Social Choice

Lecture 11

Lecture Overview

1. Recap
2. Social Choice
3. Fun Game
4. Voting Paradoxes
5. Properties

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,
- $P : Q \times A \times Q \rightarrow [0, 1]$ is the transition probability function; let $P(q, a, \hat{q})$ be the probability of transitioning from state q to state $\hat{q}$ after joint action a,
- $r_i : Q \times A \rightarrow \mathbb{R}$ 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

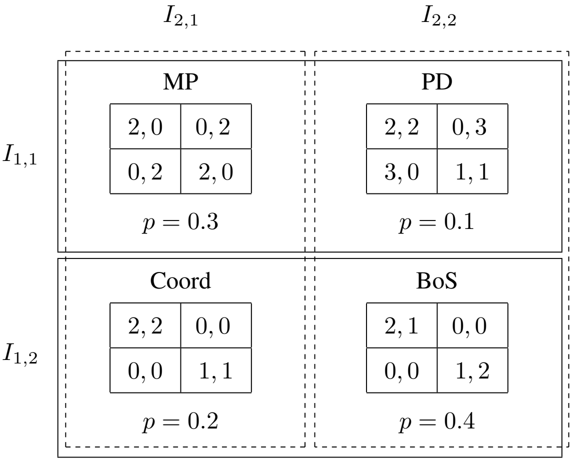

MP

$2, 0$	$0, 2$
$0, 2$	$2, 0$

$p = 0.3$

PD

$2, 2$	$0, 3$
$3, 0$	$1, 1$

$p = 0.1$

Coord

$2, 2$	$0, 0$
$0, 0$	$1, 1$

$p = 0.2$

BoS

$2, 1$	$0, 0$
$0, 0$	$1, 2$

$p = 0.4$

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

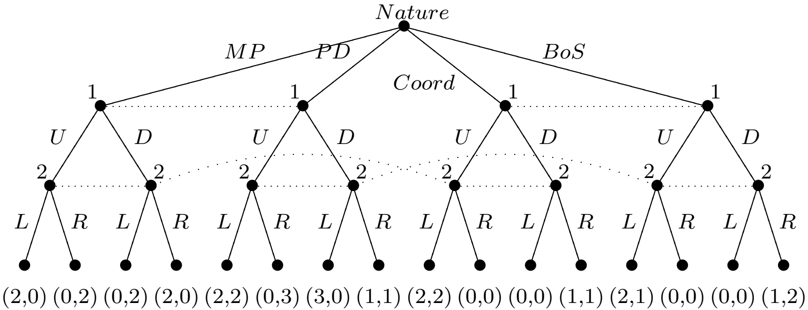

Definition 3: Epistemic Types

- 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,
- $p : \Theta \to [0, 1]$ is the common prior over types,
- $u = (u_1, \ldots, u_n)$, where $u_i : A \times \Theta \to \mathbb{R}$ is the utility function for player i.

Definition 3: Example

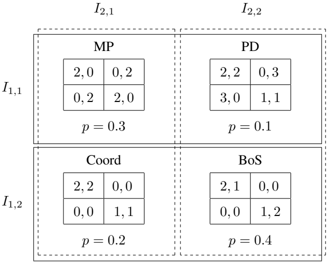

$I_{2,1}$ $I_{2,2}$

$I_{1,1}$

MP

2, 0	0, 2
0, 2	2, 0

$p = 0.3$

PD

2, 2	0, 3
3, 0	1, 1

$p = 0.1$

$I_{1,2}$

Coord

2, 2	0, 0
0, 0	1, 1

$p = 0.2$

BoS

2, 1	0, 0
0, 0	1, 2

$p = 0.4$

a_1	a_2	θ_1	θ_2	u_1	u_2
U	L	$\theta_{1,1}$	$\theta_{2,1}$	2	0
U	L	$\theta_{1,1}$	$\theta_{2,2}$	2	2
U	L	$\theta_{1,2}$	$\theta_{2,1}$	2	2
U	L	$\theta_{1,2}$	$\theta_{2,2}$	2	1
U	R	$\theta_{1,1}$	$\theta_{2,1}$	0	2
U	R	$\theta_{1,1}$	$\theta_{2,2}$	0	3
U	R	$\theta_{1,2}$	$\theta_{2,1}$	0	0
U	R	$\theta_{1,2}$	$\theta_{2,2}$	0	0

a_1	a_2	θ_1	θ_2	u_1	u_2
D	L	$\theta_{1,1}$	$\theta_{2,1}$	0	2
D	L	$\theta_{1,1}$	$\theta_{2,2}$	3	0
D	L	$\theta_{1,2}$	$\theta_{2,1}$	0	0
D	L	$\theta_{1,2}$	$\theta_{2,2}$	0	0
D	R	$\theta_{1,1}$	$\theta_{2,1}$	2	0
D	R	$\theta_{1,1}$	$\theta_{2,2}$	1	1
D	R	$\theta_{1,2}$	$\theta_{2,1}$	1	1
D	R	$\theta_{1,2}$	$\theta_{2,2}$	1	2

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.

