Lecture 22
Game Theory, Part 1

Alice Gao

November 2, 2021

Contents

1 Learning Goals

- Determine dominant-strategy equilibria of a 2-player normal form game.
- Determine pure-strategy Nash equilibria of a 2-player normal form game.

2 Applications of Game Theory

Game Theory is a branch of microeconomics. Due to the amount of time available in this course, only the basics of game theory are covered. We talk about two-person normal form games and a few solution concepts for them: dominant strategy equilibrium, Nash equilibrium, Pareto optimal outcomes, and how to derive mixed strategy Nash equilibria.

This course is about building intelligent agents for environments, but so far we've assumed that we are the only intelligent agent in the environment. What if this is not true? What if there are other intelligent agents in the environment? These other intelligent agents may have their own goals and preferences, and are also reasoning about what to do.

We shouldn't treat these other intelligent agents as noise in the environment when making decisions. We should take into account that these other intelligent agents are out there, and make our decisions based on our knowledge about how they would behave.

Game theory is an important tool we can use to reason about how agents will behave when there are multiple intelligent agents in the same environment.

Game theory is a commonly misunderstood term, as it suggests that it's only relevant to games. Game theory applies much more broadly than literally to games that we play. Game theory applies to most situations in life where we would have to act strategically. A game is a mathematical model of a strategic scenario.

For any scenario, as long as there are multiple intelligent agents interacting with each other, we can call it a game. There are many examples of applications of game theory and mechanism design that are not strictly games.

Example: An auction is a type of mechanism that is widely studied in algorithmic game theory.

At the Dutch Flower auction, every morning, fresh flowers from all over the world gets shipped to the Netherlands. The flowers are sold at the auction with big clocks on the walls displaying the current prices that the flowers are being sold at.

The buyers can look at these prices on the wall and decide whether or not they want to buy the flowers that those prices. Once the flowers are sold, they get shipped to their destination. The auction here is simply a mechanism for resource allocation. We're trying to allocate flowers through auctions in this scenario, but there are all sorts of auctions happening around the world at any given time.

Another example of auctions are the ads displayed in google search results, as those spots are sold in automatic auctions.

Example:

Matching problems deal with multiple agents having their own preferences, and trying to match these agents in some way so that they are happy in some sense.

The medical residency matching problem is concerned with matching a student that has graduated from a medical school to a hospital to be a resident for a few years before they become an actual doctor. If you've ever watched Grey's Anatomy, the main characters are new grads from medical schools going through their residencies.

The problem is that each student ranks a bunch of hospitals, then each hospital ranks a bunch of students, and we want to match students to hospitals such that the resulting matching is stable.

The matching is stable if we cannot find a pair of student-hospital pairings such that both the students and the hospitals would be happier if they switched. This is called a stable matching problem.

Similar examples of matching problems include matching students with secondary schools (school choice) or matching kidney transplant donors with patients (organ transplant).

Example: Crowd-sourcing is based on the idea that the crowd knows all the useful information, and all we need to do is find a way to extract that information from the crowd. Crowd-sourcing uses game theory.

The ESP game is an example of crowd-sourcing. The original purpose for those creating the ESP game was to help with labeling images. At the time, it was very difficult to label images automatically, so the creators of the game wanted people to label the images for free.

The researchers designed a game where two people are randomly matched, and the goal is to guess the word the other player is going to type. The name ESP comes from this, as the player is trying to connect with the other player in the "sixth sense" to guess what they're thinking. If the players match a word, the two players win the game and move on to the next game. The two players can't communicate with each other, so the only thing they have is the same image in front of them. Naturally, they'll try to match words that are relevant to the images that they see.

The players enjoy the game, but as a side effect, the researchers are able to generate labels for these images. This game is no longer available, but there are numerous other examples of crowd-sourcing in the real world.

It's believed that if you design a website in the right way. you can get a crowd to do anything for you. They can design things for you like 99 designs, write code like top coder, translate text like Duolingo, and rate courses and teachers like uwflow.com.

3 Introduction to Game Theory

Now that we've seen some examples of game theory, let's talk about two concepts: game theory and mechanism design. What kinds of problems are we trying to solve with these two concepts?

For game theory, we're usually given the rules of a game, and we're asking how would agents, or players, in the game actually play the game? How would they behave? What strategies would the agents use?

For mechanism design, the question is reversed. We want to design a game in such a way that the agents will behave in a particular way. Given that we want the agents to behave in a particular way, how should we design the rules of the game? Often, mechanism design is referred to as reverse game theory.

