Topics in Algorithmic Game Theory and Economics

Pieter Kleer

Max Planck Institute for Informatics (D1)
Saarland Informatics Campus

December 2, 2020

Lecture 4
Finite games - Existence and Computation of MNE

Finite game

Finite game $\Gamma = (N, (\mathcal{S}_i)_{i \in N}, (C_i)_{i \in N})$ consists of:

- Finite set N of players.
- Finite strategy set $\mathcal{S}_i$ for every player $i \in N$.
- Cost function $C_i : \times_j \mathcal{S}_j \to \mathbb{R}$ for every $i \in N$.

Matching pennies

Alice and Bob both choose side of a penny.

- (a, b) denotes cost for Alice (A) and Bob (B) in given profile.

		Bob	
		Head	Tails
Alice	Head	$(0, 1)$	$(1, 0)$
	Tails	$(1, 0)$	$(0, 1)$

No PNE: $(\text{Head}, \text{Head}) \xrightarrow{B} (\text{Head}, \text{Tails}) \xrightarrow{A} (\text{Tails}, \text{Tails}) \xrightarrow{B} (\text{Tails}, \text{Head}) \xrightarrow{A} (\text{Head}, \text{Head})$.

Game does have mixed Nash equilibrium (MNE).

- Both randomize over their strategies $\{\text{Head}, \text{Tails}\}$.
  - Mixed strategies $\sigma_A = (1/2, 1/2)$ and $\sigma_B = (1/2, 1/2)$.

Mixed strategies

We focus on two-player games (for sake of notation). Players are

- Row player Alice (A) with strategy set $\mathcal{S}_A = \{a_1, \ldots, a_m\}$, and
- Column player Bob (B) with strategy set $\mathcal{S}_B = \{b_1, \ldots, b_n\}$.

Definition (Mixed strategy)

A mixed strategy is a probability distribution over $\mathcal{S}_i$ for $i \in \{\text{Alice}, \text{Bob}\}$. The collection of all mixed strategies will be denoted by Δ_i, i.e.,

$$\begin{aligned} \Delta_{\text{Alice}} &= \{(x_1, \ldots, x_m) : \textstyle\sum_i x_i = 1, x_i \geq 0 \text{ for } i = 1, \ldots, m\}, \\ \Delta_{\text{Bob}} &= \{(y_1, \ldots, y_n) : \textstyle\sum_j y_j = 1, y_j \geq 0 \text{ for } j = 1, \ldots, n\}. \end{aligned}$$

- Interpretation: Alice plays strategy a_1 with prob. x_1, etc...

Example

Strategies of Alice and Bob are given by:

$$\begin{aligned} \Delta_{\text{Alice}} &= \{(x_1, x_2) : x_1 + x_2 = 1, x_1, x_2 \geq 0\}, \\ \Delta_{\text{Bob}} &= \{(y_1, y_2, y_3) : y_1 + y_2 + y_3 = 1, y_1, y_2, y_3 \geq 0\}. \end{aligned}$$

	b_1	b_2	b_3
a_1	$(0, 2)$	$(1, 0)$	$(2, 1)$
a_2	$(3, 0)$	$(0, 1)$	$(1, 4)$

$$\Delta_{\text{A(lice)}} = \{(x_1, \ldots, x_m) : \sum_i x_i = 1, x_i \geq 0 \text{ for } i = 1, \ldots, m\},$$
$$\Delta_{\text{B(ob)}} = \{(y_1, \ldots, y_n) : \sum_j y_j = 1, y_j \geq 0 \text{ for } j = 1, \ldots, n\}.$$

For $x \in \Delta_A, y \in \Delta_B$, we get product distribution $\sigma_{x,y} : \mathcal{S}_A \times \mathcal{S}_B \to [0, 1]$ over strategy profiles,

- $\sigma_{x,y}(a_k, b_\ell) = x_k y_\ell$ for $k = 1, \ldots, m$ and $\ell = 1, \ldots, n$.

