CS 217: Artificial Intelligence and Machine Learning	**Jan-Apr 2024**

Lecture 22: Simultaneous Move Games

Lecturer: Swaprava Nath — *Scribe(s): SG43, SG44*

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor.*

22.1 Expected Utility

Let us start by considering a reduced version of the penalty shootout game (players 1 and 2 are the shooter and the goalkeeper respectively).

		P2	
		L	R
P1	L	(-1, +1)	(+1, -1)
	R	(+1, -1)	(-1, +1)

Table 22.1: Simultaneous move game

We note that the utility of Player 1 is the same if he chooses L or R if $\sigma_2 = (\frac{1}{2}, \frac{1}{2})$
Similarly, Player 2 is indifferent towards L or R if $\sigma_1 = (\frac{1}{2}, \frac{1}{2})$

For any (independent) σ_1 and σ_2, the utility of Player 1 is given by

$$U_1(\sigma_1, \sigma_2) = \sum_{s_2 \in S_2} \sum_{s_1 \in S_1} \sigma_1(s_1)\ \sigma_2(s_2)\ u_1(s_1, s_2)$$

For example, if we take $\sigma_1 = (\frac{2}{3}, \frac{1}{3})$ and $\sigma_2 = (\frac{4}{5}, \frac{1}{5})$, we get

$$u_1\left(\left(\frac{2}{3}, \frac{1}{3}\right), \left(\frac{4}{5}, \frac{1}{5}\right)\right) = \frac{2}{3} \cdot \frac{4}{5} \cdot (-1) + \frac{2}{3} \cdot \frac{1}{5} \cdot (1) + \frac{1}{3} \cdot \frac{4}{5} \cdot (1) + \frac{1}{3} \cdot \frac{1}{5} \cdot (-1) = -\frac{1}{5}$$

In general, for any number of players, the expected utility of player i is given by

$$U_i(\sigma_i, \sigma_{-i}) = \sum_{s_n \in S_n} \cdots \sum_{s_1 \in S_1} \sigma_1(s_1)\ \ldots \sigma_n(s_n)\ u_1(s_1, s_2, \ldots, s_n)$$

where σ_{-i} is a shorthand notation for $(\sigma_1, \sigma_2, \ldots, \sigma_{i-1}, \sigma_{i+1}, \ldots, \sigma_n)$, which denotes the action profile of all the players except player i.

22.2 Mixed Strategy Nash Equilibrium (MSNE)

22.2.1 Definiton

A mixed strategy profile $(\sigma_i^*, \sigma_{-i}^*)$ is a *mixed strategy Nash equilibrium* (MSNE) if

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i', \sigma_{-i}^*) \quad \forall \sigma_i' \ \ \forall i \in N$$

This is referred to as **unilateral deviation**, where player i changes his strategy, but all other players continue to use their earlier strategies.

22.2.2 How to find MSNE

22.2.2.1 Support of a Probability Distribution

It is easy to see that a mixed strategy σ_i of player i is a probability distribution over the S_i, the strategies of player i. Since, we are assuming finite strategy space for all players, $|S_i| < \infty$, for all players i.

We define the **support of the probability distribution** (in our case, the support of a mixed strategy) as the subset of the state space (set of pure strategies) with positive probabilities given by the distribution (or mixed strategy).

Let the set of pure strategies of player 1 be

$$S_1 = \{s_{11}, s_{12}, s_{13}, \ldots\}$$

with the following probabilities denoting a mixed strategy

$$\sigma_1 = \{p_1, p_2, p_3, 0, 0, \ldots\}; \quad p_1, p_2, p_3 > 0$$

The support of σ_1 is given by $\delta(\sigma_1) = \{s_{11}, s_{12}, s_{13}\}$ as the probability of the player choosing them is positive.

