Lecture 7: Game Theory

- Games in Normal Form
- Equilibria
- Dominance
- Mixed Strategies

Example (1)

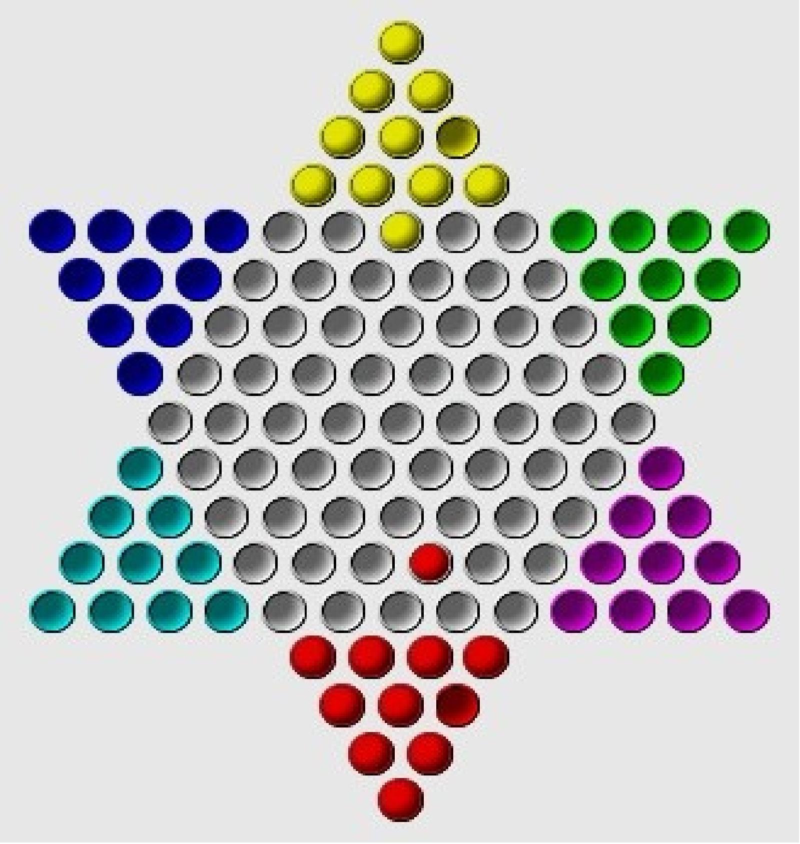

Example (2)

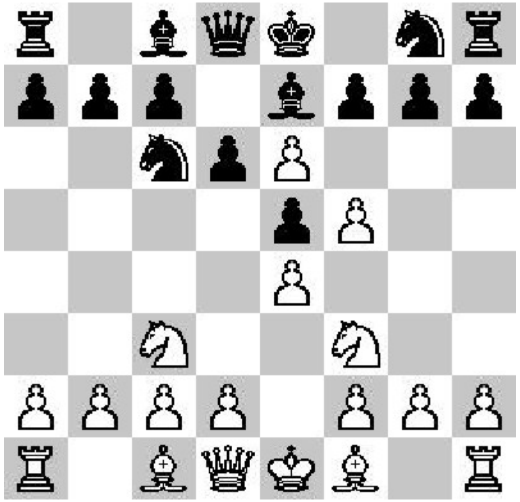

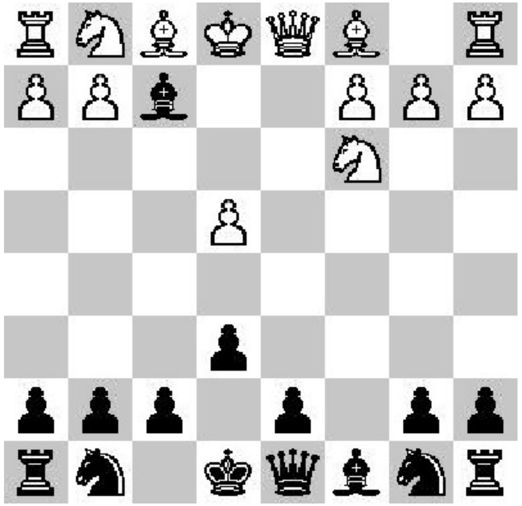

Example (3)

