CS861: Algorithmic Game Theory & Learning

Chapter 1: General Sum Games

Kirthevasan Kandasamy

UW-Madison

Outline

1. Basic definitions: normal form game, dominant strategy, safe strategy, Nash equilibrium
2. Indifference principle
3. Potential games and repeated play dynamics
4. Price of Anarchy/Stability
5. Correlated and coarse correlated equilibria

Slides are intended as teaching aids only and do not include all material discussed in class. Students are strongly encouraged to attend lectures and take their own notes.

Ch 1.1: Basic definitions

Example 1: Prisoner's dilemma. Two suspects in a robbery are questioned separately. Each must individually choose to *Remain Loyal (L)* or *Betray (B)* their partner. Their sentences depend on both decisions:

- If both remain loyal: both receive light sentences (1 year).
- If both betray: both receive moderate sentences (3 years).
- If one betrays while the other remains loyal: betrayer goes free, the loyal prisoner receives a heavy sentence (5 years).

P1 \ P2	**B**	**L**
B	$(-3, -3)$	$(0, -5)$
L	$(-5, 0)$	$(-1, -1)$

Both prisoners are better off if they both remain loyal. But individually, betraying is better *regardless of what the other chooses*.
This is called a *dominant strategy*.

Example 2: Driver vs Inspector

A driver (P1) must choose between parking in an illegal spot (I) or a legal but less convenient spot (L). Simultaneously, an inspector (P2) should decide whether to inspect (In) or not (N). If the driver parks illegally and an inspection occurs, the driver pays a substantial fine. The city bears a cost when drivers park illegally, though this is partially offset by the revenue from the fine.

P1 \ P2	**N** (Not inspect)	**In** (Inspect)
L (Legal)	$(0, 0)$	$(0, -1)$
I (Illegal)	$(10, -10)$	$(-90, -6)$

Question: Does either player have a dominant strategy?
Ans: No. If P1 chooses L, P2 should choose N. If P1 chooses I, P2 should choose In. Similarly for P1.

Question: What are “safe” strategies for P1 and P2?
Ans: P1 should choose L and P2 should choose In.

Example 3: Stag hunt

Two hunters can (individually) choose to chase either a stag (S) or hare (H). Both hunters are needed to successfully hunt the stag, but one is sufficient to catch the hare. If they both chase the same animal, they share the catch. The stag (8) is worth four times the hare (2).

P1 \ P2	**S**	**H**
S	$(4, 4)$	$(0, 2)$
H	$(2, 0)$	$(1, 1)$

1. Neither player has a dominant strategy in this game. If P1 chooses S, P2 should choose S (and vice versa). If P1 chooses H, P2 should choose H.

2. The safe strategy for both players is to choose H.

3. The action profiles (S, S) and (H, H) are called *Nash equilibria* (NE), meaning that no agent benefits from deviating (we will define it formally soon).
We also see that (S, H) and (H, S) are not NE.

Example 3: Stag hunt (cont'd)

In this example, (S, S) and (H, H) are examples of *pure* NE with,

$$u_1(S, S) = u_2(S, S) = 4, \qquad u_1(H, H) = u_2(H, H) = 1.$$

Suppose we allow agents to randomly choose either action with some probability.
- Say P1 follows $s_1 = (x, 1 - x)$, meaning P1 chooses S with probability x and H with probability H.
- Say P2 follows $s_2 = (y, 1 - y)$.

Then, the expected utility of player 1 is,

$$\begin{aligned} u_1(s_1, s_2) &= \mathbb{E}_{a_1 \sim s_1, a_2 \sim s_2} [u_1(a_1, a_2)] \\ &= xy \cdot 4 + x(1 - y) \cdot 0 + (1 - x)y \cdot 2 + (1 - x)(1 - y) \cdot 1 \\ &= 1 - x + y + 3xy. \end{aligned}$$

Similarly, you can show $u_2(s_1, s_2) = 1 - y + x + 3xy$.

Example 3: Stag hunt (cont'd)

Suppose P2 follows $(1/2, 1/2)$. What is the best strategy for P1?

$$u_1(s_1, (1/2, 1/2)) = 1 - x + \frac{1}{2} + \frac{3x}{2} = \frac{3}{2} + \frac{x}{2}.$$

Therefore, $\mathrm{argmax}_{x \in [0,1]} \frac{3}{2} + \frac{x}{2} = 1$, *i.e.*, P1 should chase the stag.

Suppose P2 follows $(1/4, 3/4)$. Then P1 should chase the hare (try at home).

Suppose P2 follows $(1/3, 2/3)$. Then, $u_1(s_1, (1/3, 2/3)) = 1 - x + \frac{1}{3} + x = \frac{4}{3}$. P1 can choose either S or H and will receive the same utility!

The same can be said about P2 if P1 follows $(1/3, 2/3)$.

As neither player has incentive to devaiate, $s_1 = s_2 = (1/3, 2/3)$ is a NE.

Next, we will rigorously define a 'game', and formalize these ideas.

Normal form games

A normal form game is an interaction between n players, denoted $[n] = \{1, \ldots, n\}$.

Each player $i \in [n]$, separately and simultaneously chooses an action $a_i \in \mathcal{A}_i$, where $\mathcal{A}_i$ is the set of available actions for player i.

Let $a = (a_1, \ldots, a_n)$ denote the action profile chosen by all players. Let $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$ denote the action space of all players.

Each player i receives utility $u_i(a)$, where $u_i : \mathcal{A} \to \mathbb{R}$. A player's utility is her degree of satisfaction (higher utility is better for player i), and it depends on the actions of all players.

A normal form game is finite if $|\mathcal{A}| < \infty$.

We will frequently refer to players as *agents* and actions as *pure strategies*.

[0] In this class, for any integer $N \in \mathbb{N}$, $[N]$ will denote the set $\{1, 2, \ldots, N\}$.

Pure and mixed strategies

A *pure strategy* for player i is an action $a_i \in \mathcal{A}_i$. A *mixed strategy*, or simply a *strategy*, for player i is a distribution over actions, *i.e.*, $s_i \in \mathcal{S}_i = \Delta(\mathcal{A}_i)$. If player i employs a mixed strategy s_i, it means they will choose an action randomly sampled from s_i.

Let $s = (s_1, \ldots, s_n)$ denote the *strategy profile* chosen by all players. Similarly, let $\mathcal{S} = \mathcal{S}_1 \times \cdots \times \mathcal{S}_n$ denote the *strategy space* of all players.

Let $s = \{s_j\}_{j \neq i}$ denote the strategy profile of all players except i. Let $\mathcal{S}_{-i} = \bigtimes_{j \neq i} \mathcal{S}_j$ denote the strategy space everyone except i.

We will overload notation to denote the expected utility $u_i : \mathcal{S} \to \mathbb{R}$ as follows:

$$u_i(s) = \mathbb{E}_{a \sim s}[u_i(a)] = \mathbb{E}_{a_i \sim s_i \, \forall \, i}[u_i(a_1, \ldots, a_n)].$$

In particular, if the action set is finite, we have

$$u_i(s) = \sum_{a \in \mathcal{A}} s(a) u_i(a) = \sum_{a \in \mathcal{A}} (s_1(a_1) \times \cdots \times s_n(a_n)) \cdot u_i(a).$$

[0]In this class, $\Delta(A)$ will denote all distributions over a set A. When A is finite, this will refer to the set $\Delta(A) = \{s : A \to [0, 1]; \sum_{a \in A} s(a) = 1\}$.

Solution concepts

How should players choose their strategies in a game?

Driver vs Inspector

P1 \ P2	**N**	**In**
L	$(0, 0)$	$(0, -1)$
I	$(10, -10)$	$(-90, -6)$

Stag hunt

P1 \ P2	**S**	**H**
S	$(4, 4)$	$(0, 2)$
H	$(2, 0)$	$(1, 1)$

We will now look at three solution concepts (and two more later in Ch 1.5):

1. Dominant strategies
2. Safe (maximin) strategies
3. Nash equilibria

Dominant strategies and Dominant Strategy Equilibria

Recall that s_{-i} and $\mathcal{S}_{-i}$ denote the strategy profile and strategy space of all players except i.

Definition (Dominant strategy). A strategy $s_i \in \mathcal{S}_i$ is a *dominant strategy* for player i if it maximizes player i's utility regardless of other's strategies, *i.e.*,

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \qquad \text{for all } s_i' \in \mathcal{S}_i,\ s_{-i} \in \mathcal{S}_{-i}.$$

Definition (Dominant strategy equilibrium). A strategy profile $s = (s_1, \ldots, s_n)$ is a *dominant strategy equilibrium* if s_i is a dominant strategy for each player i.

Dominant strategy equilibria (cont'd)

Prisoner's dilemma

P1 \ P2	**B**	**L**
B	$(-3,-3)$	$(0,-5)$
L	$(-5,0)$	$(-1,-1)$

(B, B) is a DSE.