22.2.2.2 Theorem

A mixed strategy profile $(\sigma_1^*, \sigma_2^*, \ldots, \sigma_n^*)$ is a MSNE iff

1. $u_i(s_i, \sigma_{-i}^*)$ is the same for all $s_i \in \delta(\sigma_i)$
2. $u_i(s_i, \sigma_{-i}^*) \geq u_i(s_i', \sigma_{-i}^*); \quad s_i \in \delta(\sigma_i^*)$ and $s_i' \notin \delta(\sigma_i^*)$

In the above example,

		P2	
		L	R
P1	L	(-1, +1)	(+1, -1)
	R	(+1, -1)	(-1, +1)

Table 22.2:

The above theorem doesn't give us a way to find the support. So, we have to iterate over all possible supports. Clearly the support possible above are $\{L\}, \{R\}, \{L, R\}$. Following are the possible support combinations for players to find the MSNE.

Case-1 $(\{L\}, \{L\})$ **:-** It means when both player 1 and player 2 puts all masses on the strategy choosing L. For this strategy, $U_1(L, L) < U_1(R, L)$ where L for player 1 belongs to support and R for player 1 doesn't belong to support. Clearly, this inequality fails the second condition of above theorem so this can't be an MSNE.

Case-2 $(\{R\}, \{R\})$ **:-** This is similar to the first case, just here $U_1(R, R) < U_1(L, R)$ which again violates the second condition of above theorem.

Case-3 $(\{L\}, \{R\})$ **:-** Here also $U_2(L, R) < U_2(L, L)$ which again violates the second condition of the theorem.

Case-4 $(\{R\}, \{L\})$ **:-** Here, $U_2(R, L) < U_2(R, R)$. This also violates the second condition of the theorem so this is also not an MSNE.

Case-5 $(\{L\}, \{L, R\})$ **:-** Strategy L and R both are in the support for player 2. Here $U_2(L, L) \neq U_2(L, R)$

and thus, it is violating first condition of the theorem. So, it is also not an MSNE.
Similiary we can show that $(\{L, R\}, L)$, $(\{L, R\}, R)$ and $(\{R\}, L, R)$ are also not MSNE.

Case-6 ({L,R},{L,R}) :- The supports considered for this case covers all the possible strategies and there is no strategy for any player outside the support. So, we don't need to check the second condition of the theorem. Now, for this case to be MSNE, first condition should hold i.e., $U_1(L, (L, R)) = U_1(R, (L, R))$ and $U_2((L, R), L) = U_2((L, R), R)$. For that, let us consider for player 1, $\sigma_1 = (p, 1-p)$ and for player 2, $\sigma_2 = (q, 1-q)$.
To hold:

$$U_1(L, (L, R)) = U_1(R, (L, R))$$
$$U_1(L, (q, 1-q)) = U_1(R, (q, 1-q))$$
$$\text{b q} = 1/2$$

and also,

$$U_2((L, R), L) = U_2((L, R), R)$$
$$U_2((p, 1-p), L) = U_2((p, 1-p), R)$$
$$p(1) + (1-p)(-1) = p(-1) + (1-p)(1)$$
$$\text{p} = 1/2$$

So, p and q should be equal to 1/2 for $\{L, R\}, \{L, R\}$ to be a MSNE.
Now, we shall see another example (Professor's Dilemma) in which we shall find the MSNE.

22.3 Deriving MSNE for Proffessor's Dilemma

The payoff matrix for the Professor and Student can be represented as follows:

	Listen (L)	Sleep (S)
Effort (E)	$(100, 100)$	$(-10, 0)$
No Effort (NE)	$(0, -10)$	$(0, 0)$

Analyzing Pure Strategies

No player has a dominant strategy:

- **Professor's Strategies:**
  - If the Student listens (L), the Professor prefers Effort (E) $(100 > 0)$.
  - If the Student sleeps (S), the Professor prefers No Effort (NE) $(0 > -10)$.
- **Student's Strategies:**
  - If the Professor puts in Effort (E), the Student prefers to Listen (L) $(100 > 0)$.
  - If the Professor does not put in Effort (NE), the Student prefers to Sleep (S) $(-10 > 0)$.

