CS 798: Multiagent Systems

Bayesian Games

Kate Larson

Computer Science
University of Waterloo

Outline

Bayesian Games

So far we have assumed that all players know what game they are playing

- Number of players
- Actions available to each player
- Payoffs associated with strategy profiles

	L	R
U	3,?	-2, ?
D	0, ?	6, ?

Bayesian games (games of incomplete information) are used to represent uncertainties about the game being played

Bayesian Games

So far we have assumed that all players know what game they are playing

- Number of players
- Actions available to each player
- Payoffs associated with strategy profiles

	L	R
U	3,?	-2, ?
D	0, ?	6, ?

Bayesian games (games of incomplete information) are used to represent uncertainties about the game being played

Bayesian Games

There are different possible representations.

Information Sets

- N set of agents
- G set of games
  - Same strategy sets for each game and agent
- $\Pi(G)$ is the set of all probability distributions over G
  - $P(G) \in \Pi(G)$ common prior
- $I = (I_1, \ldots, I_n)$ are information sets (partitions over games)

Example

Extensive Form With Chance Moves

A special player, Nature, makes probabilistic moves.

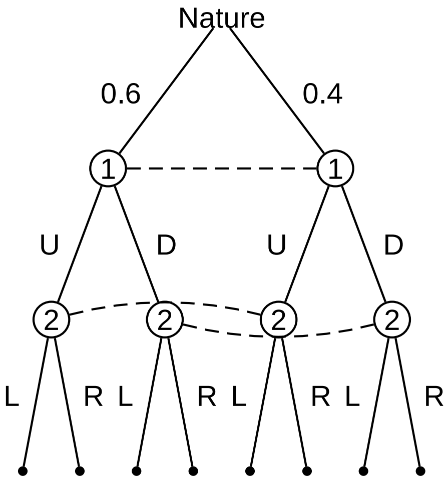

Epistemic Types

Epistemic types captures uncertainty directly over a game's utility functions.

- N set of agents
- $A = (A_1, \ldots, A_n)$ actions for each agent
- $\Theta = \Theta_1 \times \ldots \times \Theta_n$ where Θ_i is *type space* of each agent
- $p : \Theta \rightarrow [0, 1]$ is common prior over types
- Each agent has utility function $u_i : A \times \Theta \rightarrow \mathbb{R}$

Example

BoS

- 2 agents
- $A_1 = A_2 = \{\text{soccer, hockey}\}$
- $\Theta = (\Theta_1, \Theta_2)$ where $\Theta_1 = \{\text{H, S}\}$, $\Theta_2 = \{\text{H, S}\}$
- Prior: $p_1(H) = 1$, $p_2(H) = \frac{2}{3}$, $p_2(S) = \frac{1}{3}$

Utilities can be captured by matrix-form

$\theta_2 = H$

	H	S
H	2,2	0,0
S	0,0	1,1

$\theta_2 = S$

	H	S
H	2,1	0,0
S	0,0	1,2

Strategies and Utility

- A strategy $s_i(\theta_i)$ is a mapping from Θ_i to A_i. It specifies what action (or what distribution of actions) to take for each type.

Utility: $u_i(s|\theta_i)$

- *ex-ante* EU (know nothing about types)

$$EU = \sum_{\theta_i \in \Theta_i} p(\theta_i) EU_i(s_i|\theta_i)$$

- *interim* EU (know own type)

$$EU = EU_i(s|\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} p(\theta_{-i}|\theta_i) \sum_{a \in A} \Pi_{j \in N} s_j(a_j, \theta_j)) u_i(a, \theta_{-i}, \theta_i)$$

- *ex-post* EU (know everyones type)

Strategies and Utility

- A strategy $s_i(\theta_i)$ is a mapping from Θ_i to A_i. It specifies what action (or what distribution of actions) to take for each type.

Utility: $u_i(s|\theta_i)$

- *ex-ante* EU (know nothing about types)

$$EU = \sum_{\theta_i \in \Theta_i} p(\theta_i) EU_i(s_i|\theta_i)$$

- *interim* EU (know own type)

$$EU = EU_i(s|\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} p(\theta_{-i}|\theta_i) \sum_{a \in A} \Pi_{j \in N} s_j(a_j, \theta_j)) u_i(a, \theta_{-i}, \theta_i)$$

- *ex-post* EU (know everyones type)

Strategies and Utility

- A strategy $s_i(\theta_i)$ is a mapping from Θ_i to A_i. It specifies what action (or what distribution of actions) to take for each type.

Utility: $u_i(s|\theta_i)$

- *ex-ante* EU (know nothing about types)

$$EU = \sum_{\theta_i \in \Theta_i} p(\theta_i) EU_i(s_i|\theta_i)$$

- *interim* EU (know own type)

$$EU = EU_i(s|\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} p(\theta_{-i}|\theta_i) \sum_{a \in A} \Pi_{j \in N} s_j(a_j, \theta_j)) u_i(a, \theta_{-i}, \theta_i)$$

- *ex-post* EU (know everyones type)

Strategies and Utility

- A strategy $s_i(\theta_i)$ is a mapping from Θ_i to A_i. It specifies what action (or what distribution of actions) to take for each type.

