Game Theory: Lesson 1

Logan Crew and Sophie Spirkl

October 2021

1 What is a Game?

When you hear the word "game" a few particular games likely come to mind: Monopoly, Chess, Scrabble, Tic-Tac-Toe. While the rules of these games are very different, each involves multiple players who make their own decisions and try to win. They differ because for each game, its rules tell us which decisions are valid moves.

Many real-life situations can be thought of using the same framework. For example, when stores choose the price of a certain product, they want to make as much profit as possible, but the prices set by other stores also influence how well the product will sell. We want to provide a systematic way of analyzing games and the best ways to play them[1], and how this translates to both recreational games and "real-life" games.

Before we can analyze games, we need to be clear on what exactly a game even is! We need a precise, mathematical definition that we can work with. In what follows, we describe the basic properties of a mathematical game, using *Rock, Paper, Scissors (RPS)* as an illustrative example. A *game* consists of the following:

- A set of *players* who are participating. RPS is usually played with 2 people.
- A set of *strategies* for each player. When the game is played, each player chooses one of their strategies simultaneously[2]. In RPS, each of the two players chooses between three strategies: Rock, Paper, and Scissors.

[1] In other words: how to win!

[2] This may sound strange to you, because in most board games, players take turns. Our definition will be easier to analyze, and it is actually mathematically equivalent. Think of it this way, for Tic-Tac-Toe: if $\circ$ goes first and $\times$ goes second, then $\times$ could decide on a strategy that says "*In the first round, if $\circ$ doesn't take the middle square, I will take it. If $\circ$ does choose the middle square, I will pick the top left one. In the second round, . . .* " A strategy can be a complex series of "what-ifs" like that. For Chess, although such a strategy theoretically exists, you couldn't write down a good complete one even if you spent the rest of your life on it!

- A *payoff function* that takes as input one strategy from each player, and outputs to each player a number representing their *payoff*, or amount won. In a game like RPS, we typically represent a win by a payoff of 1, a loss by a payoff of -1, and a tie by a payoff of 0.

Sometimes the payoff represents an actual amount, like points in a game, or a gain or loss of money; sometimes the payoff can represent something more abstract, like happiness arising from a real-life decision, in which case numbers stand in to approximate the relative values of different outcomes[3].

Suppose that Rose and Colin are playing RPS. We represent the game by the following table. To play the game, you find Rose's choice of strategy on the left, and Colin's choice of strategy up top, and locate the corresponding table entry. To read the payoff, interpet the ordered pair as (Rose's payoff, Colin's payoff):

Rose \ Colin	Rock	Paper	Scissors
Rock	$(0,0)$	$(-1,1)$	$(1,-1)$
Paper	$(1,-1)$	$(0,0)$	$(-1,1)$
Scissors	$(-1,1)$	$(1,-1)$	$(0,0)$

Table 1: Payoffs in the game of Rock, Paper, Scissors.

For example, suppose that Rose chooses Rock, and Colin chooses Scissors. We find the corresponding entry is $(1,-1)$, indicating that Rose gets a payoff of 1 for winning, and Colin gets a payoff of -1 for losing.

Another example is the game of Chicken, a reckless game in which Rose and Colin head towards each other on bicycles and the first to swerve (to avoid a collision) "loses". To simplify things, let us assume both bikes keep heading towards each other until their last chance to swerve. Then the payoff table might roughly look like this:

Rose \ Colin	Swerve	Straight
Swerve	$(-1,-1)$	$(-3,1)$
Straight	$(1,-3)$	$(-1000,-1000)$

Table 2: The game of Chicken.

Note that for this game, some of the payoffs do not correspond to a monetary amount, nor are just winning or losing, so they instead roughly correspond to how the players feel about the outcome. For example, if both swerve, then while the game is a draw in some sense, both feel like they chickened out a bit, and so both receive a negative payoff rather

[3]Numbers are useful, because we can compare them. If the payoffs were "you win a book" and "you win a box of chocolates," it would be harder to tell mathematically which outcome we should strive for!

than 0. If one of them swerves, that person feels a bit worse, and the other "wins". If neither swerves, then they crash into each other, likely wrecking their bikes and injuring each other. Note that in fact each pair has a negative sum, so the net effect of playing this game on the players is always negative, which doesn't seem like a good game to play!

Note that in general there may be more than two players, or even infinitely many strategies (for example when ordering or using an amount of a material) in which case payoffs cannot be represented by a table. In these cases, remember that there is still a payoff function that represents the payoff to each player.

Note: Before proceeding to the next section, if you want to work through more examples of games, see Problem 1 at the end of the lesson.

