Analyzing Bayesian Games; Social Choice

CPSC 532A Lecture 14

October 26, 2006

Lecture Overview

Recap

Analyzing Bayesian games

Social Choice

Formal Definition

Definition

A stochastic game is a tuple $(Q, N, A_1, \ldots, A_n, P, r_1, \ldots, r_n)$, where

- Q is a finite set of states,
- N is a finite set of n players,
- A_i is a finite set of actions available to player i. Let $A = A_1 \times \cdots \times A_n$ be the vector of all players' actions,
- $P : Q \times A \times Q \to [0,1]$ is the transition probability function; let $P(q, a, \hat{q})$ be the probability of transitioning from state s to state $\hat{q}$ after joint action a,
- $r_i : Q \times A \to \mathbb{R}$ is a real-valued payoff function for player i.

Strategies

- What is a pure strategy?
  - pick an action conditional on every possible history
  - of course, mixtures over these pure strategies are possible too!
- Some interesting restricted classes of strategies:
  - behavioral strategy: $s_i(h_t, a_{i_j})$ returns the probability of playing action a_{i_j} for history h_t.
    - the substantive assumption here is that mixing takes place at each history independently, not once at the beginning of the game
  - Markov strategy: s_i is a behavioral strategy in which $s_i(h_t, a_{i_j}) = s_i(h'_t, a_{i_j})$ if $q_t = q'_t$, where q_t and q'_t are the final states of h_t and h'_t, respectively.
    - for a given time t, the distribution over actions only depends on the current state
  - stationary strategy: s_i is a Markov strategy in which $s_i(h_{t_1}, a_{i_j}) = s_i(h'_{t_2}, a_{i_j})$ if $q_{t_1} = q'_{t_2}$, where q_{t_1} and q'_{t_2} are the final states of h_{t_1} and h'_{t_2}, respectively.
    - no dependence even on t

Definition 1: Information Sets

- Bayesian game: a set of games that differ only in their payoffs, a common prior defined over them, and a partition structure over the games for each agent.

Definition (Bayesian Game: Information Sets)

A Bayesian game is a tuple (N, G, P, I) where

- N is a set of agents,
- G is a set of games with N agents each such that if $g, g' \in G$ then for each agent $i \in N$ the strategy space in g is identical to the strategy space in g',
- $P \in \Pi(G)$ is a common prior over games, where $\Pi(G)$ is the set of all probability distributions over G, and
- $I = (I_1, ..., I_N)$ is a set of partitions of G, one for each agent.

Definition 1: Example

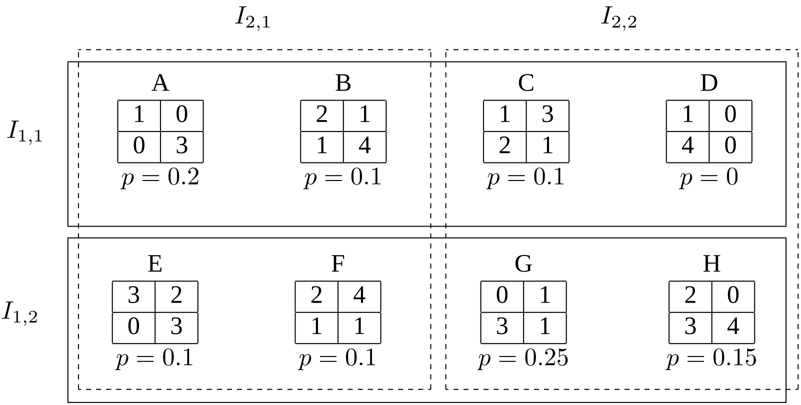

Definition 2: Extensive Form with Chance Moves

- Add an agent, "Nature," who follows a commonly known mixed strategy.
- Thus, reduce Bayesian games to extensive form games of imperfect information.
- This definition is cumbersome for the same reason that IIEF is a cumbersome way of representing matrix games like Prisoner's dilemma
  - however, it makes sense when the agents really do move sequentially, and at least occasionally observe each other's actions.

Definition 2: Example

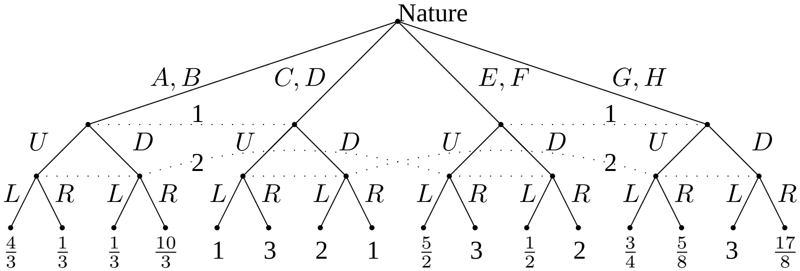

Definition 3: Epistemic Types

- Directly represent uncertainty over utility function using the notion of epistemic type.

