Game Theory

Aaron Anderson for Los Angeles Math Circle

5/31/2020

In this worksheet, we will study a few kinds of games, and try to mathematically determine what *rational* behavior is in each game, and then determine what the outcome of the game will be if each of the players is rational.

1 Tragedy of the Commons

Let's play a game! This game, called *Tragedy of the Commons*, simulates the upkeep of a common resource, and probably wouldn't guarantee hours of fun for the whole family, but understanding it can tell us a lot about human behavior.

There are several players (your breakout room), and all of the gameplay happens at once. Each of you will choose one of the following two choices: "Cooperate" or "Defect," and send either "C" or "D" to your breakout room instructor in chat, without letting anyone else know your choice. Once the instructor has heard from everybody, they announce the results. Everybody who cooperated loses 2 points (the cost of donating to maintain some shared resource). Then the number of people who cooperated is tallied up, call it c. Everybody then gets c points, because everybody profits from this shared resource, regardless of whether they contributed. The goal is to maximize your score (your ranking among the players is irrelevant - no one player wins).

Problem 1 Play one round of Tragedy of the Commons as a group. Who do you think won?

Definition. The options "Cooperate" and "Defect" are called *strategies.* We say that one strategy *dominates* another if playing the first strategy will always give you more points than the second, regardless of what the other players do. A *dominating* strategy is just a strategy that dominates another.

We can assume that no rational player will ever use a dominated strategy, because they would always benefit from switching to the strategy that dominates it.

Problem 2 In the Tragedy of the Commons, does one strategy dominate another?

Play another round as a group, using what you know about dominating strategies to try to maximize your score.

Problem 3 Play again, but now try to maximize the total score of everybody in this group. Did you get a higher score when everyone tried to maximize your own score, or when you tried to maximize the score of the group? Discuss how to reconcile this with the notion that it is never rational to play a dominated strategy.

Definition. An *outcome* is a list of one strategy for each player. For example, one outcome of 3-player Tragedy of the Commons is described by Player 1 cooperating and Players 2 and 3 defecting, which we can abbreviate as CDD. An outcome is a *Pareto alternative* to another outcome if the

Pareto alternative gives at least as good of a score to each player as the other outcome, and at least one player gets a strictly higher score. An outcome is *Pareto optimal* if it has no Pareto alternatives.

Problem 4 In Tragedy of the Commons with n players, which outcomes are Pareto optimal? Do you expect them to happen if all players rationally try to maximize their own scores?

2 Zero-Sum Games

Now let's start studying 2-player games in more detail. Let's start with *zero-sum games*, which are games where Player 1's score and Player 2's score have to add to 0. We can summarize such a game in a *payoff matrix*, where the rows correspond to the strategies for Player 1, and the columns correspond to the strategies for Player 2. Then each square in the matrix represents an outcome, and the number we record there is the score Player 1 gets in that outcome, while Player 2 gets the negative of that.

Player 2

Player 1		R	P	S
	R	0	-1	1
	P	1	0	-1
	S	-1	1	0

Figure 1: A Payoff Matrix for Rock, Paper, Scissors

Problem 5 Do either of these games have dominating strategies? Assuming that all players play rationally, rejecting dominated strategies, can you determine the outcome of each game?

Player 2

Player 1		L	R
	U	1	2
	D	1	-2

Player 2

Player 1		L	R
	U	-1	1
	D	2	-2

Figure 2: Zero-Sum Game 2 and Zero-Sum Game 3

Problem 6 In a zero-sum game, which outcomes are Pareto optimal?

2.1 Maximin and Minimax

For any game, and any choice of player, define the *minimum value* of any of their strategies to be the minimum over all the scores of all outcomes where that player takes that strategy. Then define a player's *maximin* to be the maximum of all the minimum values of their strategies.

Player 2

Player 1		L	R
	U	a	b
	D	c	d

Figure 3: A General Zero-Sum Game

For example, in this general zero-sum game, Player 1's maximin is $\max(\min(a, b), \min(c, d))$.

A player's minimax is the same expression as the maximin, except we switch the order of the maximum and minimum. In our example, Player 1's minimax is $\min(\max(a, c), \max(b, d))$.

Problem 7 For each player in Zero-Sum Games 2 and 3, and each strategy, what is the worst-case score? What scores will the players get if they both play the strategy with the highest minimum value (the *maximin strategy*)?

Problem 8 What happens in this game if both players play the strategy with the highest minimum value?

Player 2

Player 1		L	R
	U	3	-2
	D	-1	0

Figure 4: Zero-Sum Game 4

Try playing this game a few times with a partner. Can you find a way to do better than this maximin strategy?

Definition. An outcome is a *Nash equilibrium* if no player could improve their score by switching to a different strategy while everybody else keeps the same strategy.

Problem 9 Which of zero-sum games 2,3, and 4 have Nash equilibria?

Problem 10 What score can Player 2 guarantee if they cheat and find out what Player 1 is going to play?

Can you come up with an upper and a lower bound for the scores each player can hope to get if they're not sure what the other player will do?

Problem 11 In a two-player zero-sum game, what are the relationships between these four numbers?