2 Analyzing Games

The goal of studying games mathematically is to understand the best way to play them, whether the players would actually play them the way that our mathematical analysis suggests, and the implications of this. The second and third of these questions are generally more in the realm of psychology and economics respectively, so mathematically it is sensible to focus on the first and most direct question: what is the optimal way to play a game, given that you do not know what strategies the other player(s) will pick?

2.1 Pure Strategies

The simplest way to start is to compare strategies against each other directly. Consider the following game between Rose and Colin, each of whom has two possible strategies labelled A and B:

Rose \ Colin	A	B
A	$(1, -1)$	$(0, 0)$
B	$(-2, 2)$	$(-1, 1)$

Table 3: A simple game.

Let us start by looking at the game from Rose's point of view, looking at the first number in each pair since this is her payoff. If Colin chooses A, then Rose gets a payoff of 1 with her strategy A and -2 with B. On the other hand, if Colin chooses B, Rose gets payoff 0 going with A and -1 going with B. No matter what strategy Colin chooses, Rose does better choosing A than choosing B, so she never has a reason to choose B. When this happens, we say that strategy A *dominates* strategy B, meaning that the payoff from choosing A is always greater than or equal to the payoff from choosing B no matter what the other player(s) do.

Note that the same is not true for Colin: either A or B could give a better result for Colin depending on what Rose plays. However, by the above analysis Colin knows that Rose will never play B, so we can simply remove that entire row from consideration. And in what is left, namely the case where Rose chooses A, it is clear that Colin should choose B. Thus each player has a unique strategy that they should always choose, so this game has a *pure-strategy solution*.

Looking for dominant strategies is a good starting point for analyzing a game. For Colin, neither strategy dominates the other in the original game, but one *does* dominate once we remove Rose's dominated strategy from consideration. Thus, analyzing games for dominant strategies may be a multi-step process.

2.2 Mixed Strategies

Of course, many games do not have any pure-strategy solutions. Think about Rock, Paper, Scissors from before. In this game, Rose's optimal strategy may be any of Rock, Paper, or Scissors depending on what Colin chooses. In fact, it is intuitive to think that if you play this game multiple times you want to play each strategy roughly $\frac{1}{3}$ of the time rather than commit more than that to any one strategy.

To analyze these sorts of games, we assume that each player will play each of their strategies with a certain probability, and we will then analyze the *expected payoff*. This is called a *mixed strategy*. To describe a mixed strategy for a player, we assign a probability to each of their normal (pure) strategies and make sure that the probabilities add up to 100%. Note that picking one strategy 100% of the time (a pure strategy) is also a mixed strategy.

If a game doesn't have a pure-strategy solution, the next thing we can look for is a natural solution in mixed strategies for both players. What exactly constitutes an optimal mixed-strategy solution?

Let us use Rock, Paper, Scissors as an example. Suppose that Rose has decided to play Rock 60% of the time and Paper 40% of the time (and Scissors 0% of the time). Suppose that Colin is playing Rock 25% of the time and Scissors 75% of the time (and Paper 0% of the time). To calculate Rose's expected payoff, we look at each possible pair of a strategy by Rose and a strategy by Colin. We take the payoff to Rose from this choice, and multiply it by the probability that this pair of strategies occurs.

For example, in the case where Rose chooses Rock and Colin chooses Scissors, the payoff is 1 to Rose, and this pair of strategies happens with probability $(0.6)(0.75) = 0.45$, or 45% of the time. Summing over all of the possible strategy pairs, take the time to verify that Rose's expected payoff is

$$0(0.6)(0.25) + 1(0.6)(0.75) + 1(0.4)(0.25) + (-1)(0.4)(0.75) = 0.25$$

so overall Rose is expected to win slightly more than she loses.

Now of course, it is not hard to see that Colin can do better. For example, since Rose never plays Scissors, Colin can never lose if he always plays Paper, so he could switch to 100% Paper and have a better expected payoff for him.[4] Of course, Rose could switch her strategy in response to this, and so on and so forth. But is it possible to prevent this process? It would seem that truly optimal play must require a player to use a mixed strategy that their opponent cannot counter, *even if they know it*. This inspires the following definition:

Definition 1. *A choice of mixed strategies for each player in a game is called a* Nash equilibrium *(or just* equilibrium*) if no player can improve their expected payoff by choosing a different mixed strategy (while keeping the other players' mixed strategies the same).*

We have seen a Nash equilibrium in the game described by Table 3. If Rose chooses A and Colin chooses B, then both get a payoff of 0. Here, Rose doesn't want to switch to B, because if Colin sticks with his strategy B, this would give Rose payoff -1, which is worse. Colin doesn't want to switch to A, because if he switches to A and Rose sticks with her strategy A, Colin would get a worse payoff of -1. For each player, we only consider if they can improve their own payoff by only changing their own strategy. If they cannot, as is the case when Rose picks A and Colin picks B, we have a Nash equilibrium.