Stag hunt

P1 \ P2	**S**	**H**
S	$(4,4)$	$(0,2)$
H	$(2,0)$	$(1,1)$

No DSEs in this game.

Safe strategies

Definition (Safe strategies, a.k.a maximin strategies). Consider an n player normal form game. Let $g_i(s_i)$ denote the lowest possible utility agent i could achieve over the strategies of the others, *i.e.*, $g_i(s_i) = \min_{s_{-i} \in \mathcal{S}_{-i}} u_i(s_i, s_{-i})$.
A strategy $\widetilde{s}_i$ is a safe strategy for agent i if $g_i(\widetilde{s}_i) \geq g_i(s_i)$ for all $s_i \in \mathcal{S}_i$.

Intuitively, a safe strategy maximizes a player's utility among the worst possible case of other players' strategies.

Safe strategies: Example 1

Recall the driver vs inspector game.

P1 \ P2	**N** (Not inspect)	**In** (Inspect)
L (Legal)	$(0, 0)$	$(0, -1)$
I (Illegal)	$(10, -10)$	$(-90, -6)$

Suppose that the driver and inspectors strategies are respectively $s_1 = (x, 1 - x)$ and $s_2 = (y, 1 - y)$. Then, the driver's utility is:

$$\begin{aligned} u_1(s_1, s_2) &= xy \cdot 0 + x(1 - y) \cdot 0 + (1 - x)y \cdot 10 + (1 - x)(1 - y) \cdot (-90) \\ &= 10y - 10xy - 90 + 90x + 90y - 90xy \\ &= 90x - 90 + 100y(1 - x) \end{aligned}$$

Therefore,

$$g_1(s_1) = \min_{s_2} u_1(s_1, s_2) = \min_{y} (90x - 90 + 100y(1 - x)) = 90x - 90.$$

Hence, the driver's safe strategy is, $\operatorname{argmax}_{s_1} g_1(s_1) = \operatorname{argmax}_{x \in [0,1]} 90x - 90 = 1$. That is, $s_1 = (1, 0)$ (always park legally).

Safe strategies: Example 2 (Try at home)

Recall the stag hunt game.

P1 \ P2	**S**	**H**
S	$(4, 4)$	$(0, 2)$
H	$(2, 0)$	$(1, 1)$

Suppose that the players have strategies $s_1 = (x, 1-x)$ and $s_2 = (y, 1-y)$. Then, P1's utility is:

$$\begin{aligned} u_1(s_1, s_2) &= xy \cdot 4 + x(1-y) \cdot 0 + (1-x)y \cdot 2 + (1-x)(1-y) \cdot 1 \\ &= -x + 1 + y(3x+1) \end{aligned}$$

Therefore,

$$g_1(s_1) = \min_{s_2} u_1(s_1, s_2) = \min_{y} \left(-x + 1 + y(3x+1)\right) = -x + 1.$$

Hence, the P1's safe strategy is, $\mathrm{argmax}_{s_1} g_1(s_1) = \mathrm{argmax}_{x \in [0,1]} -x + 1 = 0$. That is, $s_1 = (0, 1)$ (always chase the hare).

Best responses and Nash equilibria

Definition (Best response). For player i, a strategy s_i is a *best response* to a strategy profile s_{-i} of the other players, if it maximizes i's utility, *i.e.*,

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \qquad \text{for all } s_i' \in \mathcal{S}_i$$

Definition (Nash equilibrium). A strategy profile $s^\star = (s_1^\star, \ldots, s_n^\star)$ is a *Nash equilibrium* if $s_i^\star$ is a best response to $s_{-i}^\star$ for each agent i, *i.e.*,

$$u_i(s_i^\star, s_{-i}^\star) \geq u_i(s_i', s_{-i}^\star) \qquad \text{for all } s_i' \in \mathcal{S}_i,\ i \in [n].$$

If $s^\star$ consists only of pure strategies, it is called a pure Nash equilibrium.

Nash equilibria (cont'd)

Prisoner's dilemma

P1 \ P2	**B**	**L**
B	$(-3, -3)$	$(0, -5)$
L	$(-5, 0)$	$(-1, -1)$

(B, B) is a NE.

Stag hunt

P1 \ P2	**S**	**H**
S	$(4, 4)$	$(0, 2)$
H	$(2, 0)$	$(1, 1)$

(S, S) and (H, H) are pure NE.
$((1/3, 2/3), (1/3, 2/3))$ is a mixed NE.

Quiz: state if the following statements are true or false

1. Every DSE is a NE.
2. Every NE is a DSE.
3. If a player has a mixed dominant strategy, she also has a pure dominant strategy.
4. Every normal form game has a DSE.
5. In every normal form game, each player has a safe strategy.
6. Every normal form game has a NE.
7. If a player has a dominant strategy, it is also a safe strategy for her.
8. If each player follows their safe strategy, it constitutes a NE.

Existence of NE

The following is the celebrated theorem of John Nash.

Theorem (Existence of NE, Nash 1950). Every *finite* normal form game has at least one (mixed) Nash equilibrium.

We will not prove this theorem in class. You may read chapter 5 of KP for a proof.

Symmetric Games

Definition (Symmetric game). Suppose all players in an n-player game have the same set of actions $\mathcal{A}$. Let $u_i(\widetilde{a}, \widetilde{a}')$ denote i's utility when she follows $\widetilde{a} \in \mathcal{A}$ and others follow $\widetilde{a}' \in \mathcal{A}^{n-1}$. We say a game is symmetric if,

$$u_i(\widetilde{a}, \widetilde{a}') = u_j(\widetilde{a}, \widetilde{a}') \qquad \text{for all } i, j \in [n],\ \widetilde{a} \in \mathcal{A},\ \widetilde{a}' \in \mathcal{A}^{n-1}.$$

A Nash equilibrium $s = \{s_i\}_{i \in [n]}$ is symmetric if $s_i = s_j$ for all i, j.

Theorem. Every *finite* symmetric game has at least one symmetric Nash equilibrium.

Proof in Chapter 5 of KP.

Ch 1.2: Indifference principle

The indifference principle states that in a (mixed) NE, a player will be indifferent to any action she may choose with positive probability.

Notation. We will use the following notation for the expected utility when a player i chooses action a_i while others are following mixed strategies s_{-i}:

$$u_i(a_i, s_{-i}) \triangleq \mathbb{E}_{a_{-i} \sim s_{-i}}[u_i(a)] = \sum_{a_{-i} \in \mathcal{A}_{-i}} s_1(a_1) \dots s_{i-1}(a_{i-1}) s_{i+1}(a_{i+1}) \dots s_n(a_n) \cdot u_i(a_1, \dots, a_n)$$

Theorem (Indifference principle). *(Lemma 4.3.7 in KP).* Consider an n player game with action spaces $\mathcal{A}_1, \dots, \mathcal{A}_n$. Let $s = (s_1, \dots, s_n)$, where $s_i \in \Delta(\mathcal{A}_i)$ be a mixed strategy profile. Let $\mathcal{B}_i = \{a \in \mathcal{A}_i ; s_i(a) > 0\}$ be the set of strategies for player i that will be selected with non-zero probability. Then, s is a NE iff there exists constant $c_1, \dots, c_n$ such that

$$\text{for all, } a_i \in \mathcal{B}_i, \qquad u_i(a_i, s_{-i}) = c_i,$$
$$\text{for all, } a_i \in \mathcal{A}_i \setminus \mathcal{B}_i, \qquad u_i(a_i, s_{-i}) \leq c_i.$$

Indifference principle (cont'd)

Theorem (Indifference principle). Consider an n player game with action spaces $\mathcal{A}_1, \ldots, \mathcal{A}_n$. Let $s = (s_1, \ldots, s_n)$, where $s_i \in \Delta(\mathcal{A}_i)$ be a mixed strategy profile. Let $\mathcal{B}_i = \{a \in \mathcal{A}_i; s_i(a) > 0\}$ be the set of strategies for player i that will be selected with non-zero probability. Then, s is a NE iff there exists constant $c_1, \ldots, c_n$ such that

$$\text{for all, } a_i \in \mathcal{B}_i, \qquad u_i(a_i, s_{-i}) = c_i,$$
$$\text{for all, } a_i \in \mathcal{A}_i \setminus \mathcal{B}_i, \qquad u_i(a_i, s_{-i}) \leq c_i.$$

Intuition.

- In a mixed-strategy NE, any action that a player uses with positive probability must give the player the same expected utility. Otherwise, the player would shift probability toward the better-paying strategy, contradicting Nash equilibrium.
- All other actions should yield equal or less utility. Otherwise, the player would simply choose the other action.

Our proof will simply formalize the above intuition.

