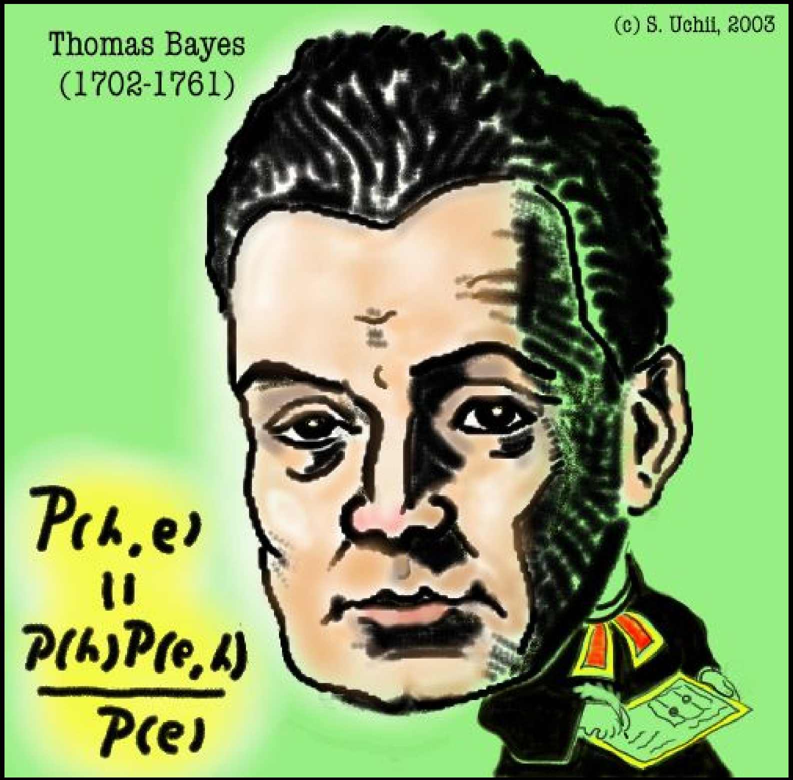

Introduction to Game Theory
Lecture 10

Disclaimer: this presentation is only a supporting material and is not sufficient to master the topics covered during the lecture. Study of relevant books is strongly recommended.

Preview

- So far:

  Games with perfect and complete information
  - structure of the game is known
  - players' payoffs are known
  - everything is a common knowledge

- Today:

  Bayesian games
  - model situations in which player is imperfectly informed about other players' preferences

Preview

Technically…

- So far:
  - Games with perfect and complete information
  - optimal behavior involves choice of actions
- Today:
  - Bayesian games
  - optimal behavior involves choice of actions AND beliefs about other player's type/preferences

Bayesian Games

- Bayesian games:
  - auctions: value of a given object to other people is unknown
  - firms: cost function of other firms is unknown

- Example: modified Battle of Sexes game
  - two possible types of one player
  - Homer is not sure whether
    - Marge wants to go out with him
    - Marge wants to avoid him
  - As before, Marge knows Homer's preferences

Modified Battle of Sexes

- Homer (based on experience):
  - with probability ½ Marge wants to meet him
    -> playing left game
  - with probability ½ Marge wants to avoid him
    -> playing right game

H \ M	B	S
B	2,1	0,0
S	0,0	1,2

H \ M	B	S
B	2,0	0,2
S	0,1	1,0

Modified Battle of Sexes

- Marge knows if she wants to meet or avoid Homer
- Homer does not know what type Marge is
- To make rational decision, Homer has to form beliefs about the action of each type of Marge (difference between games with complete and incomplete information)
- Given these beliefs he computes expected payoff of each action and chooses optimally

Expected Payoff

Example: if Homer believes that Marge who wants to meet him chooses B and Marge who wants to avoid him chooses B, then:

