15-382 Collective Intelligence – S18

Lecture 27: Game Theory 2

Instructor:
Gianni A. Di Caro

جامعة كارنيجي ميلون في قطر
Carnegie Mellon University Qatar

THE PROFESSOR'S DILEMMA

- Simultaneous move
- Non-cooperative game
- Complete information
- Imperfect information
- **Solution concept:** predict how the game will be played with rational agents
- Prediction ≡ Solution
- Nash: ***Equilibrium* concept**

Class

Professor	Listen	Sleep
Make effort	$10^6,10^6$	-10,0
Slack off	0,-10	0,0

Dominant strategies?

Nope, if Class listen, and Professor slacks off, Sleep provides a higher payoff!
No *dominant* strategy: best strategy it doesn't matter what other player's strategy

NASH EQUILIBRIUM (1951)

- Can we find an **equilibrium** also in absence of a dominant strategy?

- At equilibrium, each player's strategy is a best response to strategies of others

- Formally, a Nash equilibrium is *strategy profile* $s = (s_1 \dots, s_n) \in S^n$ such that:

$$\forall i \in N, \forall s_i' \in S, u_i(s) \geq u_i(s_i', s_{-i})$$

John F. Nash,
Nobel Prize in Economics, 1994

NASH EQUILIBRIUM

- In equilibrium, each player is playing the strategy that is a "**best response**" to the strategies of the other players. No one has an *incentive* to change strategy given the strategy choices of the others

- A NE is an equilibrium where each player's strategy is *optimal given the strategies of all other players*.

- A Nash Equilibrium exists when there is no unilateral profitable deviation from any of the players involved

- Nash Equilibria are *self-enforcing* strategies: when players are at a Nash Equilibrium they have no desire to move because they will be worse off → **Equilibrium** in the policy space

- **Dominant strategy** ⟹ **Nash equilibrium**: All solutions in dominant strategies are also Nash equilibria, but the vice versa is *not* necessarily (and not usually) true

Nash equilibrium

Equilibrium is *not*:

- *The best possible outcome of the game.*
  Equilibrium in the one-shot prisoners' dilemma is for both players to confess, which is not the best possible outcome (not Pareto optimal)

- A situation where players always choose the same action. Sometimes equilibrium will involve changing action choices (*mixed strategy* equilibrium).

Nash equilibrium

- How many Nash equilibria does the Professor's Dilemma have?

	Listen	Sleep
Make effort	$10^6,10^6$	-10,0
Slack off	0,-10	0,0

NASH EQUILIBRIA: HOW DO WE FIND THEM?

- *Nash equilibrium:* A play of the game where each strategy is a best reply to the given strategy of the other.

- Let's examine all the possible pure strategy profiles and check if for a profile (X,Y) one player could improve its payoff, given the strategy of the other

✓ (M, L)? If Prof plays M, then L is the best reply given M. Neither player can increase its the payoff by choosing a different action

○ (S,L)? If Prof plays S, S is the best reply given S, not L

○ (M, S)? If Prof plays M, then L is the best reply given M, not S

✓ (S,S)? If Prof plays S, then S is the best reply given S. Neither player can increase its the payoff by choosing a different action

NASH EQUILIBRIUM FOR PRISONER'S DILEMMA

		Prisoner B	
		Don't confess	Confess
Prisoner A	Don't Confess	-1,-1	-9,0
	Confess	0,-9	-6,-6