Proof of the indifference principle

Indifference principle. s is a NE iff there exists constant $c_1, \ldots, c_n$ such that

$$\forall\, a_i \in \mathcal{B}_i,\ \ u_i(a_i, s_{-i}) = c_i, \qquad \forall\, a_i \in \mathcal{A}_i \setminus \mathcal{B}_i,\ \ u_i(a_i, s_{-i}) \le c_i. \quad \text{where, } \mathcal{B}_i = \{a \in \mathcal{A}_i; s_i(a) > 0\}.$$

Proof. First, suppose that s is a NE. Let $c_i = \max_{a_i \in \mathcal{A}_i} u_i(a_i, s_{-i})$. As s_i is the best response to s_{-i}, for any $a_i' \in \mathcal{A}_i$,

$$\text{if } u_i(a_i', s_{-i}) < c_i, \quad \implies s_i(a_i') = 0.$$

Otherwise, we can shift the probability $s_i(a_i')$ from a_i' to a_i, for some action a_i with $u_i(a_i, s_{-i}) = c_i$ and increase the expected utility of user i.

The above statement, (along with its constrapositive, *i.e.*, $s_i(a_i') > 0 \implies u_i(a_i', s_{-i}) = c_i$) implies the condition.

Proof of the indifference principle (cont'd)

Indifference principle. s is a NE iff there exists constant $c_1, \ldots, c_n$ such that

$$\forall\, a_i \in \mathcal{B}_i,\;\; u_i(a_i, s_{-i}) = c_i, \qquad \forall\, a_i \in \mathcal{A}_i \setminus \mathcal{B}_i,\;\; u_i(a_i, s_{-i}) \leq c_i. \quad \text{where, } \mathcal{B}_i = \{a \in \mathcal{A}_i; s_i(a) > 0\}.$$

Proof (cont'd). Now suppose that the condition holds. Consider any agent i and an *alternative* strategy s_i' for i. We can expand $u_i(s_i', s_{-i})$ as follows:

$$\begin{aligned}
u_i(s_i', s_{-i}) &= \sum_{a \in \mathcal{A}} s_1(a_1) \times \cdots \times s_{i-1}(a_{i-1}) \times s_i'(a_i) \times s_{i+1}(a_{i+1}) \times \ldots s_n(a_n) u_i(a) \\
&= \sum_{a_i \in \mathcal{A}_i} s_i'(a_i) \sum_{a_{-i} \in \mathcal{A}_{-i}} s_1(a_1) \times \cdots \times s_{i-1}(a_{i-1}) \times s_{i+1}(a_{i+1}) \times \ldots s_n(a_n) u_i(a) \\
&= \sum_{a_i \in \mathcal{A}_i} s_i'(a_i) u_i(a_i, s_{-i}) \\
&\leq \sum_{a_i \in \mathcal{A}_i} s_i(a_i) u_i(a_i, s_{-i}). \qquad \texttt{via the given condition}
\end{aligned}$$

Hence, s_i is a best response to s_{-i} player i. As this is true for all players i, s is a NE.

Computing all NE using the indifference principle

Indifference principle. s is a NE iff there exists constant $c_1, \ldots, c_n$ such that

$$\text{for all, } a_i \in \mathcal{B}_i, \quad u_i(a_i, s_{-i}) = c_i, \qquad\qquad \text{for all, } a_i \in \mathcal{A}_i \setminus \mathcal{B}_i, \quad u_i(a_i, s_{-i}) \leq c_i.$$

The indifference principle suggests the following recipe for finding all NE in (small) games:

1. Enumerate all possible supports (which strategies each player plays with positive probability) for the Nash equilibria.
2. For each such support, equate expected payoffs of all actions in each player's support using the indifference principle.
3. Solve for the mixed strategies.

Eg1: Stag hunt

Recall the stag hunt game:

P1 \ P2	**S**	**H**
S	$(4, 4)$	$(0, 2)$
H	$(2, 0)$	$(1, 1)$

Let us first enumerate all possible supports for a NE:

1. Both players use a single action (pure NE)
2. One player uses a single action and the other uses both actions.
3. Both players use both actions.

1. First consider pure NE:

- $s_1 = s_2 = (1, 0)$, (S, S).
- $s_1 = s_2 = (0, 1)$, (H, H).

2. Next, consider NE where one agent follows a pure strategy and the other follows a mixed strategy. However, such NE does not exist. If one agent follows S, the best response for the other is to also follow S, and likewise if the agent follows H.

Eg1: Stag hunt (cont'd)

Recall the stag hunt game:

P1 \ P2	**S**	**H**
S	$(4, 4)$	$(0, 2)$
H	$(2, 0)$	$(1, 1)$

3. Finally, let us consider fully mixed NE. Suppose agents 1 and 2 follow $s_1 = (x, 1 - x)$ and $s_2 = (y, 1 - y)$ with $x, y \in (0, 1)$.

Let us first write P2's utility when following both actions,

If P2 chooses S: $u_2(s_1, S) = \mathbb{E}_{a_1 \sim s_1}[u_2(a_1, S)] = x \cdot 4 + (1 - x) \cdot 0 = 4x$,

If P2 chooses H: $u_2(s_1, H) = \mathbb{E}_{a_1 \sim s_1}[u_2(a_1, H)] = x \cdot 2 + (1 - x) \cdot 1 = 1 + x$.

By the indifference principle, P2 should receive the same expected utility regardless of either action (as both actions will be chosen with nonzero probability), *i.e.*, $4x = 1 + x \implies x = 1/3$.

By a similar argument, $y = 1/3$. Hence $s_1 = s_2 = (1/3, 2/3)$ is a mixed NE.

Eg2: Driver vs Inspector (try at home)

Recall the driver vs inspector game:

P1 \ P2	**N** (Not inspect)	**In** (Inspect)
L (Legal)	$(0, 0)$	$(0, -1)$
I (Illegal)	$(10, -10)$	$(-90, -6)$

It is straightforward to verify that there are no pure NE or NE where one player follows a pure strategy and the other follows a mixed strategy. To find fully mixed NE, suppose agents 1 and 2 follow $s_1 = (x, 1-x)$ and $s_2 = (y, 1-y)$ with $x, y \in (0, 1)$.

Applying the indifference principle for P2, *i.e.*, $u_2(s_1, N) = u_2(s_1, In)$, we have

$$x \cdot 0 + (1-x) \cdot (-10) = x \cdot (-1) + (1-x) \cdot (-6) \quad \Longleftrightarrow \quad x = 0.8$$

Applying the indifference principle for P1, *i.e.*, $u_1(L, s_2) = u_1(I, s_2)$, we have

$$y \cdot 0 + (1-y) \cdot 0 = y \cdot 10 + (1-y) \cdot (-90) \Longleftrightarrow \quad y = 0.9$$

Hence, the following is a NE: $s_1 = (0.8, 0.2)$, $s_2 = (0.9, 0.1)$.

Eg 3: Pollution game (E.g. 4.3.2 in KP)

Three firms (P1, P2, P3) are sharing a lake. they can choose to pollute it (P) or clean it (C). It costs 1 to clean, and 0 to pollute. If 1 or fewer pollute, the lake is still usable. If 2 or more pollute, the lake is unusable, and everyone incurs an additional cost of -3. The utilities can be written as follows:

If P3 chooses C:

P1 \ P2	**C**	**P**
C	$(-1,-1,-1)$	$(-1,\ 0,-1)$
P	$(\ 0,-1,-1)$	$(-3,-3,-4)$

If P3 chooses P:

P1 \ P2	**C**	**P**
C	$(-1,-1,\ 0)$	$(-4,-3,-3)$
P	$(-3,-4,-3)$	$(-3,-3,-3)$

Let us first enumerate all possible supports for a NE:

1. All 3 players use a single action (pure strategy NE).
2. Two players use a single action and the other uses both actions.
3. One player uses a single action and the other two use both actions.
4. All 3 players use fully mixed strategies.

Eg 3: Pollution game (cont'd)

1. Let us first consider all pure NE:
- All agents pollute (1 pure NE).
- One player pollutes while the other 2 clean (3 pure NE).

2. No NE where two players use a single action and the other uses both.

3. Next, let us consider NE where one player (say P3) uses a single action and the other two use both actions:
- If P3 chooses to pollute, the only NE is when the other choose to clean (already accounted for under pure NE).
- Suppose P3 chooses to clean. Suppose $s_1 = (x, 1-x)$, $s_2 = (y, 1-y)$, and $s_3 = (1, 0)$ is a NE where $x, y \in (0, 1)$.