Example (cont'd)

Distribution over strategy profiles is given by

$$\begin{pmatrix} x_1y_1 & x_1y_2 & x_1y_3 \\ x_2y_1 & x_2y_2 & x_2y_3 \end{pmatrix}$$

	b_1	b_2	b_3
a_1	$(0, 2)$	$(1, 0)$	$(2, 1)$
a_2	$(3, 0)$	$(0, 1)$	$(1, 4)$

Then expected cost $C_i(\sigma_{x,y}) = C_i(x, y)$, of $i \in \{\text{Alice}, \text{Bob}\}$ is

$$C_i(x, y) = \mathbb{E}_{(a_k, b_\ell) \sim \sigma_{x,y}}[C_i(a_k, b_\ell)] = \sum_{(a_k, b_\ell) \in \mathcal{S}_A \times \mathcal{S}_B} x_k y_\ell C_i(a_k, b_\ell)$$

Matrix representation

Matrix representation of cost functions $C_i : \Delta_A \times \Delta_B \to \mathbb{R}$ for $i \in \{\text{Alice}, \text{Bob}\}$ given by $A, B \in \mathbb{R}^{m \times n}$ defined as

$$A_{k\ell} = C_A(a_k, b_\ell) \text{ and } B_{k\ell} = C_B(a_k, b_\ell) \text{ for } k = 1, \ldots, m \text{ and } \ell = 1, \ldots, n.$$

Example (cont'd)

$$A = \begin{pmatrix} 0 & 1 & 2 \\ 3 & 0 & 1 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} 2 & 0 & 1 \\ 0 & 1 & 4 \end{pmatrix}.$$

	b_1	b_2	b_3
a_1	$(0, 2)$	$(1, 0)$	$(2, 1)$
a_2	$(3, 0)$	$(0, 1)$	$(1, 4)$

Expected cost under mixed strategies $x \in \Delta_A, y \in \Delta_B$ is then

$$C_{\text{Alice}}(x, y) = x^T A y = \sum_{k=1}^{m} \sum_{\ell=1}^{n} A_{k\ell} x_k y_\ell, \quad C_{\text{Bob}}(x, y) = x^T B y = \sum_{k=1}^{m} \sum_{\ell=1}^{n} B_{k\ell} x_k y_\ell$$

Short overview

Two-player game (A, B) is given by matrices $A, B \in \mathbb{R}^{m \times n}$, with player Alice choosing mixed strategy x over rows, and player Bob mixed strategy y over columns. Expected costs are given by $x^T A y$ and $x^T B y$, respectively.

Mixed Nash equilibrium

Hierarchy of equilibrium concepts

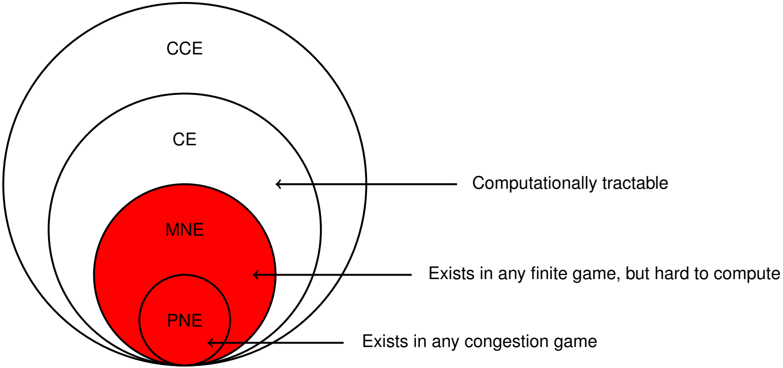

Mixed Nash equilibrium (2-player case)

