A (Brief) Introduction to Game Theory

Johanne Cohen

PRiSM/CNRS, Versailles, France.

Goal

Goal

is a Nash equilibrium.

Today

- The game of Chicken
  - Definitions
  - Nash Equilibrium
- Rock-paper-scissors Game
  - Mixed strategy
  - Mixed Nash Equilibrium
- Prisoner's dilemma
- Bonus : Toward learning equilibria

Outline

The game of Chicken (Hawk-dove game)

- A single lane bridge
- Two drivers `Bob` and `Alice` want to go cross it from opposite directions.

Each driver can `Cross` or `Stop`

The game of Chicken (Hawk-dove game)

- A single lane bridge
- Two drivers `Bob` and `Alice` want to go cross it from opposite directions.

  Each driver can `Cross` or `Stop`
- Both drivers want to minimize the time spent to reach other side.

The game of Chicken (Hawk-dove game)

- A single lane bridge
- Two drivers `Bob` and `Alice` want to go cross it from opposite directions.

Each driver can `Cross` or `Stop`

- if both attempt to cross, the result is a fatal traffic accident.

There are 4 outcomes depending on the choices made by each of the 2 drivers

Model.

1. Who play ?
2. Which actions/strategies ?

3. Which payoff according to *strategy profile* ?

Model.

1. Who play ? Alice and Bob
2. Which actions/strategies ?

   `Cross` or `Stop`
3. Which payoff according to *strategy profile* ?

   cost = transport time

Model.

1. Who play ? Alice and Bob
2. Which actions/strategies ?

   `Cross` or `Stop`
3. Which payoff according to *strategy profile* ?

   cost = transport time

		Bob	
		`Cross`	`Stop`
Alice	`Cross`		$(1, 2)$
	`Stop`	$(2, 1)$	

Model.

1. Who play ?

   Alice and Bob
2. Which actions/strategies ?

   `Cross` or `Stop`
3. Which payoff according to *strategy profile* ?

   cost = transport time

		Bob	
		`Cross`	`Stop`
Alice	`Cross`	$(60, 60)$	$(1, 2)$
	`Stop`	$(2, 1)$	

Model.

1. Who play ? Alice and Bob
2. Which actions/strategies ?

   `Cross` or `Stop`
3. Which payoff according to *strategy profile* ?

   cost = transport time

		Bob	
		`Cross`	`Stop`
Alice	`Cross`	$(60, 60)$	$(1, 2)$
	`Stop`	$(2, 1)$	$(5, 5)$

Games in standard form

		Bob	
		Cross	Stop
Alice	Cross	$(60, 60)$	$(1, 2)$
	Stop	$(2, 1)$	$(5, 5)$

- The two strategies of Bob correspond to the two columns.
- The entries of the matrix are the outcomes incurred by the players in each situation.

Rational behavior

rational behavior of player : select strategy which minimizes its cost.

		Bob	
		`Cross`	`Stop`
Alice	`Cross`	$(60, 60)$	$(1, 2)$
	`Stop`	$(2, 1)$	$(5, 5)$

For example :

1. If Bob selects to `Cross`, then Alice would select to `Stop`.
2. If Bob selects to `Stop`, then Alice would select to `Cross`.

Rational behavior

rational behavior of player : select strategy which minimizes its cost.

		Bob	
		`Cross`	`Stop`
Alice	`Cross`	$(60, 60)$	$(1, 2)$
	`Stop`	$(2, 1)$	$(5, 5)$

For example :

1. If Bob selects to `Cross`, then Alice would select to `Stop`.
2. If Bob selects to `Stop`, then Alice would select to `Cross`.

Alice has a rational behavior.