Definition

A Bayesian game is a tuple (N, A, Θ, p, u) where

- N is a set of agents,
- $A = (A_1, \ldots, A_n)$, where A_i is the set of actions available to player i,
- $\Theta = (\Theta_1, \ldots, \Theta_n)$, where Θ_i is the type space of player i,
- $p : \Theta \rightarrow [0, 1]$ is the common prior over types,
- $u = (u_1, \ldots, u_n)$, where $u_i : A \times \Theta \rightarrow \mathbb{R}$ is the utility function for player i.

Definition 3: Example

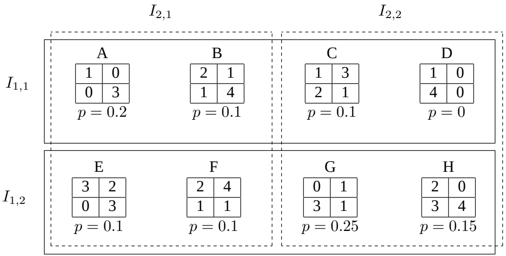

a_1	a_2	θ_1	θ_2	u_1
U	L	$\theta_{1,1}$	$\theta_{2,1}$	4/3
U	L	$\theta_{1,1}$	$\theta_{2,2}$	1
U	L	$\theta_{1,2}$	$\theta_{2,1}$	5/2
U	L	$\theta_{1,2}$	$\theta_{2,2}$	3/4
U	R	$\theta_{1,1}$	$\theta_{2,1}$	1/3
U	R	$\theta_{1,1}$	$\theta_{2,2}$	3
U	R	$\theta_{1,2}$	$\theta_{2,1}$	3
U	R	$\theta_{1,2}$	$\theta_{2,2}$	5/8

a_1	a_2	θ_1	θ_2	u_1
D	L	$\theta_{1,1}$	$\theta_{2,1}$	1/3
D	L	$\theta_{1,1}$	$\theta_{2,2}$	2
D	L	$\theta_{1,2}$	$\theta_{2,1}$	1/2
D	L	$\theta_{1,2}$	$\theta_{2,2}$	3
D	R	$\theta_{1,1}$	$\theta_{2,1}$	10/3
D	R	$\theta_{1,1}$	$\theta_{2,2}$	1
D	R	$\theta_{1,2}$	$\theta_{2,1}$	2
D	R	$\theta_{1,2}$	$\theta_{2,2}$	17/8

Lecture Overview

Strategies

- Pure strategy: $s_i : \Theta_i \rightarrow A_i$
  - a mapping from every type agent i could have to the action he would play if he had that type.
- Mixed strategy: $s_i : \Theta_i \rightarrow \Pi(A_i)$
  - a mapping from i's type to a probability distribution over his action choices.
- $s_j(a_j|\theta_j)$
  - the probability under mixed strategy s_j that agent j plays action a_j, given that j's type is θ_j.

Expected Utility

Three meaningful notions of expected utility:

- *ex-ante*
  - the agent knows nothing about anyone's actual type;
- *ex-interim*
  - an agent knows his own type but not the types of the other agents;
- *ex-post*
  - the agent knows all agents' types.

Ex-interim expected utility

Definition (*Ex-interim* expected utility)

Agent i's *ex-interim* expected utility in a Bayesian game (N, A, Θ, p, u), where i's type is θ_i and where the agents' strategies are given by the mixed strategy profile s, is defined as

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

- i must consider every θ_{-i} and every a in order to evaluate $u_i(a, \theta_i, \theta_{-i})$.
- i must weight this utility value by:
  - the probability that a would be realized given all players' mixed strategies and types;
  - the probability that the other players' types would be θ_{-i} given that his own type is θ_i.

Ex-ante expected utility

Definition (*Ex-ante* expected utility)

Agent i's *ex-ante* expected utility in a Bayesian game (N, A, Θ, p, u), where the agents' strategies are given by the mixed strategy profile s, is defined as

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

or equivalently as

$$EU_i(s) = \sum_{\theta \in \Theta} p(\theta) \sum_{a \in A} \left(\prod_{j \in N} s_j(a_j|\theta_j) \right) u_i(a, \theta).$$

Ex-post expected utility

Definition (*Ex-post* expected utility)

Agent i's *ex-post* expected utility in a Bayesian game (N, A, Θ, p, u), where the agents' strategies are given by s and the agent' types are given by θ, is defined as

$$EU_i(s, \theta) = \sum_{a \in A} \left(\prod_{j \in N} s_j(a_j | \theta_j) \right) u_i(a, \theta).$$

- The only uncertainty here concerns the other agents' mixed strategies, since i knows everyone's type.