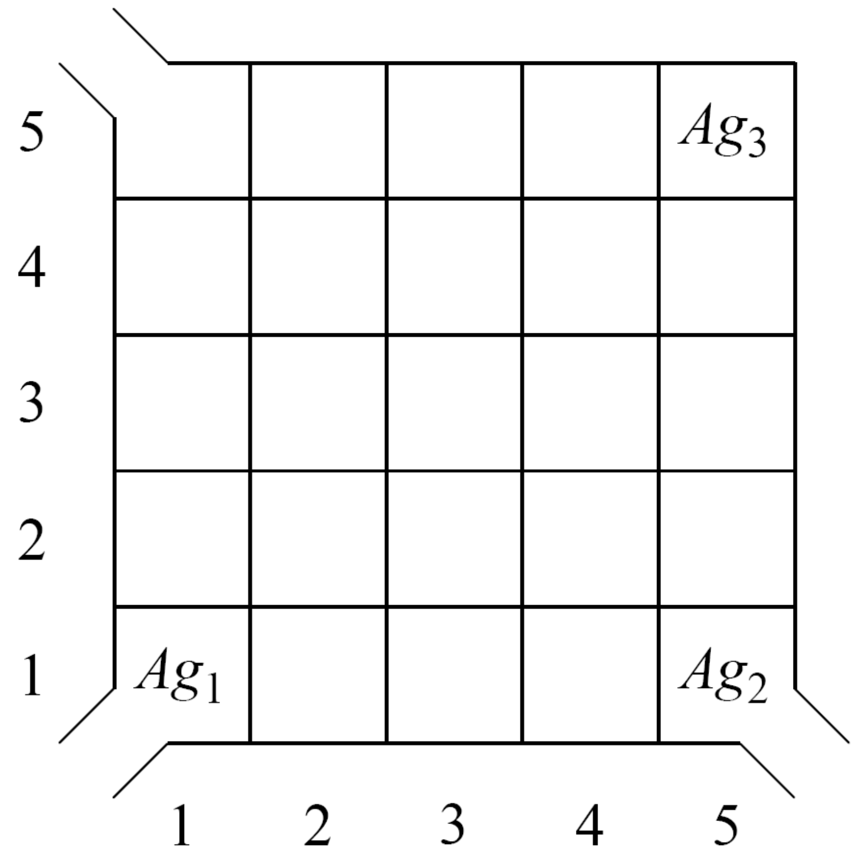

Competition and Cooperation

- The "pathological" assumption says that opponents choose the joint move that is most harmful for us.
- This is usually too pessimistic for non-zerosum games or games with $n > 2$ players. A *rational* player chooses the move that's best for him rather than the one that's worst for us.

