Learning and Optimization in Multiagent Decision-Making Systems
Lecture Notes: Static Noncooperative Games
Instructor: Rasoul Etesami

In this part, we consider static nooncooperative games. However, before we study the general case, we first consider a special case of the 2-player static noncooperative games known as zero-sum matrix games.

Zero-Sum Matrix Games

Consider two players. Let us denote the action sets of Player 1 and Player 2 by $[n] = \{1, 2, \ldots, n\}$ and $[m] = \{1, 2, \ldots, m\}$, respectively. Let A be an $n \times m$ matrix representing the payoffs to Player 2 (column player). Each entry A_{ij} represents the payoff to Player 2 when he chooses action j and Player 1 (row player) chooses action i. The payoff to Player 1 is defined to be $-A_{ij}$, hence the game is called a zero-sum game as the sum of players' payoffs is always zero.

Let x and y be two probability distributions over $[n]$ and $[m]$, respectively, i.e.,

$$x_i \geq 0 \quad \text{for all } i, \quad \sum_{i=1}^{n} x_i = 1$$

$$y_j \geq 0 \quad \text{for all } j, \quad \sum_{j=1}^{m} y_j = 1$$

If Player 1 chooses an action according to x and Player 2 chooses an action according to y, the expected payoff to Player 1 equals:

$$\mathbb{E}[\text{Payoff to Player 1}] = -x^T A y.$$

Similarly, the expected payoff to Player 2 equals:

$$\mathbb{E}[\text{Payoff to Player 2}] = x^T A y.$$

The vectors x and y are called *mixed strategies*.

Question: Does there exist x^* and y^* such that

$$x^T A y^* \geq (x^*)^T A y^* \;\; \forall x \in \Delta_n,$$
$$(x^*)^T A y \leq (x^*)^T A y^* \;\; \forall y \in \Delta_m.$$

If so, (x^*, y^*) is called a Nash equilibrium (NE), i.e., neither player has an incentive to deviate.

Player 1's problem: Choose x to minimize the maximum loss:

$$\min_{x \in \Delta_n} \max_{y \in \Delta_m} x^T A y$$

Player 2's problem: Choose y to maximize the minimum gain:

$$\max_{y \in \Delta_m} \min_{x \in \Delta_n} x^T A y$$

Theorem 81 (Minimax Theorem). *For any finite $n \times m$ payoff matrix A, we have*

$$\min_{x\in\Delta_n}\max_{y\in\Delta_m} x^T A y = \max_{y\in\Delta_m}\min_{x\in\Delta_n} x^T A y = v^*,$$

*and the common value v^* is called the **value of the game**.*

Proof: We can compute the value of the game using linear programming as follows:

$$\min_{x\in\Delta_n}\max_{y\in\Delta_m} x^T A y = \min_{x\in\Delta_n}\max_{j}\,(x^T A)_j.$$

Therefore, by introducing a new variable v_2 the right-hand side can be computed using the following linear program (LP1):

$$\begin{aligned} &\min\ v_2 \\ &(x^T A)_j \le v_2 \quad \forall j \\ &\sum_{i=1}^{n} x_i = 1 \\ &x_i \ge 0\ \forall i. \end{aligned}$$

Similarly, we can write

$$\max_{y\in\Delta_m}\min_{x\in\Delta_n} x^T A y = \max_{y\in\Delta_m}\min_{i}(Ay)_i.$$

Therefore, by introducing a new variable v_1 the right-hand side can be computed using the following linear program (LP2):

$$\begin{aligned} &\max\ v_1 \\ &(Ay)_i \ge v_1 \quad \forall i \\ &\sum_{j=1}^{m} y_j = 1 \\ &y_j \ge 0\ \forall j. \end{aligned}$$

Now, it is easy to verify that these two LPs are duals of each other, and therefore, by the strong duality theorem $\max_y \min_x x^T A y = \min_x \max_y x^T A y = v^*$. □