For two-player game (A, B), we have

$$C_{\mathrm{A}}(x,y) = x^T A y = \sum_{k=1}^{m} \sum_{\ell=1}^{n} A_{k\ell} x_k y_\ell, \quad C_{\mathrm{B}}(x,y) = x^T B y = \sum_{k=1}^{m} \sum_{\ell=1}^{n} B_{k\ell} x_k y_\ell$$

Definition (Mixed Nash equilibrium)

Pair $(x^*, y^*) \in \Delta_A \times \Delta_B$ is mixed Nash equilibrium (MNE) if neither Alice nor Bob can deviate to other mixed strategy and improve cost:

$$\begin{aligned} C_{\mathrm{A}}(x^*, y^*) &\leq C_{\mathrm{A}}(x', y^*) \quad \forall x' \in \Delta_A \\ C_{\mathrm{B}}(x^*, y^*) &\leq C_{\mathrm{B}}(x^*, y') \quad \forall y' \in \Delta_B \end{aligned}$$

For $\epsilon > 0$, pair (x^*, y^*) is ϵ-approximate MNE (or simply ϵ-MNE) if

$$\begin{aligned} C_{\mathrm{A}}(x^*, y^*) &\leq C_{\mathrm{A}}(x', y^*) + \epsilon \quad \forall x' \in \Delta_A \\ C_{\mathrm{B}}(x^*, y^*) &\leq C_{\mathrm{B}}(x^*, y') + \epsilon \quad \forall y' \in \Delta_B \end{aligned}$$

- Will see later that is suffices to have these conditions only for pure strategies: One strategy is played with probability 1.

Example

Alice has $\mathcal{S}_A = \{a_1, a_2\}$ and $\mathcal{S}_B = \{b_1, b_2, b_3\}$.

$$A = \begin{pmatrix} 2 & 1 & 2 \\ 3 & 3 & 1 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} 2 & 4 & 2 \\ 2 & 0 & 4 \end{pmatrix}.$$

Suppose that $x = (0.5, 0.5)$ and $y = (0.3, 0.4, 0.3)$, then

$$C_{\mathrm{B}}(x, y) = x^T By = \begin{pmatrix} 0.5 & 0.5 \end{pmatrix} \begin{pmatrix} 2 & 4 & 2 \\ 2 & 0 & 4 \end{pmatrix} \begin{pmatrix} 0.3 \\ 0.4 \\ 0.3 \end{pmatrix} = 2.3$$

Is (x, y) MNE? For $y' = (0.3, 0.7, 0)$, $C_B(x, y') = x^T By' = 2 < 2.3$.

(Row) vector $x^T B = (2, 2, 3)^T$ gives (expected) cost for Bob per column.

- Bob assigns positive probability to b_3: not optimal.
- Should only give positive probability to b_1, b_2 (given Alice plays x).

In MNE, players only have positive probability on rows/columns that minimize expected cost per row/column (given other's strategy).

Definition

Column b_j is best response against x for Bob if $(x^T B)_j = \min_k (x^T B)_k$.
Row a_i is best response against y for Alice if $(Ay)_i = \min_k (Ay)_k$.

(E.g., if $x^T B = (7, 1, 3)^T$, then $(x^T B)_1 = 7, (x^T B)_2 = 1, (x^T B)_3 = 3$.)

- $(x^T B)_j$ is expected cost for Bob in column j given Alice plays x.
- $(Ay)_i$ is expected cost for Alice in row i given Bob plays y.

Definition (MNE, best response version)

Mixed strategies (x^*, y^*) form MNE if Alice and Bob only assign positive probability to best responses. That is, pair (x^*, y^*) is MNE if

$$\begin{aligned} x_i^* > 0 \quad &\Rightarrow \quad (Ay^*)_i = \min_k (Ay^*)_k \qquad && \forall i = 1, \ldots, m, \\ y_j^* > 0 \quad &\Rightarrow \quad ((x^*)^T B)_j = \min_k ((x^*)^T B)_k \qquad && \forall j = 1, \ldots, n. \end{aligned}$$

Example (cont'd)

An MNE is given by $x^* = (1, 0), y^* = (0.5, 0, 0.5)$.