Both game theory and mechanism design are quite broad concepts with lots of small topics in each. Due to the limited time in this course, we only cover a tiny bit of game theory, and we don't cover mechanism design.

Let's try to make the multi-agent framework more specific. How should we think about the multi-agent framework?

There are multiple agents in an environment, and each agent is trying to decide on what strategies they should use to play the game. How does each agent make their decisions? What are their decisions based on?

The agents' decisions will be based on the information that they have about the environment. These are passive things, things that don't change in the environment.

The agents' decisions will also be based on the other agents out there. These other intelligent agents will make their own decisions, and choose their own strategies.

The agents' decisions will also be based on their utility functions. So, all agents' decisions will be based on

- their information about the world

- their information about other agents
- their utility function

The outcome of the entire game depends on the actions or strategies of all the agents in the game. All of the agents' strategies will jointly determine what will happen.

In life, there are many people acting in the world, and the outcome, which is the welfare of the world, is based on how every person jointly acts.

When thinking about a multi-agent environment, you may be tempted to think that all agents are competing with each other, but this is not necessarily the case. In a multi-agent environment, all we know is that each agent has their own interest as their goal, and are trying to act in a way to achieve their own goal to maximize their own utility.

However, there can be many relationships between the utility functions of multiple agents. It could be possible that all the agents have conflicting goals. In this case, the agents are competitive. It could also be the case that all the agents have the same goal. In this case, the agents are cooperative.

It could also be the case that it's somewhere in between. There could be multiple teams of agents, so within the teams, they are cooperative, but between teams, they are competitive.

Most board games are competitive, but a few of them like Hanabi, Ghost Stories, Pandemic, or Codenames are cooperative or somewhere in between.

4 Dominant Strategy Equilibrium

We will use a game called "Home or Dancing?" to introduce a solution concept called dominant strategy equilibrium.

The story of this game is as follows: Alice and Bob are 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.

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

This is an example of a two-person normal form game. In a normal form game, there's a set of players. In this case, there are two players which we usually called the row player and the column player.

In this game, Alice is the row player since her two actions are represented by the two rows, and Bob is the column player since his two actions are represented by the two columns. The two actions for both Alice and Bob are staying at home, or going swing dancing.

We often talk about the outcome of a game, and this game has four different outcomes. Each outcome consists of one action for each player. Outcomes for this game are $(home, home)$, $(home, dancing)$, $(dancing, home)$, and $(dancing, dancing)$.

We also have the payoff matrix which specifies how happy each player is for each outcome of the game. For every outcome, there's a tuple of two numbers. The first number denotes the row player's utility, and the second number denotes the column player's utility.

The top left outcome of the payoff matrix is $(home, home)$, and we can see that Alice's utility is 0, and Bob's utility is also 0 in this scenario where they both decide to stay home.

Since each person prefers going dancing to staying at home, the top right outcome of the payoff matrix is $(home, dancing)$, and we can see that Bob's utility is 1, since he'd rather go dancing than stay at home, regardless of whether Alice goes dancing, and Alice's utility is 0 since she stays at home.

Recall that from the story that the players can't communicate with each other. They choose their actions separately, but their actions jointly determine the outcome of the game. Players choose their actions at the same time, without communicating with each other, and without observing other players' actions. Another name for this type of game is the simultaneous-move game since players move simultaneously.

To solve a normal form game, we look for a strategy profile. By strategy profile, we mean we need to choose one strategy for each player. The most general form of strategy is called a mixed strategy.

- A mixed strategy is a probability distribution over the actions. e.g. stay at home w/ prob 80% go dancing w/ prob 20%.
- A pure-strategy plays one action w/ probability 1.

Each player i chooses a mixed strategy σ_i.

In general, a mixed strategy is a probability distribution over the actions of the player. A mixed strategy is saying that the player can choose to play the actions randomly, but this randomness is based on a fixed distribution. For example, a player in this game can decide to stay at home with 80% chance and then go dancing with 20% chance.

This means that every time the player is playing this game, they need to choose a random number, and based on that random number from 0 to 1, and based on a threshold determining whether they stay at home or whether they go dancing, in the long term, if they play this game enough times, they will roughly end up staying at home 80% of the time, and going dancing 20% of the time.

A general mixed strategy allows a player to play a game probabilistically rather than deterministically.

