PREFERENCES AND UTILITIES & NORMAL FORM GAMES

Heinrich H. Nax & Bary S. R. Pradelski

hnax@ethz.ch bpradelski@ethz.ch

March 06, 2017: Lecture 3

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Introduction

We talked in previous lectures about the difference between cooperative and non-cooperative game theory.

We now focus on non-cooperative game theory where the sets of actions of *individual* players are the primitives of the games.

We thus focus on strategic interactions between self-interested, independent agents:

- Chess: the game between two opponents
- Cold war: the game between U.S.A. and Soviet Union
- Rock-paper-scissors

The focus on a single player

To rigorously analyze such strategic interactions between separate agents (e.g., individuals, firms, countries) we need to define

Preferences: what does each individual strive for in the interaction

If we can express these preferences through a real-valued function we gain analytical tractability:

Utilities: a real-valued function expressing a player's preferences

Preferences

Let $x \in X$ be the set of decision alternatives for a player

Definition: binary relation

A *binary relation* $\succeq$ on a set X is a non-empty subset $P \subset X \times X$. We write $x \succeq y$ if and only if $(x, y) \in P$.

$x \succeq y$: "the player weakly prefers x over y"

Define two associated binary relations:

$x \succ y$: "the player strictly prefers x over y"

$x \sim y$: "the player is indifferent between x and y"

Examples

- 'I prefer a 6 over a 5' ($6 \succeq 5$) and 'I prefer a 4 over a fail' ($4 \succeq$ fail)
  $6 \succeq 4$, $6 \succeq$ fail, $5 \succeq 4$, $5 \succeq$ fail
  $6 \succeq 5$ and $5 \succeq$ fail $\Rightarrow 6 \succeq$ fail
- BMW$\succeq$ Toyota (because it is faster), Ford$\succeq$ BMW (because it is bigger), Toyota $\succeq$ Ford (because it is safer)
  BMW$\succeq$ Toyota$\succeq$ Ford **BUT NOT** BMW $\succeq$ Ford

Axiom 1: Completeness

$\forall x, y \in X : x \succeq y$ or $y \succeq x$ or both

Axioms 2: Transitivity

$\forall x, y, z \in X$: if $x \succeq y$ and $y \succeq z$, then $x \succeq z$

Worse and better alternatives

Axioms 1 and 2

Completeness. $\forall x, y \in X : x \succeq y$ or $y \succeq x$ or both
Transitivity. $\forall x, y, z \in X$: if $x \succeq y$ and $y \succeq z$, then $x \succeq z$

If a preference ordering satisfies completeness and transitivity it is called *rational*.

Let $\succeq$ be a rational preference ordering on X. For $x \in X$ define the subsets of alternatives that are (weakly) *worse/better* than x :

$$W(x) = \{y \in X : x \succeq y\}$$
$$B(x) = \{y \in X : y \succeq x\}$$

Utility function

It is very convenient for the mathematical modeling of an agent with binary relation $\succeq$ if we can find a real-valued function whose expected value the agent aims to maximize.

Definition

A **utility function** for a binary relation $\succeq$ on a set X is a function $u : X \to \mathbb{R}$ such that

$$u(x) \geq u(y) \iff x \succeq y$$

Proposition 1

There exists a utility function for every transitive and complete preference ordering on any countable set.

Proof. Exercise.

Generalizing Proposition 1 to Borel sets

Let X be a Borel set in $\mathbb{R}^n$ ($n \in \mathbb{N}$)

Axiom 3: Positive measurability

$\succeq$ is positive measurable if

- $\forall x \in X$ the sets $W(x)$ and $B(x)$ are Borel sets
- $x \prec y \Rightarrow int\{z \in X : x \prec z \preceq y\} \neq \emptyset$

Proposition 2

There exists a utility function for each complete, transitive, and positively measurable preference ordering on any Borel set.

Proposition 1/2 for continuous preferences

Let $X \subset \mathbb{R}^n$ ($n \in \mathbb{N}$) be a closed set.

Axiom 4: Continuity

$\forall x \in X : B(x)$ and $W(x)$ are closed sets.

Proposition 3

There exists a utility function for each complete, transitive, positively measurable, and continuous preference ordering on any closed set.

Let's play a game!