Proposition 82. *The following statements are true:*

- *If x^* solves LP1 and y^* solves LP2, then,* $\max_y \min_x x^T A y = \min_x \max_y x^T A y = (x^*)^T A y^*$.
- *If x^* solves LP1 and y^* solves LP2, then, (x^*, y^*) is a NE.*
- *If (x^*, y^*) is a NE, then* $\max_y \min_x x^T A y = \min_x \max_y x^T A y = (x^*)^T A y^*$.
- *If (x^*, y^*) is a NE, then x^* must solve LP1 and y^* must solve LP2.*

Proof: (a): If x^* solves LP1, we have

$$\min_x \max_y x^T A y = \max_j \big((x^*)^T A\big)_j = v_2^*,$$

which implies $v_2^* \geq \left((x^*)^T A\right)_j \ \forall j$. Therefore, for any $y \in \Delta_m$, and in particular $y = y^*$, we have $v_2^* \geq (x^*)^T A y^*$. Similarly, if y^* solves LP2, we must have $v_1^* \leq (x^*)^T A y^*$. But we know from the Minimax theorem that $v_1^* = v_2^* = v^*$, which implies $v_1^* = v_2^* = (x^*)^T A y^*$.

(b): Since x^* solves LP1, we have

$$\begin{aligned}\max_y (x^*)^T A y &= \min_x \max_y x^T A y \\ &= \max_y \min_x x^T A y \\ &= \min_x x^T A y^* \leq (x^*)^T A y^*,\end{aligned}$$

where the second equality is by the Minimax Theorem and the last equality holds because y^* solves LP2. Using a similar argument one can see that $(x^*)^T A y^* \leq \min_x x^T A y^*$, which shows that (x^*, y^*) must be a NE.

(c), (d): The proof of the last two statements are similar and we leave as an exercise. □

Remark 19. *The implication of the above proposition is that we can find a NE in zero-sum games in polynomial time by solving LP1 and LP2. Moreover, if (x_1^*, y_1^*) and (x_2^*, y_2^*) are two NE, then (x_1^*, y_2^*) and (x_2^*, y_1^*) are also NE (as they solve the corresponding LPs). This property is called* saddle point interchangeability *because in zero-sum games, the NE is also called a* saddle point.

General Noncooperative Static Games

Definition: A *normal (strategic) form game* is a triplet $G = ([n], \{A_i\}_{i\in[n]}, \{u_i\}_{i\in[n]})$, where:

- $[n] = \{1, 2, \dots, n\}$ is a finite set of players,
- $A_i \neq \emptyset$ is the action set for player $i \in [n]$,
- $u_i : \prod_{j\in[n]} A_j \to \mathbb{R}$ is the payoff function of player i. Our convention is that each player i is interested in maximizing its payoff function $u_i(\cdot)$.
- Any $a_i \in A_i$ is called an *action* or *pure strategy* for player i. We let a_{-i} denote a vector of actions for all players other than i, i.e., $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$. Then we write $A_{-i} = \prod_{j\neq i} A_j$.

Example: 2-Player Nonzero-sum Games

For a 2-player nonzero-sum game, we can represent the payoff function using a matrix. Each row corresponds to an action of the first player (row player), and each column corresponds to an action of the second player (column player). The cell indexed by row i and column j contains a pair (a, b), where:

$$a = u_1(i, j), \quad b = u_2(i, j)$$

	Heads (P2)	Tails (P2)
Heads (P1)	$(1, 1)$	$(-1, 1)$
Tails (P1)	$(1, -1)$	$(-1, -1)$

Note that this game is not a zero-sum game because the sum of the payoffs is not zero. In fact, in a sharp contrast with zero-sum games, computing a NE in 2-player nonzero-sum games is generally a computationally hard problem.

Example: Games with Infinite Action Sets (Cournot Competition Game)

Action sets can also be infinite. For instance, in a Cournot competition game, two firms (players) produce a homogeneous good and aim to maximize their profits. Formally, let $G = ([2], \{A_1, A_2\}, \{u_1, u_2\})$, where:

- $A_i = [0, \infty)$ for $i = 1, 2$ (amount of good produced by firm i),
- $u_i(a_1, a_2) = p(q) \cdot a_i - c_i a_i$, with $q = a_1 + a_2$, $p(\cdot)$ the price function, and c_i the unit cost of production for firm i.