The pure strategy Nash equilibria (PSNE) in this game occur where each player's strategy is a best response to the other's strategy:

1. (Effort, Listen): Both players get the highest payoff possible $(100, 100)$.
2. (No Effort, Sleep): Each player avoids a negative payoff $(0, 0)$.

Finding the Mixed Strategy Nash Equilibrium (MSNE)

We calculate situations where each player is indifferent to the strategies of the other player.

For the Professor:

- Suppose the Student listens with probability p and sleeps with probability $1 - p$.
- The Professor is indifferent between putting in Effort and No Effort when:

$$100p - 10(1 - p) = 0p - 10(1 - p)$$
$$110p = 20$$
$$p = \frac{20}{110} = \frac{2}{11}$$

For the Student:

- Suppose the Professor puts in Effort with probability q and no Effort with probability $1 - q$.
- The Student is indifferent between Listening and Sleeping when:

$$100q - 10(1 - q) = 0$$
$$110q = 10$$
$$q = \frac{10}{110} = \frac{1}{11}$$

Conclusion

The mixed-strategy Nash equilibrium for this game is:

- The Professor puts in Effort with probability $\frac{1}{11}$ and No Effort with probability $\frac{10}{11}$.
- The Student Listens with probability $\frac{2}{11}$ and Sleeps with probability $\frac{9}{11}$.

This analysis indicates that both players will most likely choose the less effort-intensive strategies (No Effort for the Professor and Sleep for the Student), albeit with some small probability of choosing the opposite strategies.

22.4 Mechanism Design and Social Choice

So far, we have asked "Given a game, what is the rational outcome?"

Now we ask, "Given an outcome, what game can we play such that the equilibrium of that game is the desired outcome?"

Let us look at **Voting**

We have a set of n candidates

$$N = \{1, 2, 3, \ldots, n\}$$

and the set of alternatives

$$A = \{a_1, a_2, a_3, \ldots, a_n\}$$

Each agent has **strict** preferences over A

$$p_i = \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ \cdot \\ \cdot \\ \cdot \\ a_n \end{bmatrix}$$

There are n! possible preferences for each agent.

22.4.1 Voting/ Social Choice Function

This function takes as input the preferences of all the agents and outputs the collective decision.

$$f(p_1, p_2, \ldots, p_n) \in A$$

for example,

$$f\begin{pmatrix} a_2 & a_1 & a_4 \\ a_1 & a_2 & a_3 \\ a_3 & a_3 & a_1 \\ a_4 & a_4 & a_2 \end{pmatrix} = a_4 \in A$$

22.4.2 Common Voting Rules

- Each voter votes for exactly 1 candidate (which is his most preferred candidate)
- The candidate with the most number of voters wins: **Plurality**

For example,

$$\begin{array}{c c} & \begin{array}{ccccc} 1 & 2 & 3 & 4 & 5 \end{array} \\ \begin{array}{c} \mathrm{w}_1 \\ \mathrm{w}_2 \\ \mathrm{w}_3 \\ \mathrm{w}_4 \end{array} & \begin{bmatrix} a & a & b & c & d \\ b & b & c & b & b \\ d & c & d & d & c \\ c & d & a & a & a \end{bmatrix} \end{array}$$

In this situation, ‘a’ is in the majority in the top priority row, therefore, ‘a’ is the plurality winner

22.4.2.1 Problem with Plurality and Emergence of the Borda Rule

Even though the most frequent candidate in the top row is ‘a’, he is also the last in 3 other columns. This means that 3 agents hate ‘a’ and 2 agents love ‘a’.
However, ‘b’ is loved by 1 agent and liked by 4 agents, so, ‘b’ seems to be a better candidate than ‘a’.

To handle this problem, Borda suggested that weights must be given to every row. Higher rows have a higher weight.

$$w_i \geq w_{i+1}$$

$$w_i = n - i$$

Sum of weights for alternative **a** $= 3 + 3 + 0 + 0 + 0 = 6$
Sum of weights for alternative **b** $= 2 + 2 + 3 + 2 + 2 = 11$
Sum of weights for alternative **c** $= 0 + 1 + 2 + 3 + 1 = 7$
Sum of weights for alternative **d** $= 1 + 0 + 1 + 1 + 3 = 6$

This gives ‘b’ as the Winner!!

