

Game Theory 2024

Ulle Endriss

Institute for Logic, Language and Computation

University of Amsterdam

(class taught by Federico Fioravanti)

Plan for Today


We have seen that every normal-form game has a Nash equilibrium, although not necessarily one that is pure. Pure equilibria are nicer.

Today we are going to discuss a family of games of practical interest where we can guarantee the existence of pure Nash equilibria:

- *congestion games*: examples and definition
- *potential games*: tool to analyse congestion games
- existence of *pure Nash equilibria* for both types of games
- finding those equilibria by means of *better-response dynamics*
- (briefly) *price of anarchy*: quality guarantees for equilibria

Example: Traffic Congestion

10 people need to get from A to B . Everyone can choose between the top and the bottom route. Via the top route, the trip takes 10 mins. Via the bottom route, it depends on the number of fellow travellers: it takes as many minutes x as there are people using this route.



What do you do? And: what are the pure Nash equilibria?

Example: The El Farol Bar Problem

100 people consider visiting the *El Farol* Bar on a Thursday night. They all have identical preferences:

- If 60 or more people show up, it's nicer to be at home.
- If fewer than 60 people show up, it's nicer to be at the bar.

Now what? And: what are the pure Nash equilibria of this game?



Congestion Games

A *congestion game* is a tuple $\langle N, R, \mathbf{A}, \mathbf{d} \rangle$, where

- $N = \{1, \dots, n\}$ is a finite set of *players*;
- $R = \{1, \dots, m\}$ is a finite set of *resources*;
- $\mathbf{A} = A_1 \times \dots \times A_n$ is a finite set of *action profiles* $\mathbf{a} = (a_1, \dots, a_n)$, with $A_i \subseteq 2^R$ being the set of *actions* available to player i ; and
- $\mathbf{d} = (d_1, \dots, d_m)$ is a vector of *delay functions* $d_r : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$, each of which is required to be *nondecreasing*.

Thus, every player chooses a set of resources to use (that's her action). Note that each d_r is associated with a resource (not with a player).

Let $n_r^{\mathbf{a}} = \#\{i \in N \mid r \in a_i\}$ be the number of players claiming r in \mathbf{a} . The *cost* incurred by player i is the sum of the delays she experiences due to the congestion of the resources she picks. Her *utility* then is:

$$u_i(\mathbf{a}) = -\text{cost}_i(\mathbf{a}) = -\sum_{r \in a_i} d_r(n_r^{\mathbf{a}})$$

Exercise: So why doesn't player i simply not pick any resources at all?

Formally Modelling the Examples

Our two initial examples fit this formal model:

- *Traffic Congestion*

- players $N = \{1, 2, \dots, 10\}$
- resources $R = \{\uparrow, \downarrow\}$
- action spaces $A_i = \{\{\uparrow\}, \{\downarrow\}\}$ representing the two routes
- delay functions $d_{\uparrow} : x \mapsto 10$ and $d_{\downarrow} : x \mapsto x$

- *El Farol Bar Problem*

- players $N = \{1, 2, \dots, 100\}$
- resources $R = \{\mathbf{Y}, \mathbf{H}_1, \mathbf{H}_2, \dots, \mathbf{H}_{100}\}$
- action spaces $A_i = \{\{\mathbf{Y}\}, \{\mathbf{H}_i\}\}$
- delay functions $d_{\mathbf{Y}} : x \mapsto \mathbb{1}_{x \geq 60}$ and $d_{\mathbf{H}_i} : x \mapsto \frac{1}{2}$

Remark: Neither example makes full use of the power of the model, as every player only ever claims a single resource.

Variations

Our model of congestion games has certain restrictions:

- Utility functions are *additive* (no synergies between resources).
- Delay functions are *not player-specific* (no individual tastes).

Some careful relaxations of these assumptions have been considered in the literature, but we are not going to discuss these here.

Existence of Pure Nash Equilibria

Good news:

Theorem 1 (Rosenthal, 1973) *Every congestion game has at least one pure Nash equilibrium.*

We postpone the proof and first introduce some additional machinery.

R.W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.

Potential Games

A normal-form game $\langle N, \mathbf{A}, \mathbf{u} \rangle$ is a *potential game* if there exists a function $P : \mathbf{A} \rightarrow \mathbb{R}$ such that, for all $i \in N$, $\mathbf{a} \in \mathbf{A}$, and $a'_i \in A_i$:

$$u_i(\mathbf{a}) - u_i(a'_i, \mathbf{a}_{-i}) = P(\mathbf{a}) - P(a'_i, \mathbf{a}_{-i})$$

Example: Prisoner's Dilemma

The game underlying the Prisoner's Dilemma is a potential game, because we can define a function P from action profiles (matrix cells) to the reals that correctly tracks any unilateral deviation:

