Strategic Games and Nash Equilibrium

Algorithmic Game Theory

Winter 2024/25

Prisoner's Dilemma

	S	C
S	2, 2	5, 1
C	1, 5	4, 4

- Two criminals interrogated separately.
- Strategies: (C)onfess, remain (S)ilent
- Confessing yields a smaller verdict if the other one is silent
- If both confess, the verdict is larger for both (4 years) compared to when they both remain silent (2 years).

Prisoner's Dilemma

	S	C
S	2 2	1 5
C	5 1	4 4

- If both players remain (S)ilent, the total cost is smallest.
- If both players (C)onfess, the cost is larger for both of them.
- Still, for each player confessing is always the preference!

Normal Form Games

Definition

A normal form game is a triple $(\mathcal{N}, (S_i)_{i\in N}, (c_i)_{i\in N})$ where

- $\mathcal{N}$ is the set of players, $n = |\mathcal{N}|$,
- S_i is the set of (pure) strategies of player i,
- $S = S_1 \times \ldots \times S_n$ is the set of states,
- a state is $s = (s_1, \ldots, s_n) \in S$,
- $c_i : S \to \mathbb{R}$ is the cost function of player $i \in \mathcal{N}$. In state s player i has a cost of $c_i(s)$.

We denote by $s_{-i} = (s_1, ..., s_{i-1}, s_{i+1}, ..., s_n)$ a state s without the strategy s_i.

Dominant Strategies

Definition

A pure strategy s_i is called a dominant strategy for player $i \in \mathcal{N}$ if $c_i(s_i, s_{-i}) \leq c_i(s_i', s_{-i})$ for every $s_i' \in S_i$ and every s_{-i}.

Definition

A state $s = (s_1, \ldots, s_n)$ is called a dominant strategy equilibrium if for every player $1 \leq i \leq n$ strategy $s_i \in S_i$ is a dominant strategy.

Does every game have a dominant strategy equilibrium? No!

Pareto Optimum

Definition

A state s Pareto-dominates another state s' (or: s is a Pareto improvement over s') if $c_i(s) \leq c_i(s')$ for every player $i \in \mathcal{N}$ and $c_j(s) < c_j(s')$ for at least one player $j \in \mathcal{N}$.

Definition

A state s is called a Pareto optimum or Pareto efficient if there is no state that Pareto dominates s.

In a Pareto optimum a player might be able to strictly decrease its cost by deviating – however, no player can strictly decrease its cost without strictly increasing the cost of another player.

Does every game have a Pareto optimum? Yes!

Battle of the Sexes

	(P)laza	(A)lemannia
(P)laza	1 2	6 6
(A)lemannia	5 5	2 1

- In state (P,P) the preference for both is (P)laza.
- In state (A,A) the preference for both is (A)lemannia.

$\Rightarrow$ No global preference.

What is a plausible outcome in this situation?

Pure Nash Equilibrium

Definition

A strategy s_i is called a best response against a collection of strategies s_{-i} if $c_i(s_i, s_{-i}) \leq c_i(s_i', s_{-i})$ for all $s_i' \in S_i$.

Note: s_i dominant strategy $\Leftrightarrow$ s_i best response for all s_{-i}.

Definition

A state $s = (s_1, \ldots, s_n)$ is called a pure Nash equilibrium if s_i is a best response against the other strategies s_{-i} for every player $1 \leq i \leq n$.

A Nash equilibrium

- ... is a collection of local preferences in the game.
- ... is stable against unilateral deviation.

Does every game have a pure Nash equilibrium? No!

Rock-Paper-Scissors

	R	P	S
R	0 0	-1 1	1 -1
P	1 -1	0 0	-1 1
S	-1 1	1 -1	0 0

Rock-Paper-Scissors

	R	P	S
R	0 ↓→ 0	-1 ↓ 1	1 ← -1
P	1 → -1	0 ↓→ 0	-1 ↑ 1
S	-1 ↑ 1	1 ← -1	0 ←↑ 0

Mixed Strategies

Definition

A mixed strategy x_i for player i is a probability distribution over the set of pure strategies S_i.

For finite games x_i is such that $x_{ij} \in [0,1]$ and $\sum_{j \in S_i} x_{ij} = 1$.

The cost of a mixed state for player i is

$$c_i(x) = \sum_{s \in S} p(s) \cdot c_i(s) \quad ,$$

where

$$p(s) = \prod_{i \in \mathcal{N}, j = s_i} x_{ij}$$

is the probability that the outcome is pure state s.

Mixed Nash Equilibrium

Definition

A (mixed) best response strategy x_i against a collection of mixed strategies x_{-i} is such that $c(x_i, x_{-i}) \leq c_i(x_i', x_{-i})$ for all other mixed strategies x_i'.

Definition

A mixed state x is called a (mixed) Nash equilibrium if x_i is a best response strategy against x_{-i} for every player $1 \leq i \leq n$.

Note:

- Every pure strategy is also a mixed strategy.
- Every pure Nash equilibrium is also a mixed Nash equilibrium.