- $(x^*)^T B = (2, 4, 2)^T$. We have $y_1^*, y_3^* > 0$ and $(x^T B)_1, (x^T B)_3$ are min.
- $Ay^* = (2, 2)$. We have $x_1^* > 0$ and $(Ay^*)_1$ is minimum.

Finally, we write $e^k \in \Delta_A$ for pure strategy in which Alice plays $a_k \in \mathcal{S}_A$ with probability 1. That is,

$$e^k_j = \begin{cases} 1 & \text{if } j = k \\ 0 & \text{if } j \neq k \end{cases}$$

- If Alice plays $e^k \in \mathcal{S}_A$, then $C_A(e^k, y) = (e^k)^T Ay = (Ay)_k$.
- Analogous definitions for Bob.

For Alice, one has $e^k \in \mathbb{R}^m$ and for Bob $e^\ell \in \mathbb{R}^n$. We abuse notation and do not always state the dimension of these vectors.

Definition (MNE, pure strategy version)

Mixed strategies (x^*, y^*) form MNE if

$$\begin{array}{lcll} (x^*)^T Ay^* & \leq & (e^i)^T Ay^* & i = 1, \ldots, m, \\ (x^*)^T By^* & \leq & (x^*)^T Ae^j & j = 1, \ldots, n. \end{array}$$

That is, players both have no improving move to pure strategy.

- I.e., suffices to focus on pure strategies in definition on Slide 8.
- Exercise: Prove that this definition is equivalent to that on Slide 8.

Mixed Nash equilibrium (general)

Definition (Mixed Nash equilibrium (MNE))

A mixed strategy $\sigma_i : \mathcal{S}_i \to [0, 1]$ of player $i \in N$ is a probability distribution over pure strategies in $\mathcal{S}_i$, i.e., coming from

$$\Delta_i = \left\{ \tau : \tau(t) \geq 0 \ \ \forall t \in \mathcal{S}_i \quad \text{and} \quad \sum_{t \in \mathcal{S}_i} \tau(t) = 1 \right\}.$$

A collection of mixed strategies $(\sigma_i)_{i \in N}$, with $\sigma_i \in \Delta_i$, is a mixed Nash equilibrium if

$$C_i(\sigma) := \mathbb{E}_{s \sim \sigma}\left[C_i(s)\right] \leq \mathbb{E}_{(s_{-i}) \sim (\sigma_{-i})}\left[C_i(s_i', s_{-i})\right] \ \ \forall s_i' \in \mathcal{S}_i. \qquad (1)$$

Here

- $\sigma : \times_j \mathcal{S}_j \to \mathbb{R}_{\geq 0}$ is given by $\sigma(t) = \prod_j \sigma_j(t_j)$, and
- $\sigma_{-i} : \times_{j \neq i} \mathcal{S}_j \to \mathbb{R}_{\geq 0}$ is given by $\sigma_{-i}(t_{-i}) = \prod_{j \neq i} \sigma_j(t_j)$.

Existence and computational complexity

Existence ("Nobel" Prize in Economics in 1994)

Theorem (Nash's theorem, 1950)

Any finite game Γ has a mixed Nash equilibrium.

Theorem (Brouwer's fixed point theorem)

Let $D \subseteq \mathbb{R}^m$ be compact and convex, and let $f : D \to D$ be a continuous function. Then there exists an $x^* \in D$ such that $f(x^*) = x^*$.

Brouwer's theorem says that f has a fixed point.

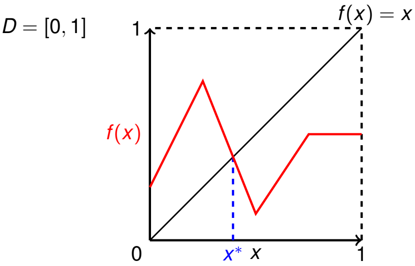

Proof of Nash's theorem

Show that MNEs correspond to fixed points of some function.
Brouwer's theorem then gives existence (proof is not constructive).