A fair coin is tossed until head shows for the first time:

- If head turns up first at 1^{st} toss you win 1 CHF
- If head turns up first at 2^{nd} toss you win 2 CHF
- If head turns up first at 3^{rd} toss you win 4 CHF
- ...
- If head turns up first at k^{th} toss you win 2^{k-1} CHF

You have a ticket for this lottery. For which price would you sell it?

https://scienceexperiment.online/gametheory/170306/vote

This game is called the St. Petersburg Paradox.

Utility $\neq$ Payoff

If you only care about expected gain:

$$\begin{aligned}\mathbb{E}[\text{lottery}] &= \frac{1}{2}\cdot 1 + \frac{1}{4}\cdot 2 + \frac{1}{8}\cdot 4 + ... \\ &= \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + ... \\ &= \infty\end{aligned}$$

- Bernoulli suggested in 1738 the theory of diminishing marginal utility of wealth.
- Further, the need for utility characterization under uncertainty arose.

This laid the foundation for ***expected utility theory***.

Expected-utility theory

Let $T = \{\tau_1, ...\tau_m\}$ be a finite set and let X consist of all probability distributions on T:

$$X = \Delta(T) = \{(x_1, ..., x_m) \in \mathbb{R}^m_+ : \sum_{k=1}^{m} x_k = 1\}$$

That is X is the unit simplex in $\mathbb{R}^m$.

Can we define a utility function in this setting?

Bernoulli function / von Neumann-Morgenstern utility function

If $\succeq$ is a binary relation on X representing the agent's preferences over lotteries over T. If there is a function $v : T \to \mathbb{R}$ such that

$$x \succeq y \iff \sum_{k=1}^{m} x_k v(\tau_k) \geq \sum_{k=1}^{m} y_k v(\tau_k)$$

then

$$u(x) = \sum_{k=1}^{m} x_k v(\tau_k)$$

defines a utility function for $\succeq$ on X.

The assumption that preferences can be expressed in this form is called the *expected utility hypothesis*. v is called a *Bernoulli function*.

Existence and independence of irrelevant alternatives

For the expected utility hypothesis to hold it is necessary that:

- $\succeq$ is complete and transitive (Axioms 1 and 2)
- $\succeq$ is continuous (note that X is closed and connected) (Axiom 4)

The representation is not only continuous but also linear we need further axioms (respecting that the choice set only contains probability distributions):

Axiom 5: Independence of irrelevant alternatives

$\forall x, y, z \in X, \forall \lambda \in (0, 1)$:

$$x \succ y \quad \Rightarrow \quad (1-\lambda)x + \lambda z \; \succ \; (1-\lambda)y + \lambda z$$

A comment on independence of irrelevant alternatives

Axiom 5: Independence of irrelevant alternatives

$\forall x, y, z \in X, \forall \lambda \in (0, 1)$:

$$x \succ y \quad \Rightarrow \quad (1-\lambda)x + \lambda z \succ (1-\lambda)y + \lambda z$$

Suppose that $x^* = (1-\lambda)x + \lambda z$ and $y^* = (1-\lambda)y + \lambda z$ are compound lottery.

Then if $x \succ y$ an agent should also have the preference $x^* \succ y^*$ independent of what z is.

Example: Suppose an agent prefers Asian food over Italian but also values an occasional Italian dish.

$$\text{sushi} \succ \text{pizza} \not\Rightarrow (1-\lambda)\cdot\text{sushi} + \lambda\cdot\text{wontons} \succ (1-\lambda)\cdot\text{pizza} + \lambda\cdot\text{wontons}$$

Allais paradox

The set of prices in CHF is $X = \{0; 1,000,000; 5,000,000\}$.

- Which probability do you prefer:
  $p_1 = (0.00; 1.00; 0.00)$ or $p_2 = (0.01; 0.089; 0.10)$?
- Which probability do you prefer:
  $p_3 = (0.90; 0.00; 0.10)$ or $p_4 = (0.89; 0.11; 0.00)$?

Most people report: $p_1 \succ p_2$ and $p_3 \succ p_4$.

Allais paradox

The set of prices in CHF is $X = \{0; 1,000,000; 5,000,000\}$.

