Nash Equilibrium

CMPUT 355: Games, Puzzles, and Algorithms

Lecture Outline

1. Logistics & Recap
2. Expected Utility
3. Best Response
4. Nash Equilibrium

Logistics

- **Practice questions #5** released last Friday
  - Solutions will be posted **tomorrow** (Mar 24)
- **Quiz #5** is **Friday** (Mar 27)
  - Covers up to the end of **last Friday's lecture** (AND/OR & rock-paper-scissors)
- **Quiz #4** marks will be posted in the next couple of days

Recap: Normal Form Games

- Rock Paper Scissors is an example of a **normal form game**
- Players choose actions **simultaneously**
- Each possible **combination** of actions maps to scores for the players ("**utilities**")
- To specify a normal form game, list:
  - **Who** the players are (need not be just 2)
  - What **actions** each player has available (need not be the same per player)
  - The **utility** to each player for each combination of actions (often a table for 2-player games)
- Unlike **perfect information alternating moves**, the optimal strategy for a normal form game may require **randomization** (i.e., **mixed strategies**)

	Rock	Paper	Scissor
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissor	-1,1	1,-1	0,0

Expected Utility

- Suppose column plays a **mixed strategy** s_c
  - i.e., a **distribution** that chooses action b with probability $s_c(b)$
- Row's aim is to maximize their **expected utility**:

$$u_r(a, s_c) = \mathbb{E}_{b \sim s_c} u_r(a, b) = \sum_b \Pr(b) u_r(a, b) = \sum_b s_c(b) u_r(a, b).$$

- Row might **themselves** be playing a mixed strategy s_r:

$$u_r(s_r, s_c) = \mathbb{E}_{\substack{a \sim s_r \\ b \sim s_c}} \left[u_r(a, b) \right] = \sum_a \Pr(a) u_r(a, s_c)$$

$$= \sum_a s_r(a) \mathbb{E}_{b \sim s_c} u_r(a, b) = \sum_a s_r(a) \sum_b s_c(b) u_r(a, b)$$

Expected Utility Example

- What is the utility for each of row's **pure strategies** (deterministic) if column plays

$$s_c(\text{R}) = 0.25 \quad s_c(\text{P}) = 0.25 \quad s_c(\text{S}) = 0.5?$$

$$u_r(\text{R}, s_c) = s_c(\text{R})u_r(\text{R}, \text{R}) + s_c(\text{P})u_r(\text{R}, \text{P}) + s_c(\text{S})u_r(\text{R}, \text{S})$$
$$= .25(0) + .25(-1) + .5(1) = \mathbf{0.25}$$

$$u_r(\text{P}, s_c) = s_c(\text{R})u_r(\text{P}, \text{R}) + s_c(\text{P})u_r(\text{P}, \text{P}) + s_c(\text{S})u_r(\text{P}, \text{S})$$
$$= .25(1) + .25(0) + .5(-1) = \mathbf{-0.25}$$

$$u_r(\text{S}, s_c) = s_c(\text{R})u_r(\text{S}, \text{R}) + s_c(\text{P})u_r(\text{S}, \text{P}) + s_c(\text{S})u_r(\text{S}, \text{S})$$
$$= .25(-1) + .25(1) + .5(0) = \mathbf{0}$$

	R	P	S
R	0,0	-1,1	1,-1
P	1,-1	0,0	-1,1
S	-1,1	1,-1	0,0

- R is row's best **pure strategy** against s_c
- **Question:** Can row do **better** by playing a **mixed strategy** (i.e., randomized)?

Theorem: Pure Best Response

Definition: A strategy s (either mixed or pure) is a **best response** for player i to strategy t (mixed or pure) of the other player if it maximizes i's expected utility among all possible strategies; that is, if $u_i(s, t) \geq u_i(s', t)$ for all possible strategies s'.

Theorem: Every strategy has at least one pure strategy best response. That is, for all strategies t, there exists at least one pure strategy a such that $u_i(a, t) \geq u_i(s, t)$ for all mixed strategies s.

Proof: Pure Best Response

Theorem: Every strategy has at least one pure strategy best response. That is, for all strategies t, there exists at least one pure strategy a such that $u_i(a, t) \geq u_i(s, t)$ for all mixed strategies s.

Proof: Suppose not; there is a best response s that has higher utility than any action a.

So $u_i(s, t) = \sum_{a'} s(a')u_i(a', t) > u_i(a, t)$ for every action a'.

If $u_i(a, t)$ is the same for every action a such that $s(a) > 0$, then $u_i(s, t) = u_i(a, t)$ for each of those a's, a contradiction.

Otherwise, there exists a, b such that $s(a) > 0$ and $s(b) > 0$ and $u_i(a, t) > u_i(b, t)$.

But consider the strategy s' that sets $s'(a) = s(a) + s(b)$ and $s'(b) = 0$.

Clearly $u_i(s', t) > u_i(s, t)$, so s is not a best response at all, and we have another contradiction. ■

Nash Equilibrium

- A **Nash equilibrium** is a profile of strategies in which every player simultaneously best responds to all the other players
- In the two-player case, s, t is a Nash equilibrium when:
  1. s is a best response to t; i.e., $u_r(s, t) \geq u_r(a, t)$ for all actions a
  2. t is a best response to s; i.e., $u_c(s, t) \geq u_c(s, b)$ for all actions b

Equilibrium Example: Rock Paper Sofa

- We can verify whether a strategy profile is an equilibrium by checking whether both players are best responding:
  - Check if any player has an **action** with a better utility (**why just actions?**)
  - Such an action is called a **deviation**

	Rock	Paper	Sofa
Rock	0,0	-1,1	0,0
Paper	1,-1	0,0	0,0
Sofa	0,0	0,0	0,0

Questions:

1. In a **normal form game**, must **every** Nash equilibrium be mixed? (**why?**)
2. Is (Rock, Rock) a Nash equilibrium?
3. Is (Sofa, Sofa) a Nash equilibrium?
4. If both players randomize uniformly between Rock and Sofa?
5. If both players randomize uniformly between Rock and Paper?

Nash's Theorem

Theorem: Every finite game has at least one (possibly mixed) Nash equilibrium.

- **Pure strategy equilibria** were studied for decades before Nash
- But not every game has one! (e.g., rock-paper-scissors)
- His main innovation was extending the concept to **mixed strategies**

Best Response & Indifference

Theorem: Suppose that s is a **best response** for player i to strategy t, and there exist actions $a \neq b$ such that $s(a) > 0$ and $s(b) > 0$ (i.e., s is a **mixed strategy**).

Then i is **indifferent** between a and b; that is,

$$u_i(s, t) = u_i(a, t) = u_i(b, t).$$

- *Implication:* If (s, t) is a Nash equilibrium, then every action that s plays with positive probability is a best response (**why?**)
- So mixed strategies only "mix over" best responses
- **Question:** If a mixed strategy never has higher utility, why would we ever play one?

Summary

- In a **normal form game**, it is sometimes necessary to consider **mixed strategies**
- When one or more players are randomizing, the goal is to maximize **expected utility**
  - A strategy s that maximizes expected utility against opponent strategy t is a **best response to t**
- In a **Nash equilibrium**, all players **simultaneously best respond** to each other
  - Can verify whether a strategy profile is a Nash equilibrium by checking for **deviations**: strategies that would yield a **higher utility** to the deviating player
- Some games have pure strategy equilibria (nobody randomizes)
- Some games have only mixed strategy equilibria (somebody randomizes)
  - Even though randomizing cannot increase the utility of the randomizing player!