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 $\forall i \;\; s_i \in BR_i(s_{-i})$.

- 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.

ex-post Equilibrium

Definition (*ex-post* equilibrium)

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

- somewhat similar to dominant strategy, but not quite
  - EP: agents do not need to have accurate beliefs about the type distribution
  - DS: agents do not need to have accurate beliefs about others' strategies

Lecture Overview

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 non-strict total orders on O. A social choice function (over N and O) is a function $C : L_-^n \mapsto O$.

Definition (Social welfare function)

Let N, O, L_- be as above. A social choice function (over N and O) is a function $C : L_-^n \mapsto L_-$.

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

Lecture Overview

1. Recap
2. Social Choice
3. **Fun Game**
4. Voting Paradoxes
5. Properties

Fun Game

- Imagine that there was an opportunity to take a one-week class trip at the end of term, to one of the following destinations:
    - (O) Orlando, FL
    - (P) Paris, France
    - (T) Tehran, Iran
    - (B) Beijing, China
- Construct your preference ordering

Fun Game

- Imagine that there was an opportunity to take a one-week class trip at the end of term, to one of the following destinations:
  - (O) Orlando, FL
  - (P) Paris, France
  - (T) Tehran, Iran
  - (B) Beijing, China
- Construct your preference ordering
- Vote (truthfully) using each of the following schemes:
  - plurality (raise hands)

Fun Game

- Imagine that there was an opportunity to take a one-week class trip at the end of term, to one of the following destinations:
    - (O) Orlando, FL
    - (P) Paris, France
    - (T) Tehran, Iran
    - (B) Beijing, China
- Construct your preference ordering
- Vote (truthfully) using each of the following schemes:
    - plurality (raise hands)
    - plurality with elimination (raise hands)

Fun Game

- Imagine that there was an opportunity to take a one-week class trip at the end of term, to one of the following destinations:
    - (O) Orlando, FL
    - (P) Paris, France
    - (T) Tehran, Iran
    - (B) Beijing, China
- Construct your preference ordering
- Vote (truthfully) using each of the following schemes:
    - plurality (raise hands)
    - plurality with elimination (raise hands)
    - Borda (volunteer to tabulate)

Fun Game

- Imagine that there was an opportunity to take a one-week class trip at the end of term, to one of the following destinations:
    - (O) Orlando, FL
    - (P) Paris, France
    - (T) Tehran, Iran
    - (B) Beijing, China
- Construct your preference ordering
- Vote (truthfully) using each of the following schemes:
    - plurality (raise hands)
    - plurality with elimination (raise hands)
    - Borda (volunteer to tabulate)
    - pairwise elimination (raise hands, I'll pick a schedule)

Lecture Overview

1. Recap
2. Social Choice
3. Fun Game
4. Voting Paradoxes
5. Properties

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?

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

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

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

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? 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

Sensitivity to Losing Candidate

35 agents: $A \succ C \succ B$
33 agents: $B \succ A \succ C$
32 agents: $C \succ B \succ A$

- What candidate wins under plurality voting?

Sensitivity to Losing Candidate

35 agents: $A \succ C \succ B$
33 agents: $B \succ A \succ C$
32 agents: $C \succ B \succ A$

- What candidate wins under plurality voting? A

Sensitivity to Losing Candidate

35 agents: $A \succ C \succ B$
33 agents: $B \succ A \succ C$
32 agents: $C \succ B \succ A$

- What candidate wins under plurality voting? A
- What candidate wins under Borda voting?

Sensitivity to Losing Candidate

35 agents:	$A \succ C \succ B$
33 agents:	$B \succ A \succ C$
32 agents:	$C \succ B \succ A$

- What candidate wins under plurality voting? A
- What candidate wins under Borda voting? A

Sensitivity to Losing Candidate

35 agents: $A \succ C \succ B$
33 agents: $B \succ A \succ C$
32 agents: $C \succ B \succ A$

- What candidate wins under plurality voting? A
- What candidate wins under Borda voting? A
- Now consider dropping C. Now what happens under both Borda and plurality?

Sensitivity to Losing Candidate

35 agents:	$A \succ C \succ B$
33 agents:	$B \succ A \succ C$
32 agents:	$C \succ B \succ A$

- What candidate wins under plurality voting? A
- What candidate wins under Borda voting? A
- Now consider dropping C. Now what happens under both Borda and plurality? B wins.