COORDINATION GAME: STAG HUNT

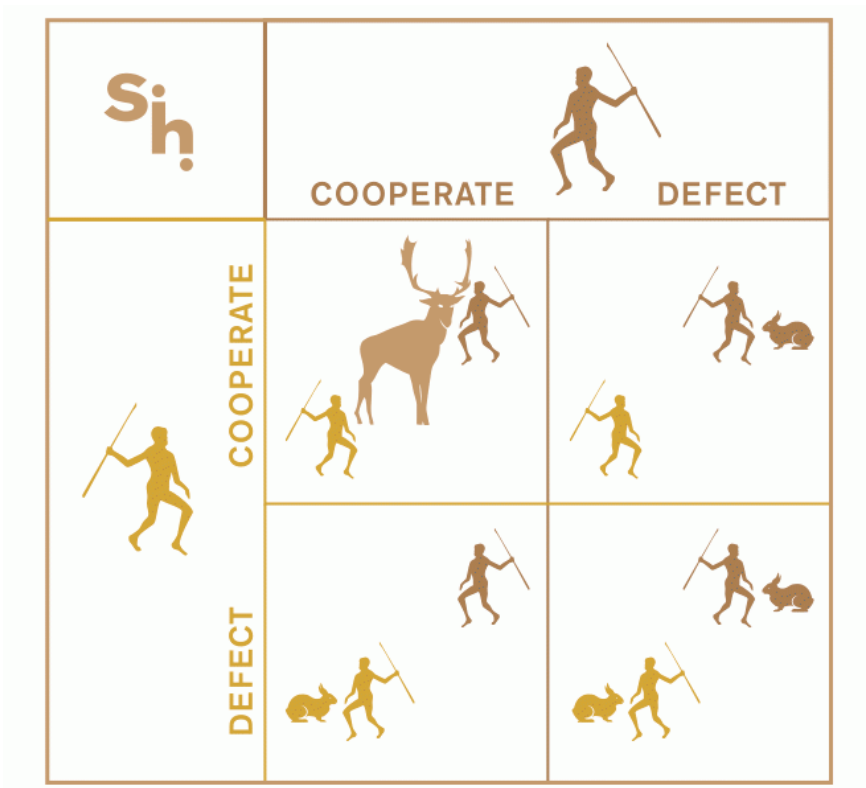

(Originally from J.J. Rousseau)

	stag	rabbit
stag	5 , 5	0 , 3
rabbit	3 , 0	3 , 3

- **Two equilibria** at (*stag, stag*) and (*rabbit, rabbit*) → Players' optimal strategy depend on their *expectation* on what the other player may do.
- This game has been used as an analogy for social cooperation, and mutual trust
- In Prisoner's dilemma, the Nash equilibrium corresponds to defect, no cooperate!

Competition Game

Player 2 / Player 1	Player 2 chooses '0'	Player 2 chooses '1'	Player 2 chooses '2'	Player 2 chooses '3'
Player 1 chooses '0'	*0, 0*	2, −2	2, −2	2, −2
Player 1 chooses '1'	−2, 2	*1, 1*	*3, −1*	3, −1
Player 1 chooses '2'	−2, 2	*−1, 3*	*2, 2*	*4, 0*
Player 1 chooses '3'	−2, 2	−1, 3	*0, 4*	*3, 3*

- Both players simultaneously choose an integer from 0 to 3
- They *both* win the smaller of the two numbers in points.
- In addition, if one player chooses a larger number than the other, then it has to give up two points to the other.

Does the (unique) NE at (0,0) make sense?

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

Nash equilibrium?

Is there a pure strategy as best response?

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

No (pure) Nash equilibria:
Best response:
randomize!

- For every pure strategy (X,Y), there is a different strategy choice that increases the payoff of a player
- E.g., for strategy (P,R), player B can get a higher payoff playing strategy S instead R
- E.g., for strategy (S,R), player A can get a higher payoff playing strategy P instead S
- No strategy equilibrium can be settled, *players have the incentive to keep switching their strategy*

Mixed strategies

- Mixed strategy: a *probability distribution* over (*pure*) strategies
- The mixed strategy of player $i \in N$ is x_i, where $x_i(s_i) = \Pr[i \text{ plays } s_i]$ (e.g., $x_i(R) = 0.3$, $x_i(P) = 0.5$, $x_i(S) = 0.2$)
- The (expected) **utility** of player $i \in N$ is

$$\underbrace{u_i(x_1, \ldots, x_n)}_{\text{Mixed strategy profile}} = \sum_{\underbrace{(s_1,\ldots,s_n)}_{\text{Pure strategy profile}} \in S^n} \underbrace{u_i(s_1, \ldots, s_n)}_{\text{Utility of pure strategy profile}} \cdot \underbrace{\prod_{j=1}^{n} x_j(s_j)}_{\text{Joint probability of the pure strategy profile given the mixed profile}}$$

