

Game Theory 2024

Ulle Endriss

Institute for Logic, Language and Computation

University of Amsterdam

Plan for Today

Today we are going to generalise our definition of “strategy” and allow players to randomise over several actions to play.

We are then going to generalise the notion of *Nash equilibrium* to this setting and discuss the following three topics:

- the (manual) *computation* of the Nash equilibria for small games
- the *existence* of Nash equilibria for arbitrary games
- (briefly) the computational *complexity* of finding Nash equilibria

Most of this is also covered in Chapters 1 and 2 of the *Essentials*.

K. Leyton-Brown and Y. Shoham. *Essentials of Game Theory: A Concise, Multi-disciplinary Introduction*. Morgan & Claypool Publishers, 2008. Chapters 1 & 2.

Reminder

A *normal-form game* is a tuple $\langle N, \mathbf{A}, \mathbf{u} \rangle$, where

- $N = \{1, \dots, n\}$ is a finite set of *players* (or *agents*);
- $\mathbf{A} = A_1 \times \dots \times A_n$ is a finite set of *action profiles* $\mathbf{a} = (a_1, \dots, a_n)$, with A_i being the set of *actions* available to player i ; and
- $\mathbf{u} = (u_1, \dots, u_n)$ is a profile of *utility functions* $u_i : \mathbf{A} \rightarrow \mathbb{R}$.

An action profile \mathbf{a} is a *pure Nash equilibrium*, if no player i wants to unilaterally deviate from her assigned action a_i : $u_i(\mathbf{a}) \geq u_i(a'_i, \mathbf{a}_{-i})$.

An action profile \mathbf{a} is *Pareto efficient*, if no other profile would be better for some player and no worse for any of the others.

Coordination Games

A (pure) *coordination game* is a normal-form game $\langle N, \mathbf{A}, \mathbf{u} \rangle$ with $u_i(\mathbf{a}) = u_j(\mathbf{a})$ for all players $i, j \in N$ and all action profiles $\mathbf{a} \in \mathbf{A}$.

Example: A world with just two drivers. Which side of the road to use?







	L	R
L	1, 1	0, 0
R	0, 0	1, 1

Remark: For this game, every pure NE is Pareto efficient. Nice.

Exercise: *Is this the case for all coordination games?*

Zero-Sum Games

A *zero-sum game* is a two-player normal-form game $\langle N, \mathbf{A}, \mathbf{u} \rangle$ with $u_1(\mathbf{a}) + u_2(\mathbf{a}) = 0$ for all action profiles $\mathbf{a} \in \mathbf{A}$. Example:

			
	0 / 0	1 / -1	-1 / 1
	-1 / 1	0 / 0	1 / -1
	1 / -1	-1 / 1	0 / 0

Exercise: What are the pure NE? Intuitively, how should you play?

Mixed Strategies and Expected Utility

So far, the space of strategies available to player i has simply been her set of actions A_i (pure strategy = action). We now generalise and allow player i to play any action in A_i with a certain probability.

For any finite set X , let $\Pi(X) = \{p : X \rightarrow [0, 1] \mid \sum_{x \in X} p(x) = 1\}$ be the set of all *probability distributions* over X .

A *mixed strategy* s_i for player i is a probability distribution in $\Pi(A_i)$.

The set of all her mixed strategies is $S_i = \Pi(A_i)$.

A *mixed-strategy profile* $\mathbf{s} = (s_1, \dots, s_n)$ is an element of $S_1 \times \dots \times S_n$.

The *expected utility* of player i for the mixed-strategy profile \mathbf{s} is:

$$u_i(\mathbf{s}) = \sum_{\mathbf{a} \in \mathbf{A}} \left[u_i(\mathbf{a}) \cdot \prod_{j \in N} s_j(a_j) \right]$$

Remark: Note the overloading of the symbol u_i (also denotes utility).

Types of Mixed Strategies

The *support* of strategy s_i is the set of actions $\{a_i \in A_i \mid s_i(a_i) > 0\}$.

A mixed strategy s_i is *pure* if its support is a singleton.

A mixed strategy s_i is *truly mixed* if it is not pure.

A mixed strategy s_i is *fully mixed* if its support is the full set A_i .

Example: Battle of the Sexes

Traditionally minded **Rowena** and **Colin** are planning a social activity. Worst of all would be not to agree on a joint activity; but if they do manage, **Colin** prefers *auto racing* and **Rowena** really prefers *ballet*.