- Homer: B -> 0.5*2+0.5*2 = 2
- Homer: S -> 0.5*0+0.5*0 = 0

H \ M	B	S
B	2,1	0,0
S	0,0	1,2

H \ M	B	S
B	2,0	0,2
S	0,1	1,0

Expected Payoff

Example 2: if Homer believes that Marge who wants to meet him chooses B and Marge who wants to avoid him chooses S, then:

- Homer: B -> 0.5*2+0.5*0 = 1
- Homer: S -> 0.5*0+0.5*1 = 0.5

H \ M	B	S
B	2,1	0,0
S	0,0	1,2

H \ M	B	S
B	2,0	0,2
S	0,1	1,0

Expected Payoff

- Similarly, we can find expected payoff of all combinations of actions for Homer and Marge (columns B,B and B,S are derived on previous slides)

H \ M	B,B	B,S	S,B	S,S
B	2	1	1	0
S	0	½	½	1

Bayesian Games - NE

- For this type of game, Nash Equilibrium is:
  - the action of player 1 is optimal, given the actions of the two types of player 2 (and player 1's belief about the state)
  - the action of each type of player 2 is optimal, given the action of player 1
- given Homer's beliefs about probability of Marge wanting to meet or avoid him, he behaves optimally
- Marge behaves optimally, whether she wants to meet Homer or avoid him

Bayesian Games - NE

- we analyze this type of Bayesian game as if there where three players: Homer, Marge - meet, Marge - avoid

- Marge knows her type, but Homer does not know Marge's type he needs to specify his optimal action given both possibilities

Bayesian Games - NE

- game can be represented in one table
- first number in each cell represents Homer's expected payoff, second number is payoff of 1st type Marge and the third one payoff of 2nd type Marge

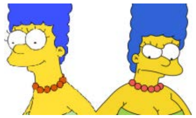
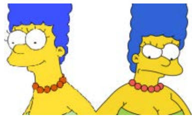
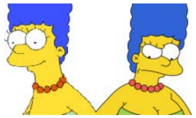
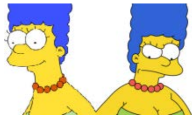

H \ M	B,B	B,S	S,B	S,S
B	2,1,0	1,1,2	1,0,0	0,0,2
S	0,0,1	½,0,0	½,2,1	1,2,0

Bayesian Games - NE

- game can be represented in one table
- first number in each cell represents Homer's expected payoff, second number is payoff of 1st type Marge and the third one payoff of 2nd type Marge

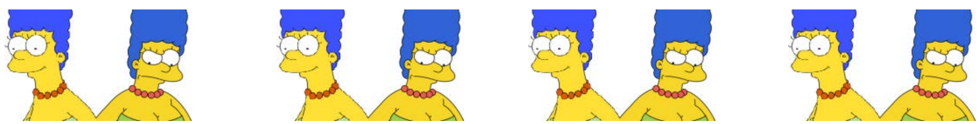

H \ M	B,B	B,S	S,B	S,S
B	2,1,0	1,1,2	1,0,0	0,0,2
S	0,0,1	½,0,0	½,2,1	1,2,0

NE (B; B,S)

Bayesian Games - NE

- interpretation of NE (B,(B,S)):
  - given Homer's beliefs and actions of both types of Marge, Homer is playing the best response
  - given Homer's action, both types of Marge are playing best response

H \ M	B,B	B,S	S,B	S,S
B	2,1,0	1,1,2	1,0,0	0,0,2
S	0,0,1	½,0,0	½,2,1	1,2,0

Bayesian Games - NE

- interpretation of NE if Marge is the 1st type:
  - Marge wants to meet Homer and chooses B
  - Homer chooses B and believes that 1st type Marge chooses B and the 2nd type Marge chooses S

