Lecture 23
Game Theory, Part 2

Alice Gao

November 2, 2021

Contents

1 Learning Goals

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

2 Pareto Optimality

Recall the "Dancing or Running" game. The following is the payoff matrix for this game:

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

Our intuition tells us that this is a coordination game. Alice and Bob want to coordinate on choosing the same activity, since they're both happier when they're doing the same thing. This intuition is confirmed by the Nash equilibrium solution concept. Both $(dancing, dancing)$ and $(running, running)$ are stable and are Nash equilibria of this game.

However, looking at the payoff matrix, Alice and Bob coordinating on dancing seems better than them coordinating on running, since their utilities are $(2,2)$ if they both go dancing, and only $(1,1)$ if they both go running.

Unfortunately, Nash equilibrium doesn't say anything about the fact that going dancing is better than going running. Whenever there's a game with multiple Nash equilibria, the concept never says which of the Nash equilibria is the best.

Is there any way to formalize this intuition, compare different Nash equilibria and determine one is better than the others?

This is where we can use the concepts called Pareto dominance and Pareto optimality to compare different outcomes of a game:

- **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.

Pareto dominance is a way to compare multiple outcomes of a game. In the definition, outcome o is being compared with outcome o'. An outcome of a game is a combination of actions of the players. In the "Dancing or Running?" game, the outcomes are $(dancing, dancing)$, $(dancing, running)$, $(running, dancing)$, and $(running, running)$.

There are two properties that need to be satisfied in order for an outcome o to Pareto dominate and outcome o'.

The first property is that every player is weakly better off in o. The logical formula to express this statement is $U_i(o) \geq U_i(o'), \forall i$.

The second property is that at least one player is strictly better off in o. The logical formula to express this statement is $U_i(o) > U_i(o'), \exists i$.

Notice that there could be situations where there are two outcomes, and outcome o does not Pareto dominate outcome o', but also outcome o' does not Pareto dominate outcome o. Therefore, a Pareto dominance relationship cannot be established between every pair of outcomes.

It seems from the name that a Pareto optimal outcome is the best outcome among all possible outcomes of a game. The definition of a Pareto optimal outcome is a bit tricky, since it may lead some to believe that if an outcome is the best or optimal, than it should Pareto dominate all other outcomes.

However, this is different from the actual definition of a Pareto optimal outcome. To understand this difference, look at the following two statements and think about the difference between them:

- An outcome o Pareto dominates all other outcomes.
- An outcome o is NOT Pareto dominated by any other outcome.

The first statement is stronger than the second statement, which is actually the definition of a Pareto optimal outcome. The difference between these two statements will become more clear as more examples are presented.

We can use Pareto optimality on the "Dancing or Running?" game.

Problem: How many of the four outcomes are **Pareto optimal**?

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

(A) 0

(B) 1

(C) 2

(D) 3

(E) 4

Solution: This is a relatively simple application of the definition, since $(dancing, dancing)$ Pareto dominates all other outcomes. If we compare $(dancing, dancing)$ and $(running, running)$, $(dancing, dancing)$ Pareto dominates

($running, running$). Therefore, ($running, running$) does not Pareto dominate ($dancing, dancing$).

Similarly, ($dancing, dancing$) Pareto dominates ($dancing, running$) and it dominates ($running, dancing$). Therefore, neither ($dancing, running$) or ($running, dancing$) Pareto dominates ($dancing, dancing$).

Since none of the other outcomes dominate ($dancing, dancing$), ($dancing, dancing$) is the only Pareto optimal outcome in this game.

The correct answer is (B) 1.

3 Prisoner's Dilemma

Prisoner's dilemma is one of the most well-known and studied games in game theory. The interesting thing about this game is that if the two prisoners are able to communicate, then they can achieve the best outcome for both of them. However, if they are not able to communicate, and they act purely out of self-interest, they both achieve the worst possible outcome.