Example

	0.3	0.7	
0.2	1 | 2	2 | 3	$0.3 \cdot 1 + 0.7 \cdot 2 = 0.3 + 1.4 = 1.7$
0.8	1 | 4	5 | 2	$0.3 \cdot 1 + 0.7 \cdot 5 = 0.3 + 3.5 = 3.8$
	$0.2 \cdot 2 + 0.8 \cdot 4 = 0.4 + 3.2 = 3.6$	$0.2 \cdot 3 + 0.8 \cdot 2 = 0.6 + 1.6 = 2.2$	

- $c_1(x) = 1.7 \cdot 0.2 + 3.8 \cdot 0.8 > 1.7$ – best response is $(1, 0)$
- $c_2(x) = 3.6 \cdot 0.3 + 2.2 \cdot 0.7 > 2.2$ – best response is $(0, 1)$
- State x with $x_1 = (0.2, 0.8)$ and $x_2 = (0.3, 0.7)$ is no mixed Nash equilibrium.

Observation

In the previous example x is not a mixed Nash equilibrium, because players play suboptimal pure strategies with positive probability.

Fact

If a mixed best response x_i against x_{-i} has $x_{ij} > 0$, then j is a pure best response against x_{-i}.

The cost of x_i is a "weighted average" of the cost of the pure strategies. It is minimal if and only if the averaging is just over pure strategies with minimum cost.

Example

	1	0	
1	2 / 1	3 / 2	$1 \cdot 1 + 0 \cdot 2 = 1$
0	4 / 1	2 / 5	$1 \cdot 1 + 0 \cdot 5 = 1$
	$1 \cdot 2 + 0 \cdot 4 = 2$	$1 \cdot 3 + 0 \cdot 2 = 3$	

- State x with $x_1 = (1,0)$ and $x_2 = (1,0)$ is a pure (and hence also a mixed) Nash equilibrium.

Example

	1	0	
$\frac{2}{3}$	2 / 1	3 / 2	$1 \cdot 1 + 0 \cdot 2 = 1$
$\frac{1}{3}$	4 / 1	2 / 5	$1 \cdot 1 + 0 \cdot 5 = 1$
	$\frac{2}{3} \cdot 2 + \frac{1}{3} \cdot 4 = \frac{8}{3}$	$\frac{2}{3} \cdot 3 + \frac{1}{3} \cdot 2 = \frac{8}{3}$	

- State x with $x_1 = (\frac{2}{3}, \frac{1}{3})$ and $x_2 = (1, 0)$ is a mixed Nash equilibrium.
- For the row player the upper strategy is a dominant strategy, but in the first column it is not *strictly* better. If it was strictly better in every column, the lower strategy would not be played in any mixed Nash equilibrium. (Why?)

Nash Theorem

Theorem (Nash Theorem)

Every finite normal form game has a mixed Nash equilibrium.

We will use Brouwer's fixed point theorem to prove it.

Theorem (Brouwer Fixed Point Theorem)

Every continuous function $f : D \to D$ mapping a compact and convex nonempty subset $D \subseteq \mathbb{R}^m$ to itself has a fixed point $x^* \in D$ with $f(x^*) = x^*$.

Brouwer's Theorem: Prerequisites and Definitions

- A set $D \subset \mathbb{R}^m$ is convex if for any $x, y \in D$ and any $\lambda \in [0,1]$ we have $\lambda x + (1-\lambda)y \in D$.
- A subset $D \subset \mathbb{R}^m$ is compact if and only if it is closed and bounded.
- A set $D \subseteq \mathbb{R}^m$ is bounded if and only if there is some integer $M \geq 0$ such that $D \subseteq [-M, M]^m$.
- Consider a set $D \subseteq \mathbb{R}^m$ and a sequence $x_0, x_1, \ldots$, where for all $i \geq 0$, $x_i \in D$ and there is $x \in \mathbb{R}^m$ such that $x = \lim_{i \to \infty} x_i$ (i.e., for all $\epsilon > 0$ there is integer $k > 0$ such that $||x - x_j||_2 < \epsilon$ for all $j > k$). A set D is closed if $x \in D$ for every such sequence.
- A function $f : D \to \mathbb{R}^m$ is continuous at a point $x \in D$ if for all $\epsilon > 0$, there exists $\delta > 0$, such that for all $y \in D$: If $||x - y||_2 < \delta$ then $||f(x) - f(y)||_2 < \epsilon$. f is called continuous if it is continuous at every point $x \in D$.

Brouwer's Theorem: Prerequisites and Examples

- Convex/Non-convex:

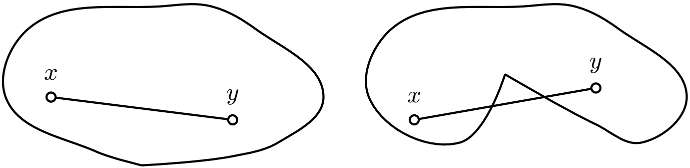

- Closed and bounded:
  $[0,1]^2$ is closed and bounded.
  $[0,1)$ is not closed but bounded.
  $[0,\infty)$ is closed and unbounded.

- Continuous: Clear.