Definition 83 (Best Response Map). *Given player i and an action profile a_{-i} of the other players, the best response map for player i is defined by:*

$$B_i(a_{-i}) = \left\{ a_i \in A_i \mid u_i(a_i, a_{-i}) \geq u_i(a'_i, a_{-i}) \text{ for all } a'_i \in A_i \right\}.$$

Example: Cournot Competition with Specific Price Function

Suppose G is a Cournot game with $c_1 = c_2 = 1$ and

$$p(q) = \max\{0, 2 - q\}. \quad \text{(linear inverse demand)}$$

Then player 1's best response is:

$$B_1(a_2) = \arg\max_{a_1 \geq 0} \left[a_1 \cdot \max\{0, 2 - (a_1 + a_2)\} - a_1 \right].$$

Solving this maximization problem, we get:

$$B_1(a_2) = \begin{cases} \frac{1-a_2}{2}, & \text{if } a_2 \leq 1 \\ 0, & \text{otherwise.} \end{cases}$$

Similarly,

$$B_2(a_1) = \begin{cases} \frac{1-a_1}{2}, & \text{if } a_1 \leq 1 \\ 0, & \text{otherwise.} \end{cases}$$

Definition 84. *A mixed strategy for player i is any probability distribution over the action set A_i.*[6] *For finite action games, we often denote a mixed strategy for player i as:*

$$\sigma_i = (p_1, p_2, \ldots, p_m),$$

where p_j is the probability that player i chooses action $j \in A_i$.

Definition 85 (Pure Nash Equilibrium). *Given a game $G = ([n], \{A_i\}, \{u_i\})$, an action profile $a^* = (a_1^*, a_2^*, \ldots, a_n^*) \in \prod_{i=1}^n A_i$ is called a pure Nash equilibrium if for all $i \in [n]$, we have:*

$$u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*) \quad \forall a_i \in A_i.$$

In other words, a pure strategy profile a^* is a Nash equilibrium if no player i can benefit by unilaterally deviating from their action a_i^*, assuming all other players stick to their actions a_{-i}^*.

Remark 20. *The definition of pure Nash equilibrium can also be stated in terms of the best response map. That is, a^* is a pure Nash equilibrium if:*

$$a_i^* \in B_i(a_{-i}^*) \quad \text{for all } i \in [n].$$

[6]For instance, in the above nonzero-sum game, the strategy of playing `Heads` with probability p and `Tails` with probability $1 - p$ forms a mixed strategy for player i.

Example: Bimatrix Games

Consider the following bimatrix game:

	C_1	C_2
R_1	$(2,1)$	$(0,0)$
R_2	$(0,0)$	$(1,2)$

This game has two pure Nash equilibria: (R_1, C_1) and (R_2, C_2). However, note that not all games admit a pure Nash equilibrium. The following is an example of a bimatrix game without any pure Nash equilibrium:

	C_1	C_2
R_1	$(-1,1)$	$(0,0)$
R_2	$(0,0)$	$(1,-1)$

Definition 86 (Mixed Nash Equilibrium). *A* mixed strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ *is called a* mixed Nash equilibrium *if for each player* i*:*

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

where Δ_i *denotes the probability simplex (set of all mixed strategies) over* A_i*, and the expected utility is defined by* $u_i(\sigma_i, \sigma_{-i}) = \sum_{a \in A} u_i(a) \cdot \sigma(a)$*, where* $\sigma(a) = \prod_{j=1}^n \sigma_j(a_j)$*.*

Remark 21. *For any fixed* σ_{-i}*, the utility function* $u_i(\sigma_i, \sigma_{-i})$ *is a linear (and hence continuous) function of* σ_i*. That is, for any* $\sigma_i^1, \sigma_i^2 \in \Delta_i$*, and any* $\lambda \in [0,1]$*, we have:*

$$u_i(\lambda \sigma_i^1 + (1-\lambda)\sigma_i^2, \sigma_{-i}) = \lambda u_i(\sigma_i^1, \sigma_{-i}) + (1-\lambda) u_i(\sigma_i^2, \sigma_{-i}).$$