		C	D	
C	<div style="display: flex; justify-content: space-between; width: 100%;"> -10 0 </div> <div style="display: flex; justify-content: space-between; width: 100%;"> -10 -25 </div>			$P(\text{C}, \text{C}) = 50$
	<div style="display: flex; justify-content: space-between; width: 100%;"> -25 -20 </div> <div style="display: flex; justify-content: space-between; width: 100%;"> 0 -20 </div>			$P(\text{C}, \text{D}) = 60$
D	<div style="display: flex; justify-content: space-between; width: 100%;"> -25 -20 </div> <div style="display: flex; justify-content: space-between; width: 100%;"> 0 -20 </div>			$P(\text{D}, \text{C}) = 60$
	<div style="display: flex; justify-content: space-between; width: 100%;"> -25 -20 </div> <div style="display: flex; justify-content: space-between; width: 100%;"> 0 -20 </div>			$P(\text{D}, \text{D}) = 65$

For example, if Colin deviates from (C, C) to (C, D) , his utility will increase by 10, and indeed: $P(\text{C}, \text{D}) - P(\text{C}, \text{C}) = 60 - 50 = 10$.

Example: Matching Pennies

Each player gets a penny and secretly displays either heads or tails.

Rowena wins if the two pennies agree; Colin wins if they don't.

	H	T
H	-1 1	1 -1
T	1 -1	-1 1

Exercise: Show that this is *not* a potential game.

Existence of Pure Nash Equilibria

Recall that a game $\langle N, \mathbf{A}, \mathbf{u} \rangle$ is called a *potential game* if there exists a function $P : \mathbf{A} \rightarrow \mathbb{R}$ such that, for all $i \in N$, $\mathbf{a} \in \mathbf{A}$, and $a'_i \in A_i$:

$$u_i(\mathbf{a}) - u_i(a'_i, \mathbf{a}_{-i}) = P(\mathbf{a}) - P(a'_i, \mathbf{a}_{-i})$$

Good news:

Theorem 2 (Monderer and Shapley, 1996) *Every potential game has at least one pure Nash equilibrium.*

Proof: Take (one of) the action profile(s) \mathbf{a} for which P is maximal. By definition, no player can benefit by deviating using a pure strategy. Thus, also not using a mixed strategy. Hence, \mathbf{a} must be a pure NE. ✓

D. Monderer and L.S. Shapley. Potential Games. *Games and Economic Behavior*, 14(1):124–143, 1996.

Back to Congestion Games

We still need to prove that also every congestion game has a pure NE.
We are done, if we can prove the following lemma:

Lemma 3 *Every congestion game is a potential game.*

Proof: Take any congestion game $\langle N, R, \mathbf{A}, \mathbf{d} \rangle$. Recall that:

$$u_i(\mathbf{a}) = - \sum_{r \in a_i} d_r(n_r^{\mathbf{a}}) \quad \text{where} \quad n_r^{\mathbf{a}} = \#\{i \in N \mid r \in a_i\}$$

Now define the function P as follows:

$$P(\mathbf{a}) = - \sum_{r \in R} \sum_{k=1}^{n_r^{\mathbf{a}}} d_r(k) \quad \text{for all } \mathbf{a} \in \mathbf{A}$$

Easy to verify that $u_i(\mathbf{a}) - u_i(a'_i, \mathbf{a}_{-i}) = P(\mathbf{a}) - P(a'_i, \mathbf{a}_{-i})$, given that both sides simplify to this: $-\sum_{r \in a_i \setminus a'_i} d_r(n_r^{\mathbf{a}}) + \sum_{r \in a'_i \setminus a_i} d_r(n_r^{(a'_i, \mathbf{a}_{-i})})$

Thus, P is a potential for our congestion game. \checkmark

Intuition: $d_r(k)$ is the cost for the k th player arriving at resource r .

Proof Detail

We just claimed that the LHS and the RHS of the crucial equation in our proof both reduce to the same term. Let's check this really is true:

- LHS: $u_i(\mathbf{a}) - u_i(\mathbf{a}'_i, \mathbf{a}_{-i})$

Use definition of u_i to get: $-\sum_{r \in a_i} d_r(n_r^{\mathbf{a}}) + \sum_{r \in a'_i} d_r(n_r^{(\mathbf{a}'_i, \mathbf{a}_{-i})})$

But for any resource $r \in a_i \cap a'_i$ the terms in the two sums cancel!

So we are left with: $-\sum_{r \in a_i \setminus a'_i} d_r(n_r^{\mathbf{a}}) + \sum_{r \in a'_i \setminus a_i} d_r(n_r^{(\mathbf{a}'_i, \mathbf{a}_{-i})}) \checkmark$

- RHS: $P(\mathbf{a}) - P(\mathbf{a}'_i, \mathbf{a}_{-i})$

Use definition of P : $-\sum_{r \in R} \left[\sum_{k=1}^{n_r^{\mathbf{a}}} d_r(k) - \sum_{k=1}^{n_r^{(\mathbf{a}'_i, \mathbf{a}_{-i})}} d_r(k) \right]$

But $n_r^{\mathbf{a}}$ and $n_r^{(\mathbf{a}'_i, \mathbf{a}_{-i})}$ can differ by at most 1, so most terms cancel!