- $p_1 = (0.00; 1.00; 0.00)$ or $p_2 = (0.01; 0.089; 0.10)$
- $p_3 = (0.90; 0.00; 0.10)$ or $p_4 = (0.89; 0.11; 0.00)$

Suppose (v_0, v_{1M}, v_{5M}) is a Bernoulli function for $\succeq$.

Then $p_1 \succ p_2$ implies:

$$v_{1M} > .01 \cdot v_0 + .89 \cdot v_{1M} + .1 \cdot v_{5M}$$

$$.11 \cdot v_{1M} - .01 \cdot v_0 > .1 \cdot v_{5M}$$

now add $.9 \cdot v_0$ to both sides:

$$.11 \cdot v_{1M} + .89 \cdot v_0 > .1 \cdot v_{5M} + .9 \cdot v_0$$

But this implies $p_4 \succ p_3$, a contradiction!

Sure thing principle

Sure thing principle (Savage)

A decision maker who would take a certain action if he knew that event B happens and also if he knew that $not - B$ happens, should also take the same action if he know nothing about B.

Lemma

Assume that everything the decision maker knows is true then *sure thing principle is equivalent to independence of irrelevant alternatives*.

Sure thing principle: Savage (1954)

"A businessman contemplates buying a certain piece of property. He considers the outcome of the next presidential election relevant. So, to clarify the matter to himself, he asks whether he would buy if he knew that the Democratic candidate were going to win, and decides that he would. Similarly, he considers whether he would buy if he knew that the Republican candidate were going to win, and again finds that he would. Seeing that he would buy in either event, he decides that he should buy, even though he does not know which event obtains, or will obtain, as we would ordinarily say. It is all too seldom that a decision can be arrived at on the basis of this principle, but except possibly for the assumption of simple ordering, I know of no other extralogical principle governing decisions that finds such ready acceptance."

Axioms

Axioms 1, 2, and 4

Completeness. $\forall x, y \in X : x \succeq y$ or $y \succeq x$ or both
Transitivity. $\forall x, y, z \in X$: if $x \succeq y$ and $y \succeq z$, then $x \succeq z$
Continuity. $\forall x \in X : B(x)$ and $W(x)$ are closed sets.

Axiom 5: Independence of irrelevant alternatives

$\forall x, y, z \in X, \forall \lambda \in (0, 1)$:

$$x \succ y \quad \Rightarrow \quad (1-\lambda)x + \lambda z \succ (1-\lambda)y + \lambda z$$

Von Neumann-Morgenstern

Theorem (von Neumann-Morgenstern)

Let $\succeq$ be a rational (complete & transitive) and continuous preference relation on $X = \Delta(T)$, for any finite set T.
Then $\succeq$ admits a utility function u of the expected-utility form if and only if $\succeq$ meets the axiom of independence of irrelevant alternatives.

Translation invariance

Given a Bernoulli function v for given preferences $\succeq$ let:

$$v' = \alpha + \beta v$$

where $\alpha \in \mathbb{R}$ and $\beta \in \mathbb{R}^+$.
Then v' is also a Bernoulli function for another utility function

$$u' = \alpha + \beta u$$

Expected utility functions are unique up to a positive affine transformation.

Ordinal, cardinal utility functions, and “utils”

Ordinal utility function. A utility function where differences between $u(x)$ and $u(y)$ are meaningless. Only the fact that, for example, $u(x) \geq u(y)$ is meaningful. An ordinal utility function can be subjected to any increasing transformation $f(u)$ which will represent the same preferences $\succeq$.

Cardinal utility function. A utility function where differences between $u(x)$ and $u(y)$ are meaningful as they reflect the *intensity* of preferences. Cardinal utility functions are only invariant to positive affine transformations.

“Utils”. An even stronger statement would be that there is a fundamental measure of utility, say one “util”. Such a utility function is not invariant to any transformation.

Comparing utility: within person

1. "She likes x less than z"
2. "She likes x over z twice as much as y over z"
3. "She likes x five times more than y"

	1.	2.	3.
Ordinal utility function	yes	no	no
Cardinal utility function	yes	yes	no
"Utils"	yes	yes	yes

Comparing utility: interpersonal

Interpersonal comparability (IC). A utility function where utility differences between players make "sense".

