CS540 Introduction to Artificial Intelligence
Lecture 20

Young Wu
Based on lecture slides by Jerry Zhu, Yingyu Liang, and Charles Dyer

July 12, 2020

Guess Average Game

Motivation

- Write down an integer between 0 and 100 that is the closest to two thirds (2/3) of the average of everyone's (including yours) integers.

Guess Average Game Derivation

Motivation

$$R^0 = [0, 1, 2, \ldots \quad \cancel{100}]$$

$$R^1 = [0, 1, 2, \ldots 66]$$

$$R^2 = [0, 1, 2, \ldots 44]$$

$$R^3 = [0, 1, 2, \ldots 30]$$

$$R^4 = [0, \ldots 20]$$

$$\vdots$$

$$R^\infty = [0, 1]$$

rationalizable

Rationalizability

Motivation

IESDS

- An action is 1-rationalizable if it is the best response to some action.
- An action is 2-rationalizable if it is the best response to some 1-rationalizable action.
- An action is 3-rationalizable if it is the best response to some 2-rationalizable action.
- An action is rationalizable if it is ∞-rationalizable.

Traveler's Dilemma Example

Motivation

$\underline{100}$

99

$99 - x$ $\quad$ $99 + x$

100

- Two identical antiques are lost. The airline only knows that its value is at most $\cancel{x}$ dollars, so the airline asks their owners (travelers) to report its value (integers larger than or equal to some integer $x > 1$). The airline tells the travelers that they will be paid the minimum of the two reported values, and the traveler who reported a strictly lower value will receive x dollars in reward from the other traveler.
- The best response of to v is $\max(v - 1, x)$, and the only mutual best response is both report x. $\quad (x, x)$

NE

- This result is inconsistent with experimental observations.

Traveler's Dilemma Example Derivation

Motivation

$x > 1$

	1	2	
100	100	100	99 - x
99+x	99	99	
99 / 98+x	98	98	

97+x

96+x

97

96-x ($\cancel{98}$, $\cancel{96}$)

96 ← 96

95+x 95

x+1

1

x ← x, x

Normal Form Games

Definition

- In a simultaneous move game, a state represents one action from each player. → list of actions
- The costs or rewards, sometimes called payoffs, are written in a payoff table.
- The players are usually called the ROW player and the COLUMN player.
- If the game is zero-sum, the convention is: ROW player is MAX and COLUMN player is MIN.

Best Response

Definition

- An action is a best response if it is optimal for the player given the opponents' actions.

payoff

$$br_{MAX}(s_{MIN}) = \arg\max_{s \in S_{MAX}} c(s, s_{MIN})$$

$$br_{MIN}(s_{MAX}) = \arg\min_{s \in S_{MIN}} c(s_{MAX}, s)$$

Strictly Dominated and Dominant Strategy

Definition

- An action s_i strictly dominates another $s_{i'}$ if it leads to a better state no matter what the opponents' actions are.

for all

$$s_i \succ_{MAX} s_{i'} \text{ if } c(s_i, s) > c(s_{i'}, s) \ \forall s \in S_{MIN}$$

$$s_i \succ_{MIN} s_{i'} \text{ if } c(s, s_i) < c(s, s_{i'}) \ \forall s \in S_{MAX}$$

- The action $s_{i'}$ is called strictly dominated.
- An action that strictly dominates all other actions is called strictly dominant.

Weakly Dominated and Dominant Strategy

Definition

- An action s_i weakly dominates another $s_{i'}$ if it leads to a better state or a state with the same payoff no matter what the opponents' actions are.

$>$

$$s_i \succ_{MAX} s_{i'} \text{ if } c(s_i, s) \geqslant c(s_{i'}, s) \;\; \forall\, s \in S_{MIN}$$
$$s_i \succ_{MIN} s_{i'} \text{ if } c(s, s_i) \leqslant c(s, s_{i'}) \;\; \forall\, s \in S_{MAX}$$

- The action $s_{i'}$ is called weakly dominated.

Nash Equilibrium

Definition

- A Nash equilibrium is a state in which all actions are best responses.

Prisoner's Dilemma

Discussion

- A simultaneous move, non-zero-sum, and symmetric game is a prisoner's dilemma game if the Nash equilibrium state is strictly worse for both players than another state.