Example

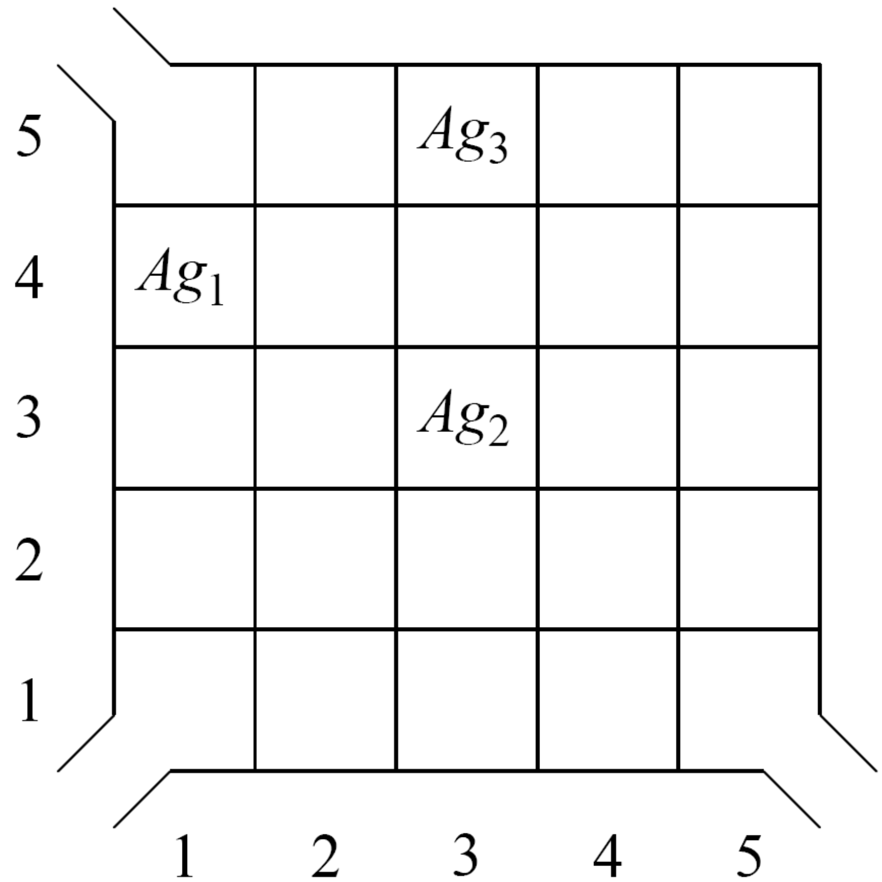

Game Theory gives the answer

Strategies

Game model:

- S – set of states
- $A_1, \ldots, A_n$ – n sets of actions, one for each player
- $l_1, \ldots, l_n$ – where $l_i \subseteq A_i \times S$, the legality relations
- $g_1, \ldots, g_n$ – where $g_i \subseteq S \times \mathbb{N}$, the goal relations

A strategy x_i for player i maps every state to a legal move for i

$$x_i : S \rightarrow A_i \qquad \text{(such that } (x_i(S),S) \in l_i \text{)}$$

(Note that even for Chess the number of different strategies is finite. They outnumber the atoms in the universe, though ...)

Strategies for Agent-Battle

Example strategy for Ag_1:

$\{At(Ag_1,1,1), At(Ag_2,5,1), At(Ag_3,5,5)\} \rightarrow Go(East)$

$\{At(Ag_1,1,1), At(Ag_2,5,1), At(Ag_3,4,5)\} \rightarrow Go(North)$

⋮

$\{At(Ag_1,1,4), At(Ag_2,3,3), At(Ag_3,3,5)\} \rightarrow Go(North)$

⋮

Similar for Ag_2, Ag_3

Towards the Normal Form of Games

Each n-tuple of strategies directly determines the outcome of a match.

Example:

Start with 7 coins. Players A and B take turn in removing one or two coins. Whoever takes the last coin wins.

Strategy Player A:

$\{7 \mapsto 2, 6 \mapsto 2, 5 \mapsto 2, 4 \mapsto 1, 3 \mapsto 1, 2 \mapsto 2, 1 \mapsto 1\}$

Strategy Player B:

$\{7 \mapsto 2, 6 \mapsto 1, 5 \mapsto 2, 4 \mapsto 1, 3 \mapsto 2, 2 \mapsto 2, 1 \mapsto 1\}$

Outcome: (0, 100)

Games in Normal Form

An *n-player game in normal form* is an n+1-tuple

$$\Gamma = (X_1, \ldots, X_n, u)$$

where X_i is the set of strategies for player i and

$$u = (u_1, \ldots, u_n) \colon \bigtimes_{i=1}^{n} X_i \rightarrow \mathbb{N}^i$$

are the utilities of the players for each n-tuple of strategies.

