From Optimality to Equilibrium

Lecture 4

Lecture Overview

1. Recap
2. Pareto Optimality
3. Best Response and Nash Equilibrium
4. Mixed Strategies

Non-Cooperative Game Theory

- What is it?
  - mathematical study of interaction between rational, self-interested agents
- Why is it called non-cooperative?
  - while it's most interested in situations where agents' interests conflict, it's not restricted to these settings
  - the key is that the individual is the basic modeling unit, and that individuals pursue their own interests
    - cooperative/coalitional game theory has teams as the central unit, rather than agents

Defining Games

- Finite, n-person game: $\langle N, A, u \rangle$:
  - N is a finite set of n players, indexed by i
  - $A = A_1 \times \ldots \times A_n$, where A_i is the action set for player i
    - $a \in A$ is an action profile, and so A is the space of action profiles
  - $u = \langle u_1, \ldots, u_n \rangle$, a utility function for each player, where $u_i : A \mapsto \mathbb{R}$

- Writing a 2-player game as a matrix:
  - row player is player 1, column player is player 2
  - rows are actions $a \in A_1$, columns are $a' \in A_2$
  - cells are outcomes, written as a tuple of utility values for each player

Prisoner's dilemma

Prisoner's dilemma is any game

	C	D
C	a, a	b, c
D	c, b	d, d

with $c > a > d > b$.

Games of Pure Competition

Players have exactly opposed interests

- There must be precisely two players (otherwise they can't have exactly opposed interests)
- For all action profiles $a \in A$, $u_1(a) + u_2(a) = c$ for some constant c
  - Special case: zero sum

	Heads	Tails
Heads	1	-1
Tails	-1	1

Games of Cooperation

Players have exactly the same interests.

- no conflict: all players want the same things
- $\forall a \in A, \forall i, j,\ u_i(a) = u_j(a)$

	Left	Right
Left	1	0
Right	0	1

General Games: Battle of the Sexes

The most interesting games combine elements of cooperation *and* competition.

	B	F
B	$2, 1$	$0, 0$
F	$0, 0$	$1, 2$

Lecture Overview

Analyzing Games

- We've defined some canonical games, and thought about how to play them. Now let's examine the games from the outside
- From the point of view of an outside observer, can some outcomes of a game be said to be better than others?

Analyzing Games

- We've defined some canonical games, and thought about how to play them. Now let's examine the games from the outside
- From the point of view of an outside observer, can some outcomes of a game be said to be better than others?
  - we have no way of saying that one agent's interests are more important than another's
  - intuition: imagine trying to find the revenue-maximizing outcome when you don't know what currency has been used to express each agent's payoff
- Are there situations where we can still prefer one outcome to another?

Pareto Optimality

- Idea: sometimes, one outcome o is at least as good for every agent as another outcome o', and there is some agent who strictly prefers o to o'
  - in this case, it seems reasonable to say that o is better than o'
  - we say that o Pareto-dominates o'.

Pareto Optimality

- Idea: sometimes, one outcome o is at least as good for every agent as another outcome o', and there is some agent who strictly prefers o to o'
  - in this case, it seems reasonable to say that o is better than o'
  - we say that o Pareto-dominates o'.

- An outcome o^* is Pareto-optimal if there is no other outcome that Pareto-dominates it.

Pareto Optimality

- Idea: sometimes, one outcome o is at least as good for every agent as another outcome o', and there is some agent who strictly prefers o to o'
  - in this case, it seems reasonable to say that o is better than o'
  - we say that o Pareto-dominates o'.

- An outcome o^* is Pareto-optimal if there is no other outcome that Pareto-dominates it.
  - can a game have more than one Pareto-optimal outcome?

Pareto Optimality

- Idea: sometimes, one outcome o is at least as good for every agent as another outcome o', and there is some agent who strictly prefers o to o'
  - in this case, it seems reasonable to say that o is better than o'
  - we say that o Pareto-dominates o'.

- An outcome o^* is Pareto-optimal if there is no other outcome that Pareto-dominates it.
  - can a game have more than one Pareto-optimal outcome?
  - does every game have at least one Pareto-optimal outcome?

Pareto Optimal Outcomes in Example Games

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

Pareto Optimal Outcomes in Example Games

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

	Left	Right
Left	1	0
Right	0	1

Pareto Optimal Outcomes in Example Games

	C	D
C	$-1,-1$	$-4,0$
D	$0,-4$	$-3,-3$

	Left	Right
Left	1	0
Right	0	1

	B	F
B	$2,1$	$0,0$
F	$0,0$	$1,2$

Pareto Optimal Outcomes in Example Games

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

	Left	Right
Left	1	0
Right	0	1

	B	F
B	$2, 1$	$0, 0$
F	$0, 0$	$1, 2$

	Heads	Tails
Heads	1	-1
Tails	-1	1

Lecture Overview

Best Response

- If you knew what everyone else was going to do, it would be easy to pick your own action

Best Response

- If you knew what everyone else was going to do, it would be easy to pick your own action
- Let $a_{-i} = \langle a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n \rangle$.
  - now $a = (a_{-i}, a_i)$

- Best response: $a_i^* \in BR(a_{-i})$ iff $\forall a_i \in A_i,\ u_i(a_i^*, a_{-i}) \geq u_i(a_i, a_{-i})$

Nash Equilibrium

- Now let's return to the setting where no agent knows anything about what the others will do
- What can we say about which actions will occur?

Nash Equilibrium

- Now let's return to the setting where no agent knows anything about what the others will do
- What can we say about which actions will occur?

- Idea: look for stable action profiles.
- $a = \langle a_1, \ldots, a_n \rangle$ is a ("pure strategy") Nash equilibrium iff $\forall i,\, a_i \in BR(a_{-i})$.