Eg 3: Pollution game (cont'd)

If P3 chooses C:

P1 \ P2	**C**	**P**
C	$(-1, -1, -1)$	$(-1, 0, -1)$
P	$(0, -1, -1)$	$(-3, -3, -4)$

If P3 chooses P:

P1 \ P2	**C**	**P**
C	$(-1, -1, 0)$	$(-4, -3, -3)$
P	$(-3, -4, -3)$	$(-3, -3, -3)$

Let us apply the indifference principle for P2:

$$\mathbb{E}_{a_1 \sim s_1}[u_2(a_1, C, C)] = \mathbb{E}_{a_1 \sim s_1}[u_2(a_1, P, C)]$$
$$\iff \quad x \cdot (-1) + (1-x) \cdot (-1) = x \cdot 0 + (1-x) \cdot (-3) \iff \quad x = 2/3.$$

By symmetry, we also get $y = 2/3$. Hence $s_1 = (2/3, 1/3)$, $s_2 = (2/3, 1/3)$, and $s_3 = (1, 0)$ is a NE.

By repeating the same argument by assuming that P1 and P2 choose a single action, we get 3 NE in total.

Eg 3: Pollution game (cont'd)

If P3 chooses C:

P1 \ P2	**C**	**P**
C	$(-1,-1,-1)$	$(-1,\ 0,-1)$
P	$(\ 0,-1,-1)$	$(-3,-3,-4)$

If P3 chooses P:

P1 \ P2	**C**	**P**
C	$(-1,-1,\ 0)$	$(-4,-3,-3)$
P	$(-3,-4,-3)$	$(-3,-3,-3)$

4. Finally, let us consider fully mixed strategies. - Suppose $s_1 = (x, 1-x)$, $s_2 = (y, 1-y)$, and $s_3 = (z, 1-x)$ is a NE where $x, y, z \in (0,1)$.

Let us apply the indifference principle for P3:

$$\begin{aligned}
&\mathbb{E}_{a_1,a_2}[u_3(a_1, a_2, C)] = \mathbb{E}_{a_1,a_2}[u_3(a_1, a_2, P)] \\
\iff\ & xy \cdot (-1) + x(1-y) \cdot (-1) + (1-x)y \cdot (-1) + (1-x)(1-y) \cdot (-4) \\
&\qquad = xy \cdot 0 + x(1-y) \cdot (-3) + (1-x)y \cdot (-3) + (1-x)(1-y) \cdot (-3)
\end{aligned}$$

Eg 3: Pollution game (cont'd)

Solving this, we get $3x + 3y = 1 + 6xy$. Via a similar argument, we can also obtain,

$$3x + 3z = 1 + 6xz, \qquad 3y + 3z = 1 + 6yz.$$

After some algebra, we get two symmetric mixed NE:

$$x = y = z = \frac{3 - \sqrt{3}}{6}, \qquad x = y = z = \frac{3 + \sqrt{3}}{6}.$$

In total, there are $1 + 3 + 3 + 2 = 9$ NE.

Computational complexity of finding NE

- Based on the above examples, what can you say about computing NE in (large) games?
  - Clearly finding all NE is difficult.
  - Even finding one NE is PPAD-complete (optional reading on the course webpage).
- Using the indifference principle to find *a* NE can be expensive even in practice.
- Next, we will study *potential games*, where a simple greedy algorithm often finds a pure NE efficiently in practice.
- A few lectures down the line, we will introduce other equilibrium concepts, which can be efficiently approximated via techniques from machine learning (online learning).
- In the next chapter, we will introduce two player zero sum games, where NE can be efficiently approximated via similar online learning techniques.

Ch 1.3: Potential games and repeated play dynamics

Recall the following two games:

Driver vs Inspector

P1 \ P2	**N**	**In**
L	$(0, 0)$	$(0, -1)$
I	$(10, -10)$	$(-90, -6)$

Prisoner's dilemma

P1 \ P2	**B**	**L**
B	$(-3, -3)$	$(0, -5)$
L	$(-5, 0)$	$(-1, -1)$

In the games above, what happens if we were to start at some arbitrary action profile, and some player changes their action to the best response given the other's action?

- In Prisoner's dilemma, we converge to the NE (DSE).
- In Driver vs Inspector, we do not.

Repeated play dynamics

Repeated play dynamics is a greedy algorithm where players start with an arbitrary action profile. Then, on each step, *exactly one player* changes their action so as to *strictly increase* their utility. We continue in this fashion until no player can strictly improve their utility.

Algorithm 1 Repeated play dynamics

1: Initialize $a = (a_1, \ldots, a_n) \in \mathcal{A}$ arbitrarily.
2: **while** there exists $i \in [n]$ such that $a_i \notin \operatorname{argmax}_{\widetilde{a}_i \in \mathcal{A}_i} u_i(\widetilde{a}_i, a_{-i})$ **do**
3:     Set $a'_i \leftarrow$ any action $\widetilde{a}_i$ such that $u_i(\widetilde{a}_i, a_{-i}) > u_i(a_i, a_{-i})$.
4:     Set $a \leftarrow (a'_i, a_{-i})$.
5: **end while**
6: **return** a

If we replace line 3 with $a'_i \leftarrow \operatorname{argmax}_{\widetilde{a}_i \in \mathcal{A}_i} u_i(\widetilde{a}_i, a_{-i})$, then the algorithm is called best response dynamics.

Repeated play dynamics (cont'd)

- **Observation 1:** By design, if RPD/BRD terminates, it terminates at a pure NE.
- But RPD/BRD may not terminate (*e.g.*, Driver vs inspector, Rock-paper-scissors).
- We are interested in studying when BRD/RPD terminates. Why?
  - An intuitive algorithm which replicates real-world behavior of agents.
  - Terminates at a pure NE.
  - Often efficient in practice, even if worst-case complexity could be large.
- In this subchapter, we will study necessary and sufficient conditions for termination of RPD.

Potential games

Definition (Potential function, Potential game). Consider an n player game. Suppose there exists a function $\psi : \mathcal{A}_1 \times \cdots \times \mathcal{A}_n \to \mathbb{R}$ such that for all $i \in [n]$ and $a_{-i} \in \mathcal{A}_{-i}$, we have that $\psi(a_i, a_{-i}) - u_i(a_i, a_{-i})$ does not depend on a_i. Then, this game is called a potential game and ψ is called the potential function of this game.

Remark. This condition can be equivalently written as follows: for all $i \in [n]$, $a_i, a_i' \in \mathcal{A}_i$, and $a_{-i} \in \mathcal{A}_{-i}$, we have

$$\psi(a_i, a_{-i}) - u_i(a_i, a_{-i}) = \psi(a_i', a_{-i}) - u_i(a_i', a_{-i}),$$

or equivalently,

$$\psi(a_i, a_{-i}) - \psi(a_i', a_{-i}) = u_i(a_i, a_{-i}) - u_i(a_i', a_{-i}).$$

Example: Congestion game (E.g. 4.4.1 in KP)

There are n players and a road network (graph) with edge set E. Player i wishes to drive from point s_i to t_i.

- Player i's action space $\mathcal{A}_i$ consists of all paths from s_i to t_i.

- Let $a_i \in \mathcal{A}_i$ be i's action and $a = \{a_i\}_{i=1}^n$ be a given action profile.

- Let $N_e(a) = |\{i \in [n]; e \in a\}|$ be the number of drivers using edge e.

- The latency to cross edge e when there are N drivers on the edge is $\ell_e(N)$, where $\ell_e : \mathbb{N} \to \mathbb{R}_+$ is an edge-specific nondecreasing function.

- Driver i's total latency (*i.e.*, cost or negative utility) under an action profile a is,

$$c_i(a) \stackrel{\Delta}{=} -u_i(a) = \sum_{e \in a_i} \ell_e(N_e(a)).$$

Example: Congestion game (cont'd)

Claim. This is a potential game with potential function

$$\psi(a) = -\sum_{e\in E}\sum_{k=1}^{N_e(a)} \ell_e(k).$$

Proof. Suppose player i switches from a_i to a_i'. Then,

$$\begin{aligned} u_i(a_i, a_{-i}) - u_i(a_i', a_{-i}) &= c_i(a_i', a_{-i}) - c_i(a_i, a_{-i}) \\ &= \sum_{e\in a_i'\setminus a_i} \ell_e(N_e(a)+1) - \sum_{e\in a_i\setminus a_i'} \ell_e(N_e(a)) \end{aligned}$$

Example: Congestion game (cont'd)

Claim. This is a potential game with potential function $\psi(a) = -\sum_{e\in E}\sum_{k=1}^{N_e(a)} \ell_e(k)$.

Now, let us consider $\psi(a_i, a_{-i}) - \psi(a_i', a_{-i})$. We have,

$$\psi(a_i, a_{-i}) - \psi(a_i', a_{-i}) = \sum_{e\in E}\sum_{k=1}^{N_e(a_i', a_{-i})} \ell_e(k) - \sum_{e\in E}\sum_{k=1}^{N_e(a)} \ell_e(k)$$

$$= \sum_{e\in a_i'\setminus a_i}\left(\sum_{k=1}^{N_e(a)+1} \ell_e(k) - \sum_{k=1}^{N_e(a)} \ell_e(k)\right) + \sum_{e\in a_i\setminus a_i'}\left(\sum_{k=1}^{N_e(a)-1} \ell_e(k) - \sum_{k=1}^{N_e(a)} \ell_e(k)\right)$$

$$= \sum_{e\in a_i'\setminus a_i} \ell_e(N_e(a)+1) \; - \sum_{e\in a_i\setminus a_i'} \ell_e(N_e(a))$$

$$= u_i(a_i, a_{-i}) - u_i(a_i', a_{-i})$$

Hence, this is a potential game.

