Learning goals:

By the end of the lecture, you should be able to

- Determine dominant-strategy equilibria of a 2-player normal form game.
- Determine pure-strategy Nash equilibria of a 2-player normal form game.
- Determine Pareto optimal outcomes of a 2-player normal form game.
- Calculate a mixed strategy Nash equilibrium of a 2-player normal form game.

Home or dancing?

Alice and Bob are best friends in grad school. They both enjoy each other's company, but neither can communicate with the other before deciding whether to stay at home (where they would not see each other) or go swing dancing this evening (where they could see each other). Each prefers going dancing to being at home. This game can be represented by the following payoff matrix.

		Bob	
		home	*dancing*
Alice	*home*	$(0, 0)$	$(0, 1)$
	dancing	$(1, 0)$	$(2, 2)$

A normal form game consists of:

- A set of player I. $I = \{Alice, Bob\}$.
- Each player $i \in I$ has a set of actions A_i. $A_{Alice} = A_{Bob} = \{home, dancing\}$.
- A payoff matrix. Once each player chooses an action, we have an outcome of the game. For example, $(home, dancing)$ is an outcome. Each agent has a utility for each outcome. For the outcome $(home, dancing)$, the utility pair $(0, 1)$ means that Alice has a utility of 0 and Bob has a utility of 1 for this outcome.

Players choose their actions

- at the same time.
- without communicating with each other.
- without knowing other players' actions.

Each player chooses a strategy, which can be pure or mixed.

- A mixed strategy is a probability distribution over all the actions.
- A pure strategy is an action. (A pure strategy is a special type of mixed strategies where one action is played with probability 1.)

A strategy profile σ contains a strategy σ_i for each player i.

For all the games before "Matching Quarters", we will focus on pure strategies, which are actions.

Pure strategy profiles for this game: (home, home), (home, dancing), (dancing, home), (dancing, dancing).

For a strategy profile σ, let σ_i be the strategy of agent i and let σ_{-i} denote the strategies of all agents except i. $\sigma_{-i} = \{\sigma_1, \sigma_2, \ldots, \sigma_{i-1}, \sigma_{i+1}, \ldots, \sigma_n\}$.

Let $U_i(\sigma) = U_i(\sigma_i, \sigma_{-i})$ denote the utility of agent i under the strategy profile σ.

What would Alice and Bob do?

Dominance and dominant strategy equilibrium

- For player i, **a strategy** σ_i **dominates strategy** σ_i' iff
  - $U_i(\sigma_i, \sigma_{-i}) \geq U_i(\sigma_i', \sigma_{-i}), \forall \sigma_{-i}$, and
  - $U_i(\sigma_i, \sigma_{-i}) > U_i(\sigma_i', \sigma_{-i}), \exists \sigma_{-i}$

  The first inequality says: regardless of the other players' strategies, player i weakly prefers σ_i to σ_i'.

  The second inequality says: there exists one strategy profile for the other players such that player i strictly prefers σ_i to σ_i'.

- **A dominant strategy** dominates all other strategies.
- When each player has a dominant strategy, the combination of those strategies is called **a dominant strategy equilibrium**.

(CQ) Which one(s) of the four outcomes is/are dominant strategy equilibria?

Alice's dominant strategy is dancing. Bob's dominant strategy is dancing as well. Thus, the only dominant strategy equilibrium is (dancing, dancing).

Alice and Bob do not need to communicate beforehand. Each pursue their own interest and the best outcome occurs for both.

Dancing or running?

Alice and Bob would like to sign up for an activity together. They both prefer dancing over running. They also prefer signing up for the same activity over signing up for two different activities.

		Bob	
		dancing	*running*
Alice	*dancing*	$(2, 2)$	$(0, 0)$
	running	$(0, 0)$	$(1, 1)$

(CQ) Which outcomes are dominant strategy equilibria of this game?

There is no dominant strategy equilibrium since each player does not have a dominant strategy.

How do we predict what the players will do in this game? We can use a different concept, called Nash equilibrium.

Best Response and Nash equilibrium

- **Best response:** Given a strategy profile (σ_i, σ_{-i}), agent i's strategy σ_i is a best response to the other agents' strategies σ_{-i} if and only if

  $$U_i(\sigma_i, \sigma_{-i}) \geq U_i(\sigma_i', \sigma_{-i}), \forall \sigma_i' \neq \sigma_i.$$

  This inequality says: Fixing the other players' strategies to be σ_{-i}, my strategy σ_i is a best response if I weakly prefer σ_i to any other strategy σ_i' of mine.