Best response

Best responses of player i :

= the strategies which produce the most favorable outcome for a player

		Bob	
		Cross	Stop
Alice	Cross	$(60, 60)$	$(1, 2)$
	Stop	$(2, 1)$	$(5, 5)$

Best response

Best responses of player i :

$=$ the strategies which produce the most favorable outcome for a player

Bob

		Cross	Stop
Alice	Cross	$(60, 60)$	$(1, 2)$
	Stop	$(2, 1)$	$(5, 5)$

Equilibrium $=$ mutual best responses

Nash Equilibrium

Consider a game with a set of n players $\{1, \dots, n\}$

- Each player has a set of possible strategies S_i
- $s = (s_1, \cdots, s_n)$ is a vector of strategies selected by the players

A pure strategy Nash Equilibrium (NE) is a vector of strategies $s = (s_1, \cdots, s_n)$ such that

$$\forall i,\ \forall s_i',\ \text{we have } c_i(s_1, \cdots, s_i, \cdots, s_n) \le c_i(s_1, \cdots, s_i', \cdots, s_n).$$

Nash Equilibrium

Consider a game with a set of n players $\{1, \ldots, n\}$

- Each player has a set of possible strategies S_i
- $s = (s_1, \cdots, s_n)$ is a vector of strategies selected by the players

A pure strategy Nash Equilibrium (NE) is a vector of strategies $s = (s_1, \cdots, s_n)$ such that

$$\forall i,\ \forall s_i' \text{ we have } c_i(s_i, s_{-i}) \leq c_i(s_i', s_{-i}).$$

Back to example

		Bob	
		Cross	Stop
Alice	Cross	(60, 60)	(1, 2)
	Stop	(2, 1)	(5, 5)

- 2 pure strategy Nash Equilibria :

  (Cross, Stop) and (Stop, Cross)

Back to example

		Bob	
		`Cross`	`Stop`
Alice	`Cross`	$(60, 60)$	$(1, 2)$
	`Stop`	$(2, 1)$	$(5, 5)$

- 2 pure strategy Nash Equilibria :

  (`Cross`, `Stop`) and (`Stop`, `Cross`)
- But which equilibrium should be selected ? Which one will be selected by the system if its converges ?

Back to example

		Bob	
		`Cross`	`Stop`
Alice	`Cross`	$(60, 60)$	$(1, 2)$
	`Stop`	$(2, 1)$	$(5, 5)$

- 2 pure strategy Nash Equilibria :

  (`Cross`, `Stop`) and (`Stop`, `Cross`)

- But which equilibrium should be selected ? Which one will be selected by the system if its converges ?

One solution

Outline

- The game of Chicken
  - Definitions
  - Nash Equilibrium
- Rock-paper-scissors Game
  - Mixed strategy
  - Mixed Nash Equilibrium
- Prisoner's dilemma
- Bonus : Toward learning equilibria

Rock-paper-scissors Game

Rules : 2 players select one strategy from `Rock/Paper/scissors`.

Standard form :

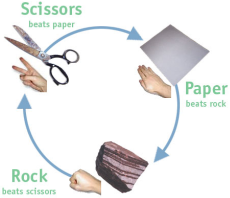

	Scissors	Paper	Rock
Scissors	$(0,0)$	$(-1,1)$	$(1,-1)$
Paper	$(1,-1)$	$(0,0)$	$(-1,1)$
Rock	$(-1,1)$	$(1,-1)$	$(0,0)$

@wikipedia

- No pure strategy Nash equilibrium.

Rock-paper-scissors Game

Rules : 2 players select one strategy from `Rock/Paper/scissors`.

Standard form :

	Scissors	Paper	Rock
Scissors	$(0,0)$	$(-1,1)$	$(1,-1)$
Paper	$(1,-1)$	$(0,0)$	$(-1,1)$
Rock	$(-1,1)$	$(1,-1)$	$(0,0)$

- No pure strategy Nash equilibrium.
- But if players select strategies at random,

$$p_i(Rock) + p_i(Paper) + p_i(Scissors) = 1$$

Rock-paper-scissors Game

Rules : 2 players select one strategy from `Rock/Paper/scissors`.