Brouwer's Theorem: Example

Every continuous $f : [0,1] \to [0,1]$ has a fixed point:

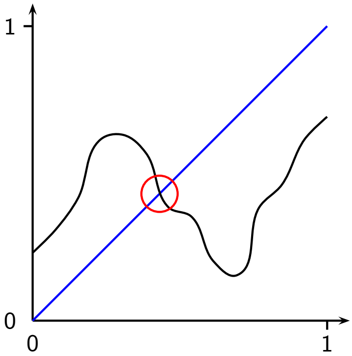

For $f : [0,1]^2 \to [0,1]^2$: Crumpled Sheet Experiment

Nash Theorem

Theorem (Nash Theorem)

Every finite normal form game has a mixed Nash equilibrium.

Proof:

First check the conditions of Brouwer's Theorem.

Fact

The set X of mixed states $x = (x_1, \ldots, x_n)$ of a finite normal form game is a convex compact subset of $\mathbb{R}^m$ with $m = \sum_{i=1}^{n} m_i$ with $m_i = |S_i|$.

We will define a function $f : X \to X$ that transforms a state into another state. The fixed points of f are shown to be the mixed Nash equilibria of the game.

Properties of Nash Equilibria

Recall:

- A mixed Nash equilibrium x is a collection of mixed best responses x_i.
- If a best response x_i against x_{-i} has $x_{ij} > 0$, then $j \in S_i$ is pure best response against x_{-i}.
- A collection of best responses (i.e. a mixed Nash equilibrium) $x = (x_1, ..., x_n)$ has

$$c_i(x) - c_i(j, x_{-i}) \leq 0 \qquad \text{for all } j \in S_i \text{ and all } i \in \mathcal{N}$$

Proof of Nash's Theorem: Definition

- For mixed state x let

$$\phi_{ij}(x) = \max\{0, c_i(x) - c_i(j, x_{-i})\} \ .$$

- Define $f : X \to X$ with $f(x) = x' = (x'_1, ..., x'_n)$ by

$$x'_{ij} = \frac{x_{ij} + \phi_{ij}(x)}{1 + \sum_{k=1}^{m_i} \phi_{ik}(x)}$$

for all $i = 1, \ldots, n$ and $j = 1, ..., m_i$.

Fact

f satisfies the prerequisites of Brouwer's Theorem: f is continuous and if $x \in X$, then $f(x) = x' \in X$ is a mixed state.

(Check as an exercise.)

Example

- Player i has 3 pure strategies
- Current mixed strategy $x_i = (0.2, 0.5, 0.3)$
- Current costs for strategies $c_i(\cdot, x_{-i}) = (2.2, 4.2, 2.2)$
- Current cost $c(x_i, x_{-i}) = 3.2$
- Under these conditions strategy x_i is mapped to x'_i as follows:

x_{ij}	$c_i(j, x_{-i})$	$\phi_{ij}(x)$	x'_{ij}
0.2	2.2	1	$\frac{0.2+1}{1+2} = 0.4$
0.5	4.2	0	$\frac{0.5+0}{1+2} \approx 0.166$
0.3	2.2	1	$\frac{0.3+1}{1+2} \approx 0.434$

Fixed Points

Brouwers Theorem tells us that there is x^* with $f(x^*) = x^*$. We need to show two directions:

$$f(x) = x \qquad \Leftrightarrow \qquad x \text{ is mixed Nash equilibrium.}$$

Easy: x is mixed Nash $\Rightarrow$ $f(x) = x$: All $\phi_{ij}(x) = 0$.

To show: $x^* = f(x^*)$ $\Rightarrow$ x^* is a mixed Nash equilibrium.

Fixed Points as Nash Equilibria

For each $i = 1, \ldots, n$ and $j = 1, \ldots, m_i$ we have

$$x^*_{ij} = \frac{x^*_{ij} + \phi_{ij}(x^*)}{1 + \sum_{k=1}^{m_i} \phi_{ik}(x^*)} \ ,$$

so

$$x^*_{ij} \cdot \left(1 + \sum_{k=1}^{m_i} \phi_{ik}(x^*)\right) = x^*_{ij} + \phi_{ij}(x^*) \ ,$$

and

$$x^*_{ij} \sum_{k=1}^{m_i} \phi_{ik}(x^*) = \phi_{ij}(x^*) \ .$$

We will show that $\sum_{k=1}^{m_i} \phi_{ik}(x^*) = 0$. This means that x^*_i chooses only pure best responses and implies that it is a mixed best response.

Fixed Points as Nash Equilibria

Claim

For every mixed state x and every player $i \in \mathcal{N}$, there is some pure strategy $j \in S_i$ such that $x_{ij} > 0$ and $\phi_{ij}(x) = 0$.

Proof of Claim:

Note that $c_i(x) = \sum_{j=1}^{m_i} x_{ij} \cdot c_i(j, x_{-i})$, so there must be some j with $x_{ij} > 0$ and cost no less than this "weighted average".