	A	B
A	2, 4	0, 0
B	0, 0	0, 3

Suppose **Rowena** chooses to go to the ballet with 75% probability, while **Colin** chooses to go to the races with certainty (pure strategy):

$$s_1 = \left(\frac{1}{4}, \frac{3}{4}\right) \quad s_2 = (1, 0)$$

Then: $u_1(s) = 2 \cdot \left(\frac{1}{4} \cdot 1\right) + 0 \cdot \left(\frac{1}{4} \cdot 0\right) + 0 \cdot \left(\frac{3}{4} \cdot 1\right) + 8 \cdot \left(\frac{3}{4} \cdot 0\right) = \frac{1}{2}$

Mixed Nash Equilibria

Consider a game $\langle N, \mathbf{A}, \mathbf{u} \rangle$ with associated (mixed) strategies $s_i \in S_i$.

We say that strategy $s_i^* \in S_i$ is a *best response* for player i to the (partial) strategy profile \mathbf{s}_{-i} if $u_i(s_i^*, \mathbf{s}_{-i}) \geq u_i(s'_i, \mathbf{s}_{-i})$ for all $s'_i \in S_i$.

We say that profile $\mathbf{s} = (s_1, \dots, s_n)$ is a *mixed Nash equilibrium*, if s_i is a best response to \mathbf{s}_{-i} for every player $i \in N$.

Thus: no player has an incentive to unilaterally change her strategy.

Remark: Note how this definition mirrors that of pure Nash equilibria.

Exercise: *Can you think of a game with infinitely many Nash equilibria?*

Example: Driving Game

Recall: A world with just two drivers. Which side of the road to use?

	L	R
L	1, 1	0, 0
R	0, 0	1, 1

For this game, it is easy to guess what the Nash equilibria are:

- (1) pure NE: both pick left with certainty: $((1, 0), (1, 0))$ [optimal!]
- (2) pure NE: both pick right with certainty: $((0, 1), (0, 1))$ [optimal!]
- (3) both choose fifty-fifty: $((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$ [anything is best response!]

There is no other NE: Suppose I pick $(\frac{1}{2} + \epsilon, \frac{1}{2} - \epsilon)$, e.g., $(0.51, 0.49)$.

Then your best response is $(1, 0)$, to which my best response is $(1, 0)$.

Computing Nash Equilibria

Suppose we have guessed (correctly) that this game has exactly one NE (s_1, s_2) and that it is fully mixed. How to compute it?

	L	R
T	<div style="display: flex; justify-content: space-between; align-items: center;"> 6 4 </div>	<div style="display: flex; justify-content: space-between; align-items: center;"> 7 5 </div>
B	<div style="display: flex; justify-content: space-between; align-items: center;"> 3 2 </div>	<div style="display: flex; justify-content: space-between; align-items: center;"> 8 1 </div>

Let $s_1 = (p, 1 - p)$ and $s_2 = (q, 1 - q)$. If your strategy is fully mixed, you must be indifferent between your two actions. Thus:

- **Player 2** is indifferent: $4p + 2(1 - p) = 5p + 1(1 - p) \Rightarrow p = \frac{1}{2}$
- **Player 1** is indifferent: $6q + 7(1 - q) = 3q + 8(1 - q) \Rightarrow q = \frac{1}{4}$

Exercise: Game of Chicken

To establish their relative levels of bravery, **Rowena** and **Colin** race their cars towards a cliff at full speed. Each can *jump out* or *wait*. If both wait, they die. If both jump, nothing happens. Otherwise, whoever jumps faces humiliation, while the other one wins.

	J	W
J	$\begin{matrix} 0 & 10 \\ 0 & -5 \end{matrix}$	$\begin{matrix} 10 & -5 \\ -5 & -20 \end{matrix}$
W	$\begin{matrix} -5 & -20 \\ 10 & -20 \end{matrix}$	$\begin{matrix} -20 & -20 \\ -20 & -20 \end{matrix}$

What are the Nash equilibria of this game?

Nash's Theorem

Recall that some games do not have pure Nash equilibria. Good news:

Theorem 1 (Nash, 1951) *Every (finite) normal-form game has at least one (truly mixed or pure) Nash equilibrium.*

We are now going to prove this seminal result. Plan:

- Definition of function f from strategy profiles to strategy profiles, simulating the *updates* players *might* use to try to improve their lot.
- Lemma showing that s is a NE iff s is a *fixed point* of f .
- Presentation of *Brouwer's Fixed-Point Theorem* (but no proof), giving sufficient conditions for a function to have a fixed point.
- Lemma showing that f meets the conditions of Brouwer's Theorem.

J.F. Nash. Non-Cooperative Games. *Annals of Mathematics*, 54(2):286–295, 1951.