- *Proof given for 2-player games. (To save on notation.)*

Proof: Consider set $D = \Delta_A \times \Delta_B$. *(Convex and compact.)*

- Remember $\begin{cases} \Delta_\text{A} & = \{(x_1, \ldots, x_m) : \sum_k x_k = 1, x_k \geq 0\}, \\ \Delta_\text{B} & = \{(y_1, \ldots, y_n) : \sum_\ell y_\ell = 1, y_\ell \geq 0\}. \end{cases}$

For $(x, y) \in \Delta_A \times \Delta_B$, define

$$\begin{array}{ll} R_{A,a_k}(x,y) = \max\{0, C_A(x,y) - C_A(e^k, y)\} & k = 1, \ldots, m \\ R_{B,b_\ell}(x,y) = \max\{0, C_B(x,y) - C_B(x, e^\ell)\} & \ell = 1, \ldots, n \end{array}$$

Note that the $R_{\cdot,\cdot}(x, y)$ encode MNE as follows:

$$R_{z,s_z}(x,y) = 0 \ \ \forall z \in \{A, B\} \, \forall s_z \in \mathcal{S}_z \quad \Leftrightarrow \quad (x, y) \text{ is MNE.}$$

Exercise: Show that $R_{z,s_z}(x, y)$ is a continuous function.

$$R_{A,a_k}(x,y) = \max\{0, C_A(x,y) - C_A(e^k, y)\} \quad k = 1,\dots,m$$
$$R_{B,b_\ell}(x,y) = \max\{0, C_B(x,y) - C_B(x, e^\ell)\} \quad \ell = 1,\dots,n$$

We use these functions to define mapping $f : \Delta_A \times \Delta_B \to \Delta_A \times \Delta_B$ by $f(x,y) = (x', y') = (x'_1, \dots, x'_m, y'_1, \dots, y'_n)$, where

$$x'_i \quad := \frac{x_i + R_{A,a_i}(x,y)}{\sum_{k=1}^m x_k + R_{A,a_k}(x,y)} = \frac{x_i + R_{A,a_i}(x,y)}{1 + \sum_{k=1}^m R_{A,a_k}(x,y)} \quad i = 1,\dots,m$$

and $y' \in \Delta_2$ by

$$y'_j \quad := \frac{y_j + R_{B,b_j}(x,y)}{\sum_{\ell=1}^n y_\ell + R_{B,b_\ell}(x,y)} = \frac{y_j + R_{B,b_j}(x,y)}{1 + \sum_{\ell=1}^n R_{B,b_\ell}(x,y)} \quad j = 1,\dots,n$$

Exercise: Show that f is a continuous function.

If (x^*, y^*) is MNE, then $R_{z,s_z}(x,y) = 0 \ \forall z \in \{A, B\} \ \forall s_z \in \mathcal{S}_z$, and so $x' = x^*$ and $y' = y^*$. In other words, (x^*, y^*) is fixed point of f.

Other direction remains: If (x^*, y^*) is fixed point of f, then it is MNE.
Suffices to show that $R_{z,s_z}(x,y) = 0 \;\; \forall z \in \{A,B\} \; \forall s_z \in \mathcal{S}_z$.

$$\begin{aligned} R_{A,a_i}(x,y) &= \max\{0, C_A(x,y) - C_A(e^i, y)\} & i = 1,\ldots,m \\ x_i' &:= \frac{x_i + R_{A,a_i}(x,y)}{1 + \sum_{k=1}^m R_{A,a_k}(x,y)} & i = 1,\ldots,m \end{aligned}$$

Note that

$$C_A(x,y) = \sum_k x_k C_A(e^k, y) \leq \max_{k:x_k>0} C_A(e^k,y) \sum_k x_k = \max_{k:x_k>0} C_A(e^k,y)$$

- There exists $\bar{i}$ with $x_{\bar{i}} > 0$ such that $R_{A,a_{\bar{i}}}(x,y) = 0$.