There's a special case of the mixed strategy called the pure strategy, which says that the player will stick to one action. This is also a probability distribution, but the distribution puts all of its probability on one action, so the player is going to play the one action for sure.

For the first few games we look at, we're only going to consider pure strategies for the players.

Now that we've described the game representation, how would we play this game? If you were Alice or Bob, would you stay home, or would you go out dancing? You would probably choose to go dancing whether you're Alice or Bob, so the outcome from this is $(dancing, dancing)$.

If you look at the payoff matrix, $(2, 2)$ will stand out as they are the biggest numbers in the matrix. If both players can get their biggest numbers through dancing, why would they not go dancing? It seems like a good choice to maximize their utility—to make them happy.

This is the right intuition, and we can formalize this intuition by introducing the solution concept called dominant strategy equilibrium. Using this solution concept, we can predict that both Alice and Bob will choose to go dancing, so the outcome of the game will be $(dancing, dancing)$.

To describe the dominant strategy equilibrium concept, we introduce some notation:

Terminologies for strategies:

- σ_i denotes the strategy of player i.
- σ_{-i} denotes the strategies of all the players except i.

Terminologies for utilities:

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

We need to define what it means for one strategy to dominate another strategy for a player. Consider a particular player i, and compare two strategies σ_i and σ'_i.

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
- $U_i(\sigma_i, \sigma_{-i}) > U_i(\sigma'_i, \sigma_{-i}), \exists \sigma_{-i}$

We compare the players' utility U_i for the two strategies, for σ_i and σ'_i with respect to some set of strategies for all the other players σ_{-i}.

The first inequality says that for any set of strategies for the other players, the current player prefers σ_i over σ'_i. Notice that "prefer" here means "weak prefer" ($\geq$). In some situations, the two may be the same, meaning we might get the same utility for both, but in other situations, I may get a higher utility for playing σ_i over σ'_i.

The second inequality is similar to the first, except that there are two differences. One difference is that it's a strict inequality ($>$) rather than a weak inequality ($\geq$). The other difference is that the quantifier is existential ($\exists$) rather than universal ($\forall$).

The second inequality says that there exists one set of strategies for the other players such that player i strictly prefers to play σ_i over σ'_i.

To summarize both inequalities, for any set of strategies for the other players σ_{-i}, player i weakly prefers σ_i over σ'_i, and for at least one set of strategies for the other players σ_{-i}, player i strictly prefers σ_i over σ_i'. If both these conditions are satisfied, we can say that for player i, strategy σ_i dominates strategy σ_i'.

We can say that a strategy is a **dominant strategy** for a player if it dominates all other strategies. Essentially, it is the best strategy for that player in some sense.

It's possible that a player does not have a dominant strategy, as sometimes strategies are not comparable using this concept. In other cases, a player does have a dominant strategy, so this strategy has to be better in the sense of the dominating relationship than any other strategy.

Finally, when each player has a dominant strategy, the combination of those strategies is called a **dominant strategy equilibrium**.

We can apply the concept of the dominant strategy equilibrium to the "Home or Dancing?" game.

Problem: Which of the following statements is correct?

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

(A) $(home, home)$ is the only dominant strategy equilibrium.

(B) $(dancing, dancing)$ is the only dominant strategy equilibrium.

(C) $(dancing, home)$ or $(home, dancing)$ is the only dominant strategy equilibrium.

(D) This game has more than one dominant strategy equilibrium.

(E) This game has no dominant strategy equilibrium

Solution: Firstly, notice that this game is completely symmetric. Alice and Bob are completely symmetric, so we only need to consider one of the players. The analysis for the other player is completely the same, so let's focus on Alice.

How would Alice play the game given a particular strategy for Bob? What would Alice do if Bob stays at home, and what would Alice do if Bob goes dancing? If Bob stays at home, we're looking at the left column of the matrix. Given this, if Alice goes dancing, her utility would be 1. If Alice stays at home, her utility would be 0. Since $1 > 0$, Alice prefers to go dancing:

$$U_{\text{Alice}}(dancing, home) > U_{\text{Alice}}(home, home)$$

If Bob goes dancing, we're looking at the right column of the matrix. Given this, if Alice goes dancing, her utility would be 2. If Alice stays at home, her utility would be 0. Since $2 > 0$, Alice again prefers to go dancing:

$$U_{\text{Alice}}(dancing, dancing) > U_{\text{Alice}}(home, dancing)$$

Recall the definition of dominant strategy. It says that if dancing is a dominant strategy for Alice, then Alice's utility for going dancing should be weakly better than her utility for any other strategy given any strategy for Bob.