–	C	D
C	(x, x)	$(0, y)$
D	$(y, 0)$	$(1, 1)$

- C stands for Cooperate and D stands for Defect (not Confess and Deny). Both players are MAX players. The game is PD if $y > x > 1$. Here, (D, D) is the only Nash equilibrium and (C, C) is strictly better than (D, D) for both players.

Prisoner's Dilemma Derivation

Discussion

		COL	
		C	D
ROW	C	x, x	0, y
	D	y, 0	$\boxed{1, 1}$

NE

$br(D) = D$

$br(C) = D$

$(\cancel{C}, \cancel{C})$ D D $\to y$

$0 \leftarrow (\cancel{C}, D)$

$1 \leftarrow D$

Public Good Game

Discussion

On Final exam

- You received one free point for this question and you have two choices.
- A: Donate the point.
- B: Keep the point.
- Your final grade is the points you keep plus twice the average donation.

2

1

0

0, 1

0

1

Properties of Nash Equilibrium

Discussion

- All Nash equilibria are rationalizable.
- No Nash equilibrium contains a strictly dominated action.
- ~~Nash equilibrium~~ can be found be iterated elimination of strictly dominated actions.

IESDS

- The above statements are not true for weakly dominated actions.

Normal Form of Sequential Games

Discussion

- Sequential games can have normal form too, but the solution concept is different.
- Nash equilibria of the normal form may not be a solution of the original sequential form game.

Fixed Point Algorithm

Description

- For small games, it is possible to find all the best responses. The states that are best responses for all players are the solutions of the game.
- For large games, start with a random action, find the best response for each player and update until the state is not changing.

Fixed Point Diagram

Definition

$\underline{x} = \underline{f}(\underline{x})$

(x, y)

$\begin{pmatrix} br_1 \\ br_2 \end{pmatrix}\begin{pmatrix} a_2 \\ a_1 \end{pmatrix} \Rightarrow \begin{pmatrix} a_1 \\ a_2 \end{pmatrix}$