Theorem 87 (Weierstrass). *Let* A *be a nonempty compact set and* $f : A \to \mathbb{R}$ *be a continuous function. Then the optimization problem* $\max_{x \in A} f(x)$ *admits a solution.*

Theorem 88 (Kakutani Fixed Point). *Let* $f : A \to 2^A$ *be a set-valued function such that:*

1. *A is a compact, convex, and nonempty subset of a finite-dimensional Euclidean space,*
2. *$f(x)$ is nonempty for all $x \in A$,*
3. *$f(x)$ is a convex set for all $x \in A$,*
4. *f has a closed graph: If $(x^k, y^k) \to (x, y)$ with $y^k \in f(x^k)$, then $y \in f(x)$.*

Then, there exists $x \in A$ *such that* $x \in f(x)$*.*

Theorem 89. *Every finite noncooperative static game admits a mixed Nash equilibrium.*

Proof: Let $\Sigma = \Delta_1 \times \cdots \times \Delta_n$ be the set of all mixed strategy profiles. Define the best response map $B : \Sigma \to 2^\Sigma$ as:

$$B(\sigma) = (B_1(\sigma_{-1}), \ldots, B_n(\sigma_{-n})),$$

where:

$$B_i(\sigma_{-i}) = \arg\max_{\sigma_i \in \Delta_i} u_i(\sigma_i, \sigma_{-i}).$$

We aim to apply Kakutani's theorem to B. Let us verify the conditions:

1. Each Δ_i is a compact, convex, nonempty subset of $\mathbb{R}^m$. Hence, Σ is compact, convex, and nonempty.

2. By Weierstrass' theorem and continuity of u_i, $B_i(\sigma_{-i})$ is nonempty for all σ_{-i}. Hence, $B(\sigma)$ is nonempty for all $\sigma \in \Sigma$.

3. For any $\sigma_i^1, \sigma_i^2 \in B_i(\sigma_{-i})$ and $\lambda \in [0,1]$, we have:

$$u_i(\lambda\sigma_i^1 + (1-\lambda)\sigma_i^2, \sigma_{-i}) = \lambda u_i(\sigma_i^1, \sigma_{-i}) + (1-\lambda)u_i(\sigma_i^2, \sigma_{-i}) \geq u_i(\sigma_i', \sigma_{-i}) \quad \forall \sigma_i' \in \Delta_i.$$

   Therefore, $\lambda\sigma_i^1 + (1-\lambda)\sigma_i^2 \in B_i(\sigma_{-i})$, so $B_i(\sigma_{-i})$ is convex. Thus, $B(\sigma)$ is convex-valued.

4. B has a closed graph: Suppose, for contradiction, $(\sigma^k, \bar{\sigma}^k) \to (\sigma, \bar{\sigma})$ with $\bar{\sigma}^k \in B(\sigma^k)$ but $\bar{\sigma} \notin B(\sigma)$. Then there exists player i, and $\hat{\sigma}_i \in \Delta_i$ such that:

$$u_i(\hat{\sigma}_i, \sigma_{-i}) > u_i(\bar{\sigma}_i, \sigma_{-i}) + 3\epsilon \quad \text{for some } \epsilon > 0.$$

   Since $\sigma_{-i}^k \to \sigma_{-i}$ and by continuity of u_i, for sufficiently large k, we have $u_i(\hat{\sigma}_i, \sigma_{-i}^k) > u_i(\hat{\sigma}_i, \sigma_{-i}) - \epsilon$. Combining these two relations and using continuity of u_i, we obtain

$$u_i(\hat{\sigma}_i, \sigma_{-i}^k) > u_i(\bar{\sigma}_i^k, \sigma_{-i}^k) + \epsilon,$$

   contradicting that $\bar{\sigma}_i^k \in B_i(\sigma_{-i}^k)$. Thus, the mapping B is closed.