Standard form :

	Scissors	Paper	Rock
Scissors	$(0,0)$	$(-1,1)$	$(1,-1)$
Paper	$(1,-1)$	$(0,0)$	$(-1,1)$
Rock	$(-1,1)$	$(1,-1)$	$(0,0)$

- No pure strategy Nash equilibrium.
- But if players select strategies at random,

  $$p_i(Rock) + p_i(Paper) + p_i(Scissors) = 1$$

  And if one player prefers one strategy (`Paper`) to the others then his opponent prefers the corresponding winning strategy (`Scissors`) :

Rock-paper-scissors Game

Rules : 2 players select one strategy from `Rock/Paper/scissors`.

Standard form :

	Scissors	Paper	Rock
Scissors	$(0, 0)$	$(-1, 1)$	$(1, -1)$
Paper	$(1, -1)$	$(0, 0)$	$(-1, 1)$
Rock	$(-1, 1)$	$(1, -1)$	$(0, 0)$

- No pure strategy Nash equilibrium.
- But if players select strategies at random,

$$p_i(Rock) + p_i(Paper) + p_i(Scissors) = 1$$

- If each player picks each of his 3 strategies with probability $1/3$,

  then nobody can improve its payoff.

Nash equilibrium of mixed strategies.

Using the random selection method.

Consider a game with a set of n players $\{1, \ldots, n\}$

- Each player has a set of possible pure strategies S_i
- a cost function $c_i : S_1 \times \cdots \times S_n \to \mathbb{N}$

A mixed strategy is a probability distribution p_i over his set of possible pure strategies (actions).

$$\forall i, \sum_{s \in S_i} p_i(s) = 1$$

A mixed profile p is a vector of n elements $(p_1, \ldots, p_n)$ such that player i selects actions using probability p_i.

Expected cost

The expected cost C_i of player i with the game profile p is

$$C_i(p) = E[c_i(p)]$$

		Bob	
		`Cross`	`Stop`
Alice	`Cross`	$(60, 60)$	$(1, 2)$
	`Stop`	$(2, 1)$	$(5, 5)$

Assume

that player Bob decides to pick `Cross` with probability $1/3$,

that player Alice decides to pick `Cross` with probability $1/2$

$$C_{Bob}(p_{Bob}, p_{Alice}) =$$

Expected cost

The expected cost C_i of player i with the game profile p is

$$C_i(p) = E[c_i(p)]$$

		Bob	
		`Cross`	`Stop`
Alice	`Cross`	$(60, 60)$	$(1, 2)$
	`Stop`	$(2, 1)$	$(5, 5)$

Assume

that player Bob decides to pick `Cross` with probability $1/3$,

that player Alice decides to pick `Cross` with probability $1/2$

$$C_{Bob}(p_{Bob}, p_{Alice}) = \frac{1}{3} \times \frac{1}{2} \times 60 +$$

Expected cost

The expected cost C_i of player i with the game profile p is

$$C_i(p) = E[c_i(p)]$$

		Bob	
		`Cross`	`Stop`
Alice	`Cross`	$(60, 60)$	$(1, 2)$
	`Stop`	$(2, 1)$	$(5, 5)$

Assume

that player Bob decides to pick `Cross` with probability $1/3$,

that player Alice decides to pick `Cross` with probability $1/2$

$$C_{Bob}(p_{Bob}, p_{Alice}) = \frac{1}{3} \times \frac{1}{2} \times 60 + \frac{1}{6} \times 1 + \frac{1}{3} \times 2 + \frac{1}{3} \times 5$$

Expected cost

The expected cost C_i of player i with the game profile p is

$$C_i(p) = E[c_i(p)]$$

		Bob	
		`Cross`	`Stop`
Alice	`Cross`	$(60, 60)$	$(1, 2)$
	`Stop`	$(2, 1)$	$(5, 5)$

Assume

that player Bob decides to pick `Cross` with probability $1/3$,

that player Alice decides to pick `Cross` with probability $1/2$

$$C_{Bob}(p_{Bob}, p_{Alice}) = \frac{75}{6}$$

Best response of a mixed strategy

		Bob	
		Cross	Stop
Alice	Cross	$(60, 60)$	$(1, 2)$
	Stop	$(2, 1)$	$(5, 5)$

$$\texttt{Bob}\text{ picks }\begin{cases} \texttt{Cross} & \text{with probability } q \\ \texttt{Stop} & \text{with probability } 1 - q \end{cases}$$

What happens for `Alice` ?