There are many versions of Prisoner's dilemma, but the game we're looking at goes as follows: 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: refuse	Bob: testify
Alice	refuse	$(-1, -1)$	$(-3, 0)$
Alice	testify	$(0, -3)$	$(-2, -2)$

We're going to look at three problems, one for each of the solution concepts: dominant strategy equilibrium, Nash equilibrium, and Pareto optimal outcomes.

Problem: Which outcome, if any, is a **dominant strategy equilibrium**?

		Bob	
		refuse	testify
Alice	refuse	$(-1, -1)$	$(-3, 0)$
	testify	$(0, -3)$	$(-2, -2)$

(A) (refuse, refuse)

(B) (refuse, testify)

(C) (testify, refuse)

(D) (testify, testify)

(E) There is no dominant strategy equilibrium.

Solution: First notice that this game is symmetric, so it's sufficient to consider only one of the two players. Let's consider Alice.

If we want to characterize a dominant strategy equilibrium, we need each player to have a dominant strategy. Recall that a dominant strategy is the best action a player can take, regardless of what actions other players take. So, we can look at the different actions Bob could take, and see if Alice has an action that is the best for all of Bob's possible actions.

If Bob decides to refuse to testify, then Alice has two possible actions. If she refuses to testify as well, than she will get a utility of -1. If Alice testifies, she gets a utility of 0. Clearly $0 > -1$, so Alice will choose to testify.

If Bob decides to testify, then Alice again has two possible actions. If she refuses to testify, then she will get a utility of -3. If Alice testifies, she will get a utility of -2. Since $-2 > -3$, Alice will choose to testify.

Therefore, regardless of what Bob does, Alice always prefers to testify, so it's a dominant strategy for Alice. Since this game is symmetric, Bob also has testify as a dominant strategy.

Since both Alice and Bob have testify as their dominant strategy, the dominant strategy equilibrium for this game is (testify, testify).

The correct answer is (D) (testify, testify).

Problem: How many of the four outcomes are **pure-strategy Nash equilibria**?

		Bob	
		refuse	testify
Alice	refuse	$(-1, -1)$	$(-3, 0)$
	testify	$(0, -3)$	$(-2, -2)$

(A) 0

(B) 1

(C) 2

(D) 3

(E) 4

Solution: Since this game is fairly small, we can use the brute force method of going through all four possible outcomes and verifying whether each outcome is a Nash equilibrium or not.

Consider (testify, testify). To verify that this is a Nash equilibrium, we need to ask whether either player can improve their utility by deviating to another action. If a player can improve their utility by deviating, then this outcome is not a Nash equilibrium.

For (testify, testify), if Alice chooses to deviate, she will get a utility of -3 instead of -2, so she would get a worse utility if she deviates. If Bob chooses to deviate, he will also get a utility of -3 instead of -2, so he would get a worse utility if he deviates as well.

Since both players will get lower utilities by deviating, they would not want to deviate. Therefore, (defect, defect) is a Nash equilibrium of this game.

For (refuse, refuse), if Alice chooses to deviate, she will get a utility of 0 instead of -1, so she would get a better utility if she deviates. Similarly for Bob, he will get a utility of 0 instead of 1, if he deviates. Since both players want to deviate, (refuse, refuse) is not a Nash equilibrium.

Similar reasoning can be used to verify that (testify, refuse) and (refuse, testify) are not Nash equilibria since one of the players in either outcome would do better by deviating from their strategy.

Therefore, we can conclude that the only pure-strategy Nash equilibrium of this game is (testify, testify). A simpler way of getting this answer requires you to realize that the dominant strategy equilibrium is a stronger concept than Nash equilibrium. In fact, if an outcome is a dominant strategy equilibrium, then it is also a Nash equilibrium.

To understand this, note that when we're characterizing a dominant strategy equilibrium, we're finding a dominant strategy for each player. The definitions of a dominant strategy is a player wants to take a specific action regardless of what other players do.