By Kakutani's theorem, B has a fixed point $\sigma^* \in B(\sigma^*)$, which is a mixed Nash equilibrium.

□

Existence of Pure Nash Equilibrium

As we saw earlier, finite noncooperative games do not necessarily admit a pure Nash equilibrium (NE). However, in many practical situations, one is only interested in pure NE. Therefore, it is important to identify subclasses of games for which the existence of a pure NE is guaranteed.

Definition (Coupled Action Sets)

We say that the action sets of players, denoted by $a = (a_1, a_2, \ldots, a_n)$, are *coupled* if $a \in R$, where $R \subseteq \prod_i A_i$ is a subset of the full (uncoupled) action space, characterized by some constraints. In the special case where $R = \prod_i A_i$, we recover the original uncoupled action space.

Remark 22. *Analyzing games with coupled action sets is generally more complex. Note that the existence of mixed NE discussed earlier only holds for finite games with uncoupled action spaces.*

Example: Coupled Cournot Game

Consider a two-player Cournot game where the actions of the players are coupled by constraints. Define the feasible set:

$$R := \left\{(a_1, a_2) \in \mathbb{R}^2 \mid 3a_1 + 2a_2 \leq 6,\ a_2 \geq 0\right\}.$$

For instance, if $a_1 = 1$, then the constraint restricts a_2 to satisfy $a_2 \leq \frac{3}{2}$.

Theorem 90 (Pure NE in Coupled Games). *Let R be a coupled action set. Suppose that:*

- *R is compact, convex, and nonempty,*
- *for each fixed $a_{-i} \in A_{-i}$, the utility function $u_i(a_i, a_{-i})$ is continuous and concave in a_i.*

Then, the game $G = ([n], R, \{u_i\})$ admits a pure Nash equilibrium.

Proof: Recall $a \in R$ is a pure NE if for all i,

$$u_i(a_i, a_{-i}) \geq u_i(a_i', a_{-i}) \quad \forall a_i' \text{ such that } (a_i', a_{-i}) \in R.$$

Define a function $L : R \times R \to \mathbb{R}$ by:

$$L(b, a) = \sum_{i=1}^{n} u_i(b_i, a_{-i}).$$

If there exists an action profile $a \in R$ such that:

$$a \in \arg\max_{b \in R} L(b, a),$$

then a must be a pure NE. Now suppose, for contradiction, that a is not a pure NE. Then there exists some player i and $a_i' \in A_i$ such that $(a_i', a_{-i}) \in R$ and:

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

Define $b = (a_i', a_{-i}) \in R$. Then:

$$L(b, a) = u_i(a_i', a_{-i}) + \sum_{j \neq i} u_j(a_j, a_{-j}) > L(a, a),$$

contradicting that a is a maximizer of $L(b, a)$.

Define the set-valued map $g : R \to 2^R$ by:

$$g(a) = \arg\max_{b \in R} L(b, a).$$

To find a pure NE, it suffices to show that g has a fixed point, i.e., some $a \in R$ such that $a \in g(a)$. We now verify that g satisfies the conditions of Kakutani's fixed point theorem:

1. R is compact, convex, and nonempty by assumption.
2. For each $a \in R$, $g(a)$ is nonempty and convex:
   - Nonemptiness follows from Weierstrass's theorem, since $L(b, a)$ is continuous in b and R is compact.
   - Convexity follows from the concavity of $L(b, a)$ in b (which holds since each $u_i(b_i, a_{-i})$ is concave in b_i).
3. The mapping g is closed: That is, if $(a^k, b^k) \to (a, b)$, with $b^k \in g(a^k)$, then $b \in g(a)$. This follows by continuity of L and upper semi-continuity of arg max.