When does `Alice` select the action `Cross` ?

Best response of a mixed strategy

Bob

		`Cross`	`Stop`
Alice	`Cross`	$(60, 60)$	$(1, 2)$
	`Stop`	$(2, 1)$	$(5, 5)$

$$\texttt{Bob}\ \text{picks} \begin{cases} \texttt{Cross} & \text{with probability } q \\ \texttt{Stop} & \text{with probability } 1 - q \end{cases}$$

What happens for `Alice` ?

if `Alice` selects the action `Cross`, then

$$\text{expected cost} = 60q + 1(1 - q)$$

if `Alice` selects the action `Stop`, then

$$\text{expected cost} = 2q + 5(1 - q)$$

When does `Alice` select the action `Cross` ?

Best response of a mixed strategy

Bob

		Cross	Stop
Alice	Cross	$(60, 60)$	$(1, 2)$
	Stop	$(2, 1)$	$(5, 5)$

$$\text{Bob picks} \begin{cases} \texttt{Cross} & \text{with probability } q \\ \texttt{Stop} & \text{with probability } 1 - q \end{cases}$$

What happens for Alice ?

if Alice selects the action Cross, then

expected cost $= 60q + 1(1 - q)$

if Alice selects the action Stop, then

expected cost $= 2q + 5(1 - q)$

When does Alice select the action Cross ?

if $60q + 1(1 - q) < 2q + 5(1 - q)$, in others words

Best response of a mixed strategy

Bob

		Cross	Stop
Alice	Cross	$(60, 60)$	$(1, 2)$
	Stop	$(2, 1)$	$(5, 5)$

$$\text{Bob picks} \begin{cases} \texttt{Cross} & \text{with probability } q \\ \texttt{Stop} & \text{with probability } 1-q \end{cases}$$

What happens for Alice ?

if Alice selects the action Cross, then

$$\text{expected cost} = 60q + 1(1-q)$$

if Alice selects the action Stop, then

$$\text{expected cost} = 2q + 5(1-q)$$

When does Alice select the action Cross ?

if $q < 2/31$

Using the random selection method...

1. **Alice** selects **Cross** if $q < 2/31$

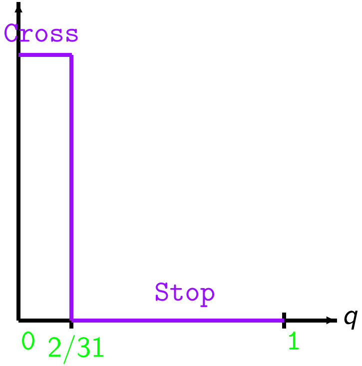

Using the random selection method...

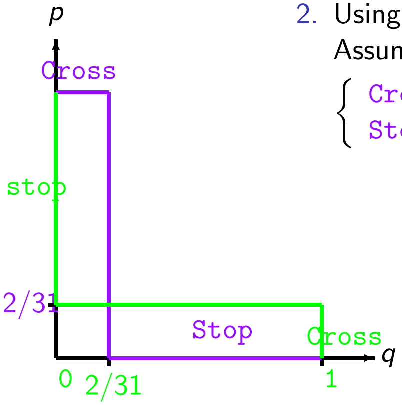

1. `Alice` selects `Cross` if $q < 2/31$
2. Using the same argument as previously :
   Assume that `Alice` selects
   $$\begin{cases} \texttt{Cross} & \text{with probability } p \\ \texttt{Stop} & \text{with probability } 1-p \end{cases},$$

   `Bob` selects `Cross` if $p < 2/31$

Using the random selection method...