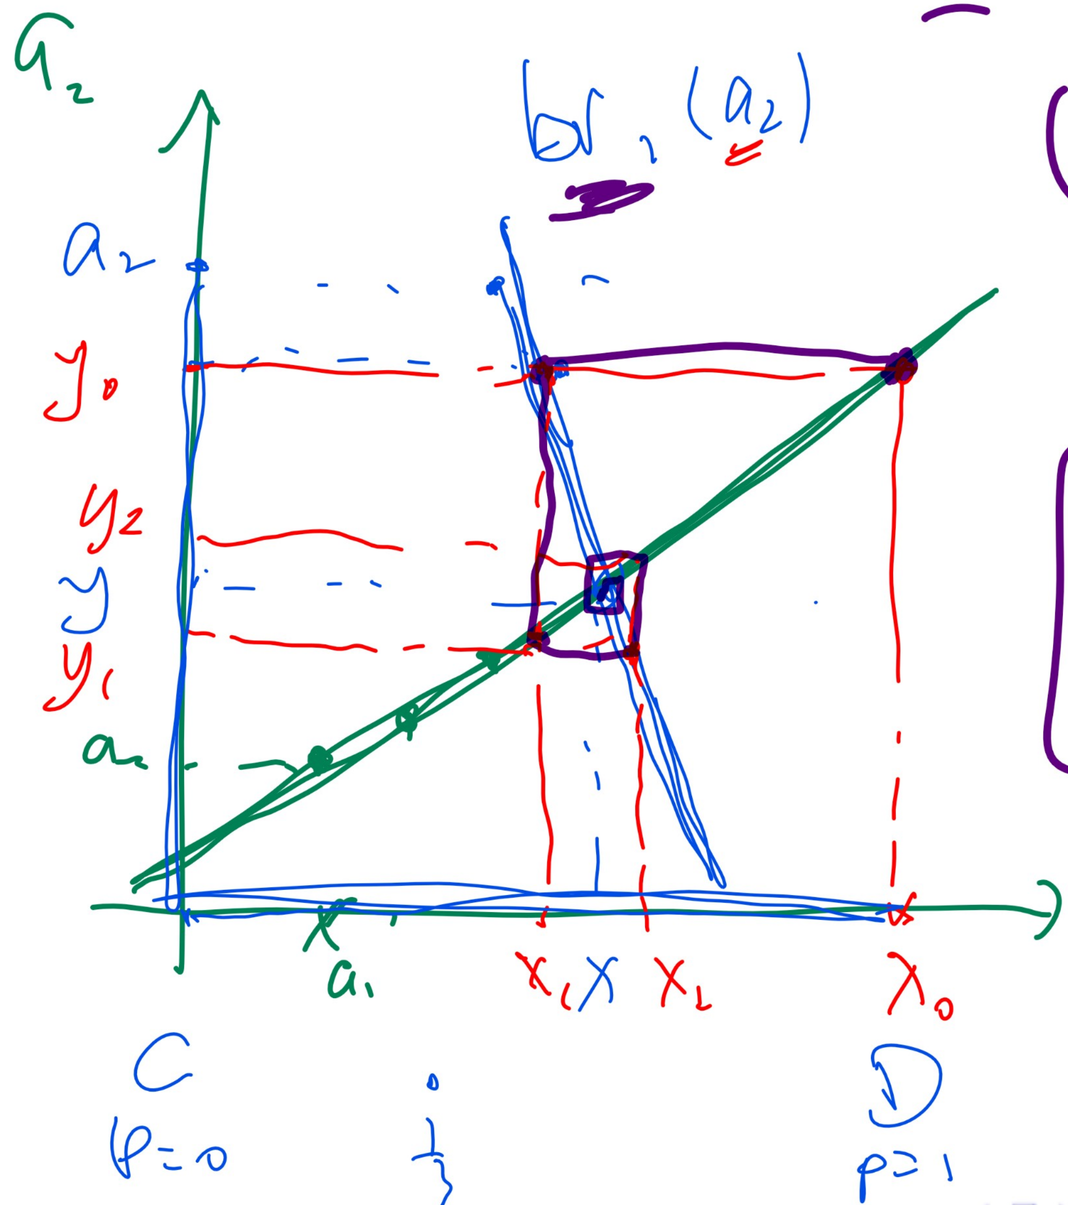

$br_2(a_1)$

$$\boxed{\begin{aligned} y &= br_2(x) \\ x &= br_1(y) \end{aligned}}$$

a_1

p

Mixed Strategy Nash Equilibrium

Definition

- A mixed strategy is a strategy in which a player randomizes between multiple actions.
- A pure strategy is a strategy in which all actions are played with probabilities either 0 or 1.
- A mixed strategy Nash equilibrium is a Nash equilibrium for the game in which mixed strategies are allowed.

Rock Paper Scissors Example

Discussion

- There are no pure strategy Nash equilibria.
- Playing each action (rock, paper, sciossors) with equal probability is a mixed strategy Nash.

Rock Paper Scissors Example Derivation

Discussion

Standard RPS

		Opponent		
		Rock	Paper	Scissors
You	Rock	0, 0	-1, 1	1, -1
	Paper	1, -1	0, 0	-1, 1
	Scissors	-1, 1	1, -1	0, 0

no pure NE

$NE = \left(\left(\frac{1}{3}_R, \frac{1}{3}_P, \frac{1}{3}_S \right), \left(\frac{1}{3}_R, \frac{1}{3}_P, \frac{1}{3}_S \right) \right)$

$\to \left(\frac{1}{2}, \frac{1}{2}, 0 \right) = P \; (0, 1, 0)$

$R \to \frac{1}{2} 0 + \frac{1}{2}(-1) = -\frac{1}{2}$

$P \to \frac{1}{2} \cdot 1 + \frac{1}{2} 0 = \frac{1}{2}$

$S \to \frac{1}{2}(-1) + \frac{1}{2}(1) = 0$

$R \to \frac{1}{3} 0 + \frac{1}{3}(-1) + \frac{1}{3}(1) = 0$

$P \to \frac{1}{3} 1 + \frac{1}{3} \cdot 0 + \frac{1}{3}(-1) = 0$

$S \to \frac{1}{3}(-1) + \frac{1}{3}(1) + \frac{1}{3} 0 = 0$

$\frac{1}{3} R \; \frac{1}{3} P \; \frac{1}{3} S$

Battle of the Sexes Example

Discussion

Romeo Juliet

- Battle of the Sexes (BoS, also called Bach or Stravinsky) is a game that models coordination in which two players have different preferences in which alternative to coordinate on.