More formally, there is j with $x_{ij} > 0$ and

$$c_i(x) - c_i(j, x_{-i}) \leq 0 \ .$$

Therefore, $\phi_{ij}(x) = \max\{0, c_i(x) - c_i(j, x_{-i})\} = 0$. □

Fixed Points as Nash Equilibria

For every player i we consider strategy j from the claim. This implies $x^*_{ij} > 0$ and

$$x^*_{ij} \cdot \sum_{k=1}^{m_i} \phi_{ik}(x^*) = \phi_{ij}(x^*) = 0 \ .$$

Since $x^*_{ij} > 0$ it must hold that

$$\sum_{k=1}^{m_i} \phi_{ik}(x^*) = 0 \ ,$$

so $\phi_{ik}(x^*) = 0$ for all $k = 1, \dots, m_i$. Therefore

$$c_i(x^*) \le c_i(j, x^*_{-i}) \qquad \text{for all } j \in S_i.$$

Hence, x^*_i is a best response. This proves Nash's Theorem. □

Normal Form Games

PPAD

Zero-Sum Games

Appendix: LP Duality

Computing Nash Equilibria

How can we compute a mixed Nash equilibrium?
What is the complexity of computing a Nash equilibrium?

This problem is different from problems we usually encounter:

- No optimization, trivial as decision problem (existence guaranteed)
- Search problem, find Nash equilibrium.
- Different Complexity Class: PPAD (polynomial parity argument, directed case)
- A notion of completeness, similar to NP: Define PPAD-complete problem, construct polynomial-time reductions

There are 3-player games with rational payoff, in which all mixed Nash equilibria have irrational probability values. Thus, we can only hope to obtain approximations to mixed Nash equilibria or Brouwer fixed points.

A PPAD-Complete Problem

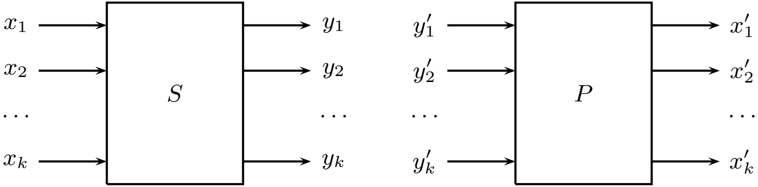

An instance of the END-OF-LINE search problem is given by

- Two circuits S and P, same number of inputs and output bits
- S and P define a directed graph:
  Vertices: k-bit vectors
  Edges: There is a directed edge (x, y) if $S(x) = y$ and $P(y) = x$
- S and P are such that the **all-0-vector** has one outgoing edge and no incoming edge!

Problem: Find a different source or sink in the graph.

END-OF-LINE

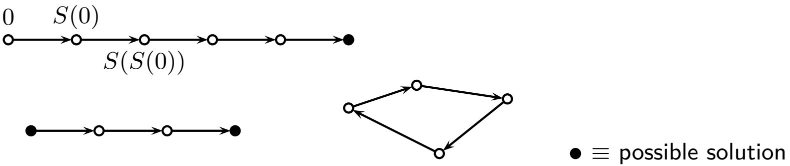

Observations:

- Every vertex in the graph has indegree and outdegree at most 1.
- By parity argument END-OF-LINE always admits a solution.
- Not necessarily the end of the line from 0, finding this specific sink is PSPACE-complete.
- Only **circuits are the input**! The graph is exponentially large in the input size. It cannot be fully enumerated in polynomial time.

Computing a solution to END-OF-LINE is PPAD-complete.
It is believed that there is no efficient algorithm for this problem.

Finding (Approximate) Brouwer Fixed Points

Lemma

Finding an (approximate) mixed Nash equilibrium is in PPAD.

Proof Sketch:

- Reduction: Finding fixed points with END-OF-LINE
- Subdivide the space into finite number of smaller areas
- Find an area close to a fixed point (Approximation)
- By continuity: Finer granularity yields more precise approximation.

Divide the space into simplices ("multidimensional triangles") and color vertices according to direction of Brouwer function

For simplicity of presentation we here consider only problems with $D \subseteq \mathbb{R}^2$, e.g., $f : [0,1]^2 \to [0,1]^2$.

Triangles

For simplicity we transform representation of $[0,1]^2$ to a triangle T. Equivalent fixed point problem with $f' : T \to T$.