  A rational player always plays the best response to all other players' strategies.

- **Nash equilibrium:** A strategy profile σ is a Nash equilibrium if and only if each agent i's strategy σ_i is a best response to the other agents' strategies σ_{-i}.

  Nash equilibrium: every agent is choosing the best strategy given the strategies of all other agents.

  Not a Nash equilibrium: at least one agent has a better strategy than their current strategy given other agents' strategies.

(CQ) Which outcomes are Nash equilibria of this game?

There are two Nash equilibria of this game: (dancing, dancing), and (running, running). At each Nash equilibrium, neither Alice nor Bob wants to change her action.

As far as Nash equilibrium is concerned, both outcomes (dancing, dancing) and (running, running) are equilibria and might be played. However, our intuition tells us that (dancing, dancing) is better for both players than (running, running). This intuition is not captured by the concept of Nash equilibrium.

How do we capture this intuition?

Which Nash equilibrium will the players choose to play?

Pareto dominance and optimality:

- **Pareto dominance:** An outcome o Pareto dominates another outcome o' iff every player is weakly better off in o and at least one player is strictly better off in o.
- **A Pareto optimal outcome:** An outcome o is Pareto optimal iff no other outcome o' Pareto dominates o.

  Notice that this definition is weaker than claiming that a Pareto optimal outcome must Pareto dominate all other outcomes. It only says that a Pareto optimal outcome cannot be Pareto dominated by any other outcome.

(CQ) Which of the four outcomes are Pareto optimal?

It is easy to see several Pareto dominance relationships: $(dancing, dancing)$ Pareto dominates all other outcomes. $(running, running)$ Pareto dominates both outcomes where the two players miscoordinate ($(running, dancing$ and $(dancing, running)$). Given these Pareto dominance relationships, the only outcome that is not Pareto dominated by any other outcome is $(dancing, dancing)$.

Thus, $(dancing, dancing)$ is the only Pareto optimal outcome.

Prisoner's dilemma

Alice and Bob have been caught by the police. Each has been offered a deal to testify against the other. They had originally agreed not to testify against each other. However, since this agreement cannot be enforced, each must choose whether to honour it. If both refuse to testify, both will be convicted of a minor charge due to lack of evidence and serve 1 year in prison. If only one testifies, the defector will go free and the other one will be convicted of a serious charge and serve 3 years in prison. If both testify, both will be convicted of a major charge and serve 2 years in prison.

		Bob	
		cooperate	defect
Alice	cooperate	$(-1,-1)$	$(-3,0)$
	defect	$(0,-3)$	$(-2,-2)$

(CQ) Dominant strategy equilibria: How many of the four outcomes are dominant strategy equilibria?

One (defect, defect).

(CQ) Nash equilibria: How many of the four outcomes are Nash equiilbria?

One (defect, defect).

(CQ) Pareto optimal outcomes: How many of the four outcomes are Pareto optimal?

Three. All but (defect, defect).

- (cooperate, cooperate): Does any other outcome Pareto dominate (cooperate, cooperate)? No.

  Does (cooperate, defect) or (-3, 0) Pareto dominate (cooperate, cooperate) (-1, -1)? No because -3 < -1. Alice's utility is worse in the former outcome, so (cooperate, defect) does NOT Pareto dominate (cooperate, cooperate)..

  Does (defect, cooperate) (0 , -3) Pareto dominate (cooperate, cooperate) (-1, -1)? No because -3 < -1. Bob's utility is worse in the former outcome, so (defect, cooperate) does NOT Pareto dominate (cooperate, cooperate)..

  Does (defect, defect) (-2, -2) Pareto dominate (cooperate, cooperate) (-1, -1)? No because -2 < -1. Both Bob and Alice's utilities are worse in the former outcome, so (defect, defect) does NOT Pareto dominate (cooperate, cooperate).

- (defect, cooperate): Does any other outcome Pareto dominate this one? No.

- (cooperate, defect): Does any other outcome Pareto dominate this one? No.

(defect, defect) is a unique Nash equilibrium, which is also a dominant strategy equilibrium. However, it is the only outcome that is not Pareto optimal.

Matching quarters

Alice and Bob are playing the game of matching quarters. They each show one side of a quarter. Alice wants the sides of the two quarters to match, whereas Bob wants the sides of the two quarters to NOT match.