—	Bach	Stravinsky
Bach	A (x, y)	B (0, 0)
Stravinsky	C (0, 0)	D (y, x)

Romeo

$y > x > 0$

Battle of the Sexes Example Diagram

Discussion

		p	$1-p$
		B	S
q	B	x, y	$0, 0$
$1-q$	S	$0, 0$	y, x

$$br_{COL}(q) \begin{cases} B & yq > x(1-q) \Rightarrow q > \frac{x}{x+y} \\ p \in [0,1] & q = \frac{x}{x+y} \\ S & q < \frac{x}{x+y} \end{cases}$$

for ROW:

$$\boxed{px + (1-p)0 = p0 + (1-p)y}$$

$$px > (1-p)y \Rightarrow p > \frac{x}{x+y}$$

$$p = \frac{y}{x+y}$$

$$p < \frac{y}{x+y}$$

$$br_{ROW}(p) = \begin{cases} B \\ B^q S^{1-q} \quad q \in [0,1] \\ S \end{cases}$$

Battle of the Sexes Example Derivation

Discussion

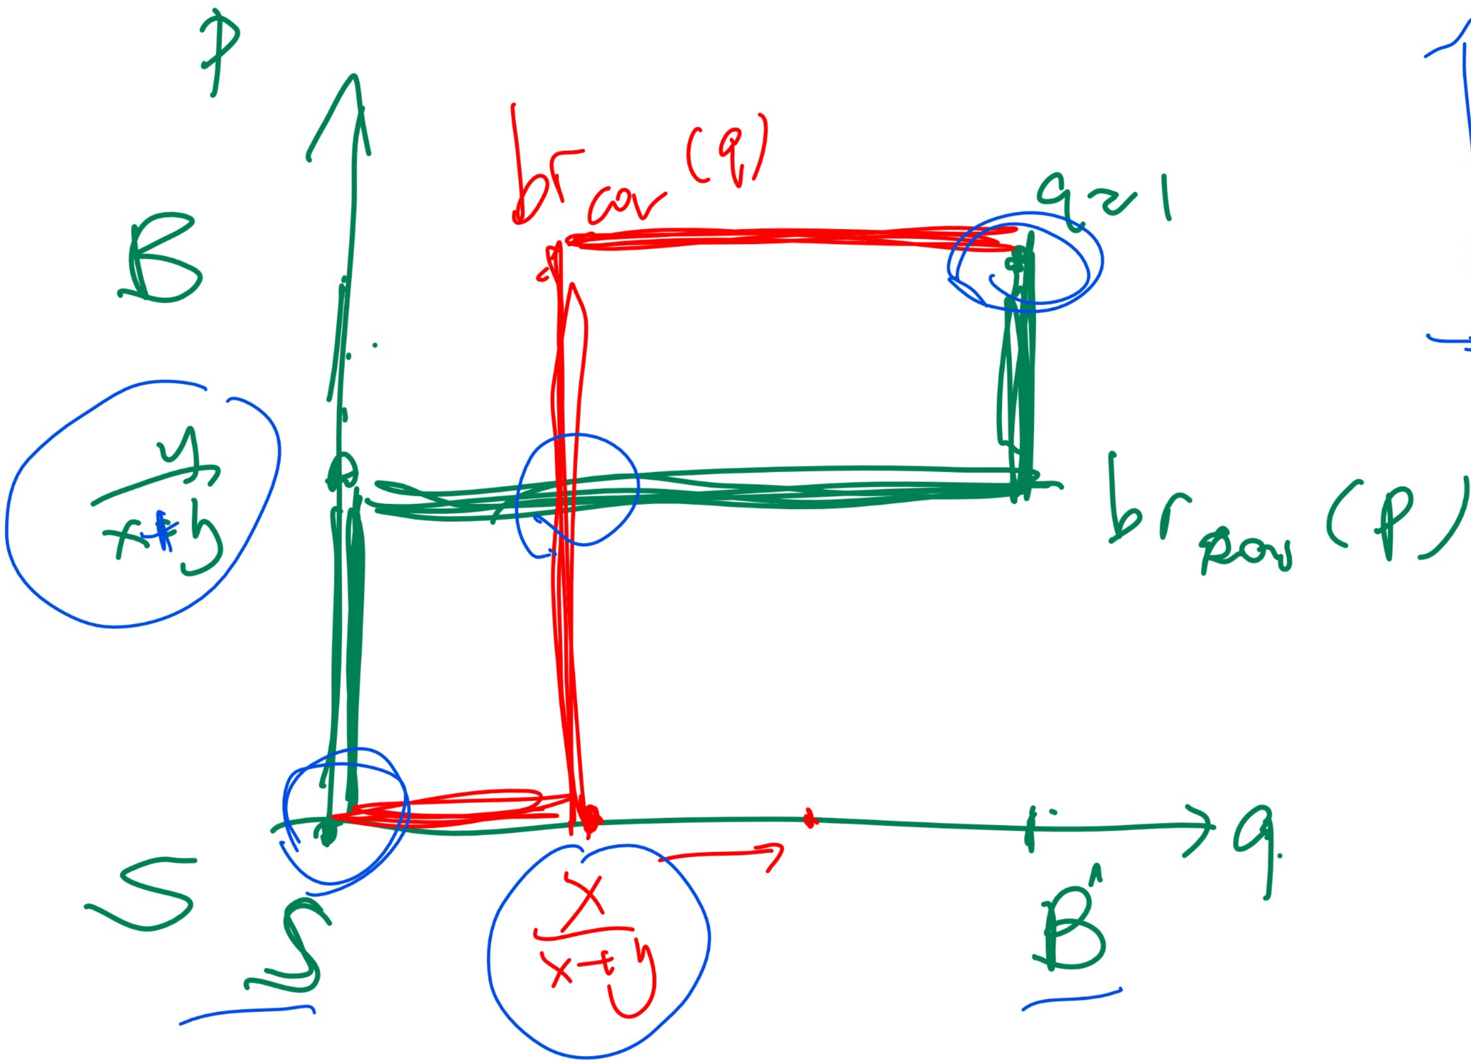

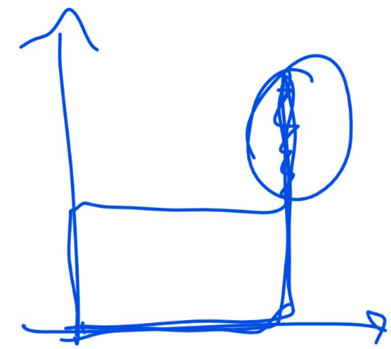

Volunteer's Dilemma

Discussion

- On March 13, 1964, Kitty Genovese was stabbed outside the apartment building. There are 38 witnesses, and no one reported. Suppose the benefit of reported crime is 1 and the cost of reporting is $c < 1$.
- Suppose every witness uses the same mixed strategy of not reporting with probability p and reporting with probability $1 - p$. Then the mixed strategy Nash equilibrium is characterized by the following expression.