Therefore, by Kakutani's theorem, there exists $a^* \in R$ such that $a^* \in g(a^*)$, and a^* will be a pure NE. □

Potential Games

Another interesting class of strategic static games that admit a pure Nash Equilibrium (NE) is the class of **potential games**, defined as follows:

Definition 91 (Weighted Potential Game). *A game $G = ([n], \{A_i\}_{i\in[n]}, \{u_i\}_{i\in[n]})$ is called a* weighted potential game *if there exists a global function $\Phi : \prod_i A_i \to \mathbb{R}$, such that for any player $i \in [n]$, and for all $a_i, a_i' \in A_i$, and $a_{-i} \in A_{-i}$, we have:*

$$\Phi(a_i', a_{-i}) - \Phi(a_i, a_{-i}) = w_i \left(u_i(a_i', a_{-i}) - u_i(a_i, a_{-i})\right),$$

where w_i is some positive scalar. In case $w_i = 1$ for all i, the game is called an exact potential game*. We refer to the function $\Phi(\cdot)$ as a weighted (or exact) potential function.*

Theorem 92. *Finite-action potential games admit a pure NE. In particular, any sequence of best-response (or better-response) moves in which players sequentially (in any order) update their actions to a strictly better action converges to a pure NE.*

Proof: Consider an exact potential game with potential function Φ. (The argument for weighted potential games is identical.) Consider the following repeated play: Each player starts with an arbitrary action. Players take turns and one player at a time chooses a new action which strictly improves their utility. Each such action increases the value of Φ. Since there are only a finite set of joint actions, we must eventually reach a local maximum of Φ, at which point no player can increase their utility by unilaterally changing their action. Thus, a pure NE is reached. □

Example: A Load Balancing Game

Consider a game with $[n] = \{1, 2, \dots, n\}$ players. Player i holds a job with weight w_i. There are $[m] = \{1, 2, \dots, m\}$ machines. An action for player i is to choose one of the machines and place their job on that machine, i.e., $a_i \in [m]$. Given an action profile $\mathbf{a} = (a_1, \dots, a_n) \in [m]^n$, the cost of player i is given by:[7]

$$c_i(\mathbf{a}) = \text{load on machine } a_i = \sum_{j:a_j=a_i} w_j.$$

Then this load balancing game admits a pure NE. To prove this, for a given action profile $\mathbf{a}$, let $L_j(\mathbf{a}) = \sum_{i:a_i=j} w_i$ be the total load on machine j under profile $\mathbf{a}$. Define:

$$\Phi(\mathbf{a}) = \sum_{j=1}^{m} L_j^2(\mathbf{a}).$$

This function captures the squared load on all machines, and can be thought of as a measure of load imbalance (larger Φ indicates worse balance). We show that Φ is a weighted potential function for the game.

Consider a unilateral deviation by player i from machine a_i to a_i'. Then:

$$\begin{aligned}\Phi(a_i', a_{-i}) - \Phi(a_i, a_{-i}) &= \left[L_{a_i'}^2(a_i', a_{-i}) + L_{a_i}^2(a_i', a_{-i})\right] - \left[L_{a_i'}^2(a_i, a_{-i}) + L_{a_i}^2(a_i, a_{-i})\right] \\ &= \left[(L_{a_i'}(\mathbf{a}) + w_i)^2 + (L_{a_i}(\mathbf{a}) - w_i)^2\right] - \left[L_{a_i'}^2(\mathbf{a}) + L_{a_i}^2(\mathbf{a})\right] \\ &= 2w_i\left(w_i + L_{a_i'}(\mathbf{a}) - L_{a_i}(\mathbf{a})\right).\end{aligned}$$

[7] Here, instead of utilities we are working with costs, i.e., $c_i(\cdot) = -u_i(\cdot)$.

On the other hand, we have

$$c_i(a_i', a_{-i}) - c_i(a_i, a_{-i}) = \left(w_i + L_{a_i'}(\mathbf{a})\right) - L_{a_i}(\mathbf{a}).$$

Hence, comparing the above two relations, we can see that the changes in the potential function and the changes in the cost functions satisfy:

$$\Phi(a_i', a_{-i}) - \Phi(a_i, a_{-i}) = 2w_i \left(c_i(a_i', a_{-i}) - c_i(a_i, a_{-i})\right).$$

Thus, the game is a *weighted potential game* with finitely many actions, and it admits a pure NE.