RPD in potential games

Algorithm 2 Repeated play dynamics

1: Initialize $a = (a_1, \ldots, a_n) \in \mathcal{A}$ arbitrarily.
2: **while** there exists $i \in [n]$ such that $a_i \notin \operatorname{argmax}_{\widetilde{a}_i \in \mathcal{A}_i} u_i(\widetilde{a}_i, a_{-i})$ **do**
3:     Set $a'_i \leftarrow$ any action $\widetilde{a}_i$ such that $u_i(\widetilde{a}_i, a_{-i}) > u_i(a_i, a_{-i})$.
4:     Set $a \leftarrow (a'_i, a_{-i})$.
5: **end while**
6: **return** a

Theorem (RPD in potential games). Repeated play dynamics terminates in a finite potential game.

Corollary. Every finite potential game has at least one pure NE.

Proof. Follows from the above theorem and observation 1 (if RPD terminates, it terminates in a pure NE.)

RPD in potential games

Proof of RPD theorem. In RPD, the utility of a player switching their action from a_i to a_i' strictly increases. Since

$$\psi(a_i', a_{-i}) - \psi(a_i, a_{-i}) = u_i(a_i', a_{-i}) - u_i(a_i, a_{-i}),$$

the value of the potential function also increases.

As the game is finite, ψ can take on only finitely many values. Hence, RPD terminates. □

Convention. If we consider agent costs (disutilities) instead of utilities, we will typically write $c_i(a) = -u_i(a)$ and $\phi(a) = -\psi(a)$. In this case, BRD can be interpreted as an algorithm which "minimizes the potential" to reach an equilibrium.

Depending on whether we use ψ, u_i or ϕ, c_i, we will say RPD increases or decreases the potential function.

Properties of potential functions

The following lemma is often useful in finding the potential function of a game.

Theorem (Separability of potential function). Consider any n-player normal form game. Let $\mathcal{A}_i$ be the action space of player i, let $\mathcal{A} = \bigtimes_{j=1}^{n} \mathcal{A}_j$ be the space of action profiles, and let $\mathcal{A}_{-i} = \bigtimes_{j \neq i} \mathcal{A}_j$ the space of action profiles of all except i. Then, the following two statements are true:

1. Let $\psi : \mathcal{A} \to \mathbb{R}$ be a function defined on the action profile of n players. Suppose that, for all $i \in [n]$, there exists $\psi_{-i} : \mathcal{A}_{-i} \to \mathbb{R}$ such that $\psi(a) = \psi_{-i}(a_{-i}) + u_i(a)$ for all $a \in \mathcal{A}$. Then ψ is a potential function for this game.

2. Suppose this game has a potential function $\psi : \mathcal{A} \to \mathbb{R}$ defined on the action profile of n players. Then, for all $i \in [n]$, there exist functions $\psi_{-i} : \mathcal{A}_{-i} \to \mathbb{R}$ such that $\psi(a) = \psi_{-i}(a_{-i}) + u_i(a)$ for all $a \in \mathcal{A}$.

Properties of potential functions (cont'd)

Proof. First, the given property implies that $\psi(a_i, a_{-i}) - u_i(a_i, a_{-i})$ does not depend on a_i, which is precisely the definition of a potential function.

We can write this out more explicitly as follows:

$$\begin{aligned}\psi(a_i, a_{-i}) - \psi(a_i', a_{-i}) &= \psi_{-i}(a_{-i}) + u_i(a_i, a_{-i}) - \big(\psi_{-i}(a_{-i}) + u_i(a_i', a_{-i})\big) \\ &= u_i(a_i, a_{-i}) - u_i(a_i', a_{-i}).\end{aligned}$$

Next, using the definition of a potential function, we have that $\psi(a_i, a_{-i}) - u_i(a_i, a_{-i})$ does not depend on a_i for all a_{-i}. Hence, we can view this quantity as a function of a_{-i} and write, $\psi(a) - u_i(a) =: \psi_{-i}(a_{-i})$. This proves the second statement.

Properties of potential functions (cont'd)

Recall: In potential games, we can write $\psi(a) = \psi_{-i}(a_{-i}) + u_i(a)$.

Another perspective of best response dynamics:

- In BRD, each player chooses $a'_i \leftarrow \mathrm{argmax}_{\widetilde{a}_i \in \mathcal{A}_i} u_i(\widetilde{a}_i, a_{-i})$ (or equivalently $a'_i \leftarrow \mathrm{argmin}_{\widetilde{a}_i \in \mathcal{A}_i} c_i(\widetilde{a}_i, a_{-i})$).

- But, $\psi(\widetilde{a}_i, a_{-i}) = \psi(a_{-i}) + u_i(\widetilde{a}_i, a_{-i})$. Hence, in effect the player is choosing

$$a'_i = \underset{\widetilde{a}_i \in \mathcal{A}_i}{\mathrm{argmax}}\, \psi(\widetilde{a}_i, a_{-i}), \qquad \left(\text{or} \quad a'_i = \underset{\widetilde{a}_i \in \mathcal{A}_i}{\mathrm{argmin}}\, \phi(\widetilde{a}_i, a_{-i}).\right)$$

- Hence, we can think of BRD as coordinate ascent (descent)[1] on the potential function. This perspective is useful when applying BRD beyond finite games.

[1] Coordinate descent is an iterative method for multivariate optimization that minimizes an objective function by repeatedly optimizing one variable (or coordinate) at a time while holding all others fixed.

Properties of potential functions (cont'd)

Theorem (Uniqueness of potential function). The potential function is unique up to an additive constant. That is, if ψ and $\widetilde{\psi}$ are potential functions of this game, then there exists some constant C such that, for all action profiles a, we have $\psi(a) = \widetilde{\psi}(a) + C$.

Proof. You will prove this in the homework.

Example 1: Congestion game revisited

Recall: In potential games, we can write $\phi(a) = \phi_{-i}(a_{-i}) + c_i(a)$.

Recall that in congestion game, player i has latency c_i and that this is a potential game with potential function ϕ, where

$$c_i(a) = -u_i(a) = \sum_{e \in a_i} \ell_e(N_e(a)). \qquad \phi(a) = -\psi(a) = \sum_{e \in E} \sum_{k=1}^{N_e(a)} \ell_e(k).$$

We can write the potential function as,

$$\phi(a) = \underbrace{\sum_{e \in E} \sum_{k=1}^{N_e(a_{-i})} \ell_e(k)}_{=\phi_{-i}(a_{-i})} + \underbrace{\sum_{e \in E} \mathbb{1}(e \in a_i)\ell_e(N_e(a))}_{=c_i(a)}$$

Example 2: Consensus game

Consider n players in an undirected graph $G = (V, E)$ where each player is a vertex.

- Let $N(i) = \{j \in V; (i, j) \in E\}$ denote the neighborhood of player i.
- Each player chooses an action $a_i \in \{0, 1\}$.
- The cost of player i is the number of disagreements among her neighbors, *i.e.*,

$$c_i(a) = \sum_{j \in N(i)} |a_i - a_j|.$$

Question: Can you guess a potential function for this game?
Hint: Use the lemma we just proved, *i.e.*, $\phi(a) = \phi_{-i}(a_{-i}) + c_i(a)$.

Example 2: Consensus game (cont'd)

One idea is to use the total number of disagreements, as we can write it as

$$\phi(a) = \underbrace{\#\text{ disgreements without } i}_{=\phi_{-i}(a_{-i})} + \underbrace{\#\text{ disgreements with } i}_{=c_i(a_i)}$$

Formally, we can try $\phi(a) = \sum_{(i,j)\in E} |a_i - a_j|$. As it can be written as,

$$\phi(a) = \underbrace{\sum_{(j,k)\in E, j,k\neq i} |a_j - a_k|}_{=\phi_{-i}(a_{-i})} + \underbrace{\sum_{j\in N(i)} |a_i - a_j|}_{=c_i(a_i)}.$$

You can also verify that ϕ is a potential function directly using the definition (try at home).

Generalized potential function

We saw that the existence of a potential function is sufficient for RPD to converge in a finite game. Next, we will show that a (slight) generalization of this concept provides a necessary and sufficient condition for convergence of RPD.