Exercise: Mixed NE

- Player 1 plays $\left(\frac{1}{2}, \frac{1}{2}, 0\right)$, player 2 plays $\left(0, \frac{1}{2}, \frac{1}{2}\right)$. What is u_1?

- Both players play $\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$. What is u_1?

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

EXERCISE: MIXED NE

Player 1 plays $\left(\frac{1}{2},\frac{1}{2},0\right)$, player 2 plays $\left(0,\frac{1}{2},\frac{1}{2}\right)$. What is u_1?

$$u_1\Big(x_1(R,P,S),x_2(R,P,S)\Big) =$$

$$u_1(R,R)p(R,R|x_1,x_2) + u_1(R,P)p(R,P|x_1,x_2) + u_1(R,S)p(R,S|x_1,x_2)+$$

$$u_1(P,R)p(P,R|x_1,x_2) + u_1(P,P)p(P,P|x_1,x_2) + u_1(P,S)p(P,S|x_1,x_2)+$$

$$u_1(S,R)p(S,R|x_1,x_2) + u_1(S,P)p(S,P|x_1,x_2) + u_1(S,S)p(S,S|x_1,x_2)$$

$$= 0 \cdot (\tfrac{1}{2} \cdot 0) + (-1) \cdot (\tfrac{1}{2} \cdot \tfrac{1}{2}) + 1 \cdot (\tfrac{1}{2} \cdot \tfrac{1}{2})$$

$$+1 \cdot (\tfrac{1}{2} \cdot 0) + 0 \cdot (\tfrac{1}{2} \cdot \tfrac{1}{2}) + (-1) \cdot (\tfrac{1}{2} \cdot \tfrac{1}{2})$$

$$+(-1) \cdot (0 \cdot 0) + 1 \cdot (0 \cdot \tfrac{1}{2}) + 0 \cdot (0 \cdot \tfrac{1}{2})$$

$$= -\tfrac{1}{4}$$

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

In the second case, because of symmetry, the utility is zero: It's a ***zero-sum game***

Mixed Strategies Equilibrium is Nash

- The mixed strategy profile x^* in a strategic game is a **mixed strategy Nash equilibrium** if

$$u_i(x_i^*, x_{-i}^*) \geq u_i(x_i, x_{-i}^*) \quad \forall\, x_i \text{ and } i$$

- $u_i(x)$ is player i's expected utility with mixed strategy profile x

- → Same definition as in the case f pure strategies, where u_i was the utility of a pure strategy instead of a mixed strategy

Mixed Strategies Nash Equilibrium

- Using *best response* functions, x^* is a mixed strategy NE iff x_i^* is the best response for every player i.
- If a mixed strategy x^* is a best response, then each of the *pure strategies in the mix must be best response*: they must yield the same expected payoff (otherwise it would just make sense to choose the one with the better payoff)
- → If a mixed strategy is a best response for player i, then the player must be indifferent among the pure strategies in the mix
- E.g., in the RPS game, if the mixed strategy of player i assigns non-zero probabilities p_R for playing R and p_P for playing P, then i's **expected utility for playing R or P has to be the same**

Exercise: Mixed NE

- Which is a NE?

1. $\left(\left(\frac{1}{2},\frac{1}{2},0\right),\left(\frac{1}{2},\frac{1}{2},0\right)\right)$
2. $\left(\left(\frac{1}{2},\frac{1}{2},0\right),\left(\frac{1}{2},0,\frac{1}{2}\right)\right)$
3. $\left(\left(\frac{1}{3},\frac{1}{3},\frac{1}{3}\right),\left(\frac{1}{3},\frac{1}{3},\frac{1}{3}\right)\right)$
4. $\left(\left(\frac{1}{3},\frac{2}{3},0\right),\left(\frac{2}{3},0,\frac{1}{3}\right)\right)$

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

Any other NE?