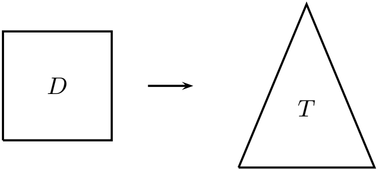

Subdivision and Coloring

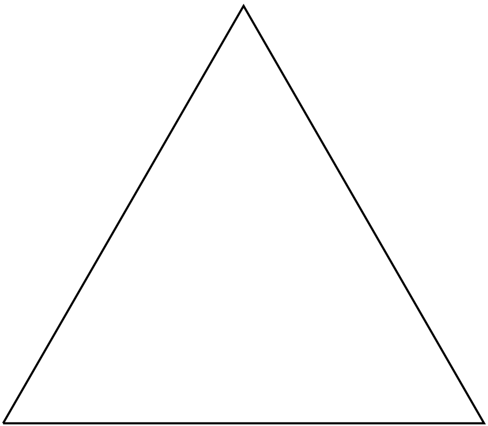

- The trianlge space T is subdivided into smaller triangles

Subdivision and Coloring

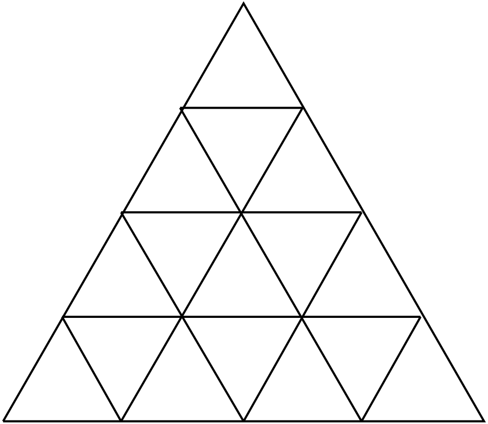

- The trianlge space T is subdivided into smaller triangles
- For each vertex consider the direction, in which f' maps the point

Subdivision and Coloring

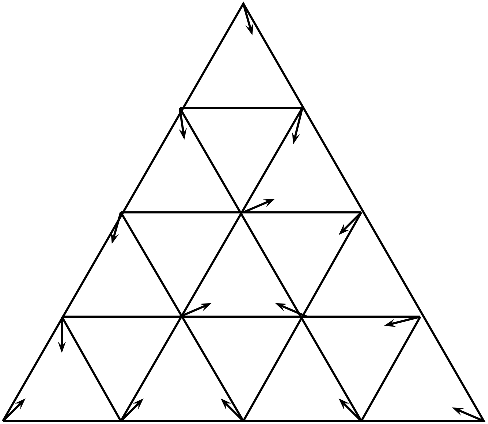

- The trianlge space T is subdivided into smaller triangles
- For each vertex consider the direction, in which f' maps the point
- Depending on the direction the vertex receives a color.

Subdivision and Coloring

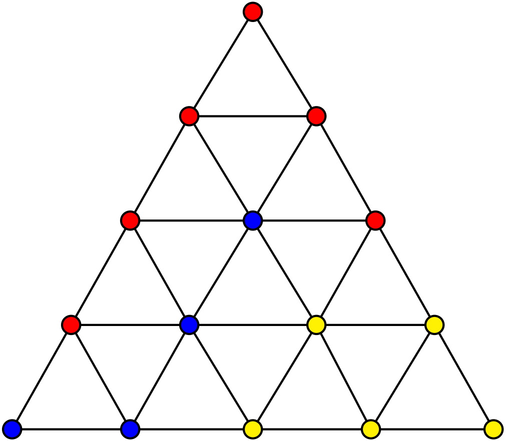

- The trianlge space T is subdivided into smaller triangles
- For each vertex consider the direction, in which f' maps the point
- Depending on the direction the vertex receives a color.
- With increasing granularity trichromatic triangles become the fixed points of f'.

Sperner Coloring

Definition

A subdivided triangle is a division of a triangle into smaller triangles.

Definition

A Sperner coloring of the vertices of a subdivided triangle satisfies:

- Each extremal vertex gets a different color.
- A vertex on a side of the largest triangle gets a color of one of the corresponding endpoints.
- Other vertices are colored arbitrarily.

Verify that our coloring based on directions of f' yields a Sperner coloring.

Sperner's Lemma

Lemma (Sperner's Lemma)

Every Sperner coloring of a subdivided triangle contains a trichromatic triangle.

Proof:

- Connect all vertices on the outer blue/red edge to the blue vertex. Start at the outside face and move over lines connecting a red and a blue vertex. There are at most 2 such lines in each triangle, never visit a triangle twice.
- This implies an instance of END-OF-LINE:
  Vertices: Small triangles
  Edges: There is an edge if two triangles share a line between a red and blue vertex.
- By construction indegree and outdegree at most 1
- There is a starting point by creation, other sources/sinks are the trichromatic triangles. □

Proof by END-OF-LINE

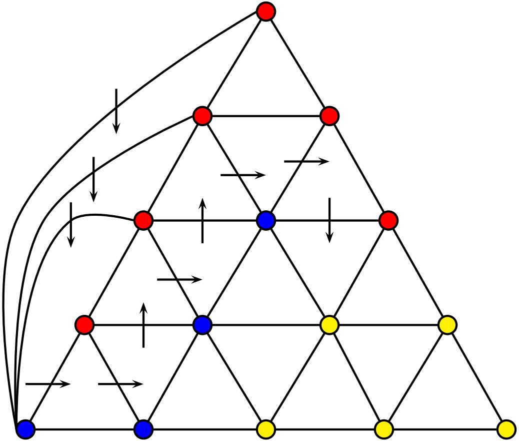

Implications and Results

Sperner's Lemma is a *discretized version of Brouwer's fixed point theorem*. The proof of the lemma...

- shows that Sperner colorings create an instance of END-OF-LINE.
- can be generalized to more dimensions and simplicies instead of triangles. Then trichromatic triangles correspond to simplicies with maximum number of colors.
- with "infinite granularity" implies maximally colored simplicies are Brouwer fixed points.

