$

3 Computational complexity

Now that we know what a Nash equilibrium is, we can ask how hard is it to find one. The most known classification of computer algorithms is P and NP. A decision problem is a question whose answer is yes or no. Such a problem is said to be in P if there exists a deterministic polynomial time algorithm that can determine whether a solution exists or not. Conversely, NP is the collection of problems for which a non-deterministic polynomial time algorithm exists instead. However, NASH (the algorithm for finding a Nash equilibrium) doesn't fit in this classification because the solution always exists, in other words, it is not a decision problem. The answer is always yes. Instead, we use another class introduced by [1] called TFNP. This is the collection of total function problems that can be solved by a non-deterministic polynomial time algorithm. An example of such a problem is factorization of integers. TFNP is further subdivided by the mathematical argument used to find the solution. NASH is in the class PPAD, which stands for Polynomial Parity Arguments on Directed Graphs. The defining problem of PPAD is called End Of the Line (EOL).

Definition 11. EOL Suppose we have a directed graph G with possibly exponentially many vertices. G satisfies:

- It has at least one source (vertex with no predecessor).
- It has no isolated vertices.
- Every vertex has at most one predecessor and one successor.
- There exists polynomial time algorithms P and S that given a vertex v, they output the predecessor $P(v)$ and the successor $S(v)$ (provided they exist).

The End of the Line (EOL) algorithm takes as input this graph G and a source vertex. The output is another source or sink (vertex with no successor). A problem is said to be in PPAD if we can find a solution by using EOL.

At first glance, EOL doesn't look like a hard problem. We can explain why by being more precise as to what it means for G to have exponentially many vertices. The vertices are encoded as an integer between 0 and M in binary form and this requires $n = \lceil \log_2(M) \rceil$ bits of storage. As a side note, P and V are defined to be polynomial time algorithms with respect to n. However, M is at least 2^{n-1}, an exponential quantity in n. This should clarify the difficulty of EOL: while it is true that we can start at the source vertex and follow the edges until we hit a sink, there are around 2^n vertices to search. The fact that P and V are relatively fast is overshadowed by the number of times we have to use them.

As it turns out, NASH is PPAD-complete. This means that we can solve NASH by constructing a suitable graph and using EOL to find a sink. Conversely, we can encode a directed graph as a game in which a Nash equilibrium represents a sink or source. One more step is required in this construction, shown in the diagram:

$$\text{NASH} \rightleftarrows \text{BROUWER} \rightleftarrows \text{EOL}$$

In other words, we can solve NASH by constructing a continuous function that has a fixed point guaranteed by the Brouwer fixed point theorem. Such a point encodes a Nash equilibrium, and we can find it by using EOL. We will show how to do this, and give a sketch of the other direction. For more details, see [1] and [2].

3.1 BROUWER $\rightarrow$ EOL

The Brouwer fixed point theorem says that any continuous function $F : [0,1]^m \to [0,1]^m$ has a fixed point. As it turns out, we can build a suitable graph to use EOL on. We will show how when $m = 2$.

First, subdivide $[0,1]^2$ into an very fine $M \times M$ grid of triangles (see picture). As before, let $n = \lceil \log_2(M) \rceil$. Suppose we are given a polynomial time algorithm (with respect to n) that calculates F in the vertices of these squares. We color each vertex according to where it is sent by F. Namely, if the vector $F(x) - x$ has an angle between $0°$ and $90°$, color it yellow; color it blue if the angle is between $90°$ and $225°$ and red if it is between $-135°$ and $0°$.

Notice that the left side of the square has no blue, the bottom has no red and the top has no yellow. Also, assume that all red vertices on the left side are above all the yellow ones. We can achieve this by extending the square a little on the left, and slightly modifying F in a way that it is still continuous. As we saw a couple of talks ago, Sperner's lemma guarantees that this grid will have a triangle with all three colors. This means that F is sending all the vertices in different directions. They are either moving into the interior of the triangle or away from it. In any case, there is a fixed point in the interior. Thus, the vertices are approximate fixed points, which are sufficient for our purposes.

To find it using EOL, we need a directed graph. Each triangle in the grid will be a vertex of the graph. We add an edge between two triangles T_1 and T_2 if they share a side with a red and a yellow vertex. The edge goes from T_1 to T_2 if, when we walk on the boundary of T_1 clockwise, we go from the red to the yellow vertex. Notice that we have $M^2 \sim 2^{2n}$ triangles. Given a triangle T, we can calculate if it has a red-yellow edge by calculating F on its three vertices, each polynomially in n. Lastly, notice that the single change from red to black makes its triangle a source vertex. In summary, we have a directed graph with exponential number of vertices that has at least one source, and calculating the predecessor and successor of each vertex can be done in polynomial time. Thus, we can find a sink with EOL. The sink has to be a trichromatic triangle, so EOL found the promised approximate fixed points.

3.2 NASH $\rightarrow$ BROUWER

We have to encode NASH as a continuous function on a cube whose fixed point is a Nash equilibrium. Suppose that all the players picked a mixed strategy, that is, a point $x \in S \simeq \Delta^M \simeq [0,1]^M$. Most often than not, the players will not be satisfied with this choice, so let them choose another strategy y close to x and set $F(x) = y$. Then F is a continuous function $[0,1]^M \rightarrow [0,1]^M$ that has a fixed point. This is a Nash equilibrium because players do not want to change their strategy in that situation. Since we can transform BROUWER into EOL, this shows NASH $\in$ PPAD.

3.3 EOL $\rightarrow$ BROUWER $\rightarrow$ NASH

This will be a very short summary of the idea. For a more detailed explanation, see [1], and [2] for the complete details. The EOL to BROUWER transformation is done by embedding the directed graph into $[0,1]^3$ and carefully constructing a continuous function $F : [0,1]^3 \rightarrow [0,1]^3$ so that the sinks are fixed points. In turn, the fixed points can be found by building a game with three "input" and three "output" players that choose between 1 and 0. Let $x_1, x_2, x_3 \in [0,1]$ be the probabilities that each input player chooses 1. It is possible to add more players to the game so that the probabilities that the three output players pick 1 are the three coordinates of $F(x_1, x_2, x_3) \in [0,1]^3$. A fixed point of this function turns out to be a Nash equilibrium, so if we have an algorithm for finding Nash equilibria, we can also use it to find the sink of the original EOL graph. This shows that NASH is PPAD-complete.

References

[1] Daskalakis, Goldberg, and Papadimitriou. "The Complexity of Computing a Nash Equilibrium". *Communications of the ACM*, 52(2), 8997, 2009.

[2] Daskalakis, Goldberg, and Papadimitriou. "The Complexity of Computing a Nash Equilibrium". *SIAM Journal on Computing*, 39(1), 195259, 2009.

[3] Nash, J. "Equilibrium points in n-person games". *Proceedings of the National Academy of Sciences* 36(1), 48-49, 1950.

[4] Racz, Yang. "Lecture: Complexity of Finding a Nash Equilibrium". Recovered from https://people.cs.pitt.edu/∼kirk/CS1699Fall2014/lect4.pdf