CSC2556

Lecture 9

Game Theory 1: Nash Equilibria

Game Theory

Game Theory

- How do rational self-interested agents act in a given environment?
- Environment modeled as a game
  - Each agent or player has a set of possible actions
  - Rules of the game dictate the rewards for the agents as a function of the actions taken by all the players
    - My reward also depends on what action you take
    - Therefore, I must reason about what action you'll take as well
- Non-cooperative games
  - No external trusted agency, no legally binding agreements

Normal Form Games

- A set of players $\mathrm{N} = \{1, \dots, n\}$
- Each player i chooses an action $a_i \in A_i$
  - Action profile $\vec{a} = (a_1, \dots, a_n) \in \mathcal{A} = A_1 \times \cdots \times A_n$
  - $\vec{a}_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$
- Each player i has a utility function $u_i : \mathcal{A} \to \mathbb{R}$
  - Given the action profile $\vec{a} = (a_1, \dots, a_n)$, each player i gets reward $u_i(a_1, \dots, a_n)$
- Note that the utility to player i depends on the action chosen by the other players too

Normal Form Games

Prisoner's dilemma

$S = \{\text{Silent}, \text{Betray}\}$

John's Actions / Sam's Actions	Stay Silent	Betray
Stay Silent	(-1 , -1)	(-3 , 0)
Betray	(0 , -3)	(-2 , -2)

$u_{Sam}(Betray, Silent)$ $u_{John}(Betray, Silent)$

s_{Sam} s_{John}

Strategies

- Pure strategy
  - Choose an action deterministically, e.g., "*betray*"
- Mixed strategy
  - Choose an action in a randomized fashion, e.g., "*stay silent* with probability 0.3, and *betray* with probability 0.7" (call this s^*)
  - We compute expected utilities when each player's action is sampled from her mixed strategy independently of the other players
  - Example: Say both Sam and John adopt s^*:

$$
\begin{aligned}
E[u_{Sam}(s^*, s^*)] = \; & 0.3 \times 0.3 \times u_{Sam}(Silent, Silent) \\
& + 0.3 \times 0.7 \times u_{Sam}(Silent, Betray) \\
& + 0.7 \times 0.3 \times u_{Sam}(Betray, Silent) \\
& + 0.7 \times 0.7 \times u_{Sam}(Betray, Betray)
\end{aligned}
$$

Domination Among Strategies

- Consider two strategies s_i, s_i' of player i

- Informally, s_i "dominates" s_i' if s_i is "better than" s_i', *irrespective of the other players' strategies*

- Weak vs strict domination
  - Both require: $u_i(s_i, \vec{s}_{-i}) \geq u_i(s_i', \vec{s}_{-i}), \forall \vec{s}_{-i}$
  - Weak domination requires: Strict inequality for some $\vec{s}_{-i}$
  - Strict domination requires: Strict inequality for all $\vec{s}_{-i}$

Dominant Strategies

- Dominant strategies
  - s_i is a strictly (resp. weakly) dominant strategy for player i if it strictly (resp. weakly) dominates every other strategy
  - Strictly/weakly dominating every other *pure* strategy is sufficient (Why?)
  - Can a player have two strictly/weakly dominant strategies?

- How does this relate to strategyproofness?
  - "Truth-telling should be at least as good as any other strategy, regardless of what the other players do"
  - Basically, truth-telling should be weakly dominant except we don't require that it be strictly better for *some* combination of strategies of the other players

Example: Prisoner's Dilemma

- Recap:

John's Actions / Sam's Actions	Stay Silent	Betray
Stay Silent	**(-1 , -1)**	**(-3 , 0)**
Betray	**(0 , -3)**	**(-2 , -2)**

- Each player strictly wants to
  - Betray if the other player will stay silent
  - Betray if the other player will betray
- Betraying strictly dominates staying silent
  - So betraying is a strictly dominant strategy for each player

Solution Concept 1:

- If each player i has a strictly/weakly dominant strategy s_i^*, then the realized strategy profile would be $(s_1^*, \dots, s_n^*)$

Iterated Elimination

- What if there are no dominant strategies?
  - No single strategy dominates every other strategy
  - But some strategies might still be dominated

- Assuming everyone knows everyone is rational…
  - Can remove their dominated strategies
  - Might reveal a newly dominant strategy

- Eliminating only strictly dominated vs eliminating weakly dominated

Iterated Elimination

- Toy example:
  - Microsoft vs Startup
  - Enter the market or stay out?

Microsoft \ Startup	Enter	Stay Out
Enter	(2 , -2)	(4 , 0)
Stay Out	(0 , 4)	(0 , 0)

- Q: Is there a dominant strategy for startup?
- Q: Do you see a rational outcome of the game?