Ordinal utility function	possibly IC
Cardinal utility function	possibly IC
"Utils"	$\Rightarrow$ IC

Suppose we have cardinal utility functions that are IC for agent 1 and 2, u_1, u_2. Transform them by some non-affine increasing transformation f resulting in $v_1 = f(u_1), v_2 = f(u_2)$.
Then v_1, v_2 are no longer cardinal but are IC.

Note: Utility functions that are ordinal and do not allow for interpersonal comparisons do not contain more information than preference relations.

Comparing utility: interpersonal

1. "Warren Buffet values 1000 CHF less than a starving child values 1000 CHF"
2. "Eve would pay 10 CHF (utils) for the chocolate, Sarah would pay 5 CHF (utils)"
3. "Mother loves d1 more than d2. Father loves d2 more than d1"

	1.	2.	3.
Ordinal utility function	yes (?)	?	?
Cardinal utility function	yes (?)	?	?
"Utils"	yes (?)	yes (?)	?!?

Comparing utilities between agents implies some welfare statement / judgment.

Utility and risk

Define a lottery:

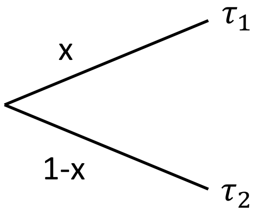

The lottery is a fair gamble if and only if $x \cdot v(\tau_1) = (1 - x) \cdot v(\tau_2)$.

Risk neutrality

Definition: risk neutral

An agent is *risk-neutral* if and only if he is indifferent between accepting and rejecting all fair gambles, that is for all x, τ_1, τ_2:

$$\mathbb{E}[u(\text{lottery})] = x \cdot v(\tau_1) + (1-x) \cdot v(\tau_2)$$
$$= u(x \cdot \tau_1 + (1-x) \cdot \tau_2)$$

An agent is risk-neutral if and only if he has a linear von Neumann-Morgenstern utility function.

Risk aversion

Definition: risk averse

An agent is *risk averse* if and only if he rejects all fair gambles, that is for all x, τ_1, τ_2:

$$\begin{aligned}\mathbb{E}[u(\text{lottery})] &= x \cdot v(\tau_1) + (1-x) \cdot v(\tau_2) \\ &< u(x \cdot \tau_1 + (1-x) \cdot \tau_2)\end{aligned}$$

Recall that a function $g(\cdot)$ is strictly concave if and only if

$$g(\lambda x + (1-\lambda)y) > \lambda g(x) + (1-\lambda)g(y)$$

An agent is risk averse if and only if he has a strictly concave utility function.

Risk seekingness

Definition: risk seeking

An agent is *risk seeking* if and only if he strictly prefers all fair gambles, that is for all x, τ_1, τ_2:

$$\begin{aligned}\mathbb{E}[u(\text{lottery})] &= x \cdot v(\tau_1) + (1-x) \cdot v(\tau_2) \\ &> u(x \cdot \tau_1 + (1-x) \cdot \tau_2)\end{aligned}$$

An agent is risk seeking if and only if he has a strictly convex utility function.

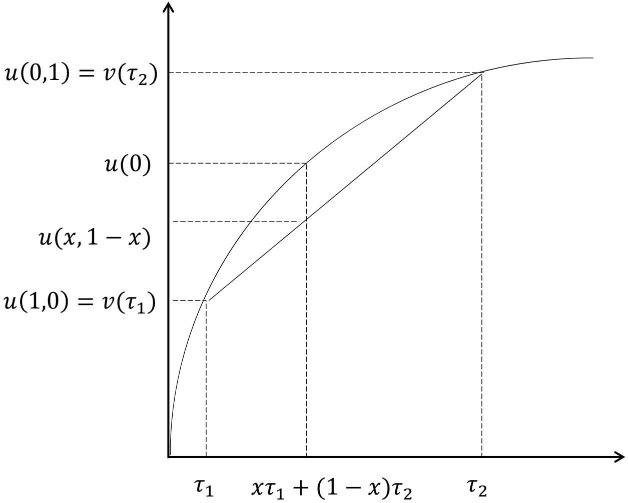
$u(0,1) = v(\tau_2)$
$u(0)$
$u(x, 1-x)$
$u(1,0) = v(\tau_1)$
τ_1
$x\tau_1 + (1-x)\tau_2$
τ_2

Some final remarks

- If you believe that people have preferences, under reasonable axioms we can translate them into utility functions.
- Money is not equal to utility (recall diminishing marginal utility).
- Preferences do not have to be self-regarding ("homo oeconomicus").