Roshambo

2-player games are often depicted as matrices

	Rock	Scissors	Paper
Rock	50 50	0 100	100 0
Scissors	100 0	50 50	0 100
Paper	0 100	100 0	50 50

2-Finger-Morra

	1 Finger	2 Fingers
1 Finger	30 70	90 10
2 Fingers	90 10	0 100

Battle of the Sexes

	Ballgame	Opera
Ballgame	3 4	2 2
Opera	1 1	4 3

Prisoner's Dilemma

	Cooperate	Defect
Cooperate	3 3	4 1
Defect	1 4	2 2

Equilibria

Let $\Gamma = (X_1, \ldots, X_n, u)$ be an n-player game.

$(x_1^*, \ldots, x_n^*) \in \mathrm{X}_1 \times \ldots \times \mathrm{X}_n$ equilibrium

if for all $i = 1, \ldots, n$ and all $x_i \in \mathrm{X}_i$

$$u_i(x_1^*, \ldots, x_{i-1}^*, x_i, x_{i+1}^*, \ldots, x_n^*) \leq u_i(x_1^*, \ldots, x_n^*)$$

An equilibrium is a tuple of optimal strategies: No player has a reason to deviate from his strategy, given the opponent's strategies.

Full Cooperation

	a	b
a	4 4	1 2
b	2 3	3 1

Battle of the Sexes

	Ballgame	Opera
Ballgame	3 4	2 2
Opera	1 1	4 3

(Note that the outcome for both players is bad if they choose to play different equilibria.)

Cooperation

	a	b
a	4, 4	2, 2
b	1, 1	3, 3

(Note that the concept of an equilibrium doesn't suffice to achieve the best possible outcome for both players.)

Prisoner's Dilemma

	Cooperate	Defect
Cooperate	3, 3	1, 4
Defect	4, 1	2, 2

(Note that the outcome which is better for both players isn't even an equilibrium!)

Dominance

A strategy $x \in X_i$ dominates a strategy $y \in X_i$ if

$$u_i(x_1, \ldots, x_{i-1}, x, x_{i+1}, \ldots, x_n) \geq u_i(x_1, \ldots, x_{i-1}, y, x_{i+1}, \ldots, x_n)$$

for all $(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n) \in X_1 \times \ldots \times X_{i-1} \times X_{i+1} \times \ldots \times X_n$.

A strategy $x \in X_i$ strongly dominates a strategy $y \in X_i$ if x dominates y and y does not dominate x.

Assume that opponents are rational: They don't choose a strongly dominated strategy.

Removing Strongly Dominated Strategies

	a	b
a	4 4	1 2
~~b~~	~~2 3~~	~~3 1~~

Iterated Dominance

Let a zero-sum game be given by

	a	b	c	d	e
a	10	7	6	9	8
b	10	4	6	9	5
c	9	7	9	8	8
d	2	6	4	3	7

Iterated Dominance (2)

	a	b	c	d	e
a	10	7	6	9	8
~~b~~	~~10~~	~~4~~	~~6~~	~~9~~	~~5~~
c	9	7	9	8	8
~~d~~	~~2~~	~~6~~	~~4~~	~~3~~	~~7~~

Iterated Dominance (3)

	a	b	c	d	e
a	10	7	6	9	8
c	9	7	9	8	8

Iterated Dominance (4)