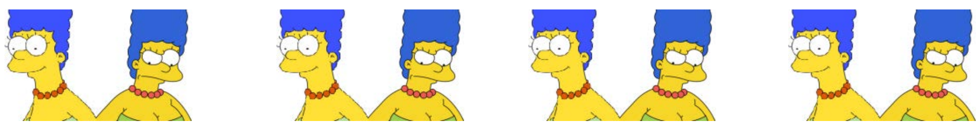

H \ M	B,B	B,S	S,B	S,S
B	2,1,0	1,1,2	1,0,0	0,0,2
S	0,0,1	½,0,0	½,2,1	1,2,0

Bayesian Games - NE

- interpretation of NE if Marge is the 2nd type:
  - Marge wants to avoid Homer and chooses S
  - Homer chooses B and believes that 1st type Marge chooses B and the 2nd type Marge chooses S

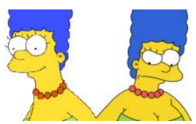
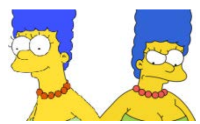
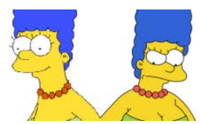
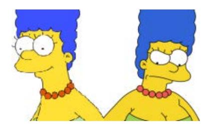

H \ M	B,B	B,S	S,B	S,S
B	2,1,0	1,1,2	1,0,0	0,0,2
S	0,0,1	½,0,0	½,2,1	1,2,0

Too Much Information Hurts

- single-person decision problem – player cannot be worse off if she has more information: if she wishes, she can ignore the information
- strategic game – if a player has more information and the other players know that she has it she may be worse off
- following game has two possible states, each player believes that states S1 and S2 are equally likely

S1	L	M	R
T	1,4	1,0	1,6
B	2,16	0,0	0,24

S2	L	M	R
T	1,4	1,6	1,0
B	2,16	0,24	0,0

Too Much Information Hurts

If P1 (row) believes that P2 (column) will choose L:

EP(T) = ½*1 + ½*1 = 1 EP(B) = ½*2 + ½*2 = 2

If P1 believes that P2 will choose M:

EP(T) = ½*1 + ½*1 = 1 EP(B) = ½*0 + ½*0 = 0

If P1 believes that P2 will choose R:

EP(T) = ½*1 + ½*1 = 1 EP(B) = ½*0 + ½*0 = 0

probability = ½

S1	L	M	R
T	1,4	1,0	1,6
B	2,16	0,0	0,24

probability = ½

S2	L	M	R
T	1,4	1,6	1,0
B	2,16	0,24	0,0

Too Much Information Hurts

If P2 believes that P1 will choose T:

EP(L) = ½*4 + ½*4 = 4 EP(M) = ½*0 + ½*6 = 3

EP(R) = ½*6 + ½*0 = 3

If P2 believes that P1 will choose B:

EP(L) = ½*16 + ½*16 = 16 EP(M) = ½*0 + ½*24 = 12

EP(R) = ½*24 + ½*0 = 12

probability = ½

S1	L	M	R
T	1,4	1,0	1,6
B	2,16	0,0	0,24

probability = ½

S2	L	M	R
T	1,4	1,6	1,0
B	2,16	0,24	0,0

Too Much Information Hurts

S1	L	M	R
T	1,4	1,0	1,6
B	2,16	0,0	0,24

S2	L	M	R
T	1,4	1,6	1,0
B	2,16	0,24	0,0

No player can distinguish the state, they believe that there is just one uninformed type of the other player

EPs	L	M	R
T	1,4	1,3	1,3
B	2,16	0,12	0,12

Too Much Information Hurts

- When we look at the table with expected payoffs, we can see, that player 2 has dominant strategy to play L, no matter what is the player 1’s action
- If no player has information about the state, there is a single NE in this game: (B,L)

