CS 396: Online Markets

Lecture 8: Game Theory

Last Time:

- multi-armed bandit learning
- reduction to online learning

Today:

- game theory
- bimatrix games
- Nash equilbrium
- dominant strategy equilibrium

Exercise: Best Response

Setup:

- you observe weather forecast, probability q of rain.
- you choose bus or bike to school.
- nature realizes weather: rain w. prob. q, sunny w. prob. $(1-q)$
- payoffs:

	it rains	it's sunny
Bus	1	2
Bike	-2	6

Questions: What is your best action given:

- Forecast of $q = 0$ chance of rain?
- Forecast of $q = 1/2$ chance of rain?
- Forecast of $q = 1$ chance of rain?

Game Theory

"predict outcomes of strategic scenarios"

Def: bimatrix game

- row player with n actions
- column player with m actions
- payoff matrices $R, C \in \mathbb{R}^{n\times m}$.
- on row action i, column action j:
  - row player payoff R_{ij}
  - column player payoff C_{ij}

Example: Chicken

	Stay	Swerve
Stay	**-5**,-5	**1**,-1
Swerve	**-1**,1	**-1**,-1

Equilibria

"stable outcomes in a game"

Def: best response

- if column player plays j
- row player's best responce is $i^* = \text{argmax}_i\, R_{ij}$

Q: in Chicken

- What should R do if C stays? **A:** Swerve
- What should C do if R swerves? **A:** Stay

Conclusion: (Swerve,Stay) is a mutual best response

A.k.a., a **pure Nash equilibrium**

Symmetry: (Swerve,Stay) is also pure Nash equilibium.

Randomized Equilibria

"a.k.a., mixed"

Q: are there other equilibria?

A: yes

- suppose C stays with probability q
- plot R's payoff:
  - for swerving: $-5\,q + 1\,(1-q) = 1 - 6q$
  - for staying: $-1\,q + -1\,(1-q) = -1$

[PICTURE]

- R is indifferent at $q = 1/3$.

Def: n dimensional probability simplex

$\Delta_n = \{\mathbf{p} \in [0,1]^n : \sum_i p_i = 1\}$

Def: A **mixed Nash equilibrium** is

- row player plays $\mathbf{p}^* \in \Delta_n$
- column player plays $\mathbf{q}^* \in \Delta_m$
- $\mathbf{p}^*$ is best response: $\mathbf{p}^* \in \text{argmax}_{\mathbf{p}}\, \mathbf{p}^T R \mathbf{q}^*$
- $\mathbf{q}^*$ is best response: $\mathbf{q}^* \in \text{argmax}_{\mathbf{q}}\, \mathbf{p}^{*T} C \mathbf{q}$

Fact: $\mathbf{p}^*, \mathbf{q}^*$) is Nash iff

- $\mathbf{p}^*_i > 0 \Rightarrow$ "action i is best response to $\mathbf{q}^*$"
- $\mathbf{q}^*_j > 0 \Rightarrow$ "action j is best response to $\mathbf{p}^*$"

Exercise: "Battle of the Sexes"

Setup:

- you are the row player.
- payoffs:

	Opera	Football
Opera	**4**, 2	**0**, 0
Football	**0**, 0	**2**, 4

Questions:

- Identify the pure Nash equilibria.
- Is there a mixed Nash; if so, what is your probability of selecting Opera in it?
- Let's play! (Imagine the column player is a random classmate.)

Existence of Equilibria

Example: Hide and Seek

	Seek at A	Seek at B
Hide at A	**-1**,1	**1**,-1
Hide at B	**1**,-1	**-1**,1

A.k.a. Matching Pennies

Q: pure Nash equilbrium? **A:** no!

Q: mixed Nash? **A:** yes!

- R mixes $(1/2, 1/2) \Rightarrow$ C is indifferent.
- C mixes $(1/2, 1/2) \Rightarrow$ R is indifferent.

Thm: [Nash '51] Every game with finite action space has a (possibly) mixed (Nash) equilibrium.

Dominant Strategy Equilibria

Example: Prisoner's Dilemma

	Defect	Cooperate
Defect	**-2**,-2	**0**,-3
Cooperate	**-3**,0	**-1**,-1

Q: in Prisoner's Dilemma

- What should R do if C Defects? **A:** Defect
- What should R do if C Cooperates? **A:** Defect

Conclusion:

- Defect is a dominant strategy in Prisoner's Dilemma
- (Defect,Defect) is a dominant strategy equilbrium

Def:

- a **dominant strategy** is one that is a best response to any action.
- a **dominant strategy equilbrium** is one where all players play dominant strategies.

Notes:

- DSE $\subseteq$ pure Nash $\subseteq$ (possibly mixed) Nash
- only (possibly mixed) Nash always guaranteed to exist.
- equilibria can be unique, e.g., Hide and Seek, Prisoner's Dilemma
- or not, e.g., Chicken

Challenge: predictions in games with multiple equilibria.