Introduction to Game Theory

Isabel Amigo

IMT Atlantique

October 2017

Notes:

Notes:

What are games with incomplete information?

In a game of complete information, all players know the rules of the game. Otherwise the game is one of incomplete information.

hmm? Old definition

Notes:

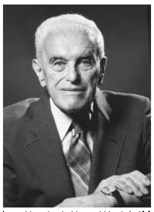

John Charles Harsanyi (Hungarian: Harsányi János Károly) (May 29, 1920 – August 9, 2000), Hungarian-Australian-American economist (wikipedia no-free media

Notes:

- Harsanyi fromalized games with uncertainty
- Introduced a new player called Nature, does not have an utility function (or has a constant one)
- Nature selects "a state of the world"

What peace of information is missing in games with incomplete information?

- Players do not exactly know which game is being played, but they know all the possible games
- All possible games have the same set of players
- and the same set of actions
- Games differ in the payoffs
- Hypothesis: common priori (belief)

Notes:

- What do they believe is possible? We talk about aposteriori beliefs, they start with a common apriori, with an individual private known information and applying Bayes they have an aposteriori belief.
- There are different definitions, which are mathematically equivalent, we shall see one
- Typically a cards games where each players see only their cards can be modeled as a Bayesian game.
- Note, however, that we will focus on simultaneous games.

Gaining Intuition, Example

Bob (Player 1) is unsure about Alice (player 2), wanting to go out with him, or avoid it. Alice knows Bob's preferences.

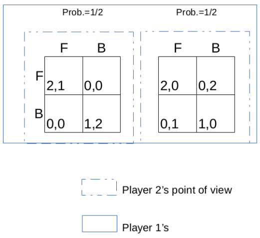

Notes:

- The example is taken from An Introduction to Game Theory, by Osborne
- In the example, we can think of there being two states, one in which the players' payoffs are given in the left table and one in which they are given in the right
- To find the equilibrium of this game, we can first find the expected payoff Bob has for each possible combination of actions taken by Alice (one action at each game) taking into account his own action, and the probability of having one or other game being played
- Then we can determine Bob's best responses given Alice's combination of actions
- We can finally look for NE such that all players are best responding the other one

Definition based on types

Bayesian Game

$\cdot G = (N, A, \Theta, p, u)$
$\cdot \Theta$ set of type spaces for each player, private information for each player
$\cdot$ p is still common prior over types
$\cdot$ u set of utility functions for each player

Notes:

- This examples follows one possible definition of Bayesian games:
  - This definition is based on defining a set of games, where the set of players and actions are the same across all games
  - We also have a common joint probability distribution, which says the probability of having each game played
  - And we have a set of partitions of games, one for each player
- There are three equivalent definitions, we shall following work with the one based in types:
- Formally $G = (N, A, \Theta, p, u)$ with $\Theta = \{\Theta_i\}_{i \in N}$ where Θ_i is the type space of player i
- The type captures any uncertainty
- $A = \{A_i\}_{i \in N}$ where A_i is the set of available actions for player i
- $p : \Theta \to [0, 1]$ common prior over types
- $u = \{u_i\}_{i \in N}$ where $u_i : A \times \Theta \to \mathbb{R}$ is the utility function of player i
- this formulation can be less intuitive than the one based on games, but is mathematically simpler

Definitions and notations

Pure strategy

For player i is a function mapping i's types into an action

Mixed Strategy

A function mapping i's types into a distribution probability of i's available actions

Notes:

- Pure strategy $\alpha_i : \Theta_i \to A_i$
- Mixed strategy $s_i : \Theta_i \to \Pi(A_i)$
- We shall also use the notation $s_i(a_i|\theta_i)$ to refer to the proba. under mixed strategy s_i that player i choses a_i given that his type is θ_i
- We shall consider finite sets (players, strategies, types).

Notions of expected utility

Ex-ante
Only the prior is known (p).

Interim
Player i knows his own type θ_i, and has a prior over others' types.

Ex-post
All types are known by everyone (so complete information game). Only uncertainty: other agents' strategies.