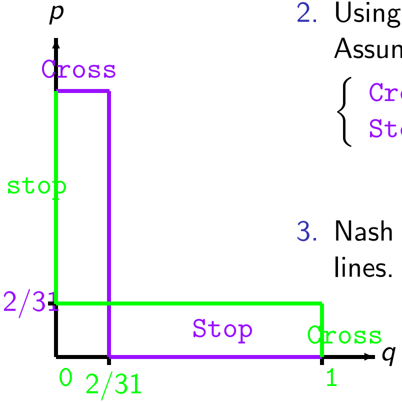

1. Alice selects Cross if $q < 2/31$
2. Using the same argument as previously : Assume that Alice selects
   $$\begin{cases} \texttt{Cross} & \text{with probability } p \\ \texttt{Stop} & \text{with probability } 1-p \end{cases},$$
   Bob selects Cross if $p < 2/31$
3. Nash Equilibrium = intersection of the both lines.

Using the random selection method...

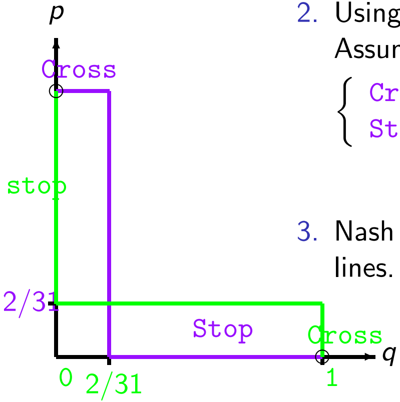

1. **Alice** selects **Cross** if $q < 2/31$
2. Using the same argument as previously : Assume that **Alice** selects

   $$\begin{cases} \texttt{Cross} & \text{with probability } p \\ \texttt{Stop} & \text{with probability } 1-p \end{cases},$$

   Bob selects **Cross** if $p < 2/31$
3. Nash Equilibrium = intersection of the both lines.

   Two Nash Equilibria of pure strategies

Using the random selection method...

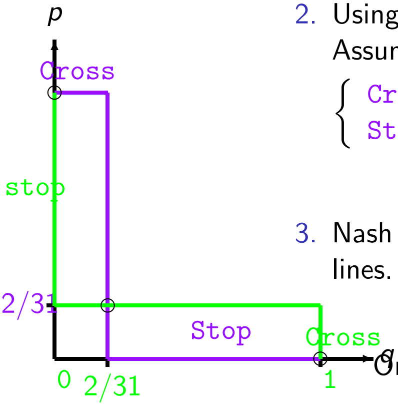

1. **Alice** selects **Cross** if $q < 2/31$
2. Using the same argument as previously : Assume that **Alice** selects

$$\begin{cases} \texttt{Cross} & \text{with probability } p \\ \texttt{Stop} & \text{with probability } 1 - p \end{cases},$$

Bob selects **Cross** if $p < 2/31$

3. Nash Equilibrium $=$ intersection of the both lines.

Two Nash Equilibria of pure strategies

One Nash Equilibrium of mixed strategies

Mixed Nash equilibrium

Consider a game with a set of n players $\{1, \ldots, n\}$

- Each player has a set of possible pure strategies S_i
- a cost function $c_i : S_1 \times \cdots \times S_n \to \mathbb{N}$

A mixed Nash equilibrium is a profile $p^* = (p^*_1, \cdots, p^*_n)$ such that

$$\forall i,\ \forall p'_i \in \mathcal{P}_i \text{ we have } C_i(p^*_i, p^*_{-i}) \leq C_i(p'_i, p^*_{-i}).$$

$\mathcal{P}_i$ the set of mixed strategies of i.

back to example

Nash's Theorem

Théorème [Nash51]

Every finite game (with a finite set of players and finite set of strategies) has a mixed strategy Nash equilibrium

Recall : there is a game without
pure strategy Nash equilibrium.