If you would like some practice with this concept before continuing, you can try Exercises 3 and 4.

The concept of a Nash equilibrium was introduced and popularized by John Nash, a Princeton mathematician who was a pioneer in game theory. He famously[5] proved that Nash equilibria exist:

Theorem 2. *Every game with a finite number of players and a finite number of strategies has at least one Nash equilibrium (possibly in mixed strategies).*

What is a Nash equilibrum of the above game? It is not hard to guess that both players choosing each strategy 1/3 of the time is an equilibrium, and indeed we can prove it. Suppose Rose knows that Colin is playing this strategy. What is her expected payoff if she plays Rock with probability r, Paper with probability p, and Scissors with probability s (so that $r + p + s = 1$)? Then her expected payoff is

$$\frac{1}{3}(r(0) + p(1) + s(-1)) + \frac{1}{3}(r(-1) + p(0) + s(1)) + \frac{1}{3}(r(1) + p(-1) + s(0)) = 0$$

which is in fact independent of the choice of r, p, and s. So if both players are playing with this mixed strategy, neither can improve their expected payoff by switching to another strategy.

[4] Calculate this and verify that Colin's expected payoff increases from -0.25 with the current set of strategies to 0.6 if Rose continues with the same mixed strategy, but Colin switches to playing Paper 100% of the time.

[5] He proved this in his doctoral dissertation, which has only 26 pages!

However, if either player is not playing the equal probability mixed strategy, then some player can change their strategy to improve their expected payoff, so the Nash equilibrium for this game is unique. As an exercise, try to prove this formally![6]

For Rock, Paper, Scissors, we could make an educated guess as to the Nash equilibrium and prove that it works. But how do we find one when it isn't obvious? Let us work through the game of Chicken mentioned before:

Rose \ Colin	Swerve	Straight
Swerve	$(-1, -1)$	$(-3, 1)$
Straight	$(1, -3)$	$(-1000, -1000)$

Table 4: The game of Chicken again.

Let r be the probability that Rose selects Swerve, and likewise let c be the probability that Colin selects Swerve (and thus $1-r$ and $1-c$ are the respective probabilities they stay going straight). The expected payoff to Rose is

$$-rc + (1-r)c - 3r(1-c) - 1000(1-r)(1-c) = 997r - 999rc + 1001c - 1000.$$

Now, Rose can't control the portion of her payoff given by $1001c - 1000$, so she can only try to maximize $997r - 999rc = r(997 - 999c)$. Thus, whether Rose wants r to be high or low depends on whether $997 - 999c$ is positive, negative, or zero. Solving, we may verify that the quantity is zero when $c = \frac{997}{999}$. So there are three cases:

- $c < \frac{997}{999}$. Then Rose wants r as high as possible, so in this case we can't have equilibrium unless $r = 1$.
- $c > \frac{997}{999}$. Then Rose wants r as low as possible, so in this case we can't have equilibrum unless $r = 0$.
- $c = \frac{997}{999}$. In this case the expected payoff is independent of r.

Note that the strategies and payoffs are symmetric[7], so we may apply the same analysis to Colin. From this, we may conclude that no equilibrium is possible unless at least one of r and c is one of the critical numbers 0, 1, or $\frac{997}{999}$.

From here we may proceed by cases. Suppose first that r is a critical number. If $r = 0$, then by the above rules and symmetry we must have $c = 1$, and clearly this is an equilibrium (in pure strategies!). Likewise, if $r = 1$ then we must have $c = 0$. If $r = \frac{997}{999}$, then we may

[6]Hint: You can assume that one player has r, p, s not all equal to $1/3$, so look at which one is largest.

[7]It may not always be clear when we can assume something by symmetry, so if you're not sure, you can write out the argument for Colin and check that it actually does come out the same as for Rose, only with r and c switched.

check that we cannot have equilibrium unless also $c = \frac{997}{999}$. If indeed $r = c = \frac{997}{999}$, then Rose's payoff does not depend on r, and Colin's payoff does not depend on c, so it is a Nash equilibrium. Since in all three cases c is also a critical number, it follows that for equilibrium r and c must both be critical, and so our analysis is actually complete.