	b	c
a	7	6
c	7	9

Game Tree Search with Dominance

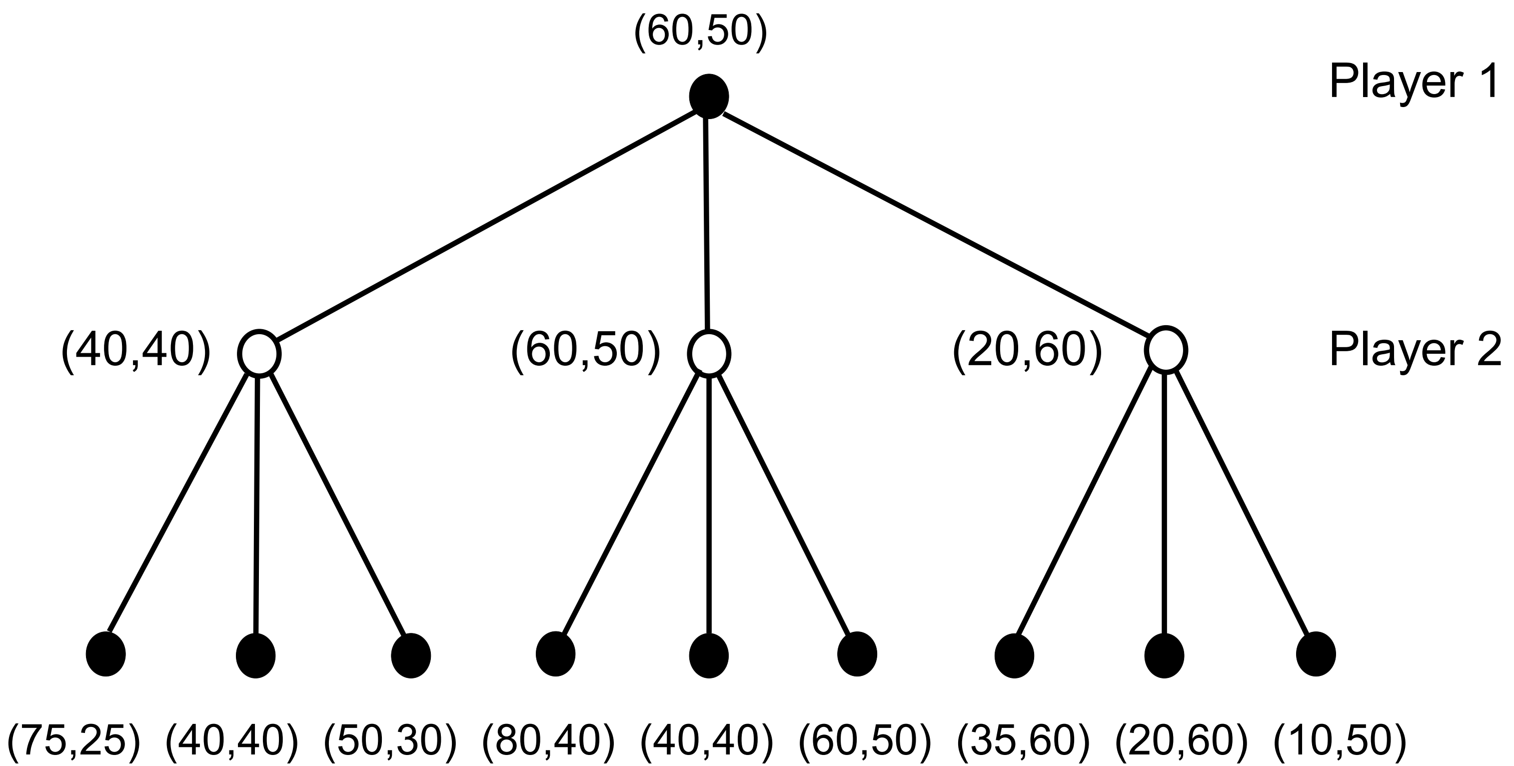

The α-β-Principle does not Apply

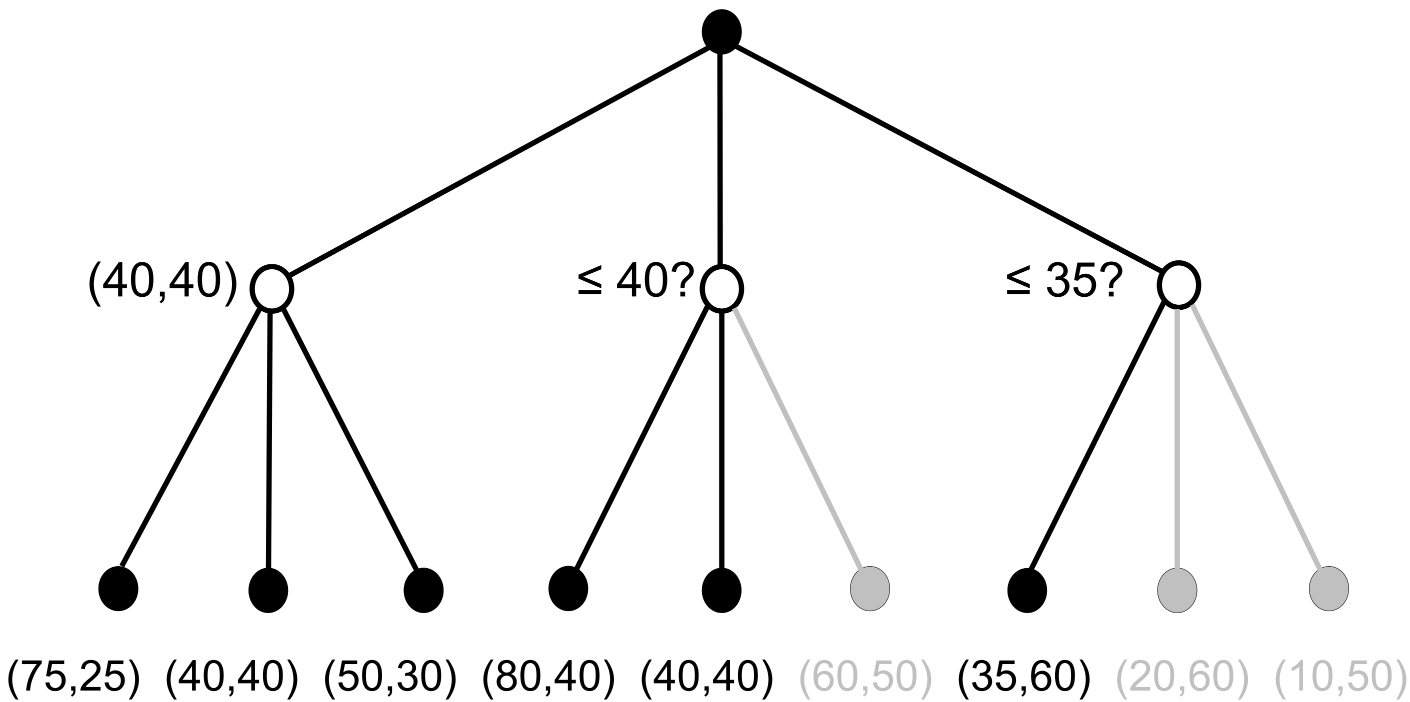

The Need to Randomize: Roshambo

	Rock	Scissors	Paper
Rock	50 50	0 100	100 0
Scissors	100 0	50 50	0 100
Paper	0 100	100 0	50 50

This game has no equilibrium

2-Finger-Morra

	1 Finger	2 Fingers
1 Finger	30 70	90 10
2 Fingers	90 10	0 100

This game, too, has no equilibrium

Mixed Strategies

Let $(X_1, \ldots, X_n, u)$ be an n-player game, then its mixed extension is

$$\Gamma = (P_1, \ldots, P_n, (e_1, \ldots, e_n))$$

where for each i=1, ..., n

$$P_i = \{p_i\!: p_i \text{ probability measure over } X_i\}$$

and for each $(p_1, \ldots, p_n) \in P_1 \times \ldots \times P_n$

$$e_i(p_1, \ldots, p_n) = \sum_{x_1 \in X_1} \ldots \sum_{x_n \in X_n} u_i(x_1, \ldots, x_n) * p_1(x_1) * \ldots * p_n(x_n)$$

Existence of Equilibria

Nash's Theorem.

Every mixed extension of an *n*-player game has at least one equilibrium.