This proves that finding a Brouwer fixed point and, hence, a mixed Nash equilibrium in a finite game is in PPAD. □

Fundamental result in the literature by Daskalakis/Papadimitriou and Chen/Deng/Teng:

Theorem

Finding a mixed Nash equilibrium in a finite 2-player game is PPAD-complete.

Definition

For consistency with literature we here consider utility functions instead of cost.

Definition

The utility of player i in a state s of a normal form game is $u_i(s) = -c_i(s)$.

Definition

A zero-sum game is a strategic game, in which for every state s we have $\sum_{i \in \mathcal{N}} u_i(s) = 0$.

In a zero-sum game every utility gain of one player results in a utility loss of another player. For instance, this can be used to model situations in which players must divide a common good.

2-player Zero Sum Games

Two players, player I (row player), player II (column player)

Representation as a matrix $A \in \mathbb{R}^{k \times \ell}$ with $k = |S_{\text{I}}|$ rows and $\ell = |S_{\text{II}}|$ columns:

$$\begin{pmatrix} a_{11} & a_{12} & . & . & . & a_{1\ell} \\ a_{21} & a_{22} & . & . & . & a_{2\ell} \\ . & . & . & & & . \\ . & . & & . & & . \\ . & . & & & . & . \\ a_{k1} & a_{k2} & . & . & . & a_{k\ell} \end{pmatrix}$$

a_{ij} is utility for player I
$-a_{ij}$ is utility for player II

Examples

Matching Pennies

$$\begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$$

Rock-Paper-Scissors

$$\begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix}$$

A game with $k \neq \ell$:

$$\begin{pmatrix} 0 & 2 & 4 \\ 1 & 2 & 3 \end{pmatrix}$$

Utility by Matrix Multiplication

We denote mixed strategies by x for I and y for II.

Computing the utility $u_{\mathcal{I}}(x,y)$:

	0.1	0.4	0.5
0.8	0	2	4
0.2	1	2	3

$$\longrightarrow \quad \begin{aligned} & 0.8 \cdot (0.1 \cdot 0 + 0.4 \cdot 2 + 0.5 \cdot 4) \\ & + 0.2 \cdot (0.1 \cdot 1 + 0.4 \cdot 2 + 0.5 \cdot 3) \\ & = 2.48 \end{aligned}$$

$$\begin{aligned} u_{\mathcal{I}}(x,y) = -u_{\text{II}}(x,y) &= \sum_{i=1}^{k} \sum_{j=1}^{\ell} x_i a_{ij} y_j \\ &= \begin{pmatrix} x_1 & x_2 \end{pmatrix} \cdot \begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{pmatrix} \cdot \begin{pmatrix} y_1 \\ y_2 \\ y_3 \end{pmatrix} \\ &= x^T A y \end{aligned}$$

Public Strategy Choices

Suppose player I has to decide first. He must pick a public strategy before player II makes his choice. How should I choose his public strategy?

$$\begin{pmatrix} 0 & 2 & 4 \\ 1 & 2 & 3 \end{pmatrix}$$

Player II will hurt player I as much as possible.

In this game II will always answer with column 1. Hence, optimal choice for I is pure strategy 2 or $x = (0, 1)^T$.

Maximin Strategies

Player I picks x, then player II best responds with y. II solves the problem $\max_y u_{\mathtt{II}}(x,y) = \max_y -x^T Ay = \min_y x^T Ay$.

Hence, player I searches for x that maximizes $\min_y x^T Ay$.

Definition

The gain-floor of a 2-player zero-sum game is

$$v_{\mathtt{I}}^* = \max_x \min_y x^T Ay \quad .$$

A strategy x^* that yields the gain-floor is an optimal strategy for I, called maximin strategy.

Example Maximin

$$\begin{pmatrix} 5 & 1 & 2 \\ 1 & 4 & 3 \end{pmatrix}$$

Player $\texttt{II}$ will hurt player $\texttt{I}$ as much as possible.

- $\texttt{I}$ picks row 1 $\Rightarrow$ $\texttt{II}$ picks column 2 $\Rightarrow$ $\texttt{I}$ gets utility 1
- $\texttt{I}$ picks row 2 $\Rightarrow$ $\texttt{II}$ picks column 1 $\Rightarrow$ $\texttt{I}$ gets utility 1

- $\texttt{I}$ picks $x = (0.5, 0.5)$, minimum loss for $\texttt{II}$ is 2.5 in columns 2 and 3

$\Rightarrow$ $\texttt{I}$ gets utility 2.5!

What is x^*, how large can $v_{\texttt{I}}^*$ be?

Dual Perspective: Minimax

Now suppose player II first picks y, then player I picks x optimally with $\max_x u_{\mathtt{I}}(x, y) = \max_x x^T Ay$.

Hence, II searches for y that minimizes $\max_x x^T Ay$.

Definition

The loss-ceiling of a 2-player zero-sum game is

$$v^*_{\mathtt{II}} = \min_y \max_x x^T Ay \quad .$$

A strategy y^* that yields the loss-ceiling is an optimal strategy for II, called minimax strategy.