22.4.2.2 Single Transferable Vote (STV)

Run multiple rounds of evaluation. In each round, eliminate the candidate with the lowest plurality.
$=>$ Sequential Elimination

Steps:

- Pick the candidate with the lowest plurality. If multiple candidates have the same and minimum plurality, randomly pick any one of them.
- Replace the eliminated candidate in each row with the candidate below it.
- Repeat this on the next row.

In our example,

Step-1: Eliminate any one of b,c, or d. Randomly pick 'b'

$$\begin{array}{ccccc} 1 & 2 & 3 & 4 & 5 \end{array}$$
$$\begin{bmatrix} a & a & c & c & d \\ d & c & d & d & c \\ c & d & a & a & a \end{bmatrix}$$

Step-2: Eliminate 'd'

$$\begin{array}{ccccc} 1 & 2 & 3 & 4 & 5 \end{array}$$
$$\begin{bmatrix} a & a & c & c & c \\ c & c & a & a & a \end{bmatrix}$$

Step-3: Eliminate 'a'

'c' is the Winner!!

22.4.3 Condorcet Consistency

A Condorcet winner is a candidate who beats every other candidate in pairwise elections. For example,

$$\begin{array}{ccc} 1 & 2 & 3 \end{array}$$
$$\begin{bmatrix} a & b & c \\ b & c & a \\ c & a & b \end{bmatrix}$$

- a $\longleftrightarrow$ b : a
- b $\longleftrightarrow$ c : b
- c $\longleftrightarrow$ a : c

We see that there may not exist a Condorcet Consistent Winner

-> A Condorcet consistent voting rule always outputs a Condorcet winner if it exists.

22.4.3.1 Is Plurality Condorcet Consistent?

Consistency is split across 3 parts

$$\begin{array}{ccc} 1 & 2 & 3 \end{array}$$
$$\begin{bmatrix} a & b & c \\ b & c & a \\ c & a & b \\ 30\% & 30\% & 40\% \end{bmatrix}$$

- a $\longleftrightarrow$ b : a (70%)
- b $\longleftrightarrow$ c : b (60%)
- c $\longleftrightarrow$ a : a (60%)

We pick 'a' over 'b' and 'c'
However, plurality will give 'c' as the winner as it has 40% weightage in the first row

$\implies$ Plurality is not Condorcet consistent

22.4.4 Copeland Rule

Copeland Score of candidate 'a' = Number of pairwise wins of 'a'

Copeland Rule: Highest Copeland scorer wins

Example: Candidate **a** wins the most in pairwise election. Following the Copeland Rule, candidate **a** is the winner.

30%	30%	40%
a	b	c
b	a	a
c	c	b

pairwise election	counts	winner
(a, b)	(70, 30)	a
(a, c)	(60, 40)	a
(b, c)	(60, 40)	b

	a	b	c
winner count	2	1	0

Table 22.3: Result of Copeland Rule

NOTE: Copeland rule is Condorcet consistent

Q: Can there exist a Copeland winner who loses to another candidate in pairwise elections?
Ans:

- If the preference profile has a Condorcet winner, then the answer is no. Because there exists one candidate who beats everyone else in pairwise election by definition, Copeland will certainly output that candidate.
- If the preference profile does not have a Condorcet winner, then the answer is yes.
  In the above example (without weights), every candidate has a Copeland score of 1. No matter which of them is selected, there exists another candidate who beats it in a pairwise election.

Q: Is Copeland a good rule in that case?
Ans: The point is that Copeland ensures Condorcet consistency, i.e., if there exists a clear winner, who beats everyone else pairwise, that candidate should be chosen. The Condorcet condition does not say when there does not exist any such clear winner – in a way, all possible choices may be bad to a certain extent – and Copeland chooses one such 'bad' outcome.