CMSC 474, Introduction to Game Theory

4. Analyzing Normal-Form Games

Mohammad T. Hajiaghayi
University of Maryland

Some Comments about Normal-Form Games

- Only two kinds of strategies in the normal-form game representation:
  - **Pure strategy**: just a single action
  - **Mixed strategy**: probability distribution over pure strategies
    - i.e., choose an action at random from the probability distribution
- The normal-form game representation may see very restricted
  - No such thing as a conditional strategy (e.g., cross the bay if the temperature is above 70)
  - No temperature or anything else to observe

	C	*D*
C	3, 3	0, 5
D	5, 0	1, 1

- However much more complicated games can be mapped into normal-form games
  - Each pure strategy is a description of what you'll do in *every* situation you might ever encounter in the game
- In later sessions, we see more examples

How to reason about games?

- In single-agent decision theory, look at an **optimal strategy**
  - Maximize the agent’s expected payoff in its environment

- With multiple agents, the best strategy depends on others’ choices
- Deal with this by identifying certain subsets of outcomes called **solution concepts**

- This second chapter of the book discusses two solution concepts:
  - Pareto optimality
  - Nash equilibrium
- Chapter 3 will discuss several others

Pareto Optimality

- A strategy profile **s** **Pareto dominates** a strategy profile **s′** if
  - no agent gets a worse payoff with **s** than with **s′**,
    i.e., $u_i(\mathbf{s}) \geq u_i(\mathbf{s}')$ for all i ,
  - at least one agent gets a better payoff with **s** than with **s′**,
    i.e., $u_i(\mathbf{s}) > u_i(\mathbf{s}')$ for at least one i

- A strategy profile **s** is **Pareto optimal** (or **Pareto efficient**) if there's no strategy profile **s′** that Pareto dominates **s**
  - Every game has at least one Pareto optimal profile
  - Always at least one Pareto optimal profile in which the strategies are pure

Examples

	C	*D*
C	3, 3	0, 5
D	5, 0	1, 1

The Prisoner's Dilemma

- (*D*,*C*) is Pareto optimal: no profile gives player 1 a higher payoff
- (*C*, *D*) is Pareto optimal: no profile gives player 2 a higher payoff
- (*C*,*C*) is Pareto optimal: no profile gives both players a higher payoff
- (*D*,*D*) isn't Pareto optimal: (*C*,*C*) Pareto dominates it

Which Side of the Road

- (Left,Left) and (Right,Right) are Pareto optimal
- In common-payoff games, all Pareto optimal strategy profiles have the same payoffs
  - If (Left,Left) had payoffs (2,2), then (Right,Right) wouldn't be Pareto optimal

	Left	Right
Left	1, 1	0, 0
Right	0, 0	1, 1

Best Response

- Suppose agent i knows how the others are going to play
  - Then i has an ordinary optimization problem: maximize expected utility

- We'll use $\mathbf{s}_{-i}$ to mean a strategy profile for all of the agents except i

$$\mathbf{s}_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$$

- Let s_i be any strategy for agent i. Then

$$(s_i, \mathbf{s}_{-i}) = (s_1, \ldots, s_{i-1}, s_i, s_{i+1}, \ldots, s_n)$$

- s_i is a **best response** to $\mathbf{s}_{-i}$ if for every strategy s_i' available to agent i,

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

- There is always at least one best response
- A best response s_i is **unique** if $u_i(s_i, \mathbf{s}_{-i}) > u_i(s_i', \mathbf{s}_{-i})$ for every $s_i' \neq s_i$

Best Response

- Given $\mathbf{s}_{-i}$, there are only two possibilities:

  (1) i has a pure strategy s_i that is a unique best response to $\mathbf{s}_{-i}$

  (2) i has infinitely many best responses to $\mathbf{s}_{-i}$

Proof. Suppose (1) is false. Then there are two possibilities:

- **Case 1**: s_i isn't unique, i.e., ≥ 2 strategies are best responses to $\mathbf{s}_{-i}$
  - Then they all must have the same expected utility
  - Otherwise, they aren't all "best"
  - Thus any mixture of them is also a best response
- **Case 2**: s_i isn't pure, i.e., it's a mixture of $k > 2$ actions
  - The actions correspond to pure strategies, so this reduces to Case 1
- **Theorem**: Always there exists a pure best response s_i to $\mathbf{s}_{-i}$

Proof. In both (1) and (2) above, there should be one pure best response.

Example

- Suppose we modify the Prisoner's Dilemma to give Agent 1 another possible action:
  - Suppose 2's strategy is to play action *C*
  - What are 1's best responses?
  - Suppose 2's strategy is to play action *D*
  - What are 1's best responses?

	C	*D*
C	3, 3	0, 5
D	5, 0	1, 1
E	3, 3	1, 3

Nash Equilibrium

- **s** = (s_1, …, s_n) is a **Nash equilibrium** if for every i, s_i is a best response to $\mathbf{s}_{-i}$
  - Every agent's strategy is a best response to the other agents' strategies
  - No agent can do better by *unilaterally* changing his/her strategy

- **Theorem (Nash, 1951)**: Every game with a finite number of agents and actions has at least one Nash equilibrium

- In Which Side of the Road, (Left,Left) and (Right,Right) are Nash equilibria

	Left	Right
Left	1, 1	0, 0
Right	0, 0	1, 1

- In the Prisoner's Dilemma, (*D*,*D*) is a Nash equilibrium
  - Ironically, it's the only pure-strategy profile that *isn't* Pareto optimal

	C	*D*
C	3, 3	0, 5
D	5, 0	1, 1