Roshambo

	Rock	Scissors	Paper
Rock	50 / 50	0 / 100	100 / 0
Scissors	100 / 0	50 / 50	0 / 100
Paper	0 / 100	100 / 0	50 / 50

The unique equilibrium is

$$\left(\left(\frac{1}{3},\frac{1}{3},\frac{1}{3}\right),\left(\frac{1}{3},\frac{1}{3},\frac{1}{3}\right)\right)$$

2-Finger-Morra

	1 Finger	2 Fingers
1 Finger	70, 30	10, 90
2 Fingers	10, 90	100, 0

The unique equilibrium is

$$(p_1^*, p_2^*) = \left(\left(\frac{3}{5}, \frac{2}{5}\right), \left(\frac{3}{5}, \frac{2}{5}\right)\right)$$

with $e_1(p_1^*, p_2^*) = 46$ and $e_2(p_1^*, p_2^*) = 54$

Computing the Equilibrium

Solution known for 2-player, zero-sum games:

Let a_{ij} be the goal value for player 1 for the pair (i, j) of pure strategies.

For player 1, minimize $\sum_i s_i$ under conditions $\sum_i a_{ij} \cdot s_i \geq 1 \quad (\forall_j)$

$$s_i \geq 0 \quad (\forall_i)$$

For player 2, maximize $\sum_j t_j$ under conditions $\sum_j a_{ij} \cdot t_j \leq 1 \quad (\forall_i)$

$$t_j \geq 0 \quad (\forall_j)$$

Normalization gives the final answer.

Example: 2-Finger-Morra

$a_{11} = 70$ $a_{12} = 10$ $a_{21} = 10$ $a_{22} = 100$

minimize $s_1 + s_2$ under conditions
$$70 \cdot s_1 + 10 \cdot s_2 \geq 1$$
$$10 \cdot s_1 + 100 \cdot s_2 \geq 1$$
$$s_1 \geq 0 \quad s_2 \geq 0$$

maximize $t_1 + t_2$ under conditions
$$70 \cdot t_1 + 10 \cdot t_2 \leq 1$$
$$10 \cdot t_1 + 100 \cdot t_2 \leq 1$$
$$t_1 \geq 0 \quad t_2 \geq 0$$

Solutions: $s_1 = 0.01304$ $s_2 = 0.00870$

$t_1 = 0.01304$ $t_2 = 0.00870$

Normalization gives $p_1^* = \left(\frac{s_1}{s_1+s_2}, \frac{s_2}{s_1+s_2}\right) = (0.6, 0.4)$

$p_2^* = \left(\frac{t_1}{t_1+t_2}, \frac{t_2}{t_1+t_2}\right) = (0.6, 0.4)$

Iterated Row Dominance for Mixed Strategies

Let a zero-sum game be given by

	a	b	c
a	10	0	8
b	6	4	4
c	2	8	7

Then $p_1 = \left(\frac{1}{2}, 0, \frac{1}{2}\right)$ strongly dominates $p_1' = (0,1,0)$.

Hence, for all $(p_a', p_b', p_c') \in P_1$ with $p_b' > 0$ there exists a dominating strategy $(p_a, 0, p_c) \in P_1$.

Iterated Row Dominance for Mixed Strategies (2)

	a	b	c
a	10	0	8
~~b~~	~~6~~	~~4~~	~~4~~
c	2	8	7

Now $p_2 = \left(\frac{1}{2}, \frac{1}{2}, 0\right)$ dominates $p_2' = (0,0,1)$.

Iterated Row Dominance for Mixed Strategies (3)

	a	b	c
a	10	0	8
c	2	8	7

The unique equilibrium is $\left(\left(\frac{3}{8}, 0, \frac{5}{8}\right), \left(\frac{1}{2}, \frac{1}{2}, 0\right)\right)$.

Challenges

- From a game theoretic point of view, modeling simultaneous moves as a sequence of our move followed by the joint moves of our opponents is incorrect.
- How to modify the node expansion?
- How to compute equilibria and mixed strategies?
- How to model (and coin against) “stupid” opponents, e.g. who always choose Rock in Roshambo?

The Floor is Yours!