Definition (Generalized potential function, Ordinal potential function). Consider a finite n-player game with utilities $u_i : \mathcal{A} \to \mathbb{R}$. Say that there exists a function $\psi : \mathcal{A} \to \mathbb{R}$ such that for all $i \in [n]$, $a_i, a_i' \in \mathcal{A}_i$, and $a_{-i} \in \mathcal{A}_{-i}$ we have,

$$\mathrm{sign}(\psi(a_i, a_{-i}) - \psi(a_i', a_{-i})) = \mathrm{sign}(u_i(a_i, a_{-i}) - \psi(a_i', a_{-i})).$$

Intuition. If i increases her utility in RPD, it also increases the potential function, so local improvements can lead to a NE.

Example: Load balancing game

The following is an example of a game which has a generalized potential function, but not a potential function.

There are n players sharing m computers. Player i has a job of size $w_i > 0$ and must choose which of the m machines to run it on. Hence $a_i = [m]$ for all i. The load on machine k under an action profile a is $\ell_k(a) = \sum_i \mathbb{1}(a_i = k) w_i$. The cost of player i is the load of the machine she is running her job on, *i.e.*, $c_i(a) = \ell_{a_i}(a)$.

We will show that $\phi(a) = \frac{1}{2}\sum_{k=1}^{m} \ell_k^2(a)$ is a generalized potential function. Supose $a_i = k$ and $a_i' = k'$. Let

$$x \triangleq \ell_k(a_{-i}) = \sum_{j \neq i} \mathbb{1}(a_j = k) w_j, \qquad x' \triangleq \ell_{k'}(a_{-i}) = \sum_{j \neq i} \mathbb{1}(a_j = k') w_j,$$

denote the loads on k, k' without i. Then, we can write

$$c_i(a_i, a_{-i}) - c_i(a_i', a_{-i}) = \ell_k(a_i, a_{-i}) - \ell_{k'}(a_i', a_{-i}) = (x + w_i) - (x' + w_i) = x - x'.$$

Example: Load balancing game (cont'd)

Let us now consider, $\phi(a_i, a_{-i}) - \phi(a'_i, a_{-i})$. We have,

$$\begin{aligned}
\phi(a_i, a_{-i}) - \phi(a'_i, a_{-i}) &= \frac{1}{2}\sum_j \ell_j^2(a_i, a_{-i}) - \frac{1}{2}\sum_j \ell_j^2(a'_i, a_{-i}) \\
&= \frac{1}{2}\left(\ell_k^2(a_i, a_{-i}) + \ell_{k'}^2(a_i, a_{-i}) - \ell_k^2(a'_i, a_{-i}) - \ell_{k'}^2(a'_i, a_{-i})\right) \\
&= \frac{1}{2}\left((x + w_i)^2 + x'^2 - x^2 - (x' + w_i)^2\right) \\
&= w_i(x - x')
\end{aligned}$$

Hence, ϕ is a generalized potential function.

Necessary and sufficient conditions for convergence of RPD

Theorem (Remark 4.4.8 in KP without proof). RPD converges in a finite game *iff* the game has a generalized potential function.

Proof. The if condition is straightforward and identical to the proof for potential function.

To prove the only if direction, let us construct a graph $G = (V, E)$ as follows:

- Let the vertices be each action profile, $V = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$.

- For each pair of vertices $a = (a_1, \ldots, a_i, \ldots, a_n)$ and $a' = (a_1, \ldots, a_i', \ldots, a_n)$ differing in only one agent's action, let us add an edge from a to a' if $u_i(a') > u_i(a)$ and an edge from a' to a if $u_i(a) > u_i(a')$. In particular, no edge if $u_i(a) = u_i(a')$.

In the above construction, an edge exists from a to a' only if a' can be reached from a via one iteration of RPD.

Necessary and sufficient conditions for convergence of RPD

RPD can be viewed as traversing this graph until you reach a sink vertex. Hence, RPD converges from any initial action profile *iff* there are no cycles in this graph.

To construct a generalized potential function, we need to show that there exists a function ψ which satisfies $\mathrm{sign}(\psi(a_i, a_{-i}) - \psi(a_i', a_{-i})) = \mathrm{sign}(u_i(a_i, a_{-i}) - \psi(a_i', a_{-i}))$ for all a_i, a_i', a_{-i}. It is sufficient to show that for any (a_i', a_i) which is a child of (a_i, a_{-i}), we have $\psi(a_i, a_{-i}) < \psi(a_i', a_{-i})$.

Let us construct a function ψ in this DAG as follows: Let $\phi(a) = -\psi(a)$ be the length of the longest path from a to *any* sink vertex.

Here, for any a, a' such that a' is a child of a, we have $\phi(a) \geq \phi(a') + 1 > \phi(a)$. Hence, $\psi(a) < \psi(a')$. □

Question: : Why did we define it to be the *longest* path?

Ch 1.4: Price of Anarchy/Stability

The *price of anarchy* (PoA) measures the degradation in efficiency due to selfish behavior, comparing the best achievable welfare (or cost) relative to the *worst* NE.

The *price of stability* (PoS) similarly compares the best achievable welfare (or cost) relative to the *best* NE.

Example 1: Stag hunt

Define the welfare under an action profile a and strategy profile s as follows:

$$W(a) = \sum_{i=1}^{n} u_i(a) = u_1(a) + u_2(a), \qquad W(s) = \mathbb{E}_{a \sim s}\left[W(a)\right].$$

Recall the stag hunt game.

P1 \ P2	**S**	**H**
S	$(4, 4)$	$(0, 2)$
H	$(2, 0)$	$(1, 1)$

Recall that this game has two pure NE, (S, S) and (H, H), and one mixed NE $s = ((1/3, 2/3), (1/3, 2/3))$.

We have, $W(S, S) = 8$, $W(H, H) = 2$, and $W(s) = \frac{8}{3}$. Therefore,

$$\mathrm{PoA} = \frac{\text{maximum welfare}}{\text{welfare at worst NE}} = \frac{\max_{s \in \mathcal{S}} W(s)}{\min_{s; s \text{ is a NE}} W(s)} = \frac{8}{2} = 4.$$

$$\mathrm{PoS} = \frac{\text{maximum welfare}}{\text{welfare at best NE}} = \frac{\max_{s \in \mathcal{S}} W(s)}{\max_{s; s \text{ is a NE}} W(s)} = \frac{8}{8} = 1.$$

Example 2: Prisoner's Dilemma

Define the total cost under an action profile a and strategy profile s as follows:

$$C(a) = \sum_{i=1}^{n} c_i(a) = c_1(a) + c_2(a), \qquad C(s) = \mathbb{E}_{a\sim s}\left[C(a)\right].$$

Recall the prisoner's dilemma with costs:

P1 \ P2	**B**	**L**
B	$(3,3)$	$(0,5)$
L	$(5,0)$	$(1,1)$

This game has just one NE (B, B).

We have, $C(B,B) = 6$, and $C(L,L) = 2$. Therefore,

$$\mathrm{PoA} = \frac{\text{cost at worst NE}}{\text{minimum cost}} = \frac{\max_{s;s \text{ is a NE}} C(s)}{\min_{s\in\mathcal{S}} C(s)} = \frac{6}{2} = 3.$$

$$\mathrm{PoS} = \frac{\text{cost at best NE}}{\text{minimum cost}} = \frac{\min_{s;s \text{ is a NE}} C(s)}{\min_{s\in\mathcal{S}} C(s)} = \frac{6}{2} = 3.$$

Example 3: Market sharing game (Sec 8.3 in KP)

There are n sports teams and K cities, with $K \geq n$. The population of city $k \in [K]$ is v_k. We will assume, w.l.o.g, that $v_1 \geq v_2 \geq \cdots \geq v_K$.

- Each team should choose a city, *i.e.*, the action space is $\mathcal{A}_i = [K]$ for each $i \in [n]$.

- Let $N_k(a) = \{i \in [n]; a_i = k\}$ be the number of teams in city k under an action profile a.

- If N' teams select city k, the utility for each is v_k/N'. Hence, $u_i(a) = v_{a_i}/N_{a_i}(a)$.

- Let $J(a) = \{k \in [K]; \exists i \in [n] \text{ s.t } a_i = k\}$ be the set of cities selected under an action profile a.

- For any $J' \subset [K]$, let $V(J) = \sum_{j \in J'} v_j$ be the total population of the cities in J.

- The welfare under an action profile a is

$$W(a) = \sum_{i=1}^{n} u_i(a) = \sum_{j \in J(a)} v_j(a) = V(J(a)).$$

- The optimal welfare is $W^{\text{OPT}} = V(J^{\text{OPT}})$, where $J^{\text{OPT}} = \{1, 2, \ldots, n\}$.

Example 3: Market sharing game (cont'd)

The market sharing game is a potential game (you will prove this in HW2). Therefore, it has at least one pure Nash equilibrium (PNE).