Regardless of what Bob does, Alice prefers to go dancing. Notice that both inequalities are strict inequalities, so this means that dancing is a dominant strategy for Alice as it satisfies the definition.

Since the game is completely symmetric, dancing is also a dominant strategy for Bob. Therefore, $(dancing, dancing)$ is a dominant strategy equilibrium of this game, and it the only dominant strategy equilibrium of this game.

The correct answer is (B) $(dancing, dancing)$ is the only dominant strategy equilibrium.

The second game we're going to look at is called "Dancing or running?". We'll use this game to introduce another solution concept called the Nash equilibrium.

The story of this game is as follows: 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	$(0, 0)$	$(1, 1)$

Let's think about what the players will do, since the goal here is to predict how they would behave.

Problem: What do you think the players will do?

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

(A) $(dancing, dancing)$

(B) $(dancing, running)$

(C) $(running, dancing)$

(D) $(running, running)$

Solution: Just looking at the numbers in the matrix, most would say that Alice and Bob would both prefer to go dancing, but some would say that they would both prefer to go running. Both of these assumptions are correct, but why is that?

In some sense, $(dancing, dancing)$ and $(running, running)$ are the two stable outcomes of this game. By stable, we mean if Alice sticks to dancing, then Bob doesn't want to go running instead. Similarly, if Bob sticks to running, Alice doesn't want to go dancing instead. Deviating from this will cause either player to get a smaller utility.

We can try to apply the dominant strategy equilibrium to help us formalize the intuition from the last problem.

Problem: Which of the following statements is correct?

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

(A) $(dancing, dancing)$ is the only dominant strategy equilibrium.

(B) $(running, running)$ is the only dominant strategy equilibrium.

(C) $(dancing, running)$ or $(running, dancing)$ is the only dominant strategy equilibrium.

(D) This game has more than one dominant strategy equilibrium.

(E) This game has no dominant strategy equilibrium.

Solution: Firstly, notice that this game is again symmetric, which means we only have to consider one of the players, and the analysis for the other player will be exactly the same, so let's focus on Alice.

How would Alice play the game given a particular strategy for Bob? What would Alice do if Bob goes dancing, and what would Alice do if Bob goes running?

If we can find one action for Alice such that no matter what Bob does, Alice prefers that action, then that action will be a dominant strategy for Alice.

If Bob goes dancing, then the utility for Alice going dancing is 2, and the utility for Alice going running is 0. Since $2 > 0$, Alice prefers to go dancing in this case:

$$U_{\text{Alice}}(dancing, dancing) > U_{\text{Alice}}(running, dancing)$$

If Bob goes running, then the utility for Alice going dancing is 0, and the utility for Alice going running is 1. Since $1 > 0$, Alice prefers to go running in this case:

$$U_{\text{Alice}}(running, running) > U_{\text{Alice}}(dancing, running)$$

Depending on Bob's action, Alice has different preferences. If Bob goes dancing, then Alice prefers dancing, but if Bob goes running, Alice prefers running.

Since Alice prefers different actions depending on Bob's action, there is no single action that's best for Alice regardless of what Bob does. Therefore, Alice doesn't have a dominant strategy. Since Alice doesn't have a dominant strategy, the entire game doesn't have a dominant strategy equilibrium.

The correct answer is (E) This game has no dominant strategy equilibrium.

The previous problem shows us that the dominant strategy equilibrium solution concept is not sufficient to capture the intuition of this game. Intuitively, we believe that $(dancing, dancing)$ and $(running, running)$ are both reasonable outcomes of the game, but the dominant strategy equilibrium doesn't give us a prediction of how the players will behave.

As such, we need to look at another solution concept called the Nash equilibrium to predict the players' behaviour.

5 Nash Equilibrium

The creator of the Nash equilibrium is John Nash, an American mathematician who made a lot of important contributions to game theory and various branches of mathematics. As part of his PhD thesis at Princeton, Nash worked out the idea for the Nash equilibrium. This idea, along with other contributions to economics and mathematics eventually won Nash the Nobel prize in Economics. "A Beautiful Mind" is a biographical film based on John Nash's life, although it is not an entirely accurate depiction of Nash.