What is y^*, how small can $v^*_{\mathtt{II}}$ be?

How do $v^*_{\mathtt{I}}$ and $v^*_{\mathtt{II}}$ compare?

Minimax Theorem

Intuitively, if both players play optimally, player I should gain at least $v_{\mathtt{I}}^*$, player II should not lose more than $v_{\mathtt{II}}^*$. It is easy to show that

Lemma

It holds that $v_{\mathtt{I}}^* \leq v_{\mathtt{II}}^*$.

Perhaps surprisingly, von Neumann and Morgenstern proved

Theorem (Minimax Theorem)

In every 2-player zero-sum game it holds that $v = v_{\mathtt{I}}^* = v_{\mathtt{II}}^*$. *The value* v *is called the* ***value*** *of the game.*

Minimax Theorem by Linear Programming Duality

Consider the optimization problem to find x^* and $v_{\mathtt{I}}^* = \max_x \min_y x^T Ay$.

Observe:

- $\mathtt{II}$ plays a best response y against x.
- For a given x, player $\mathtt{II}$ sets $y_j > 0$ if and only if his expected loss $\sum_{i=1}^k x_i a_{ij}$ in column j is minimal.
- Hence,

$$v_{\mathtt{I}} = \sum_{j=1}^{\ell} \sum_{i=1}^{k} x_i a_{ij} y_j = \min_{j=1}^{\ell} \sum_{i=1}^{k} x_i a_{ij}$$

- For any x and the resulting utility $v_{\mathtt{I}}$ obtained by $\mathtt{I}$ we thus know

$$v_{\mathtt{I}} \leq \sum_{i=1}^{k} x_i a_{ij} \qquad \text{for all } j = 1, \ldots, \ell.$$

Gain-Floor Optimization as a Linear Program

$$
\begin{aligned}
\text{Maximize} \quad & v_{\mathtt{I}} & & \\
\text{subject to} \quad & v_{\mathtt{I}} - \sum_{i=1}^{k} x_i a_{ij} \leq 0 & & \text{for all } j = 1, \ldots, \ell \\
& \sum_{i=1}^{k} x_i = 1 & & \\
& x_i \geq 0 & & \text{for all } i = 1, \ldots, k \\
& v_{\mathtt{I}} \in \mathbb{R} & &
\end{aligned} \tag{1}
$$

Loss-Ceiling Optimization as a Linear Program

Similar arguments yield a linear program for loss-ceiling minimization.

$$
\begin{array}{llcll}
\text{Minimize} & v_{\text{II}} & & & \\
\text{subject to} & v_{\text{II}} - \sum_{j=1}^{\ell} a_{ij} y_j & \geq & 0 & \text{for all } i = 1, \ldots, k \\
& \sum_{j=1}^{\ell} y_j & = & 1 & \\
& y_j & \geq & 0 & \text{for all } j = 1, \ldots, \ell \\
& v_{\text{II}} & \in & \mathbb{R} &
\end{array}
\qquad (2)
$$

In the appendix we show that this represents the LP-dual of the Gain-Floor LP (1).

Implications

Finding optimal strategies for players I and II can be formulated as dual linear programs.

Strong duality in Linear Programming:

- Consider a linear program with a feasible optimum solution
- Let f^* be the optimal objective function value
- Then the dual has a feasible optimum solution, objective function value g^*
- Strong Duality: It holds that $f^* = g^*$.

Thus, strong duality proves the minimax theorem. □

Example

$$\begin{pmatrix} 5 & 1 & 2 \\ 1 & 4 & 3 \end{pmatrix}$$

$$\begin{array}{llrcl} \text{Max.} & v_{\mathrm{I}} & & & \\ \text{s.t.} & v_{\mathrm{I}} - 5x_1 - 1x_2 & \le & 0 \\ & v_{\mathrm{I}} - 1x_1 - 4x_2 & \le & 0 \\ & v_{\mathrm{I}} - 2x_1 - 3x_2 & \le & 0 \\ & x_1 + x_2 & = & 1 \\ & x_1, x_2 & \ge & 0 \\ & v_{\mathrm{I}} & \in & \mathbb{R} \end{array}$$

$$x^* = (0.4, 0.6)$$
$$v_{\mathrm{I}}^* = 2.6$$

$$\begin{array}{llrcl} \text{Min.} & v_{\mathrm{II}} & & & \\ \text{s.t.} & v_{\mathrm{II}} - 5y_1 - 1y_2 - 2y_3 & \ge & 0 \\ & v_{\mathrm{II}} - 1y_1 - 4y_2 - 3y_3 & \ge & 0 \\ & y_1 + y_2 + y_3 & = & 1 \\ & y_1, y_2, y_3 & \ge & 0 \\ & v_{\mathrm{II}} & \in & \mathbb{R} \end{array}$$

$$y^* = (0.2, 0, 0.8)$$
$$v_{\mathrm{II}}^* = 2.6$$

Is (x^*, y^*) a mixed Nash equilibrium?