- Player 1's maximin
- Player 1's minimax
- Player 2's maximin
- Player 2's minimax

Problem 12 Show that an outcome of a two-player zero-sum game is a Nash equilibrium if and only if both players are playing a maximin strategy and getting their maximin score. Conclude that all Nash equilibria in a particular game of this sort give the same scores.

Can you find a criterion on the four numbers from the previous problem, to determine when a Nash equilibrium exists?

2.2 Mixed Strategies

We now have an upper and lower bound for everybody's scores if they play rationally, but can we find exactly what everyone's scores will be under rational play when there is no stable outcome?

We can find a stable "outcome," where we can calculate everyone's scores and call those the real value of the game, but we will need to use probability.

Definition. If a player in a particular game has n strategies, $S_1, \ldots, S_n$, define a *mixed strategy* to be a list of nonnegative real numbers $(p_1, \ldots, p_n)$ such that $\sum_{i=1}^{n} p_i = 1$. To play a mixed strategy, play strategy S_i with probability p_i.

To contrast with mixed strategies, the strategies we're already familiar with can be called *pure strategies.*

Problem 13 For each pure strategy, explain how to construct a mixed strategy which is essentially the same.

Problem 14 Say that in Zero-Sum Game 3, Player 1 plays a mixed strategy of U with probability 2/3 and D with probability 1/3. Show that the expected value of Player 1's score when Player 2 plays L is the same as their expected score when Player 2 plays R.

Find a mixed strategy for Player 2 such that Player 2's expected score does not depend on Player 1.

Problem 15 For a general two-player zero-sum game where each player has two pure strategies, neither of which dominates the other, show that there is a Nash equilibrium with mixed strategies. (By this we mean a pair of mixed strategies, one for each player, such that neither player can improve their expected value by changing to a different mixed strategy, while the other player's mixed strategy stays fixed.)

Problem 16 Find the maximins, minimaxes, and a (mixed) Nash equilibrium for rock, paper, scissors above.

Theorem. *In fact, one can use Brouwer's Fixed Point Theorem to show that all games of this sort with finitely many pure strategies have a Nash Equilibrium, even if they are not zero-sum.*

3 Non-Zero-Sum Games

We can write a payoff matrix for a non-zero sum game as follows:

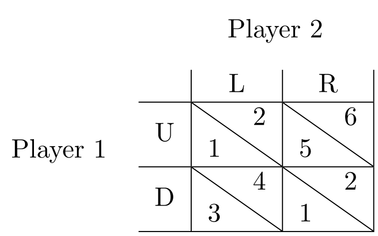

Figure 5: A Non-Zero-Sum Game

We divide each outcome's square into two pieces, where the lower left is Player 1's score and the upper right is Player 2's. For instance, if Player 1 plays U and Player 2 plays R, then Player 1 gets 5 point and Player 2 gets 6.

Problem 17 Given an n-player non-zero-sum game, explain how you can construct an $n+1$-player zero-sum game which is essentially the same.

Problem 18 Here are some famous non-zero-sum games, with some flavortext to explain their names. See if you can identify the Pareto optimal outcomes and (mixed) Nash equilibria:

Stravinsky-lover

		B	S
Bach-lover	B	10, 5	3, 3
	S	0, 0	5, 10

Figure 6: Bach or Stravinsky

Two people with strong preferences on classical music want to go to a concert together, but they have failed to coordinate and choose which concert to go to. They have to just pick one and go to it, and hope that their friend also attends.

Player 2

		D	H
Player 1	D	5, 5	10, 0
	H	10, 0	-2, -2

Figure 7: Hawks vs. Doves

This game is used in studying evolution of cooperation and aggressive behavior: the pure strategies are to be hawklike, and fight over resources, or dovelike, and share. Two doves will share their resources, a hawk will steal all the resources from a dove, and two hawks will fight each other, to the detriment of both.

Player 2

		S	H
Player 1	S	10, 10	0, 5
	H	5, 0	5, 5

Figure 8: Stag Hunt

Two hunters must choose whether to hunt a stag or a hare. They're guaranteed to catch a hare if they try, but if they try to hunt the larger, tastier, stag, they will only succeed as a team.

		Player 2	
		C	D
Player 1	C	-2, -2	-5, 0
	D	0, -5	-4, -4

Figure 9: Prisoner's Dilemma

Two criminals have been captured, and can either cooperate with each other, telling the police nothing, or defect, reporting their co-conspirator. If they both cooperate, they get a small sentence. If only one defects, they get a plea deal for reporting the co-conspirator, and serve no time. If they report each other, then they are both easily convicted and serve a long sentence.

Problem 19 Is a Nash equilibrium necessarily Pareto optimal?

Problem 20 Which of these games is most similar to 2-player Tragedy of the Commons?

4 Promises and Iterations

Everything we've said so far assumes that there is no communication prior to the game.

Problem 21 How do the non-zero-sum games above change if one player is able to declare their move ahead of time, and fully commit to it?

Problem 22 The games can also change if you play them against the same opponent repeatedly. Try playing Prisoner's Dilemma with a partner for a streak of 5 or 10 games at a time, and see if you can find strategies that give you scores better than predicted by the Nash equilibria. How do these strategies violate the rules for single games?