Heuristic Improvement Dynamics

Let $\mathcal{S} = S_1 \times \cdots \times S_n$ be the space of all mixed-strategy profiles.

In any given profile $\mathbf{s} \in \mathcal{S}$, player i will look for a new strategy that improves her payoff. Let $f_i : \mathcal{S} \rightarrow S_i$ describe how she *updates*.

Player i might use a heuristic. First, for every pure strategy $a \in A_i$, she computes her *gain* if she switches to a (and nobody else moves):

$$g_i(\mathbf{s}, a) = \max\{u_i(a, \mathbf{s}_{-i}) - u_i(\mathbf{s}), 0\}$$

Then, she updates the probability of every action $a \in A_i$, in line with these expected gains (ensuring they still add up to 1), from $s_i(a)$ to:

$$f_i(\mathbf{s})(a) = \frac{s_i(a) + g_i(\mathbf{s}, a)}{\sum_{a' \in A_i} s_i(a') + g_i(\mathbf{s}, a')} = \frac{s_i(a) + g_i(\mathbf{s}, a)}{1 + \sum_{a' \in A_i} g_i(\mathbf{s}, a')}$$

If everyone does this, we get a *global* update function $f : \mathcal{S} \rightarrow \mathcal{S}$ with $f(\mathbf{s}) = (f_1(\mathbf{s}), \dots, f_n(\mathbf{s}))$. Exercise: Suppose \mathbf{s} is a *NE*. What is $f(\mathbf{s})$?

Nash Equilibria and Fixed Points

Recall: $f : \mathcal{S} \rightarrow \mathcal{S}$ with $f(\mathbf{s}) = (f_1(\mathbf{s}), \dots, f_n(\mathbf{s}))$, where:

$$f_i(\mathbf{s})(a) = \frac{s_i(a) + g_i(\mathbf{s}, a)}{1 + \sum_{a' \in A_i} g_i(\mathbf{s}, a')} \quad \left[\begin{array}{l} g_i(\mathbf{s}, a) = \\ \max\{u_i(a, \mathbf{s}_{-i}) - u_i(\mathbf{s}), 0\} \end{array} \right]$$

Lemma 2 A strategy profile $\mathbf{s} \in \mathcal{S}$ is a *Nash equilibrium* iff \mathbf{s} is a *fixed point* of the update function f (meaning that $f(\mathbf{s}) = \mathbf{s}$).

Proof: (\Rightarrow) If \mathbf{s} is a NE, no strategy, certainly no pure strategy a , can increase i 's payoff, i.e., $g_i(\mathbf{s}, a) = 0$. Thus: $f_i(\mathbf{s})(a) = s_i(a)$ ✓

(\Leftarrow) Suppose \mathbf{s} is a fixed point of f , i.e., $f_i(\mathbf{s})(a) = s_i(a)$. Cases:

- $g_i(\mathbf{s}, a) = 0$ for all $i \in N$ and $a \in A_i$: If it is impossible to improve via a pure strategy, then also via a mixed strategy. Thus, \mathbf{s} is a NE. ✓
- *Otherwise*, use definition of f_i to get $s_i(a) = g_i(\mathbf{s}, a) / \sum_{a' \in A_i} g_i(\mathbf{s}, a')$. Thus $s_i(a) > 0 \Leftrightarrow g_i(\mathbf{s}, a) > 0$, meaning $s_i(a) > 0 \Leftrightarrow u_i(a, \mathbf{s}_{-i}) > u_i(\mathbf{s})$. But this contradicts $u_i(\mathbf{s}) = \sum_{a \in A_i} u_i(a, \mathbf{s}_{-i}) \cdot s_i(a)$. ✓

Brouwer's Fixed-Point Theorem

Theorem 3 (Brouwer, 1911) Let $X \subseteq \mathbb{R}^m$ ($X \neq \emptyset$) be *compact* and *convex*. Then every *continuous* function $f : X \rightarrow X$ has a *fixed point*.

Recall that this means that there exists an $x \in X$ such that $f(x) = x$.

Explanation of the terminology used:

- *Compactness*. X is compact if it is *closed* (contains its limit points) and *bounded* (any two elements have distance $\leq K$, for some K).
- *Convexity*. X is convex if any point “between” x and y in X is also in X , i.e., if $\lambda \cdot x + (1 - \lambda) \cdot y \in X$ for any $\lambda \in [0, 1]$.
- *Continuity*. You know this one (“no sudden jumps”).