Notes:

- Since players know different things at different moments, we should re visit the concept of expected utility
- Let us denote $EU_i(s)$ as the expected utility of function u_i under mixed strategy s, since we do not include the realizations of the types this is also known as the Exante expected utility
- Exante $EU_i(s) = \sum_{\theta \in \Theta} p(\theta) \sum_{a \in A} u_i(a, \theta) \prod_{j \in N} s_j(a_j|\theta_j)$
- Interim expected utility $EU_i(s|\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} p(\theta_{-i}|\theta_i) \sum_{a \in A} u_i(a, \theta_{-i}, \theta_i) \prod_{j \in N} s_j(a_j|\theta_j)$
- Ex-post $EU_i(s, \theta) = \sum_{a \in A} u_i(a, \theta) \prod_{j \in N} s_j(a_j|\theta_j)$

Solution concept: Bayes-Nash equilibrium (BNE)

Best Response
for player i is the set of strategies for player i that maximize the ex-pected ex-ante utility, given other player's strategies

Bayes NE
Is a strategy profile where all player is best responding the other players

Notes:

- Equivalently, is a strategy profile s such that $s_i \in$ argmax $EU_i(s_i, s_{-i}|\theta_i)$ $\forall i \in N$ and all $\theta_i \in \Theta_i$
- Notion introduced by Harsanyi
- Formally, the set best response for player i is $BR_i(s_i) = \text{argmax}_{s_i} EU_i(s_i, s_{-i})$
- BNE is then defined as a strategy profile s such that $\forall$ i that $s_i \in BR_i(s_i)$
- Note that if all types have positive probability, then we can equivalently define best response and NE considering the interim expected utility

Example: Finding BNE

θ_T, θ_L:

	L	R
T	2,0	0,2
B	0,2	2,0

θ_T, θ_R:

	L	R
T	2,2	0,3
B	3,0	1,1

θ_B, θ_L:

	L	R
T	2,2	0,0
B	0,0	1,1

θ_B, θ_R:

	L	R
T	2,1	0,0
B	0,0	1,2

$p(\theta_T, \theta_R) = 0.1$
$p(\theta_T, \theta_L) = 0.3$
$p(\theta_B, \theta_R) = 0.4$
$p(\theta_B, \theta_L) = 0.2$

Notes:

- As for extensive form games, we can find an equivalent Strategic game for a Bayesian game
- We can find the BNE in the strategic way, as we already know how to do it
- Lets consider the provided example, originally proposed in Multiagent Systems Algorithmic, Game-Theoretic, and Logical Foundations, by Shoham and Leyton-Brown
- We could also calculate the conditional probabilities using Bayes' rule
- We will find the equivalent normal form game, by listing all possible pure strategies for each player and computing the exante expected payoff for each strategy profile
- the strategies will be the actions of the normal form game

Expost NE

Expost NE
Is a strategy profile such that no player would ever want to deviate from his strategy even if he knew the complete type vector (θ).

Notes:

- Mathematically, a expost NE is a strategy profile s verifying for all θ and for all i s_i=argmax $EU_i(s_i, s_{-i}, \theta)$
- This kind of equilibrium is not guaranteed to exist

Quizz Lecture 3

The questions proposed here are taken from: the MOOC Game Theory on Coursera platform, created by Matthew Jackson, Kevin Leyton-Brown and Yoav Shoham and from the Book An Introduction to Game Theory, by Martin J. Osborne.

Exercise 1 (Bayesian Games) *In the following two-player Bayesian game, the payoffs to player 2 depend on whether 2 is a friendly player (with probability p) or a enemy (with probability 1p). See the following payoff matrices for details.*

With probability p, the payoff matrix is:

Friend		*Player 2*	
		Left	*Right*
Player 1	*Left*	*(3,1)*	*(0,0)*
	Right	*(2,1)*	*(1,0)*

while with probability 1p, the payoff matrix is:

Enemy		*Player 2*	
		Left	*Right*
Player 1	*Left*	*(3,0)*	*(0,1)*
	Right	*(2,0)*	*(1,1)*

Player 2 knows if he/she is a friend or a enemy, but player 1 doesn't know. If player 2 uses a strategy of Left when a friend and Right when a enemy, what is true about player 1's expected utility?

a. *It is 3 when 1 chooses Left;*

b. *It is 3p when 1 chooses Left;;*

c. *It is 2p when 1 chooses Right;*

d. *It is 1 when 1 chooses Right;*

Exercise 2 (Bayesian Games) *Consider the conflict game:*

With probability p, the payoff matrix is:

Strong		*Player 2*	
		Fight	*Not*
Player 1	*Fight*	*(1,-2)*	*(2,-1)*
	Not	*(-1,2)*	*(0,0)*

and with probability $1-p$, the payoff matrix is:

Weak		*Player 2*	
		Fight	*Not*
Player 1	*Fight*	*(-2,1)*	*(2,-1)*
	Not	*(-1,2)*	*(0,0)*

Assume that player 1 plays fight when strong and not when weak. Given this strategy of player 1, there is a certain p^* such that player 2 will prefer 'fight' when $p < p^*$, and 'not' when $p > p^*$.

What is p^* in this modified game? (Hint: Write down the payoff of 2 when choosing Fight and Not Fight. Equalize these two payoffs to get p^*):

a. 1/3

b. 2/3

c. 1/2

d. 3/4