If an action is already a dominant strategy, then it also would satisfy the requirements for Nash equilibrium. The requirement for Nash equilibrium is an action that the player prefers to take given what the other players are doing. This is just a weaker form of dominant strategy.

In other words, if a player prefers a particular action regardless of what other players are doing, then the player definitely also prefers this action given what other players are doing.

The correct answer is (B) 1.

Problem: How many of the four outcomes are **Pareto optimal**?

		Bob: refuse	Bob: testify
Alice	refuse	$(-1,-1)$	$(-3,0)$
	testify	$(0,-3)$	$(-2,-2)$

(A) 0

(B) 1

(C) 2

(D) 3

(E) 4

Solution: We can go through all four outcomes and verify whether each outcome is Pareto optimal or not.

Consider (testify, testify). To verify that an outcome is Pareto optimal, we need to check if any other outcome Pareto dominates the outcome in question. We can easily see that (refuse, refuse) Pareto dominates (testify, testify). Both players strictly prefers (refuse, refuse) over (testify, testify) since they both get a utility of -1 instead of -2.

Since we've found an outcome that Pareto dominates (testify, testify), we can say that (testify, testify) is not Pareto optimal.

Now consider (refuse, refuse). We need to verify that (refuse, refuse) is not Pareto dominated by any of the other three outcomes: (refuse, testify), (testify, refuse), and (testify, testify).

Comparing (refuse, refuse) to (testify, refuse), we can see that Alice prefers (testify, refuse) since she gets a utility of 0 rather than -1, but Bob prefers (refuse, refuse) since he gets a utility of -1 instead of -3. Given this, we know that (testify, refuse) does not Pareto dominate (refuse, refuse).

Using the same argument as (testify, refuse), we can also verify that (refuse, testify) does not Pareto dominate (refuse, refuse), since this game is symmetric.

For (testify, testify), we've already established that (refuse, refuse) Pareto dominates (testify, testify), so (testify, testify) cannot Pareto dominate (refuse, refuse). Both players prefer (refuse, refuse) to (testify, testify).

Since we've established that no other outcome dominates (refuse, refuse), we can say that (refuse, refuse) is Pareto optimal.

Using the same strategy as above, we can establish that (refuse, testify) and (testify, refuse) are also Pareto optimal.

Notice two things here: The first is that the definition of Pareto optimality does not require Pareto dominance over all other outcomes. That is a much stronger condition. Instead, it only requires that the outcome is not dominated by any other outcome.

The second is that there can be two outcomes where neither outcome Pareto dominates the other. Comparing (refuse, refuse) and (refuse, testify), Alice prefers (refuse, refuse), but Bob prefers (refuse, testify). In this case, neither outcome Pareto dominates the other.

The correct answer is (D) 3.

After thinking about these three problems, we can start to think about this game strategy a little more. Looking rationally at this game, (refuse, refuse) seems to be the best outcome for both players. If both players testify, they both serve 2 years in prison, but if they both refuse to testify, then they both only serve 1 year in prison.

It turns out, though, that there is only one dominant strategy equilibrium and only one pure-strategy Nash equilibrium, and they are both (testify, testify). Even more interesting, the only dominant strategy equilibrium and only Nash equilibrium is also the only outcome that is not Pareto optimal.

(testify, testify) is the worst outcome in multiple senses, yet it's the only stable one. There's quite a bit of research on how people can be motivated to achieving the (refuse, refuse) outcome when playing this kind of game in practice.

From numerous economic experiments on the Prisoner's Dilemma game, it was found that if the two players are only going to play the game once, or a small number of times, then it's very difficult to sustain or motivate cooperation.

However, if the two players know that they're going to repeatedly interact with each other multiple times, there is a lot more cooperation between players since they can build a trust between each other which will motivate them to cooperate and achieve the (refuse, refuse) outcome.

4 Mixed-Strategy Nash Equilibrium

4.1 Matching quarters

We will use a game called "Matching quarters" to introduce the solution concept called mixed-strategy Nash equilibrium.