Plan

- Introduction normal form games
- Dominance in pure strategies
- Nash equilibrium in pure strategies
- Best replies
- Dominance, Nash, best replies in mixed strategies
- Nash's theorem and proof via Brouwer
- Equilibrium refinements

The Prisoner's Dilemma

"Two suspects are arrested and interviewed separately. If they both keep quiet (i.e., *cooperate*) they go to prison for one year. If one suspect supplies evidence (*defects*) then that one is freed, and the other one is imprisoned for eight years. If both defect then they are imprisoned for five years."

PLAYERS The players are the two suspects $N = \{1, 2\}$.

STRATEGIES The strategy set for player 1 us $S_1 = \{C, D\}$, and for player 2 is $S_2 = \{C, D\}$.

PAYOFFS For example, $u_1(C, D) = -8$ and $u_2(C, D) = 0$. All payoffs are represented in this matrix:

	Cooperate	*Defect*
Cooperate	$-1, -1$	$-8, 0$
Defect	$0, -8$	$-5, -5$

Definition: Normal form game

A normal form (or strategic form) game consists of three object:

1. *Players:* $N = \{1, \ldots, n\}$, with typical player $i \in N$.
2. *Strategies:* For every player i, a finite set of strategies, S_i, with typical strategy $s_i \in S_i$.
3. *Payoffs:* A function $u_i : (s_1, \ldots, s_n) \rightarrow \mathbb{R}$ mapping strategy profiles to a payoff for each player i. $u : S \rightarrow \mathbb{R}^n$.

Thus a normal form game is represented by the triplet:

$$G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$$

Strategies

Definition: strategy profile

$s = (s_1, \dots, s_n)$ is called a *strategy profile*.
It is a collection of strategies, one for each player. If s is played, player i receives $u_i(s)$.

Definition: opponents strategies

Write s_{-i} for all strategies except for the one of player i. So a strategy profile may be written as $s = (s_i, s_{-i})$.

Dominance

A strategy strictly dominates another if it is always better whatever others do.

STRICT DOMINANCE A strategy s_i *strictly dominates* s'_i if $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$ for all s_{-i}.

	Cooperate	*Defect*
~~*Cooperate*~~	~~$-1, -1$~~	~~$-8, 0$~~
Defect	$0, -8$	$-5, -5$

Dominance

A strategy strictly dominates another if it is always better whatever others do.

Strict Dominance A strategy s_i *strictly dominates* s'_i if $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$ for all s_{-i}.

Dominated Strategy A strategy s'_i is *strictly dominated* if there is an s_i that strictly dominates it.

Dominant Strategy A strategy s_i is *strictly dominant* if it strictly dominates all $s'_i \neq s_i$.

If players are rational they should never play a strictly dominated strategy, no matter what others are doing, they may play weakly dominated strategies:

Weak Dominance A strategy s_i *weakly dominates* s'_i if $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$ for all s_{-i}.

Dominant-Strategy Equilibrium

Definition: Dominant-Strategy Equilibrium

The strategy profile s^* is a *dominant-strategy equilibrium* if, for every player i, $u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$ for all strategy profiles $s = (s_i, s_{-i})$.

Example: Prisoner's dilemma

	Cooperate	*Defect*
Cooperate	$-1, -1$	$-8, \quad 0$
Defect	$0, -8$	$-5, -5$

(D, D) is the (unique) *dominant-strategy equilibrium*.

Common knowledge of rationality and the game

Suppose that players are rational decision makers and that mutual rationality is common knowledge, that is:

- I know that she knows that I will play rational
- She knows that “I know that she knows that I will play rational”
- I know that “She knows that “I know that she knows that I will play rational””
- ...

Further suppose that all players know the game and that again is common knowledge.

Iterative deletion of strictly dominated strategies

If the game and rationality of players are common knowledge, iterative deletion of strictly dominated strategies yields the set of "rational" outcomes.

	L	C	R
T	1, 1	2, 0	2, 2
M	0, 3	1, 5	4, 4
B	2, 4	3, 6	3, 0

Note: Iteratively deletion of strictly dominated strategies is independent of the order of deletion.

Battle of the Sexes

PLAYERS The players are the two students $N = \{row, column\}$.