Best response

Definition (Best response in a Bayesian game)

The set of agent i's best responses to mixed strategy profile s_{-i} are given by

$$BR_i(s_{-i}) = \arg\max_{s'_i \in S_i} EU_i(s'_i, s_{-i}).$$

- it may seem odd that BR is calculated based on i's *ex-ante* expected utility.
- However, write $EU_i(s)$ as $\sum_{\theta_i \in \Theta_i} p(\theta_i) EU_i(s|\theta_i)$ and observe that $EU_i(s'_i, s_{-i}|\theta_i)$ does not depend on strategies that i would play if his type were not θ_i.
- Thus, we are in fact performing independent maximization of i's *ex-interim* expected utility conditioned on each type that he could have.

Nash equilibrium

Definition (Bayes-Nash equilibrium)

A Bayes-Nash equilibrium is a mixed strategy profile s that satisfies $\forall i \;\; s_i \in BR_i(s_{-i})$.

Definition (*ex-post* Bayes-Nash equilibrium)

A *ex-post* Bayes-Nash equilibrium is a mixed strategy profile s that satisfies $\forall \theta,\ \forall i,\ s_i \in \arg\max_{s_i' \in S_i} EU_i(s_i', s_{-i}, \theta)$.

- almost a dominant strategy, but not quite

Induced normal form

- we can also construct an induced normal form for Bayesian games
- the numbers in the cells will correspond to *ex-ante* expected utilities
  - however as argued above, as long as the strategy space is unchanged, best responses don't change between the *ex-ante* and *ex-interim* cases.

Lecture Overview

Recap

Analyzing Bayesian games

Social Choice

Introduction

Our setting now:

- a set of outcomes
- agents have preferences across them
- for the moment, we won't consider incentive issues:
  - center knows agents' preferences, or they declare truthfully
- the goal: a social choice function: a mapping from everyone's preferences to a particular outcome, which is enforced
  - how to pick such functions with desirable properties?

Formal model

Definition (Social choice function)

Assume a set of agents $N = \{1, 2, \ldots, n\}$, and a set of outcomes (or alternatives, or candidates) O. Let L be the set of strict total orders on O. A social choice function (over N and O) is a function $C : L^n \rightarrow O$.

Definition (Social welfare function)

Let N, O, L be as above. A social welfare function (over N and O) is a function $W : L^n \rightarrow L^-$, where L^- is the set of weak total orderings (that is, total preorders) on O.

Non-Ranking Voting Schemes

- Plurality
  - pick the outcome which is preferred by the most people
- Cumulative voting
  - distribute e.g., 5 votes each
  - possible to vote for the same outcome multiple times
- Approval voting
  - accept as many outcomes as you "like"

Ranking Voting Schemes

- Plurality with elimination ("instant runoff")
  - everyone selects their favorite outcome
  - the outcome with the fewest votes is eliminated
  - repeat until one outcome remains
- Borda
  - assign each outcome a number.
  - The most preferred outcome gets a score of $n-1$, the next most preferred gets $n-2$, down to the n^{th} outcome which gets 0.
  - Then sum the numbers for each outcome, and choose the one that has the highest score
- Pairwise elimination
  - in advance, decide a schedule for the order in which pairs will be compared.
  - given two outcomes, have everyone determine the one that they prefer
  - eliminate the outcome that was not preferred, and continue with the schedule

Condorcet Condition

- If there is a candidate who is preferred to every other candidate in pairwise runoffs, that candidate should be the winner
- While the Condorcet condition is considered an important property for a voting system to satisfy, there is not always a Condorcet winner
- sometimes, there's a cycle where A defeats B, B defeats C, and C defeats A in their pairwise runoffs

Condorcet example

$$
\begin{aligned}
499 \text{ agents:} &\quad A \succ B \succ C \\
3 \text{ agents:} &\quad B \succ C \succ A \\
498 \text{ agents:} &\quad C \succ B \succ A
\end{aligned}
$$

- What is the Condorcet winner?

Condorcet example

499 agents:	$A \succ B \succ C$
3 agents:	$B \succ C \succ A$
498 agents:	$C \succ B \succ A$

- What is the Condorcet winner? B

Condorcet example

$$\begin{aligned} \text{499 agents:} \quad & A \succ B \succ C \\ \text{3 agents:} \quad & B \succ C \succ A \\ \text{498 agents:} \quad & C \succ B \succ A \end{aligned}$$

- What is the Condorcet winner? B
- What would win under plurality voting?

Condorcet example

499 agents:	$A \succ B \succ C$
3 agents:	$B \succ C \succ A$
498 agents:	$C \succ B \succ A$

- What is the Condorcet winner? B
- What would win under plurality voting? A

Condorcet example

$$\begin{aligned} 499 \text{ agents: } & A \succ B \succ C \\ 3 \text{ agents: } & B \succ C \succ A \\ 498 \text{ agents: } & C \succ B \succ A \end{aligned}$$

- What is the Condorcet winner? B
- What would win under plurality voting? A
- What would win under plurality with elimination?

Condorcet example

499 agents:	$A \succ B \succ C$
3 agents:	$B \succ C \succ A$
498 agents:	$C \succ B \succ A$

- What is the Condorcet winner? B
- What would win under plurality voting? A
- What would win under plurality with elimination? C