Nash's Theorem

- **Theorem [Nash, 1950]:** In any game with finite number of strategies there exists at least one (possibly mixed) Nash equilibrium

Player B

Player A		Left	Right
	Up	1,2	0,4
	Down	0,5	3,2

This game has no pure strategy Nash equilibria but it does have a Nash equilibrium in **mixed strategies**. How is it computed?

Computation of MS NE

		Player B	
		Left	Right
Player A	Up	1,2	0,4
	Down	0,5	3,2

Player A plays Up with probability p_U and plays Down with probability $1-p_U$

Player B plays Left with probability p_L and plays Right with probability $1-p_{L}$.

Computation of MS NE

Player B

		L, π_L	$R, 1-\pi_L$
Player A	U, π_U	1,2	0,4
	$D, 1-\pi_U$	0,5	3,2

Computation of MS NE

Player B

Player A		L, π_L	$R, 1-\pi_L$
	U, π_U	1,2	0,4
	$D, 1-\pi_U$	0,5	3,2

If B plays Left, its expected utility is

$$2\pi_U + 5(1 - \pi_U)$$

Computation of MS NE

Player B

Player A		L, π_L	R, $1-\pi_L$
	U, π_U	1,2	0,4
	D, $1-\pi_U$	0,5	3,2

If B plays Right, its expected utility is

$$4\pi_U + 2(1-\pi_U)$$

Computation of MS NE

		Player B	
		L, π_L	$R, 1-\pi_L$
Player A	U, π_U	1,2	0,4
	$D, 1-\pi_U$	0,5	3,2

If $2\pi_U + 5(1-\pi_U) > 4\pi_U + 2(1-\pi_U)$ then

B would play only Left, which would be a pure strategy.
But there are no (pure) Nash equilibria in which B plays only Left.

Computation of MS NE

Player B

		L,π_L	$R,1-\pi_L$
Player A	U,π_U	1,2	0,4
	$D,1-\pi_U$	0,5	3,2

If $2\pi_U + 5(1-\pi_U) < 4\pi_U + 2(1-\pi_U)$ then

B would play only Right.
But there are no (pure) Nash equilibria in which B plays only Right

COMPUTATION OF MS NE

		Player B	
		L, π_L	R, 1-π_L
Player A	U, π_U	1,2	0,4
	D, 1-π_U	0,5	3,2

For there to exist a MS Nash equilibrium, B must be indifferent between playing Left or Right:

$$2\pi_U + 5(1 - \pi_U) = 4\pi_U + 2(1 - \pi_U)$$

Computation of MS NE

		Player B	
		L, π_L	R, 1-π_L
Player A	U, π_U	1,2	0,4
	D, 1-π_U	0,5	3,2

$$2\pi_U + 5(1-\pi_U) = 4\pi_U + 2(1-\pi_U)$$

$$\Rightarrow \quad \pi_U = 3/5$$

Computation of MS NE

Player B

Player A

	L, π_L	R, $1-\pi_L$
U, $\frac{3}{5}$	1,2	0,4
D, $\frac{2}{5}$	0,5	3,2

$$\pi_U = \frac{3}{5} \qquad 1 - \pi_U = \frac{2}{5}$$

COMPUTATION OF MS NE

Player B

		L, π_L	$R, 1-\pi_L$
Player A	$U, \frac{3}{5}$	1,2	0,4
	$D, \frac{2}{5}$	0,5	3,2

If A plays Up its expected payoff is

$$1 \times \pi_L + 0 \times (1 - \pi_L) = \pi_L$$

Computation of MS NE

Player B

		L, π_L	$R, 1-\pi_L$
Player A	U, $\frac{3}{5}$	1,2	0,4
	D, $\frac{2}{5}$	0,5	3,2

If A plays Down his expected payoff is

$$0 \times \pi_L + 3 \times (1 - \pi_L) = 3(1 - \pi_L)$$

COMPUTATION OF MS NE

Player B

			L, π_L	$R, 1-\pi_L$
Player A	U,	$\frac{3}{5}$	1,2	0,4
	D,	$\frac{2}{5}$	0,5	3,2

If $\boldsymbol{\pi_L > 3(1-\pi_L)}$ then

A would play only Up.
But there are no pure Nash equilibria in which A plays only Up.