EPs	L	M	R
T	1,4	1,3	1,3
B	2,16	0,12	0,12

Too Much Information Hurts

Now, consider that Player 2 can distinguish between two states

Player 1 now believes that there are two types of Player 2 (Left, and Right table) with equal probabilities

probability = ½

S1	L	M	R
T	1,4	1,0	1,6
B	2,16	0,0	0,24

probability = ½

S2	L	M	R
T	1,4	1,6	1,0
B	2,16	0,24	0,0

Too Much Information Hurts

Each type of Player 2 has a dominant action:

- For first type of player 2 it is best to play R, no matter of what is P1's action
- For second type of player 2 it is best to play M, no matter what is P1's action

probability = ½

S1	L	M	R
T	1,4	1,0	1,6
B	2,16	0,0	0,24

probability = ½

S2	L	M	R
T	1,4	1,6	1,0
B	2,16	0,24	0,0

Too Much Information Hurts

If Player 1 believes that first type Player 2 chooses R and the second type of Player 2 chooses M:

- EP(T) = ½*1 + ½*1 = 1
- EP(B) = ½*0 + ½*0 = 0
- And the NE in this game is: (T,(R,M))

probability = ½

S1	L	M	R
T	1,4	1,0	1,6
B	2,16	0,0	0,24

probability = ½

S2	L	M	R
T	1,4	1,6	1,0
B	2,16	0,24	0,0

Too Much Information Hurts

Comparison of outcomes

No information about state:

- NE = (B,L)
- Equilibrium payoff:
  - Player 1: 2
  - Player 2: 16 (both types)

Information about state:

- NE = (T,(R,M))
- Equilibrium payoff:
  - Player 1: 1
  - Player 2: 6 (both types)

Too Much Information Hurts

Comparison of outcomes

When Player 2 knows her type (knows the state), she optimally tailors her actions to the state which induces Player 1 to choose T rather than B and both players are worse off.

Note that this result is not general and depends on choice of payoffs of both players.

Modified Prisoner's dilemma

- Player 1 knows that:
  - with probability 2/3 Prisoner 2 is rational
    -> playing left game
  - with probability 1/3 Prisoner 2 is super nice
    -> playing right game

1 \ 2	C	RS
C	1,1	3,0
RS	0,3	2,2

1 \ 2	C	RS
C	1,1	3,3
RS	0,0	2,2

Expected Payoff

Example: if Prisoner 1 believes that rational Prisoner 2 chooses C and super nice Prisoner 2 chooses RS, then:

- Prisoner 1: C -> 2/3*1+1/3*3 = 5/3
- Prisoner 2: RS -> 2/3*0+1/3*2 = 2/3

1 \ 2	C	RS
C	1,1	3,0
RS	0,3	2,2

1 \ 2	C	RS
C	1,1	3,3
RS	0,0	2,2

Expected Payoff

- Similarly, we can find expected payoff of all action profiles:

1 \ 2	C	RS
C	1,1	3,0
RS	0,3	2,2

1 \ 2	C	RS
C	1,1	3,3
RS	0,0	2,2

1 \ 2	C,C	C,RS	RS,C	RS,RS
C	1,1,1	5/3,1,3	7/3,0,1	3,0,3
RS	0,3,0	2/3,3,2	4/3,2,0	2,2,2

Bayesian NE

- Similarly, we can find expected payoff of all action profiles:

1 \ 2	C	RS
C	1,1	3,0
RS	0,3	2,2

1 \ 2	C	RS
C	1,1	3,3
RS	0,0	2,2

1 \ 2	C,C	C,RS	RS,C	RS,RS
C	<u>1</u>,<u>1</u>,1	<u>5/3</u>,<u>1</u>,<u>3</u>	<u>7/3</u>,0,1	<u>3</u>,0,<u>3</u>
RS	0,<u>3</u>,0	2/3,<u>3</u>,<u>2</u>	4/3,2,0	2,2,<u>2</u>

Summary

- Bayesian games – information is incomplete (several possible states, types)

- How to find NE in Bayesian games:
  - consider each type as an individual player
  - given the beliefs, compute expected payoffs
  - find NE in this game