Claim. For the market sharing game,

$$\mathrm{PoA}_{\mathsf{PNE}} = \frac{W^{\mathrm{OPT}}}{\min_{a;a \text{ is a pure NE}} W(J(a))} \leq 2.$$

Proof. Let a, a' be two action profiles in the market sharing game and let the cities selected by $J(a), J(a')$ respectively. If $a'_i \in J(a'_i, a_{-i}) \setminus J(a_i, a_{-i})$, *i.e.*, if a new city was added by i switching to a'_i, we have $u_i(a'_i, a_{-i}) = v_{a'_i}$. Otherwise

$u_i(a'_i, a_{-i}) = \frac{v_{a'_i}}{N_{a'_i}(a'_i, a_{-i})} \geq 0$. Therefore,

$$\sum_{i=1}^{m} u_i(a'_i, a_{-i}) \geq V(J(a') \setminus J(a)) = V(J(a')) - V(J(a') \cap J(a)) \geq V(J(a')) - V(J(a)).$$

The first inequality follows from the observation that we are adding every city in $J(a') \setminus J(a)$ at least once.

Example 3: Market sharing game (cont'd)

Let $a^\star$ be any pure NE and let a^{OPT} be an optimal action profile in which $J(a^{\mathrm{OPT}}) = \{1, \ldots, n\}$. Then,

$$V(J(a^\star)) = \sum_{i=1}^{n} u_i(a^\star) \overset{(a)}{\geq} \sum_{i=1}^{n} u_i(a_i^{\mathrm{OPT}}, a_{-i}^\star) \overset{(b)}{\geq} V(J(a^{\mathrm{OPT}})) - V(J(a^\star)).$$

Here, (a) uses the fact that $a^\star$ is a NE, while (b) uses the previously derived result. Therefore, $W(a^{\mathrm{OPT}}) = V(J(a^{\mathrm{OPT}})) \leq 2V(J(a^\star)) = 2W(a^\star)$. □

Remark. This can be improved to $2 - 1/n$ (in HW3).

Example 4: Fair network formation games (§8.2 in KP)

A set of n players wish to construct a road network across cities. We are given a weighted directed graph $G = (V, E, b)$, where V represents the cities, E represents links that can be constructed between cities, and $b = \{b\}_{e \in E}$ represents the costs, with b_e being the cost to construct link e. Each player i wishes to construct a path between city $s_i \in V$ and city $t_i \in V$. If N players decide to construct a link along e, they share the cost b_e. Player i's action space $\mathcal{A}_i$ is all paths from s_i to t_i.

- Let $N_e(a) = |\{i \in [n]; e \in a_i\}|$ denote the number of players who have chosen a path which includes edge i under action profile a.

- The cost of player i is

$$c_i(a) = \sum_{e \in a_i} \frac{b_e}{N_e(a)}.$$

- Letting $E(a)$ denote the edges constructed under a, the total cost is,

$$C(a) = \sum_{i=1}^{n} c_i(a) = \sum_{e \in E(a)} b_e.$$

Example 4: Fair network formation games (cont'd)

This is a potential game with potential function, (Try at home)

$$\phi(a) = \sum_{e \in E} \sum_{j=1}^{N_e(a)} \frac{b_e}{j}.$$

Therefore, it has at least one pure Nash equilibrium (PNE).

Let us define the price of anarchy and stability as follows:

$$\mathrm{PoA}_{\mathrm{PNE}} = \frac{\max_{a;a \text{ is a pure NE}} C(a)}{\min_a C(a)} \qquad \mathrm{PoS}_{\mathrm{PNE}} = \frac{\min_{a;a \text{ is a pure NE}} C(a)}{\min_a C(a)}$$

Question: How large can the price on anarchy be on a given graph?

Example 4: Fair network formation games (cont'd)

Consider the following graph, where all agents have the same source s_i and sink t_i, and $\varepsilon > 0$. What is the PoA and PoS on this graph?

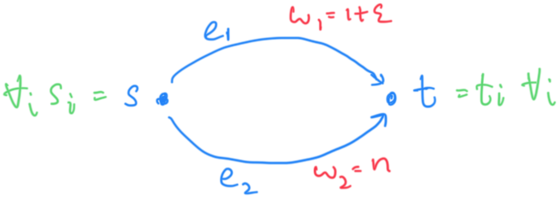

$$\mathrm{PoA}_{\mathrm{PNE}} = \frac{n}{1+\varepsilon} \qquad \mathrm{PoS}_{\mathrm{PNE}} = 1$$

Question: In this example, the PoA is large, but the PoS is small. Is it possible to bound the price of stability for an arbitrary network?

Example 4: Fair network formation games (cont'd)

Claim. For any fair network formation game with n players, $\mathrm{PoS} \leq H_n$.

Proof. First note that

$$\underbrace{\sum_{e \in E(a)} b_e}_{=C(a)} \overset{(a)}{\leq} \underbrace{\sum_{e \in E(a)} b_e \sum_{j=1}^{N_e(a)} \frac{1}{j}}_{=\phi(a)} = \sum_{e \in E(a)} b_e H_{N_e(a)} \overset{(b)}{\leq} H_n \underbrace{\sum_{e \in E(a)} b_e}_{=C(a)}$$

Let $a^{\mathrm{OPT}} = \mathrm{argmin}_a C(a)$ be the cost-optimal action profile. Suppose we execute RPD starting at a^{OPT}. We will terminate at a NE $a^\star$ where $\phi(a^\star) \leq \phi(a^{\mathrm{OPT}})$. Therefore,

$$C(a^\star) \leq \phi(a^\star) \leq \phi(a^{\mathrm{OPT}}) \leq H_n C(a^{\mathrm{OPT}}).$$

Above, the first inequality is by (a) and the last is by (b). Therefore, there exists a pure NE $a^\star$ (specifically, the one obtained by running RPD starting from a^{OPT}) such that $\frac{c(a^\star)}{c(a^{\mathrm{OPT}})} \leq H_n$. $\square$

Example 4: Fair network formation games (cont'd)

This bound is tight, in that, $\forall \varepsilon > 0$, there exists a network with $\mathrm{PoS} = \frac{H_n}{1+\varepsilon}$.

Consider the following network:

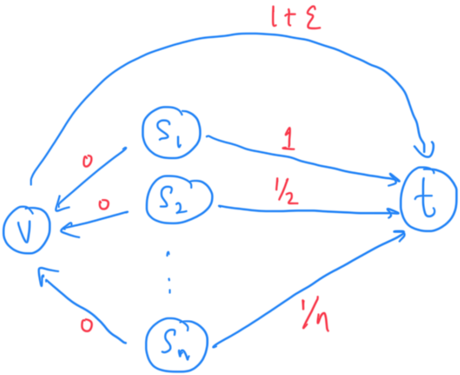

- $C(a^{\mathrm{OPT}}) = 1 + \varepsilon$, where all agents choose $s_i \to v \to t$.
- Here, each agent incurs cost $\frac{1+\varepsilon}{n}$.
- But from this action profile, player n can benefit by switching to the edge $s_n \to t$.
- Then, player $n-1$ can benefit by switching to $s_{n-1} \to t$.

...

Here, $\mathrm{PoS} = \frac{H_n}{1+\varepsilon}$.

(Non-atomic) Selfish Routing

Consider the following road network, where a large number of drivers wish to go from A to B. They can choose to go via $A \to C \to B$ or $A \to D \to B$. Each driver will choose a path to minimize their own latency.

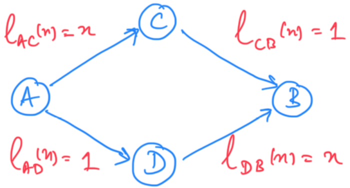

- The latency on a link if a fraction x of the drivers are on that link is,

$$\ell_{AC}(x) = \ell_{DB}(x) = x, \qquad \ell_{AD}(x) = \ell_{CB}(x) = 1.$$

- At equilibrium, half of the drivers will choose $A \to C \to B$ and the other half will choose $A \to D \to B$.

- The average latency at equilibrium is $3/2$.

- This is also the optimal flow as it minimizes the average latency.

(Non-atomic) Selfish Routing (cont'd)

Braess's paradox. Suppose we add a zero latency path from $C \to D$.

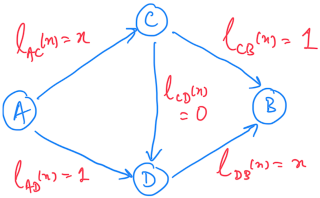

The optimal flow does not change, but this is not an equilibrium anymore!

At equilibrium, all drivers will choose $A \to C \to D \to B$.
The average latency at this equilibrium is 2.

$$\mathrm{PoA} = \frac{\text{latency at worst equilibrium}}{\text{minimum average latency}} = \frac{4}{3}.$$

(Non-atomic) Selfish Routing (cont'd)

Theorem ((Theorem 8.18 in KP)). In any (non-atomic) selfish routing network with affine latency functions, *i.e.*, $\ell_e(x) = a_e x + b_e$, we have $\mathrm{PoA} \leq 4/3$.