L.E.J. Brouwer. Über Abbildung von Mannigfaltigkeiten. *Mathematische Annalen*, 71(1):97–115, 1911.

Examples

On $X = [0, 1]$, the function $f : x \mapsto x^2$ has the fixed points 0 and 1.

On $X = \{(x, y) \in [0, 1]^2 \mid x^2 + y^2 \leq 1\}$, the “mirroring” function $f : (x, y) \mapsto (-x, -y)$ has the fixed point $(0, 0)$.

Put a map of Amsterdam on a table somewhere in A’dam. Then some point on the map will be directly above the location it represents.

(X is the set of locations in A’dam; f is the projection to the map.)

On $X = [0, 1)$, the function $f : x \mapsto \frac{x+1}{2}$ has no fixed point, because X is not closed. But for $X = [0, 1]$ we get the fixed point 1.

On $X = \mathbb{R}$, the function $f : x \mapsto x + 1$ has no fixed point, because X is not bounded (even though it is closed).

On $X = \{-1, 1\}$, the function $f : x \mapsto -x$ has no fixed point, because X is not convex.

On $X = [0, 9]$, the function $f : x \mapsto (\lfloor x \rfloor + 1) \bmod 10$ has no fixed point, because f is not continuous.

Applying Brouwer's Fixed-Point Theorem

Recall: $f : \mathcal{S} \rightarrow \mathcal{S}$ with $f(\mathbf{s}) = (f_1(\mathbf{s}), \dots, f_n(\mathbf{s}))$, where:

$$f_i(\mathbf{s})(a) = \frac{s_i(a) + g_i(\mathbf{s}, a)}{1 + \sum_{a' \in A_i} g_i(\mathbf{s}, a')} \quad \left[\begin{array}{l} g_i(\mathbf{s}, a) = \\ \max\{u_i(a, \mathbf{s}_{-i}) - u_i(\mathbf{s}), 0\} \end{array} \right]$$

We can think of every s_i as a *stochastic vector* of length $|A_i|$, so $S_i \subseteq [0, 1]^{|A_i|}$ and similarly $\mathcal{S} \subseteq [0, 1]^m$ for $m = |A_1| + \dots + |A_n|$.

This \mathcal{S} and $f : \mathcal{S} \rightarrow \mathcal{S}$ satisfy the conditions of Brouwer's Theorem:

- *Compactness* of \mathcal{S} : $\mathcal{S} \subseteq [0, 1]^m$ is *bounded*, because $[0, 1]^m$ is. $\mathcal{S} = S_1 \times \dots \times S_n$ is *closed*, because all S_i are (the limit of any converging sequence of stochastic vectors is itself stochastic). ✓
- *Convexity* of \mathcal{S} : Follows from fact that every convex combination of stochastic vectors is itself a stochastic vector. ✓
- *Continuity* of f : As all $u_i : \mathcal{S} \rightarrow \mathbb{R}$ are continuous, so are all $g_i : \mathcal{S} \times A_i \rightarrow \mathbb{R}_{\geq 0}$, all $f_i(\mathbf{s}) : A_i \rightarrow [0, 1]$, and all $f_i : \mathcal{S} \rightarrow S_i$. ✓

Complexity of Computing Nash Equilibria

Note: The proof of Nash's Theorem does *not* provide us with a method to actually compute Nash equilibria, because it is not constructive.

Both the design of *algorithms* for computing Nash equilibria and the analysis of the *computational complexity* of this task are important research topics in *algorithmic game theory*.

Regarding the complexity:

- Finding a NE is *in NP*: if you guess a NE, I can easily verify.
- But likely not *not NP-hard*: guaranteed existence would be atypical.
- Complete for *PPAD* (“polynomial parity argument for directed graphs”), which lies “between” P and NP. Believed to be *intractable*.

Discussion: NE as model of rational behaviour vs. high complexity.

N. Nisan, T. Roughgarden, E. Tardos, and V.V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.

Summary

We have introduced the notion of a *mixed strategy*, where a player randomises over several pure strategies (i.e., actions). And:

- *Nash's Theorem*: every normal-form game has a Nash equilibrium
- technique for *computing* NE's for small (two-player) games

The NE, although not perfect, is the most important *solution concept* in game theory, and we'll return to it frequently.

Points of concern regarding the notion of (mixed) NE:

- Is it reasonable to assume players are perfectly *rational*?
- Is it reasonable to assume players work with *probabilities*?
- Is it reasonable to assume players can handle the high *complexity*?
- What if there are *many* Nash equilibria? How do you choose?

What next? Focus on a specific class of games with nice properties, the so-called *congestion games*.