Mixed Nash equilibrium

Corollary

A state (x, y) in a 2-player zero-sum game is a mixed Nash equilibrium
$\Leftrightarrow$ *x and y are optimal strategies for the players.*

Proof ($\Rightarrow$):

- Consider (x, y) and assume x is suboptimal (similar for y suboptimal)
- There is y' that achieves $u_{\texttt{II}}(x, y') > -v$, thus $u_{\mathcal{I}}(x, y') < v$.
- To be NE we must have $u_{\texttt{II}}(x, y) \geq u_{\texttt{II}}(x, y')$, so $u_{\texttt{I}}(x, y) < v$.
- If $u_{\texttt{I}}(x, y) < v$, $\texttt{I}$ can improve by optimal strategy $\Rightarrow$ (x, y) no mixed NE.

($\Leftarrow$):

- Suppose both play optimal, but player $\texttt{I}$ has a better strategy x'
- This means $u_{\texttt{I}}(x', y) > v$, but then y is suboptimal for $\texttt{II}$
- Same argument for player $\texttt{II}$ having a better strategy.

Mixed Nash equilibrium

Corollary

All mixed Nash equilibria in a 2-player zero-sum game yield an expected utility of v $(-v)$ *for player* I *(*II*).*

We can find optimal strategies by solving the linear programs (1) and (2). There are efficient algorithms for solving linear programs, which proves the following result:

Theorem

In 2-player zero-sum games a mixed Nash equilibrium can be computed in polynomial time.

Literature

- G. Owen. Game Theory. Academic Press, 2001. (Chapters 1 + 2)
- Chapters 1 and 2 in the AGT book
- J. Nash. Non-cooperative Games. Annals of Mathematics 54, pp. 286–295, 1951.
- P. Goldberg, C. Daskalakis, C. Papadimitriou. The Complexity of Computing a Nash Equilibrium. SIAM Journal on Computing, 39(1), pp. 195–259, 2009.
- X. Chen, X. Deng, S.-H. Teng. Settling the Complexity of Computing Two-Player Nash Equilibria. Journal of the ACM, 56(3), 2009.
- For background on linear programming, duality, and algorithms see: Cormen, Leiserson, Rivest, Stein. Introduction to Algorithms, 3rd edition. MIT Press, 2009. (Chapter 29)

Constructing the Dual

We construct an upper bound on v_{I} for every solution of (1).

- Consider a solution (v_{I}, x) of (1).
- We take a linear combination of the constraints to construct an upper bound. In particular, we use multipliers z_j and $w_{\mathcal{I}}$:

$$\begin{aligned} z_j \cdot \left(v_{\mathrm{I}} - \sum_{i=1}^k x_i a_{ij}\right) &\leq z_j \cdot 0 \quad \text{for each } j \text{ and} \\ w_{\mathrm{I}} \cdot \sum_{i=1}^k x_i &= w_{\mathcal{I}} \cdot 1 \end{aligned}$$

Here $z_j \geq 0$ to keep the correct inequality.

Constructing the Dual

- Now we try to get an upper bound by using the linear combination:

$$\begin{aligned} v_{\mathrm{I}} &\leq \sum_{j=1}^{\ell} z_j \left(v_{\mathrm{I}} - \sum_{i=1}^{k} x_i a_{ij} \right) + w_{\mathrm{I}} \cdot \sum_{i=1}^{k} x_i \\ &= \left(\sum_{j=1}^{\ell} z_j \right) \cdot v_{\mathrm{I}} + \sum_{i=1}^{k} \left(w_{\mathrm{I}} - \sum_{j=1}^{\ell} a_{ij} z_j \right) \cdot x_i \\ &\leq \sum_{j=1}^{\ell} z_j \cdot 0 + w_{\mathrm{I}} \cdot 1 \ = \ w_{\mathrm{I}} \end{aligned}$$

- This works if the first inequality is fulfilled, and this holds if the following conditions for coefficients for the v_{I} and x_i on l.h.s. and r.h.s. are true:

$$1 = \sum_{j=1}^{\ell} z_j \qquad \text{(Same ones because } v_{\mathrm{I}} \in \mathbb{R}.\text{)}$$

$$0 \leq w_{\mathrm{I}} - \sum_{j=1}^{\ell} a_{ij} z_j \qquad \text{(Possibly larger ones, because } x_i \geq 0.\text{)}$$

What is the best upper bound w_{I} that can be obtained in this way?

Finding the best Upper Bound

$$\begin{array}{llcll} \text{Minimize} & w_{\mathtt{I}} & & & \\ \text{subject to} & w_{\mathtt{I}} - \sum_{j=1}^{\ell} a_{ij} z_j & \geq & 0 & \text{for all } i = 1, \ldots, k \\ & \sum_{j=1}^{\ell} z_j & = & 1 & \\ & z_j & \geq & 0 & \text{for all } j = 1, \ldots, \ell \\ & w_{\mathtt{I}} & \in & \mathbb{R} & \end{array} \quad (3)$$

This linear program is called the dual program of (1), $w_{\mathtt{I}}$ and z_j are the dual variables.

Note that this represents exactly the optimization problem to find the loss-ceiling and an optimal strategy of player II (with $w_{\mathtt{I}} = v_{\mathtt{II}}$ and $z_j = y_j$)!