Utility: $u_i(s|\theta_i)$

- *ex-ante* EU (know nothing about types)

$$EU = \sum_{\theta_i \in \Theta_i} p(\theta_i) EU_i(s_i|\theta_i)$$

- *interim* EU (know own type)

$$EU = EU_i(s|\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} p(\theta_{-i}|\theta_i) \sum_{a \in A} \Pi_{j \in N} s_j(a_j, \theta_j)) u_i(a, \theta_{-i}, \theta_i)$$

- *ex-post* EU (know everyones type)

Example

- 2 firms, 1 and 2, competing to create some product.
- If one makes the product then it has to share with the other.
- Product development cost is $c \in (0, 1)$
- Benefit of having the product is known only to each firm
  - Type θ_i drawn uniformly from $[0, 1]$
  - Benefit of having product is θ_i^2

Bayes Nash Equilibrium

Definition (BNE)

Strategy profile s^* *is a Bayes Nash equilibrium if* $\forall i$, $\forall \theta_i$

$$EU(s_i^*, s_{-i}^* | \theta_i) \geq EU(s_i', s_{-i}^* | \theta_i) \forall s_i' \neq s_i^*$$

Example Continued

- Let $s_i(\theta_i) = 1$ if i develops product, and 0 otherwise.
- If i develops product

$$u_i = \theta_i^2 - c$$

If it does not then

$$u_i = \theta_i^2 Pr(s_j(\theta_j) = 1)$$

- Thus, develop product if and only if

$$\theta_i^2 - c \geq \theta_i^2 Pr(s_j(\theta_j) = 1) \Rightarrow \theta_i \geq \sqrt{\frac{c}{1 - Pr(s_j(\theta_j) = 1)}}$$

Example Continued

- Let $s_i(\theta_i) = 1$ if i develops product, and 0 otherwise.
- If i develops product

$$u_i = \theta_i^2 - c$$

If it does not then

$$u_i = \theta_i^2 Pr(s_j(\theta_j) = 1)$$

- Thus, develop product if and only if

$$\theta_i^2 - c \geq \theta_i^2 Pr(s_j(\theta_j) = 1) \Rightarrow \theta_i \geq \sqrt{\frac{c}{1 - Pr(s_j(\theta_j) = 1)}}$$

Example Continued

- Let $s_i(\theta_i) = 1$ if i develops product, and 0 otherwise.
- If i develops product

$$u_i = \theta_i^2 - c$$

If it does not then

$$u_i = \theta_i^2 Pr(s_j(\theta_j) = 1)$$

- Thus, develop product if and only if

$$\theta_i^2 - c \geq \theta_i^2 Pr(s_j(\theta_j) = 1) \Rightarrow \theta_i \geq \sqrt{\frac{c}{1 - Pr(s_j(\theta_j) = 1)}}$$

Example Continued

Suppose $\hat{\theta}_1$, $\hat{\theta}_2 \in (0,1)$ are cutoff values in BNE.

- If so, then $Pr(s_j(\theta_j) = 1) = 1 - \hat{\theta}_j$
- We must have

$$\hat{\theta}_i \geq \sqrt{\frac{c}{\hat{\theta}_j}} \Rightarrow \hat{\theta}_i^2 \hat{\theta}_j = c$$

  and

$$\hat{\theta}_j^2 \hat{\theta}_i = c$$

- Therefore

$$\hat{\theta}_i^2 \hat{\theta}_j = \hat{\theta}_j^2 \hat{\theta}_i$$

  and so

$$\hat{\theta}_i = \hat{\theta}_j = \theta^* = c^{\frac{1}{3}}$$

Example Continued

Suppose $\hat{\theta}_1$, $\hat{\theta}_2 \in (0,1)$ are cutoff values in BNE.

- If so, then $Pr(s_j(\theta_j) = 1) = 1 - \hat{\theta}_j$
- We must have

$$\hat{\theta}_i \geq \sqrt{\frac{c}{\hat{\theta}_j}} \Rightarrow \hat{\theta}_i^2 \hat{\theta}_j = c$$

  and

$$\hat{\theta}_j^2 \hat{\theta}_i = c$$

- Therefore

$$\hat{\theta}_i^2 \hat{\theta}_j = \hat{\theta}_j^2 \hat{\theta}_i$$

  and so

$$\hat{\theta}_i = \hat{\theta}_j = \theta^* = c^{\frac{1}{3}}$$

Example Continued

Suppose $\hat{\theta}_1$, $\hat{\theta}_2 \in (0, 1)$ are cutoff values in BNE.

- If so, then $Pr(s_j(\theta_j) = 1) = 1 - \hat{\theta}_j$
- We must have

$$\hat{\theta}_i \geq \sqrt{\frac{c}{\hat{\theta}_j}} \Rightarrow \hat{\theta}_i^2 \hat{\theta}_j = c$$

  and

$$\hat{\theta}_j^2 \hat{\theta}_i = c$$

- Therefore

$$\hat{\theta}_i^2 \hat{\theta}_j = \hat{\theta}_j^2 \hat{\theta}_i$$

  and so

$$\hat{\theta}_i = \hat{\theta}_j = \theta^* = c^{\frac{1}{3}}$$

Example Continued

Suppose $\hat{\theta}_1$, $\hat{\theta}_2 \in (0,1)$ are cutoff values in BNE.

- If so, then $Pr(s_j(\theta_j) = 1) = 1 - \hat{\theta}_j$
- We must have
$$\hat{\theta}_i \geq \sqrt{\frac{c}{\hat{\theta}_j}} \Rightarrow \hat{\theta}_i^2 \hat{\theta}_j = c$$
and
$$\hat{\theta}_j^2 \hat{\theta}_i = c$$
- Therefore
$$\hat{\theta}_i^2 \hat{\theta}_j = \hat{\theta}_j^2 \hat{\theta}_i$$
and so
$$\hat{\theta}_i = \hat{\theta}_j = \theta^* = c^{\frac{1}{3}}$$