Nash not only proposed the concept of the Nash equilibrium, but he was also able to prove that every finite game has at least one Nash equilibrium. This means given any finite game, we're guaranteed to be able to derive a Nash equilibrium, and we could use it to predict how players would behave in any finite game.

This makes the Nash equilibrium a really applicable and useful solution concept compared to the dominant strategy equilibrium, where it's possible to not have a dominant strategy equilibrium, which means not being able to make predictions about some games.

The concept of Nash equilibrium is based on the idea of best response:

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

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

In short, this definition says that given what the other agents are doing σ_{-i}, σ_i is the best strategy compared to any other strategy σ_i' for agent i.

Given the concept, the Nash equilibrium is easy to define:

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

The Nash equilibrium is trying to characterize a stable set of strategies. This set of strategies is stable because if the actions of all the other agents are fixed, then agent i wouldn't want to change strategies, since the current strategy is already the "weakly" best option. If every agent thinks they're playing their best option, no agent would want to switch to another strategy.

Another way of characterizing the Nash equilibrium is that no single agent wants to deviate to another strategy if the strategy profile is at the Nash equilibrium.

We can apply the Nash equilibrium solution concept to the "Dancing or running?" game.

Problem: Which of the following is correct? Consider only pure-strategy Nash equilibria.

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

(A) $(dancing, dancing)$ is the only Nash equilibrium.

(B) $(running, running)$ is the only Nash equilibrium.

(C) $(dancing, dancing)$ and $(running, running)$ are both Nash equilibria.

(D) This game has more than two Nash equilibria.

Solution: There are a couple of general approaches that can be used to derive a pure-strategy Nash equilibrium for any normal form game.

One approach is to use the best responses, and follow a chain of best responses until a stable outcome of the game is reached. To use this approach, we start with choosing an arbitrary outcome. Let's choose $(dancing, running)$.

From this outcome, we pick either player and try to determine whether the player is playing the best response to the other players' strategy. Basically, we pick a player, try to improve their strategy to be a best response to other players, and continue to do this until every player is playing a best response to other players' strategies, and then we get a Nash equilibrium. Let's consider Bob in this case.

In this case, we're fixing Alice's action to be dancing. Bob's current strategy is running, which gives him a utility of 0. If Bob chooses to go dancing, he gets a utility of 2, so dancing is clearly better than running, which means the current strategy profile for Bob is not his best response. We switch Bob's strategy from running to dancing which means we switch the profile from $(dancing, running)$ to $(dancing, dancing)$.

Now we have a new profile after improving Bob's strategy to his best response to Alice's strategy. Next, we try to improve Alice's strategy in a similar way. In this case, we're fixing Bob's action to be dancing. Alice's current strategy is dancing, which gives her a utility of 2. If Alice goes running, she gets a utility of 0, so dancing is clearly better than running for Alice, which means the current strategy profile for Alice is her best response.

At this point, we've reached a stable point, since both Alice and Bob are playing their best responses to each others' strategies. Therefore, this is a Nash equilibrium.

By using similar reasoning, we can also derive that $(running, running)$ is a Nash equilibrium. We can start with $(dancing, running)$ or $(running, dancing)$, and choosing the correct player and following the strategy we used above will get $(running, running)$ as a stable outcome as well, and therefore a Nash equilibrium.

Another approach is to look at all possible strategies for each player and figure out the best responses for all of them. Given that the game is symmetric, we only need to consider one of the players. Let's consider changing Bob's strategy, and determining what Alice's best response is with respect to Bob's strategy.

If Bob goes dancing, Alice's best response is to go dancing. If Bob goes running, Alice's best response is to go running. This can be determined by looking at the utility values for the different outcomes on the matrix.

Since the game is symmetric, Bob's best responses are the same as Alice's, which means this is a coordination game, which means it's best if players coordinate on taking the same actions. This leads us to the conclusion that there are two stable outcomes (($dancing, dancing$) and ($running, running$)) that are Nash equilibria for this game.

The difference between the two approaches is that the first approach is somewhat ad hoc. We start with a random outcome, and then we randomly choose a player to try to improve their strategy, so we'll likely be able to find at least one pure-strategy Nash equilibrium. However, it's not guaranteed that following the first strategy will lead to finding all the pure-strategy Nash equilibria.

The second approach is more systematic. We end up analysing all the possible best responses, and we should be able to find all the pure-strategy Nash euqilibia of a game.

The correct answer is (C) ($dancing, dancing$) and ($running, running$) are both Nash equilibria.