Nash Equilibria of Example Games

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

Nash Equilibria of Example Games

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

	Left	Right
Left	1	0
Right	0	1

Nash Equilibria of Example Games

	C	D
C	$-1,-1$	$-4,0$
D	$0,-4$	$-3,-3$

	Left	Right
Left	1	0
Right	0	1

	B	F
B	$2,1$	$0,0$
F	$0,0$	$1,2$

Nash Equilibria of Example Games

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

	Left	Right
Left	1	0
Right	0	1

	B	F
B	$2, 1$	$0, 0$
F	$0, 0$	$1, 2$

	Heads	Tails
Heads	1	-1
Tails	-1	1

Nash Equilibria of Example Games

	C	D
C	$-1,-1$	$-4,0$
D	$0,-4$	$-3,-3$

	Left	Right
Left	1	0
Right	0	1

	B	F
B	$2,1$	$0,0$
F	$0,0$	$1,2$

	Heads	Tails
Heads	1	-1
Tails	-1	1

The paradox of Prisoner's dilemma: the Nash equilibrium is the only non-Pareto-optimal outcome!

Lecture Overview

Mixed Strategies

- It would be a pretty bad idea to play any deterministic strategy in matching pennies
- Idea: confuse the opponent by playing randomly
- Define a strategy s_i for agent i as any probability distribution over the actions A_i.
  - pure strategy: only one action is played with positive probability
  - mixed strategy: more than one action is played with positive probability
    - these actions are called the support of the mixed strategy
- Let the set of all strategies for i be S_i
- Let the set of all strategy profiles be $S = S_1 \times \ldots \times S_n$.

Utility under Mixed Strategies

- What is your payoff if all the players follow mixed strategy profile $s \in S$?
    - We can't just read this number from the game matrix anymore: we won't always end up in the same cell

Utility under Mixed Strategies

- What is your payoff if all the players follow mixed strategy profile $s \in S$?
  - We can't just read this number from the game matrix anymore: we won't always end up in the same cell
- Instead, use the idea of expected utility from decision theory:

$$u_i(s) = \sum_{a \in A} u_i(a) Pr(a|s)$$

$$Pr(a|s) = \prod_{j \in N} s_j(a_j)$$

Best Response and Nash Equilibrium

Our definitions of best response and Nash equilibrium generalize from actions to strategies.

- Best response:
  - $s_i^* \in BR(s_{-i})$ iff $\forall s_i \in S_i,\ u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$

Best Response and Nash Equilibrium

Our definitions of best response and Nash equilibrium generalize from actions to strategies.

- Best response:
  - $s_i^* \in BR(s_{-i})$ iff $\forall s_i \in S_i,\ u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$
- Nash equilibrium:
  - $s = \langle s_1, \ldots, s_n \rangle$ is a Nash equilibrium iff $\forall i,\ s_i \in BR(s_{-i})$

Best Response and Nash Equilibrium

Our definitions of best response and Nash equilibrium generalize from actions to strategies.

- Best response:
  - $s_i^* \in BR(s_{-i})$ iff $\forall s_i \in S_i$, $u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$
- Nash equilibrium:
  - $s = \langle s_1, \ldots, s_n \rangle$ is a Nash equilibrium iff $\forall i$, $s_i \in BR(s_{-i})$
- Every finite game has a Nash equilibrium! [Nash, 1950]
  - e.g., matching pennies: both players play heads/tails 50%/50%

Computing Mixed Nash Equilibria: Battle of the Sexes

	B	F
B	$2, 1$	$0, 0$
F	$0, 0$	$1, 2$

- It's hard in general to compute Nash equilibria, but it's easy when you can guess the support
- For BoS, let's look for an equilibrium where all actions are part of the support

Computing Mixed Nash Equilibria: Battle of the Sexes

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

- Let player 2 play B with p, F with $1 - p$.
- If player 1 best-responds with a mixed strategy, player 2 must make him indifferent between F and B (why?)

Computing Mixed Nash Equilibria: Battle of the Sexes

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

- Let player 2 play B with p, F with $1-p$.
- If player 1 best-responds with a mixed strategy, player 2 must make him indifferent between F and B (why?)

$$u_1(B) = u_1(F)$$
$$2p + 0(1-p) = 0p + 1(1-p)$$
$$p = \frac{1}{3}$$

Computing Mixed Nash Equilibria: Battle of the Sexes

	B	F
B	$2, 1$	$0, 0$
F	$0, 0$	$1, 2$

- Likewise, player 1 must randomize to make player 2 indifferent.
    - Why is player 1 willing to randomize?

Computing Mixed Nash Equilibria: Battle of the Sexes

	B	F
B	$2, 1$	$0, 0$
F	$0, 0$	$1, 2$

- Likewise, player 1 must randomize to make player 2 indifferent.
  - Why is player 1 willing to randomize?
- Let player 1 play B with q, F with $1 - q$.

$$u_2(B) = u_2(F)$$
$$q + 0(1 - q) = 0q + 2(1 - q)$$
$$q = \frac{2}{3}$$

- Thus the mixed strategies $(\frac{2}{3}, \frac{1}{3})$, $(\frac{1}{3}, \frac{2}{3})$ are a Nash equilibrium.

Interpreting Mixed Strategy Equilibria

What does it mean to play a mixed strategy? Different interpretations:

- Randomize to confuse your opponent
  - consider the matching pennies example
- Players randomize when they are uncertain about the other's action
  - consider battle of the sexes
- Mixed strategies are a concise description of what might happen in repeated play: count of pure strategies in the limit
- Mixed strategies describe population dynamics: 2 agents chosen from a population, all having deterministic strategies. MS is the probability of getting an agent who will play one PS or another.