COMPUTATION OF MS NE

Player B

		L, π_L	$R, 1-\pi_L$
Player A	U, $\frac{3}{5}$	1,2	0,4
	D, $\frac{2}{5}$	0,5	3,2

If $\pi_L < 3(1 - \pi_L)$ then

A would play only Down.
But there are no Nash equilibria in which A plays only Down.

Computation of MS NE

Player B

		L, π_L	R, $1-\pi_L$
Player A	U, $\frac{3}{5}$	1,2	0,4
	D, $\frac{2}{5}$	0,5	3,2

For there to exist a Mixed Nash equilibrium, A must be indifferent between playing Up or Down:

$$\pi_L = 3(1-\pi_L)$$

Computation of MS NE

Player B

		L, π_L	$R, 1-\pi_L$
Player A	$U, \frac{3}{5}$	1,2	0,4
	$D, \frac{2}{5}$	0,5	3,2

$$\pi_L = 3(1-\pi_L) \quad \Rightarrow \quad \pi_L = 3/4$$

Computation of MS NE

Player B

		$L, \frac{3}{4}$	$R, \frac{1}{4}$
Player A	$U, \frac{3}{5}$	1,2	0,4
	$D, \frac{2}{5}$	0,5	3,2

$$\pi_L = \frac{3}{4} \qquad 1 - \pi_L = \frac{1}{4}$$

Computation of MS NE

Player B

		L, $\frac{3}{4}$	R, $\frac{1}{4}$
Player A	U, $\frac{3}{5}$	1,2	0,4
	D, $\frac{2}{5}$	0,5	3,2

Game's only Nash equilibrium has A playing the mixed strategy $\left(\frac{3}{5}, \frac{2}{5}\right)$ and B playing the mixed strategy $\left(\frac{3}{4}, \frac{1}{4}\right)$

Computation of MS NE

Player B

		L, $\frac{3}{4}$	R, $\frac{1}{4}$
Player A	U, $\frac{3}{5}$	1,2	0,4
	D, $\frac{2}{5}$	0,5	3,2

Payoffs:

- (1,2) with probability $(\frac{3}{5} \times \frac{3}{4}) = \frac{9}{20}$
- (0,4) with probability $(\frac{3}{5} \times \frac{1}{4}) = \frac{3}{20}$
- (0,5) with probability $(\frac{2}{5} \times \frac{3}{4}) = \frac{6}{20}$
- (3,2) with probability $(\frac{2}{5} \times \frac{1}{4}) = \frac{2}{20}$

Computation of MS NE

Player B

Player A

		L, $\frac{3}{4}$	R, $\frac{1}{4}$
Player A	U, $\frac{3}{5}$	1,2	0,4
	D, $\frac{2}{5}$	0,5	3,2

A's expected Nash equilibrium payoff:

$$1 \times \frac{9}{20} + 0 \times \frac{3}{20} + 0 \times \frac{6}{20} + 3 \times \frac{2}{20} = \frac{3}{4}$$

COMPUTATION OF MS NE

		Player B	
		L, $\frac{3}{4}$	R, $\frac{1}{4}$
Player A	U, $\frac{3}{5}$	1,2	0,4
	D, $\frac{2}{5}$	0,5	3,2

B's expected Nash equilibrium payoff:

$$2 \times \frac{9}{20} + 4 \times \frac{3}{20} + 5 \times \frac{6}{20} + 2 \times \frac{2}{20} = \frac{16}{5}$$