Outline

Prisoner's dilemma : statement

`Bob` and `Alice` that committed a crime
are interviewed separately by the police.

The offer of the police is the following :

1. If only one of them confess, then he/she will be relaxed and the other will get a sentense of 10 years.
2. If they both remain silent, then they both will have to serve prison sentences of 1 year.
3. if they both confess then they both will get a sentense of 8 years.

They have two strategies : `Confess` or `Silent`.

Standard form

Two strategies : `Confess` or `Silent`.

		Bob	
		`Confess`	`Silent`
Alice	`Confess`	$(8, 8)$	$(0, 10)$
	`Silent`	$(10, 0)$	$(1, 1)$

- The strategy `Confess` dominates strategy `Silent`.

$$\forall s \in S_i \;\; c_i(s, s_{-i}) \geq c_i(\texttt{Confess}, s_{-i})$$

Standard form

Two strategies : Confess or Silent.

		Bob	
		Confess	Silent
Alice	Confess	(8, 8)	(0, 10)
	Silent	(10, 0)	(1, 1)

- The strategy Confess dominates strategy Silent.

$$\forall s \in S_i \;\; c_i(s, s_{-i}) \geq c_i(\texttt{Confess}, s_{-i})$$

- (Confess, Confess) is a Nash equilibrium
- (Silent, Silent) is more favorable than (Confess, Confess)

Domination in the sense of Pareto

Bob

		`Confess`	`Silent`
Alice	`Confess`	$(8, 8)$	$(0, 10)$
	`Silent`	$(10, 0)$	$(1, 1)$

Definition : Profile $\hat{s}$ Pareto-dominates profile s if

1. $\forall i, c_i(\hat{s}) \leq c_i(s)$,
2. $\exists j, c_j(\hat{s}) < c_j(s)$,

Remark : (`Silent`, `Silent`) Pareto-dominates (`Confess`, `Confess`).

Domination in the sense of Pareto

Notion of cooperation

Definition : Profile $\hat{s}$ Pareto-dominates profile s if

1. $\forall i, c_i(\hat{s}) \leq c_i(s)$,
2. $\exists j, c_j(\hat{s}) < c_j(s)$,

Remark : (`Silent`, `Silent`) Pareto-dominates (`Confess`, `Confess`).

And if games are repeated ?

- a fixed number k of times,
  - `Confess` dominates `Silent` at step k of the repeated game ; the two players hence play `Confess`.
  - same reasoning for the last but one step.
  - players play `Confess` at time $k, k-1, \cdots, 1$

And if games are repeated ?

- a fixed number k of times,
  - `Confess` dominates `Silent` at step k of the repeated game ;
    the two players hence play `Confess`.
  - same reasoning for the last but one step.
  - players play `Confess` at time $k, k-1, \cdots, 1$

Introduction of a probability δ
that the game continues for one more step

And if games are repeated ?

- a fixed number k of times,
  - `Confess` dominates `Silent` at step k of the repeated game ; the two players hence play `Confess`.
  - same reasoning for the last but one step.
  - players play `Confess` at time $k, k-1, \cdots, 1$

Introduction of a probability δ
that the game continues for one more step

- in a infinite number of steps,
  - Strategies $=$ mixed strategies in the static game.
  - Construction of a strategy of behaviors that correspond to a simulation of mixed strategy S and if a player i deviates, then it is punished

Outline

- The game of Chicken
  - Definitions
  - Nash Equilibrium
- Rock-paper-scissors Game
  - Mixed strategy
  - Mixed Nash Equilibrium
- Prisoner's dilemma
- Bonus : Toward learning equilibria

Toward learning equilibria

- The same game is repeated at each step.
- At every step t, player i has to solve the following problem :

Which action to play at time t, given the past history of the game ?

that is to say

for all players `i`, $v_i(t) = f(Q)$,

where f_i is a function that gives the behavior of `i` in function of history Q.