Iterated Elimination

- "Guess 2/3 of average"
  - Each student guesses a real number between 0 and 100 (inclusive)
  - The student whose number is the closest to 2/3 of the average of all numbers wins!

- Poll: What would you do?

Solution Concept 2:

- If iterated elimination of strictly dominated strategies leads to a single strategy profile, then that would be the realized strategy profile

Nash Equilibrium

- What if not all players have a dominant strategy and iterated elimination does not help predict the outcome of the game either?

Students \ Professor	Attend	Be Absent
Attend	(3 , 1)	(-1 , -3)
Be Absent	(-1 , -1)	(0 , 0)

Nash Equilibrium

- Instead of hoping to find strategies that players would play *irrespective of what other players play,* we find strategies that players would play *given what the other players are playing*

- Nash Equilibrium
  - A strategy profile $\vec{s}$ is in Nash equilibrium if s_i is the best action for player i given that other players are playing $\vec{s}_{-i}$

$$u_i(s_i, \vec{s}_{-i}) \geq u_i(s_i', \vec{s}_{-i}), \forall s_i'$$

  - Pure NE: All strategies are pure
  - Mixed NE: At least one strategy is mixed

Recap: Prisoner's Dilemma

Sam's Actions \ John's Actions	Stay Silent	Betray
Stay Silent	(-1 , -1)	(-3 , 0)
Betray	(0 , -3)	(-2 , -2)

- Nash equilibrium?
- (Dominant strategies)

Recap: Microsoft vs Startup

Microsoft \ Startup	Enter	Stay Out
Enter	(2 , -2)	(4 , 0)
Stay Out	(0 , 4)	(0 , 0)

- Nash equilibrium?
- (Iterated elimination of strongly dominated strategies)

Recap: Attend or Not

Students \ Professor	Attend	Be Absent
Attend	(3 , 1)	(-1 , -3)
Be Absent	(-1 , -1)	(0 , 0)

- Nash equilibria?

- Lack of predictability

Example: Rock-Paper-Scissor

P1 / P2	Rock	Paper	Scissor
Rock	**(0 , 0)**	**(-1 , 1)**	**(1 , -1)**
Paper	**(1 , -1)**	**(0 , 0)**	**(-1 , 1)**
Scissor	**(-1 , 1)**	**(1 , -1)**	**(0 , 0)**

- Pure Nash equilibrium?

Nash's Beautiful Result

- Theorem: Every normal form game admits a mixed-strategy Nash equilibrium.

- What about Rock-Paper-Scissor?

P1 / P2	Rock	Paper	Scissor
Rock	(0 , 0)	(-1 , 1)	(1 , -1)
Paper	(1 , -1)	(0 , 0)	(-1 , 1)
Scissor	(-1 , 1)	(1 , -1)	(0 , 0)

Indifference Principle

- Let $\vec{s}$ be a Nash equilibrium
- Let s_i be a mixed strategy with support T_i
- Then, the expected payoff of player i from each $a_i \in T_i$ must be identical and at least as much as the expected payoff from any $a_i' \notin T_i$

- Derivation of rock-paper-scissor on the board.

Complexity

- Theorem [DGP'06, CD'06]
  - The problem of computing a Nash equilibrium of a given game is PPAD-complete even with two players.

Stag-Hunt

Hunter 2 \ Hunter 1	Stag	Hare
Stag	**(4 , 4)**	**(0 , 2)**
Hare	**(2 , 0)**	**(1 , 1)**

- Game
  - Stag requires both hunters, food is good for 4 days for each hunter.
  - Hare requires a single hunter, food is good for 2 days
  - If they both catch the same hare, they share.

- Two pure Nash equilibria: (Stag,Stag), (Hare,Hare)

Stag-Hunt

Hunter 2 \ Hunter 1	Stag	Hare
Stag	(4 , 4)	(0 , 2)
Hare	(2 , 0)	(1 , 1)

- Two pure Nash equilibria: (Stag,Stag), (Hare,Hare)
  - Other hunter plays “Stag” → “Stag” is best response
  - Other hunter plays “Hare” → “Hare” is best reponse

- What about mixed Nash equilibria?

Stag-Hunt

Hunter 1 / Hunter 2	Stag	Hare
Stag	(4 , 4)	(0 , 2)
Hare	(2 , 0)	(1 , 1)

- Symmetric: $s \rightarrow \{\text{Stag w.p. } p, \text{ Hare w.p. } 1-p\}$

- Indifference principle:
  - *Given the other hunter plays s, equal $\mathbb{E}$[reward] for Stag and Hare*
  - $\mathbb{E}[\text{Stag}] = p * 4 + (1-p) * 0$
  - $\mathbb{E}[\text{Hare}] = p * 2 + (1-p) * 1$
  - Equate the two $\Rightarrow p = 1/3$

Extra Fun 1: Cunning Airlines

- Two travelers lose their luggage.
- Airline agrees to refund up to $100 to each.
- Policy: Both travelers would submit a number between 2 and 99 (inclusive).
  - ➢ If both report the same number, each gets this value.
  - ➢ If one reports a lower number (s) than the other (t), the former gets s+2, the latter gets s-2.