		Bob	
		heads	tails
Alice	heads	$(1, 0)$	$(0, 1)$
	tails	$(0, 1)$	$(1, 0)$

Does this game have a pure strategy Nash equilibrium?

No.

But every finite game has a Nash equilibrium. Was Nash wrong? No. This game has a mixed strategy Nash equilibrium. At this equilibrium, each player plays heads with 50% probability.

A mixed strategy is a probability distribution over the actions. Here are some examples of mixed strategies for Alice or Bob.

- Play heads with probability 0.3 and tails with probability 0.7.
- Play Heads with probability 0.1 and tails with probability 0.9.
- Play Heads with probability 1 and tails with probability 0.

A mixed strategy profile consists of a mixed strategy for each player. Here are some examples of mixed strategy profiles.

- Alice's strategy is to play heads with probability 0.8 and to play tails with probability 0.2. Bob's strategy is to play heads with probability 0.1 and to play tails with probability 0.9.
- Alice's strategy is to play heads with probability 0.4 and to play tails with probability 0.6. Bob's strategy is to play heads with probability 0.4 and to play tails with probability 0.6.

For the matching quarters game, there is a mixed strategy equilibrium where Alice's strategy is to play heads with probability 0.5 and to play tails with probability 0.5, and Bob's strategy is to play heads with probability 0.5 and to play tails with probability 0.5.

How do we derive this mixed strategy equilibrium?

- Suppose that Alice plays heads with probability p and Bob plays heads with probability q.
- When a player is mixing between two actions, it means that the two actions have the same expected utility for the player — the player is indifferent between the actions.
- Alice will choose p such that Bob is indifferent between his two actions.

  If Bob plays heads, his expected utility is $p * 0 + (1 - p) * 1 = 1 - p$.

  If Bob plays tails, his expected utility is $p * 1 + (1 - p) * 0 = p$.

  Bob is indifferent between heads and tails, thus $1 - p = p \Rightarrow p = 0.5$.
- Bob will choose q such that Alice is indifferent between her two actions.

  If Alice plays heads, her expected utility is $q * 1 + (1 - q) * 0 = q$.

  If Alice plays tails, her expected utility is $q * 0 + (1 - q) * 1 = 1 - q$.

  Alice is indifferent between heads and tails, thus $q = 1 - q \Rightarrow q = 0.5$.

Why does this mixed strategy equilibrium make sense? For pure strategy equilibrium, we said that each player must be playing a best response to the strategies of all other players. Is this still the case?

At this mixed strategy equilibrium, Alice is indifferent between her two actions. Thus, both actions are best responses to Bob's strategy. Alice indeed plays both actions, each with a probability of 0.5.

The same goes for Bob. Bob is indifferent between his two actions. Thus, both of Bob's actions are best responses to Alice's strategy.

Dancing or concert?

Alice and Bob would like to sign up for an activity together. Alice prefers going dancing whereas Bob prefers going to a concert. They also prefer signing up for the same activity over signing up for two different activities.

		Bob	
		dancing	*concert*
Alice	*dancing*	$(2, 1)$	$(0, 0)$
	concert	$(0, 0)$	$(1, 2)$

How many pure strategy Nash equilibria are there?

2 (dancing, dancing) and (concert, concert).

Is there a mixed strategy Nash equilibrium? If so, at this equilibrium, Alice goes dancing with what probability? Bob goes to the concert with what probability?

Yes.

- Suppose that Alice goes dancing with probability p and Bob goes dancing with probability q.
- Alice wants to make Bob indifferent between the two actions.

  If Bob goes dancing, his expected utility is $p * 1 + (1 - p) * 0 = p$.

  If Bob goes to a concert, his expected utility is $p * 0 + (1 - p) * 2 = 2 - 2p$.

  Bob is indifferent between the two actions. So $p = 2 - 2p \Rightarrow p = 2/3$.
- Bob wants to make Alice indifferent between the two actions.

  If Alice goes dancing, her expected utility is $q * 2 + (1 - q) * 0 = 2q$.

  If Alice goes to a concert, her expected utility is $q * 0 + (1 - q) * 1 = 1 - q$.

  Alice is indifferent between the two actions. So $2q = 1 - q \Rightarrow q = 1/3$.
- At the equilibrium, Alice goes dancing with probability 2/3 and Bob goes to a concert with probability 2/3. Each goes to their preferred activity with a higher probability.