A dynamic : fictitious player

A player is a fictitious player if the player has the following behavior :

the player will play a best response
in function of the past statistic of
of strategies of his/her adversary,

That is to say

If player 2 used n_j times the action `j` between step 1 and $t-1$, then player 1 will estimate

that player 2 will play the action `i`
with probability $q_{2,j}(t) = \frac{n_j}{t-1}$ at time t.

A dynamic : fictitious player

		player 2	
		1	2
player 1	1	(3,1)	(0,3)
	2	(1,2)	(2,0)

- Going from a discrete time to a continuous time
- The system = the couple $(q_{1,1}, q_{2,1})$
  with $q_{i,1}$ = probability that player `i` plays strategy 1.

A dynamic : fictitious player

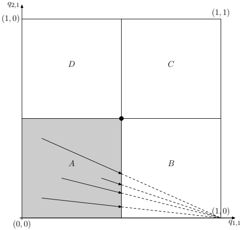

direction of the dynamic
in zone A

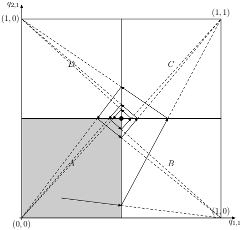

example of behavior
of the dynamic

For zone A :

- player 1 will be willing to use pure strategy 2, and player 2 pure strategy 1.
- the dynamic $(q_{1,1}, q_{2,1})$ will stay in A up to time $t+\tau$ for small $\tau > 0$.

  So $q_{2,1}(t+\tau) = \frac{tq_{2,1}(t)}{t+\tau}$. By making converging $\tau \to 0$, we obtain

$$q'_{2,1}(t) = \frac{q_{2,1}(t)}{t}.$$

Questions ?

Prisoner's dilemma (other interpretation)

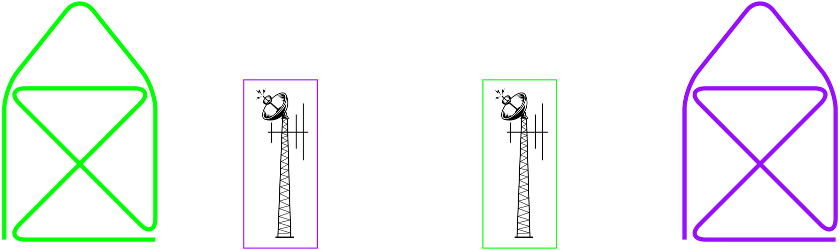

J 2

		transmit	$\overline{\text{transmit}}$
J 1	transmit	(1-c , 1-c)	(-c , 1)
	$\overline{\text{transmit}}$	(1 , -c)	(0 , 0)

- $c > 0$ is the cost of traffic,
- "1" represents the fact that packets reach the destination.

Domination in the sense of Pareto

		J 2	
		transmit	$\overline{\text{transmit}}$
J 1	transmit	(1-c , 1-c)	(-c , 1)
	$\overline{\text{transmit}}$	(1 , -c)	(0 , 0)

Remark :

(transmit,transmit) is more favorable than ($\overline{\text{transmit}}$,$\overline{\text{transmit}}$).

Definition : The profile $\hat{s}$ Pareto-domine the profil s si

1. $\forall i, u_i(\hat{s}) \geq u_i(s)$,
2. $\exists j, u_j(\hat{s}) > u_j(s)$,

Domination in the sense of Pareto

Notion of cooperation

Remark :

($\texttt{transmit}$,$\texttt{transmit}$) is more favorable than ($\overline{\texttt{transmit}}$,$\overline{\texttt{transmit}}$).

Definition : The profile $\hat{s}$ Pareto-domine the profil s si

1. $\forall i, u_i(\hat{s}) \geq u_i(s)$,
2. $\exists j, u_j(\hat{s}) > u_j(s)$,

Questions ?