When $n_r^{\mathbf{a}} > n_r^{(\mathbf{a}'_i, \mathbf{a}_{-i})}$ only $d_r(n_r^{\mathbf{a}})$ is left, and *vice versa*.

So we are left with: $-\sum_{r \in a_i \setminus a'_i} d_r(n_r^{\mathbf{a}}) + \sum_{r \in a'_i \setminus a_i} d_r(n_r^{(\mathbf{a}'_i, \mathbf{a}_{-i})}) \checkmark$

Better-Response Dynamics

We start in some action profile \mathbf{a}^0 . Then, at every step, some player i unilaterally deviates to achieve an outcome that is better for her:

- $a_i^k \in A_i$ such that $u_i(a_i^k, \mathbf{a}_{-i}^{k-1}) > u_i(\mathbf{a}^{k-1})$
- $a_{i'}^k = a_{i'}^{k-1}$ for all other players $i' \in N \setminus \{i\}$

This leads to a sequence $\mathbf{a}^0 \rightsquigarrow \mathbf{a}^1 \rightsquigarrow \mathbf{a}^2 \rightsquigarrow \mathbf{a}^3 \rightsquigarrow \dots$

A game has the *finite improvement property* (FIP) if it does not permit an infinite sequence of better responses of this kind.

By definition, if a profile \mathbf{a} does *not* admit a *better response*, then that profile \mathbf{a} is a *pure Nash equilibrium*. Thus:

Observation 4 *Every game with the FIP has a pure Nash equilibrium.*

The converse is *not* true (as we are going to see next).

Exercise: Better-Response Dynamics

For the games below, all pure Nash equilibria are shown in boldface:

		C		D										
C	<table style="border-collapse: collapse; width: 100%; height: 100%;"> <tr> <td style="border: none;"></td> <td style="border: none; color: teal;">-10</td> <td style="border: none; color: teal;">0</td> </tr> <tr> <td style="border: none; color: red;">-10</td> <td style="border: 1px solid black; padding: 5px;"></td> <td style="border: 1px solid black; padding: 5px; color: red;">-25</td> </tr> <tr> <td style="border: none;"></td> <td style="border: none; color: teal;">-25</td> <td style="border: none; color: teal;">-20</td> </tr> </table>		-10	0	-10		-25		-25	-20				
	-10	0												
-10		-25												
	-25	-20												
D		0												

		L			M			R		
T		-1	1	-5						
C		1	-1	-5						
B		-5	-5	5						

		H		T	
H		-1	1		
T		1	-1		

Suppose we start in the upper lefthand cell and players keep playing better (or best) responses. *What will happen?*

Finite Improvement Property

Potential and congestion games not only all have pure Nash equilibria, but it also is natural to believe players will actually find them ...

Theorem 5 (Monderer and Shapley, 1996) *Every potential game has the FIP. Thus, also every congestion game has the FIP.*

Proof: By definition of the potential P , we get $P(\mathbf{a}^k) > P(\mathbf{a}^{k-1})$ for any two consecutive action profiles in a better-response sequence.


The claim then follows from finiteness. ✓

D. Monderer and L.S. Shapley. Potential Games. *Games and Economic Behavior*, 14(1):124–143, 1996.

Price of Anarchy

So: in a congestion game, the natural better-response dynamics will always lead us to a pure NE. Nice. But: how good is that equilibrium?

Recall our traffic congestion example:



10 people overall
 top delay = 10 minutes
 bottom delay = # on route

If $x \leq 10$ players use bottom route, the *sum of costs* is:

$$x \cdot x + (10 - x) \cdot 10 = x^2 - 10x + 100$$

This term is minimal for $x = 5$ and maximal for $x = 0$ and $x = 10$.


In equilibrium, 9 or 10 people will use the bottom route (10 is worse).

The so-called the *price of anarchy* of this game is: $\frac{100}{75} = \frac{4}{3}$.

Thus: not perfect, but not too bad either (for this example).

Braess' Paradox

Something to think about. 10 people have to get from A to D :



If the delay-free link from B to C is not present:

- In equilibrium, 5 people will use the top route $A-B-D$ and 5 people the bottom route $A-C-D$. Everyone will take **15 minutes**.

Now, if we add the delay-free link from B to C , this happens:

- In the worst equilibrium, everyone will take the route $A-B-C-D$ and take **20 minutes!** (Other equilibria are only slightly better.)

Summary

We have analysed a specific class of games, the *congestion games*:

- natural model for *applications* (but: note *restrictions*)
- *Rosenthal's Theorem*: every congestion game has a *pure* NE
- analysis via the more general *potential games*
- *finite improvement property*: can find NE via *better responses*
- *price of anarchy*: how much worse if we don't impose outcome?
- beware of counter-intuitive effects (*Braess' paradox*)

What next? Other solution concepts, besides the Nash equilibrium.