Let us look at $x_{\bar{i}}'$ for fixed point $(\mathbf{x}^*, \mathbf{y}^*)$:

- $x_{\bar{i}}^* = \dfrac{x_{\bar{i}}^*}{1 + \sum_{k=1}^m R_{A,a_k}(x^*,y^*)} \overset{x_{\bar{i}}>0}{\Longleftrightarrow} 1 = \dfrac{1}{1 + \sum_{k=1}^m R_{A,a_k}(x^*,y^*)}$
- This gives $\sum_{k=1}^m R_{A,a_k}(x^*,y^*) = 0$.
- R_{A,a_k} is always non-negative $\Rightarrow R_{A,a_k}(x^*,y^*) = 0$ for $k = 1,\ldots,m$.

Computation of MNE

Theorem (Nash's theorem, 1950)

Any finite game Γ has a mixed Nash equilibrium.

Can we compute an MNE efficiently?

Assuming cost functions are rational (think of $A, B \in \mathbb{Q}^{m \times n}$),

- MNE is always rational when $n = 2$, but
- MNE can be irrational when $n \geq 3$.
  - Irrational numbers are, e.g., π, e (Euler's number), etc.

(Context: Suppose $f(z) = z^2 + z - 2$, then $f(z) = z$ is solved by $z^* = \pm\sqrt{2}$.)

For $n \geq 3$,

- Rational ϵ-approximate MNE still exists for any $\epsilon > 0$.
- Algorithms are known to compute approx. equilibrium.
  - E.g., Scarf's algorithm (1967) for approximating fixed points.
  - Probably hard to compute in general (similar to upcoming discussion for $n = 2$).

Complexity of computing MNE ($n = 2$)

For $n = 2$, proof(s) of Brouwer's theorem give no algorithm.

- (Combinatorial) algorithms are known, e.g., Lemke-Howson algorithm.
  - Worst-case running time is exponential (in #strategies).

How to study computational complexity of MNE in 2-player games?
Computing MNE will be referred to as problem NASH.

Some (informal) intuition

Consider function/search problem version of NP:

- For problem X, decide whether solution exists. If YES, output one.

Is NASH NP-complete? Not likely.

- "Deciding" whether Nash equilibrium exists is trivial.

NASH is complete for complexity class PPAD (already for $n = 2$).

- "Polynomial Parity Arguments on Directed graphs"
- See Chapter 20 [R2016] for this class, and more..

Theorem (Chen and Deng, 2006)

Computing MNE in 2-player games is PPAD-complete

- Same is true for approximate equilibria when $n \geq 3$.

What about approximate equilibria in 2-player games?

Assuming game is normalized ($0 \leq A_{ij}, B_{ij} \leq 1$) and $m = n$, we have:

Theorem (Lipton, Markakis and Mehta, 2003)

There is an $O^*(n^{24\log(n)/\epsilon^2})$ algorithm known for computing ϵ-approximate MNE in 2-player game.

- Quasi-polynomial in n.

Theorem (Rubinstein, 2016)

There exists a constant $\epsilon > 0$ such that, assuming the "Exponential Time Hypothesis for PPAD", computing ϵ-approximate MNE in 2-player game requires time at least $n^{\log^{1-o(1)}(n)}$.

Two-player zero-sum games

Two-player zero-sum game

Two-player game is called zero-sum if $A + B = 0$, i.e, $A = -B$.

- Minimizing cost under A is same as maximizing cost under B.

Viewpoint that we take: Given is $m \times n$ matrix C.

- Row player (Alice) tries to maximize utility $x^T Cy$;
- Column player (Bob) tries to minimize cost $x^T Cy$.