$$p^{37} \cdot 0 + (1 - p^{37}) \cdot 1 = 1 - c \Rightarrow p = c^{\frac{1}{37}}$$

Volunteer's Dilemma Derivation

Discussion

$1-p$ $\quad\quad\quad\quad$ p

R $\quad\quad\quad\quad$ not R

$$\boxed{1 - c = 0 \cdot p^{37} + 1 \cdot (1 - p^{37})}$$

$$c = p^{37} \qquad\qquad p = c^{\frac{1}{37}}$$

$$\Pr(\text{no repose}) = p^{38} = c^{\frac{38}{37}}$$

Nash Theorem

Definition

- Every finite game has a Nash equilibrium. (mixed or pure)
- The Nash equilibria are fixed points of the best response functions.

Fixed Point Nash Equilibrium

Algorithm

- Input: the payoff table $c(s_i, s_j)$ for $s_i \in S_{MAX}, s_j \in S_{MIN}$.
- Output: the Nash equilibria.
- Start with random state $s = (s_{MAX}, s_{MIN})$.
- Update the state by computing the best response of one of the players.

$$\begin{aligned} \text{either } s' &= \left(br_{MAX}\left(s_{MIN}\right), br_{MIN}\left(br_{MAX}\left(s_{MIN}\right)\right)\right) \\ \text{or } s' &= \left(br_{MAX}\left(br_{MIN}\left(s_{MAX}\right)\right), br_{MIN}\left(s_{MAX}\right)\right) \end{aligned}$$

- Stop when $s' = s$.