The story for this game goes as follows: 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	$(0, 1)$	$(1, 0)$

Problem:

How many of the four outcomes are **pure-strategy Nash equilibria**?

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

(A) 0

(B) 1

(C) 2

(D) 3

(E) 4

Solution: The correct answer is (A).

Recall that every finite game has a Nash equilibrium. So why doesn't the "Matching quarters" game have a Nash equilibrium? This is because the exact statement was that every finite game has a mixed-strategy Nash equilibrium, and a pure-strategy Nash equilibrium is just a special case of a mixed-strategy Nash equilibrium.

"Matching quarters" doesn't have a pure-strategy Nash equilibrium, but it is guaranteed to have a mixed-strategy Nash equilibrium according to the statement. We can derive this mixed-strategy Nash equilibrium.

Example: To derive a mixed-strategy Nash equilibrium for this game, recall that a mixed-strategy is a probability distribution over all the actions for all other players. In this case, there are two actions per player, so Alice needs to play heads with some probability, and tails with the rest of the probability. Bob also needs to play heads with some probability, and tails with the rest of the probability.

In order to derive the mixed-strategy Nash equilibrium, we need to derive the mixed-strategies for each of the players. Assume that Alice plays heads with probability p. Similarly, Bob plays heads with probability q. We only need one probability for each player to define this distribution, since there's only 2 actions for each player.

To derive these probabilities, the general principle states that a player should choose their probability such that the other players are indifferent between their actions. Specifically for the "Matching quarters" game, Alice must choose a value for p such that Bob is indifferent between his two actions. Similarly, Bob must choose a value for q such that Alice is indifferent between her two actions.

Indifferent means that the expected utility of playing either action is the same for that player. Bob being indifferent between his actions means that both of his actions give him the same expected utility, so he doesn't really care about which action he takes.

We can use this general principle to derive the probability for Alice. So, Alice needs to choose a value for p such that Bob is indifferent between his actions. If Bob is indifferent, than we can calculate his expected utility for each action and equate them together to solve for p.

The expected utility if Bob plays heads is

$$U_{Bob}(heads) = p * 0 + (1 - p) * 1 = 1 - p.$$

When Bob plays heads, Alice plays heads with probability p, and the expected utility for Bob in that case is 0, so we have $p * 0$. When Bob plays heads, Alice plays tails with probability $1-p$, and the expected utility for Bob in that case is 1, so we have $(1-p)*1$.

Similarly, the expected utility if Bob plays tails is

$$U_{Bob}(tails) = p*1 + (1-p)*0 = p.$$

We equate the expected utility of Bob playing heads to the expected utility of Bob playing tails to get

$$\begin{aligned} U_{Bob}(heads) &= U_{Bob}(tails) \\ 1-p &= p \\ p &= 0.5. \end{aligned}$$

The calculations for Bob are very similar:

$$U_{Alice}(heads) = q*1 + (1-q)*0 = q$$

$$U_{Alice}(tails) = q*0 + (1-q)*1 = 1-q$$

$$\begin{aligned} U_{Alice}(heads) &= U_{Alice}(tails) \\ q &= 1-q \\ q &= 0.5. \end{aligned}$$

Therefore, there is one mixed-strategy Nash equilibrium where Alice plays heads with probability $p = 0.5$ and Bob plays heads with probability $q = 0.5$.

Now that we know how to derive a mixed-strategy Nash equilibrium, we can answer a second question: What does it mean if a player is mixing between two actions?

The answer to this question is that it means that the two actions have the same expected utilities for the player. It means that the player is indifferent between their actions, since both actions have the same expected utility.

There's a dynamic here where each player is choosing their mixing probability to make the other player indifferent, but once they're indifferent between the two actions, the mixing probability doesn't really matter to the player anymore.

Problem: Consider a 2-player normal form game and fix Bob's strategy. Alice is willing to play heads 60% of the time and tails 40% of the time. Which of the following statements is true?