Sensitivity to Agenda Setter

35 agents: $A \succ C \succ B$
33 agents: $B \succ A \succ C$
32 agents: $C \succ B \succ A$

- Who wins pairwise elimination, with the ordering A, B, C?

Sensitivity to Agenda Setter

35 agents: $A \succ C \succ B$
33 agents: $B \succ A \succ C$
32 agents: $C \succ B \succ A$

- Who wins pairwise elimination, with the ordering A, B, C? C

Sensitivity to Agenda Setter

35 agents: $A \succ C \succ B$
33 agents: $B \succ A \succ C$
32 agents: $C \succ B \succ A$

- Who wins pairwise elimination, with the ordering A, B, C? C
- Who wins with the ordering A, C, B?

Sensitivity to Agenda Setter

35 agents: $A \succ C \succ B$
33 agents: $B \succ A \succ C$
32 agents: $C \succ B \succ A$

- Who wins pairwise elimination, with the ordering A, B, C? C
- Who wins with the ordering A, C, B? B

Sensitivity to Agenda Setter

35 agents: $A \succ C \succ B$
33 agents: $B \succ A \succ C$
32 agents: $C \succ B \succ A$

- Who wins pairwise elimination, with the ordering A, B, C? C
- Who wins with the ordering A, C, B? B
- Who wins with the ordering B, C, A?

Sensitivity to Agenda Setter

35 agents: $A \succ C \succ B$
33 agents: $B \succ A \succ C$
32 agents: $C \succ B \succ A$

- Who wins pairwise elimination, with the ordering A, B, C? C
- Who wins with the ordering A, C, B? B
- Who wins with the ordering B, C, A? A

Another Pairwise Elimination Problem

1 agent: $B \succ D \succ C \succ A$
1 agent: $A \succ B \succ D \succ C$
1 agent: $C \succ A \succ B \succ D$

- Who wins under pairwise elimination with the ordering A, B, C, D?

Another Pairwise Elimination Problem

1 agent: $B \succ D \succ C \succ A$
1 agent: $A \succ B \succ D \succ C$
1 agent: $C \succ A \succ B \succ D$

- Who wins under pairwise elimination with the ordering A, B, C, D? D.

Another Pairwise Elimination Problem

1 agent: $B \succ D \succ C \succ A$
1 agent: $A \succ B \succ D \succ C$
1 agent: $C \succ A \succ B \succ D$

- Who wins under pairwise elimination with the ordering A, B, C, D? D.
- What is the problem with this?

Another Pairwise Elimination Problem

1 agent: $B \succ D \succ C \succ A$
1 agent: $A \succ B \succ D \succ C$
1 agent: $C \succ A \succ B \succ D$

- Who wins under pairwise elimination with the ordering A, B, C, D? D.
- What is the problem with this?
  - *all* of the agents prefer B to D—the selected candidate is Pareto-dominated!

Lecture Overview

1. Recap
2. Social Choice
3. Fun Game
4. Voting Paradoxes
5. Properties

Notation

- N is the set of agents
- O is a finite set of outcomes with $|O| \geq 3$
- L is the set of all possible strict preference orderings over O.
  - for ease of exposition we switch to strict orderings
  - we will end up showing that desirable SWFs cannot be found *even if* preferences are restricted to strict orderings
- $[\succ]$ is an element of the set L^n (a preference ordering for every agent; the input to our social welfare function)
- $\succ_W$ is the preference ordering selected by the social welfare function W.
  - When the input to W is ambiguous we write it in the subscript; thus, the social order selected by W given the input $[\succ']$ is denoted as $\succ_{W([\succ'])}$.

Pareto Efficiency

Definition (Pareto Efficiency (PE))

W is Pareto efficient if for any $o_1, o_2 \in O$, $\forall i \, o_1 \succ_i o_2$ implies that $o_1 \succ_W o_2$.

- when all agents agree on the ordering of two outcomes, the social welfare function must select that ordering.

Independence of Irrelevant Alternatives

Definition (Independence of Irrelevant Alternatives (IIA))

W is independent of irrelevant alternatives if, for any $o_1, o_2 \in O$ and any two preference profiles $[\succ'], [\succ''] \in L^n$, $\forall i\,(o_1 \succ'_i o_2$ if and only if $o_1 \succ''_i o_2)$ implies that $(o_1 \succ_{W([\succ'])} o_2$ if and only if $o_1 \succ_{W([\succ''])} o_2)$.

- the selected ordering between two outcomes should depend only on the relative orderings they are given by the agents.

Nondictatorship

Definition (Non-dictatorship)

W does not have a dictator if $\neg \exists i \, \forall o_1, o_2 (o_1 \succ_i o_2 \Rightarrow o_1 \succ_W o_2)$.

- there does not exist a single agent whose preferences always determine the social ordering.
- We say that W is dictatorial if it fails to satisfy this property.