To summarize, the game of Chicken has three equilibria:

$$(1,0), (0,1), \left(\frac{997}{999}, \frac{997}{999}\right).$$

What do these equilibria actually tell us about the game? The first two tell us that if you're playing a game of Chicken, and you know for a fact that your opponent will go straight, you should always swerve. If you know that your opponent will definitely swerve, you may as well go straight. If you're pretty sure that your opponent will swerve, but not entirely sure, then both of your strategies look about equally good.

Here, you can also see that not all Nash equilibria are created equal: for each of the Nash equilibria above, compute Rose's expected payoff and Colin's expected payoff. You will see that one of the Nash equilibria is better for Rose, one is better for Colin, and one is okay but not great for both.

2.3 Lesson 1 Exercises

1. Write each of the following situations as a game, explicitly describing the players, strategies, and what the payoffs represent.

    (a) Rose and Colin both pick either 1 or 2. If the sum is odd then Rose wins, and if the sum is even then Colin wins.

    (b) Rose and Colin both choose to cultivate 10, 20, 30, 40, or 50 % of a shared field, and get back a harvest with value equal to the percent chosen. However, if they use at least 80% of the field combined, then it will not grow back the following year, and both players suffer a penalty of value 100.

    (c) Rose, Colin, and Larry attend an auction for a single item. Rose values the item at \$20, Colin at \$25, and Larry at \$30. Each one submits a bid secretly, and the highest bid wins the item and pays their bid for it (with ties broken in favor of Larry, then Colin, then Rose).

2. Simplify the two games below by crossing out a strategy if there is another that dominates it. How can you verify that you have simplified a game as much as possible using this method?

Rose \ Colin	A	B	C
A	$(3,-1)$	$(2,0)$	$(-1,4)$
B	$(2,2)$	$(1,6)$	$(1,0)$
C	$(4,1)$	$(-1,3)$	$(0,-1)$

Rose \ Colin	A	B	C
A	$(2,1)$	$(3,0)$	$(2,4)$
B	$(4,2)$	$(1,6)$	$(1,3)$
C	$(3,3)$	$(-1,1)$	$(0,-1)$

Table 5: Two games.

3. (a) Consider the game given by the table below. You can think of it as a bidding game: Rose and Colin both bid 1, 2, or 3 dollars; and the higher bid wins (and gets the money from both players). However, if they bid the same amount of money, they both lose 1 dollar.[8] Which pure strategy Nash equilibria[9] does this game have?

Rose \ Colin	1	2	3
1	$(-1,-1)$	$(-1,1)$	$(-1,1)$
2	$(1,-1)$	$(-1,-1)$	$(-2,2)$
3	$(1,-1)$	$(2,-2)$	$(-1,-1)$

Table 6: Bidding game.

[8] Maybe they'll donate the money to a charity in this case.

[9] That is, Rose and Colin each choose one of their strategies with probability 100%.

(b) Here is another game. Rose and Colin are giving class presentations. They both prefer topic A. However, if both talk about the same topic, then their classmates will get bored, so they could each switch to topic B instead. This game shows how

Rose \ Colin	A	B
A	$(-1,-1)$	$(2,0)$
B	$(0,2)$	$(-2,-2)$

Table 7: Presentations.

vastly different Nash equilibria can be. Show that these are all Nash equilibria, and compare the payoffs:

- Rose picks A, Colin picks B.
- Colin picks A, Rose picks B.
- Rose and Colin each pick A with probability 80% and B with probability 20%.

4. Find at least one Nash equilibrium for each of the games in Exercise 1 (you do not need to try to find all of them), and explain how you know it is a Nash equilibrium.

5. In doing the previous problem, you may have noted for the game described in 1(c), each player bidding how much they value the item is *not* an equilibrium (try checking this if you have not already done so!).

   Consider a variant in which the player making the largest bid wins the item, but pays the *second-largest* bid for it. Show that for this modified game (known as a *Vickrey auction*), each player bidding how much they value the item is now an equilibrium.

6. A Nash equilibrium is called *stable* if when one player deviates slightly from the best probabilities, no other player has an incentive to change their strategies. Otherwise it is called *unstable*.

   Is the Nash equilibrium for Rock, Paper, Scissors stable or unstable? Which of the three Nash equilibria for the game of Chicken are stable and which are unstable?[10]

7. The following game is called the *Prisoners' Dilemma*.

   Criminals Bonnie and Clyde have finally been caught running from the police after robbing a bank. Unfortunately for the cops, they can't prove that Bonnie and Clyde are the serial bank robbers, only that they ran from and shot at police officers, which will come with a penalty of one year in prison. So they offer each of them a deal: to betray their partner and tell the police everything. If one of them takes the deal and the other doesn't, the one who tells the police everything will be released, and the other will be imprisoned for ten years. If both take the deal, both will be imprisoned for the robberies, but with a lighter sentence of six years for cooperating.