Think of it as that Bob has to pay $x^T Cy$ to Alice.

Algorithmic aspects of MNE:

- Can be modeled as optimal solution of linear program (LP).
  - Solvable in polynomial time.
  - *(Any LP can be written as zero-sum game as well.)*
- Certain player dynamics can "learn" it: Fictitious Play
  - Holds for more classes of games, but not in general.

Value of zero-sum game

What can Alice guarantee to get from Bob?

- Suppose Alice plays mixed strategy x. What should Bob do?
- Choose y such that $x^T Cy$ is minimal, i.e., strategy attaining

$$\min_{y \in \Delta_B} x^T Cy.$$

- So what should Alice do? Choose x maximizing $\min_{y \in \Delta_B} x^T Cy$.
- Alice can guarantee to get $v_A = \max_x \min_y x^T Cy$.

Similarly, Bob can guarantee to pay no more than $v_B = \min_y \max_x x^T Cy$.

- Exercise: Show that $v_A \leq v_B$

Theorem (Von Neumann, 1928)

Consider a two-player zero-sum game given by matrix C. Then

$$v_A = \max_x \min_y x^T Cy = \min_y \max_x x^T Cy = v_B.$$

The number $v = v_A = v_B$ is called the value of the game.

- Often referred to as the “Minimax theorem”

Theorem (Minimax)

Consider a two-player zero-sum game given by matrix C. Then

$$v_A = \max_x \min_y x^T C y = \min_y \max_x x^T C y = v_B.$$

We say that x^* is optimal for Alice if v_A is attained for x^*, i.e.,

$$\max_x \min_y x^T C y = \min_y (x^*)^T C y,$$

and, similarly, y^* is optimal for Bob if v_B is attained for y^*, i.e.,

$$\min_y \max_x x^T C y = \max_x x^T C y^*.$$

Corollary

(x^*, y^*) is MNE if and only x^* optimal for Alice and y^* optimal for Bob.

- Computing MNE comes down to computing optimal strategies.

Corollary

Any MNE yield the same utility/loss for Alice/Bob, namely $v = v_A = v_B$.

- Exercise: Prove these corollaries

Two-player zero-sum games

Computing MNE using linear programming

LP formulation for optimal strategy

Optimal strategy x^* for Alice is solution to optimization problem.

- We assume that the C is $m \times n$ matrix, i.e., m rows, n columns.

$$
\begin{array}{lll}
\max & w & \\
\text{subject to} & w \leq \sum_{i=1}^{m} C_{ik} x_i & k = 1, \ldots, n \\
& \sum_{i=1}^{m} x_i = 1 & \\
& x_i \geq 0 & i = 1, \ldots, m \\
& w \in \mathbb{R} &
\end{array}
$$

Problem above is indeed LP, with variables $(x_1, \ldots, x_m, w)$.

- First m variables of optimum give optimal strategy x^*.
- Variable w of optimum gives value $v = v_A$ of the game.

The dual of this program precisely computes optimal strategy for Bob!
In fact, strong duality can be used to prove the minimax theorem.

Theorem

MNE can be computed in polynomial time in 2-player zero-sum game.

Two-player zero-sum games

Fictitious play

Simultaneous fictitious play (Brown, 1951)

Introduced as algorithm for approximating value of zero-sum game.

Game is played repeatedly. In every round:

- Alice (A) and Bob (B) play a pure strategy.
- They base their decision on history of the other player.
    - Choose best response w.r.t. empirical distribution (so far) of strategies chosen by the other.

Informally speaking, empirical distributions "converge" to MNE.

Let $\mathcal{S}_A = \{a_1, \ldots, a_m\}$ (rows) and $\mathcal{S}_B = \{b_1, \ldots, b_n\}$ (columns).