(A) Alice's expected utility of playing heads is greater than her expected utility of playing tails.

(B) Alice's expected utility of playing heads is less than her expected utility of playing tails.

(C) Alice's expected utility of playing heads is same as her expected utility of playing tails.

Solution: If Alice is willing to mix between two actions, if she's willing to play either action with positive probability, then that must mean that both actions give her the same expected utility.

If one action gave her strictly better expected utility than the other, why would she bother with playing the other action at all? She would always choose to play the action with the better expected utility. If two actions give her different expected utilities, she will always prefer to play the better one.

It's only in the case where the two actions give the same expected utilities that Alice is willing to mix between the two.

This may be confusing since the mixing probabilities have nothing to do with the expected utilities. Whether she plays heads with a probability of 60% and tails with 40%, or she plays heads with 90% and tails with 10% doesn't change the fact that the expected utilities for both heads and tails must be the same.

The correct answer is (C) Alice's expected utility of playing heads is the same as her expected utility of playing tails.

4.2 Dancing or Concert?

We can look at another game called "Dancing or concert?" to calculate a mixed-strategy Nash equilibrium.

The story for this game is as follows: Alice and Bob want to sign up for an activity together. They both prefer to sign up for the same activity. However, Alice prefers dancing over going to a concert whereas Bob prefers going to a concert over dancing.

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

The next two problems are deriving the mixed-strategy Nash equilibrium.

Problem: At the mixed-strategy Nash equilibrium, with what probability does **Alice go dancing**?

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

(A) 0

(B) 1/3

(C) 2/3

(D) 1

Solution: To answer this question, we must first derive the mixed-strategy Nash equilibrium. There are two probability distributions: one for the actions of Alice, and another for the actions of Bob. Since there are two actions for each player, we only need 1 probability per player to specify each distribution.

Suppose that Alice goes dancing with probability p and Bob goes dancing with probability q.

Recall that the general principle states that since each player is mixing between two actions, the two actions must give the player the same expected utility. As such, we need to set p such that Bob gets the same expected utility regardless of whether he goes dancing or goes to the concert.

From the payoff matrix, we can write the expected utilities of Bob for dancing and going to the concert in terms of p:

$$U_{Bob}(dancing) = p * 1 + (1 - p) * 0 = p$$

$$U_{Bob}(concert) = p * 0 + (1 - p) * 2 = 2 - 2p$$

In order for Bob to be indifferent, we need to equate $U_{Bob}(dancing)$ and $U_{Bob}(concert)$ to get p:

$$\begin{aligned} U_{Bob}(dancing) &= U_{Bob}(concert) \\ p &= 2 - 2p \\ p &= 2/3. \end{aligned}$$

Since we're only looking for Alice's probability, and we've derived $p = 2/3$, we can say that Alice goes dancing with a probability of 2/3.

The correct answer is (C) 2/3.

Problem: At the mixed-strategy Nash equilibrium, with what probability does **Bob go dancing**?

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

(A) 0

(B) 1/3

(C) 2/3

(D) 1

Solution: The process for finding q is very similar to the process for finding p.

From the payoff matrix, we can write the expected utilities of Alice for dancing and going to the concert in terms of q:

$$U_{Alice}(dancing) = q * 2 + (1 - q) * 0 = 2q$$

$$U_{Alice}(concert) = q * 0 + (1 - q) * 1 = 1 - q$$

In order for Alice to be indifferent, we need to equate $U_{Alice}(dancing)$ and $U_{Alice}(concert)$ to get q:

$$\begin{aligned} U_{Alice}(dancing) &= U_{Alice}(concert) \\ 2q &= 1 - q \\ q &= 1/3. \end{aligned}$$

Since we're only looking for Bob's probability, and we've derived $q = 1/3$, we can say that Bob goes dancing with a probability of 1/3.

The correct answer is (B) 1/3.