Ch 1.5: Correlated and Coarse-correlated equilibria

Example: Coordination game (a.k.a battle of the sexes). Two friends should choose between attending a Ballet (B) or a Lecture (L). Both players wish to spend time with each other, but disagree on which event they prefer.

P1 \ P2	**B**	**L**
B	$(10, 7)$	$(0, 0)$
L	$(0, 0)$	$(7, 10)$

This game has 3 NE:

1) (B, B), $u_1 = 10, u_2 = 7$.

2) (L, L), $u_1 = 7, u_2 = 10$.

3) $((10/17, 7/17), (7/17, 10/17))$, $u_1 = 70/17, u_2 = 70/17$.

But these three NE are “unsatisfying”:

- The two pure NE are welfare-optimal ($u_1 + u_2 = 17$), but not fair (one user has lower utility).

- The mixed NE is fairer, but has lower welfare $u_1 + u_2 < 17$.

Coordination game (cont'd)

Can we consider other solutions concepts, besides NE?

P1 \ P2	**B**	**L**
B	$(10, 7)$	$(0, 0)$
L	$(0, 0)$	$(7, 10)$

Consider the following suggestion:

- P1 and P2 flip a coin. If 'Heads', they both go to the ballet, and if 'Tails', they both go to the lecture.

- The expected welfare under this scheme is 17, ane expected utility for each player is 8.5.

- This is an example of a "correlated equilibrium".

Correlated equilibrium

Recall the following definitions: There are n players in a normal form game. The action space of player i is $\mathcal{A}_i$. A player's (mixed) strategy is an element of $\Delta(\mathcal{A}_i)$.
Let $\mathcal{A} = \bigtimes_{i=1}^{n} \mathcal{A}_i$. We can write the utility of player i as $u_i(a) = u_i(a_i, a_{-i})$.

Before we define CEs, let us first recall the definition of a NE.

Definition (Nash equilibrium). A set of n distributions (mixed strategies) $s_1, \ldots, s_n$, where $s_i \in \Delta(\mathcal{A}_i)$ for all $i \in [n]$ is a Nash equilibrium if, for all $i \in [n]$, and any $a'_i \in \mathcal{A}_i$, we have

$$\mathbb{E}_{a\sim s}\left[u_i(a)\right] \geq \mathbb{E}_{a\sim s}\left[u_i(a'_i, a_{-i})\right].$$

Here, $s = s_1 \times \cdots \times s_n$ is the product distribution.

Correlated equilibrium

Recall the following definitions: There are n players in a normal form game. The action space of player i is $\mathcal{A}_i$. Let $\mathcal{A} = \bigtimes_{i=1}^{n} \mathcal{A}_i$. We can write the utility of player i as $u_i(a) = u_i(a_i, a_{-i})$.

Definition (Correlated equilibrium). A joint distribution $s \in \Delta(\mathcal{A})$ is a correlated equilibirum if, for all $i \in [n]$, and any $a_i, a_i' \in \mathcal{A}_i$, we have

$$\mathbb{E}_{a \sim s}\left[u_i(a)\middle|a_i\right] \geq \mathbb{E}_{a \sim s}\left[u_i(a_i', a_{-i})\middle|a_i\right].$$

Interpretation of a correlated equilibrium.

- The distrbution s is known ahead of time to the players.

- A trusted third party draws $a \sim s$ and reveals a_i and only a_i to player i. Importantly, she does not reveal other's actions to i.

- Player i can follow a_i or choose any other action a_i' (or any mixed strategy s_i').

- If s is a CE, then following a_i maximizes player i's utility, provided that others are following a_{-i}.

Example 1: Coordination game

Recall the coordination game (left). Suppose the friends agree to flip a fair coin. If it is 'Heads', they both go to the ballet, and if 'Tails', they both go to the lecture. This results in the following joint distribution s (right).

Utilities

P1 \ P2	**B**	**L**
B	$(10, 7)$	$(0, 0)$
L	$(0, 0)$	$(7, 10)$

Joint distribution s

P1 \ P2	**B**	**L**
B	0.5	0
L	0	0.5

Which of the following statements are true?

1. s is NE.
2. s is CE.

Example 1: Coordination game (cont'd)

1. s **is not a NE.**

2. s **is a CE.**

- Suppose we reveal B to P1. P1 knows that (B, B) was chosen and by following B, her utility is 10, whereas if she switches to L, her utility is 0 (given that P2 is following B).

- Suppose we reveal L to P1. P1 knows that (L, L) was chosen, and by following L, her utility is 7, whereas if she switches to B, her utility is 0.

Example 2: Traffic lights

Two drivers, P1 and P2 arrive at an intersection. They should decide whether to go (G), or stop (S) and let the other pass. Consider the following joint distribution s. Show that s is a correlated equilibrium. (Try at home)

Utilities

P1 \ P2	**S**	**G**
S	$(-1,-1)$	$(-1,+1)$
G	$(+1,-1)$	$(-5,-5)$

Joint distribution s

P1 \ P2	**S**	**G**
S	0	0.5
G	0.5	0

Example 2: Traffic lights (cont'd)

Consider instead the following joint distribution, which may represent a faulty traffic light. We will show that this is also a CE.

Utilities

P1 \ P2	**S**	**G**
S	$(-1,-1)$	$(-1,+1)$
G	$(+1,-1)$	$(-5,-5)$

Joint distribution s

P1 \ P2	**S**	**G**
S	0.2	0.4
G	0.4	0

- Suppose P1 was told G. Then, P1 knows P2 will stop, so P1 will go.
- Suppose P1 was told S. The expected utilities of choosing G and S are:

$$\mathbb{E}_{a\sim s}\left[u_1(G,a_2)\middle|S\right] = \frac{0.2}{0.2+0.4}\times 1 + \frac{0.4}{0.2+0.4}\times(-5) = -3.$$

$$\mathbb{E}_{a\sim s}\left[u_1(S,a_2)\middle|S\right] = -1.$$

Therefore, s is a CE.

Coarse correlated equilibrium

Recall the following definitions: There are n players in a normal form game. The action space of player i is $\mathcal{A}_i$. Let $\mathcal{A} = \bigtimes_{i=1}^{n} \mathcal{A}_i$. We can write the utility of player i as $u_i(a) = u_i(a_i, a_{-i})$.

Definition (Coarse correlated equilibrium). A joint distribution $s \in \Delta(\mathcal{A})$ is a coarse correlated equilibirum if, for all $i \in [n]$, and any $a_i' \in \mathcal{A}_i$, we have

$$\mathbb{E}_{a \sim s}\left[u_i(a)\right] \geq \mathbb{E}_{a \sim s}\left[u_i(a_i', a_{-i})\right].$$

Interpretation of a coarse correlated equilibrium.

- The distrbution s is known ahead of time to the players.

- Each player i can choose their own alternative action a_i' (or a mixed strategy s_i'), or agree to a contract where a trusted third party will draw $a \sim s$ and play a_i on behalf of player i.

- Unlike a correlated equilibrium, the player cannot change her action after her action is revealed to her.

Quiz: Draw a Venn diagram illustrating DSE, NE, CE, and CCE

Example 3

Consider the following game.

Utilities

P1 \ P2	**A**	**B**	**C**
A	$(0, 0)$	$(0, 0)$	$(0, 0)$
B	$(0, 0)$	$(-100, -100)$	$(10, -1)$
C	$(0, 0)$	$(-1, 10)$	$(-1, -1)$

Joint distribution s

P1 \ P2	**A**	**B**	**C**
A	0	0	0
B	0	0	0.5
C	0	0.5	0

Which of the following statements are true?

1. s is NE.
2. s is CE.
3. s is CCE.

Example 3 (cont'd)

1. s **is not a NE.**

2. s **is not a CE.** Suppose P1 was told C. Then, it is better to deviate to A, as her utility improves from -10 to 0.

3. s **is a CCE.** Suppose P2 has agreed to the contract. P1's expected utility under s is $0.5 \times 10 + 0.5 \times (-1) = 4.5$. Let us compute P1's expected utility for independently choosing any action:

$$\mathbb{E}_{a \sim s}\left[u_1(A, a_2)\right] = 0$$
$$\mathbb{E}_{a \sim s}\left[u_1(B, a_2)\right] = 0.5 \times (-100) + 0.5 \times 10 = -45.$$
$$\mathbb{E}_{a \sim s}\left[u_1(C, a_2)\right] = -10$$

As all three are lower than 4.5, this is a CCE.

Why are we interested in CE and CCE?

- Often, a social planner can provide a signal to induce socially desirable outcomes (*e.g.*, traffic lights).
- CE or CCE may exist, when NE may not.
- Easier to compute than NE via linear programming (chapter 3) or online (machine) learning (chapter 5).
- No-regret learning converges to CE/CCE (also in chapter 5). Hence, CE/CCE better explain long-run behavior of adaptive agents, than NE.