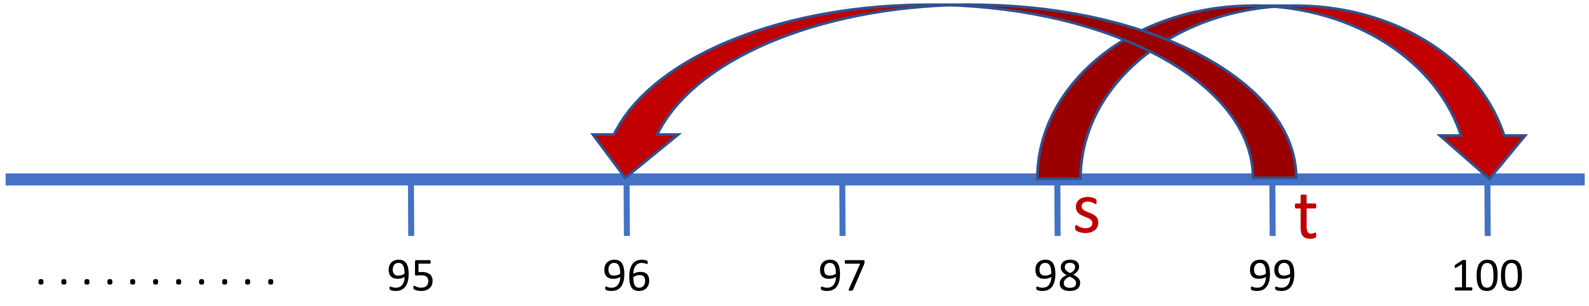

Extra Fun 2: Ice Cream Shop

- Two brothers, each wants to set up an ice cream shop on the beach ([0,1]).
- If the shops are at s, t (with $s \leq t$)
  - The brother at s gets $\left[0, \frac{s+t}{2}\right]$, the other gets $\left[\frac{s+t}{2}, 1\right]$

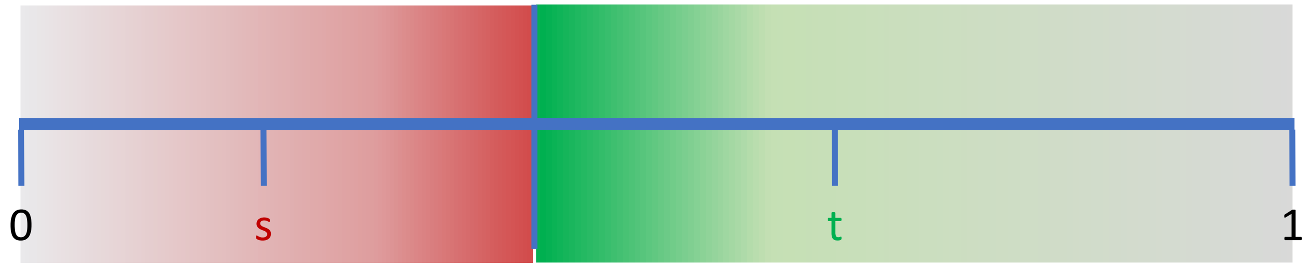

Nash Equilibria: Critique

- Noncooperative game theory provides a framework for analyzing rational behavior.
- But it relies on many assumptions that are often violated in the real world.
- Due to this, human actors are observed to play Nash equilibria in some settings, but play something far different in other settings.

Nash Equilibria: Critique

- Assumptions:
  - Rationality is common knowledge.
    - All players are rational.
    - All players know that all players are rational.
    - All players know that all players know that all players are rational.
    - … [Aumann, 1976]
    - Behavioral economics
  - Rationality is perfect = "infinite wisdom"
    - Computationally bounded agents
  - Full information about what other players are doing.
    - Bayes-Nash equilibria

Nash Equilibria: Critique

- Assumptions:
  - No binding contracts.
    - Cooperative game theory
  - No player can commit first.
    - Stackelberg games (will study this in a few lectures)
  - No external help.
    - Correlated equilibria
  - Humans reason about randomization using expectations.
    - Prospect theory

Nash Equilibria: Critique

- Also, there are often multiple equilibria, and no clear way of "choosing" one over another.

- For many classes of games, finding a single equilibrium is provably hard.
  - Cannot expect humans to find it if your computer cannot.

Nash Equilibria: Critique

- Conclusion:
  - For human agents, take it with a grain of salt.
  - For AI agents playing against AI agents, perfect!