   The payoff table is given on the next page:

[10] Hint: For the game of Chicken, there should be two stable ones and one unstable one.

Bonnie \ Clyde	Say nothing	Take deal
Say nothing	$(-1,-1)$	$(-10,0)$
Take deal	$(0,-10)$	$(-6,-6)$

Table 8: The Prisoners' Dilemma.

(a) Find the unique Nash equilibrium of this game.

(b) Find strategies for Bonnie and Clyde that give both of them a strictly better payoff than the Nash equilibrium.[11]

(c) Suppose that four players Jessie, James, Butch, and Cassidy each play the game once against each other player (so there are a total of six games, one for each pair). Suppose that Jessie and James both always pick "say nothing" and that Butch and Cassidy both always pick "take deal". For each player, state what their overall payoff is after all of their games (that is, the sum of their payoffs from each game). What if each player plays a 10-game series against each other player?

(d) Suppose that these four players again play this game 10 times against each other player, but Jessie and James have come back with a new strategy that changes from game to game within a series: in the 10-game series against a given player, they always choose "say nothing" in the first game, and then in each subsequent game they choose whatever their opponent chose the previous game. This multi-game strategy is called *tit-for-tat*. Explain how the games go, and give the total payoff to all players at the end.

8. Suppose that Rose and Colin play a game where each chooses a positive integer. Suppose that if Rose chooses r and Colin chooses c the payoff to Rose is $r - c$ and the payoff to Colin is $c - r$. Explain why this game has no Nash equilibrium (Hint: there is a lowest number Rose chooses with positive probability. Could she do better?). Why does this not violate Theorem 2?

9. Challenge Problem: Let $1, 2, \dots$ be infinitely many climbers, each of whom is considering climbing Mount Everest. They each decide whether or not to climb the mountain in order, first 1, then 2, and so on.[12]

   Now, climbing Mount Everest is a lot of work, and is not really worth it unless you are either the first or the last person to ever climb the mountain. Since the first to climb already occurred (shoutouts to Edmund Hillary and Tenzing Norgay), it is only worthwhile to be the last.

   Thus, each climber i may choose to ascend Everest or to stay off of the mountain. The payoff to a given i is then given as follows:

[11] This is why they always put people in different rooms when offering them a deal: if Bonnie and Clyde could agree to play these better strategies beforehand, they would!

[12] For example, perhaps 1 was considering this in 2001, and 2 in 2002, and so on. Assume for this problem that we expect the Earth and humanity to last forever.

Other choices \ i's choice	Nobody after i ascends	At least one person after i ascends
Ascend	100	-100
Stay off	0	0

Table 9: The Everest Climbing Game.

Two examples are given to illustrate. First, suppose that only climbers 1, 4, and 9 choose to ascend. Then only these climbers get nonzero payoff. Climbers 1 and 4 each get -100, and 9 gets payoff 100.

Second, suppose that all even numbered climbers choose to ascend. Then all odd numbered climbers get payoff 0, and all even numbered climbers get payoff -100, since for each even climber, there is always a higher-numbered climber who also climbs.

We will show, step-by-step, that this game has no Nash equilibrium.

(a) Let us first consider pure strategies, that is, each climber chooses "ascend" or "stay off." Show that if finitely many climbers choose to ascend Mount Everest, it is not a Nash equilibrium.[13] Now show that if infinitely many climbers choose to ascend Mount Everest, it is not a Nash equilibrium.[14]

(b) Now let us assume that there is a Nash equilibrium with mixed strategies. In this strategy, let x_i be the probability that i will decide to ascend Mount Everest. By (a), we can assume there is a player i with $0 < x_i < 1$. Show that this means that the probability that at least one of the climbers after i chooses "ascend" is exactly 50%.

(c) Continuing from (b), since the probability that i is the last to ascend has to be exactly 50%, conclude that there is someone after i, say j with $j > i$, who also played a mixed strategy $0 < x_j < 1$.

(d) Continuing again, this means that the probability that at least one of the climbers after j ascends is also exactly 50%. Why is this impossible?

(e) Conclude from (a)-(d) that there is no Nash equilibrium for this game. Why does this not violate Theorem 2?

[13]The advantage of having finitely many people ascending Mount Everest is that there is a last person, say i, who chooses to ascend. But, knowing this, who should change their strategy?

[14]If infinitely many people all choose "ascend", then none of them the is last one to do so!