Definition (Empirical distribution)

Let r_t be row chosen by Alice in step $t = 1, \ldots, T-1$. Empirical distribution over $\mathcal{S}_A$ in round t is given by

$$\bar{x}_i(t) = \frac{|\{j : r_j = a_i, 1 \leq j \leq t-1\}|}{t-1}$$

for $i = 1, \ldots, m$. *(Fraction of rounds in which Alice chose row i.)*

- Analogous definition for Bob (with chosen column c_t in round t).

Example

Suppose the matrix C has $n = 6$ rows, and that Alice plays $(a_1, a_1, a_4, a_6, a_4, a_5, a_2, a_3, a_4)$ in first $t - 1 = 9$ rounds. Then

$$\bar{x}(t) = \bar{x}(10) = \frac{1}{9}(2, 1, 1, 3, 1, 1) = \left(\frac{2}{9}, \frac{1}{9}, \frac{1}{9}, \frac{3}{9}, \frac{1}{9}, \frac{1}{9}\right).$$

The idea of fictitious play is that Alice believes Bob plays every round according to some (unknown to her) probability distribution y.

- She uses empirical distribution $\bar{y}(t)$ as guess for y in step t.
- Alice chooses best response row $r_t \in \mathcal{S}_A$ with respect to $\bar{y}(t)$:

$$r_t \in \text{argmax}_j\{(e^i)^T C\bar{y}(t) : i = 1, \ldots, m\}.$$

Bob is doing the same w.r.t Alice (for unknown distribution x).

- He uses empirical distribution $\bar{x}(t)$ as guess for x in step t.
- Bob chooses best response column $c_t \in \mathcal{S}_B$ with respect to $\bar{x}(t)$:

$$c_t \in \text{argmin}_j\{\bar{x}(t)^T Ce^j : j = 1, \ldots, n\}.$$

Fictitious play algorithm

ALGORITHM 1: Fictitious play (with index tie-breaking rule)

Input : $m \times n$ matrix C; initial row r, column c; round total $T \in \mathbb{N}$.
Output: Empirical distributions $\bar{x}(T), \bar{y}(T)$.

$\bar{x}(1) = e_r$ and $\bar{y}(1) = e_c$.
for $t = 2, \ldots, T$ **do**
  Choose $r_t \in \text{argmax}\{(e^i)^T C\bar{y}(t) : i = 1, \ldots, m\}$
  Choose $c_t \in \text{argmin}\{\bar{x}(t)^T Ce^j : j = 1, \ldots, n\}$
  (Choose lowest indexed row/column in case of multiple best responses.)
  Update empirical distributions $(\bar{x}(t), \bar{y}(t))$ to $(\bar{x}(t+1), \bar{y}(t+1))$
end
return $\bar{x}(T), \bar{y}(T)$

- Observe that we specify a tie-breaking rule that decides which column/row to choose, in case there are multiple best responses.

Theorem (Robinson, 1951)

Utility/cost of Alice/Bob converges to value v of the game. That is, as $t \to \infty$, it holds that

- $\max_i (e^i) C\bar{y}(t) \to v$, $\min_j \bar{x}(t)^T Ce_j \to v$, *and* $\bar{x}(t)^T C\bar{y}(t) \to v$.

Empirical distributions $(\bar{x}(t), \bar{y}(t))$ "converge" to MNE as $t \to \infty$.

- Convergence in the sense that $(\bar{x}(t), \bar{y}(t))$ is $\epsilon(t)$-approximate equilibrium, where $\epsilon(t) \to 0$ as $t \to \infty$.

Convergence time of Fictitious Play still not fully understood!

Some notes on fictitious play

- Simple way to compute value and ϵ-MNE.
  - Avoiding the need to solve LPs.
- Players do not need to know each other's empirical distribution.
  - Alice only needs to know vector $(C\bar{y}(t))$ in round t.
  - Bob only needs to know (row) vector $(\bar{x}(t)^T C)$ in round t.
- Fictitious play can be defined for any two-player game (A, B).
  - Convergence fails beyond zero-sum games.