STRATEGIES Row chooses from $S_{row} = \{Cafe, Pub\}$
Column chooses from $S_{column} = \{Cafe, Pub\}$.

PAYOFFS For example, $u_{row}(Cafe, Cafe) = 4$. The following matrix summarises:

	Cafe	*Pub*
Cafe	4, 3	1, 1
Pub	0, 0	3, 4

Battle of the Sexes

In this game, nothing is dominated, so profiles like (Cafe, Pub) are not eliminated. Should they be?

- Column player would play Cafe if row player played Cafe!
- Row player would play Pub if column player played Pub!

In other words, after the game, both players may "regret" having played their strategies.

This a truly interactive game – best responses depend on what other players do ... next slides!

Nash Equilibrium

Definition: Nash Equilibrium

A *Nash equilibrium* is a strategy profiles s^* such that for every player i,

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*) \text{ for all } s_i$$

At s^*, no i regrets playing s_i^*. Given all the other players' actions, i could not have done better

In words: If no player has an incentive to deviate from their part in a particular strategy profile, then it is Nash equilibrium.

Best-Reply functions

What should each player do given the choices of their opponents? They should "best reply".

Definition: best-reply function

The *best-reply function* for player i is a function B_i such that:

$$B_i(s_{-i}) = \{s_i | u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \text{ for all } s_i'\}$$

Best-reply functions in Nash

Nash equilibrium can be redefined using best-reply functions:

Definition: Nash equilibrium

s^* is a *Nash equilibrium* if and only if $s_i^* \in B_i(s_{-i}^*)$ for all i.

In words: a Nash equilibrium is a strategy profile of mutual best responses each player picks a best response to the combination of strategies the other players pick.

Example

For the Battle of the Sexes:

- $B_{row}(Cafe) = Cafe$
- $B_{row}(Pub) = Pub$
- $B_{column}(Cafe) = Cafe$
- $B_{column}(Pub) = Pub$

So (Cafe, Cafe) is a Nash equilibrium and so is (Pub, Pub) …

Cook book: how to find pure-strategy Nash equilibria

The best way to find (pure-strategy) Nash equilibria is to underline the best replies for each player:

	L	C	R
T	$\underline{5}, 1$	$2, 0$	$2, \underline{2}$
M	$0, 4$	$1, \underline{5}$	$\underline{4}, \underline{5}$
B	$2, 4$	$\underline{3}, \underline{6}$	$1, 0$

Matching Pennies

"Each player has a penny. They simultaneously choose whether to put their pennies down heads up (H) or tails up (T). If the pennies match, column receives row's penny, if they don't match, row receives columns' penny."

PLAYERS The players are $N = \{row, column\}$.

STRATEGIES Row chooses from $\{H, T\}$; Column from $\{H, T\}$.

PAYOFFS Represented in the strategic-form matrix:

	H	T
H	$\underline{1}, -1$	$-1, \underline{1}$
T	$-1, \underline{1}$	$\underline{1}, -1$

- Best replies are: $B_{row}(H) = H, B_{row}(T) = T, B_{column}(T) = H$, and $B_{column}(H) = T$
- There is no pure-strategy Nash equilibrium in this game

Plan

- Introduction normal form games
- Dominance in pure strategies
- Nash equilibrium in pure strategies
- Best replies
- Dominance, Nash, best replies in mixed strategies
- Nash's theorem and proof via Brouwer
- Equilibrium refinements

Hawk-dove game

Player 2

Player 1		Hawk	Dove
	Hawk	-2,-2	4,0
	Dove	0,4	2,2

Harmony game

		Company B	
		Cooperate	Not Cooperate
Company A	Cooperate	9,9	4,7
	Not Cooperate	7,4	3,3

A three player game

L

	l	r
T	$0, \underline{21}, 0$	$-10, 11, \underline{1}$
B	$\underline{10}, 0, -10$	$\underline{0}, \underline{10}, \underline{11}$

R

	l	l
T	$\underline{1}, \underline{11}, \underline{10}$	$\underline{11}, 1, -9$
B	$-9, 10, \underline{0}$	$1, \underline{20}, 1$

THANKS EVERYBODY
See you next week!
and keep checking the website for new materials as we progress:
http://www.coss.ethz.ch/education/GT.html