Game Theory: Lesson 2

Logan Crew and Sophie Spirkl

October 2021

1 Extensive-Form Games

So far all of the games we have studied have consisted of each player choosing a single strategy simultaneously, and then computing corresponding payoffs. However, we have already noted that this does not capture the usual notion we have when we hear the word "game". We have already discussed one issue, which is that while it is possible to do so, it is weird to write Tic-Tac-Toe, Chess, or other games in which players alternate turns in this form. In these cases, it is more natural to write a game in a way that naturally illustrates the sequence of events.

We will show how to do so using a simple example. Suppose that Rose is given the opportunity to take an item of payoff r to her. If she accepts, the game ends. If she declines, the same item is offered to Colin, and he can accept it with payoff c to him. Regardless of whether Colin accepts or declines, the game ends.

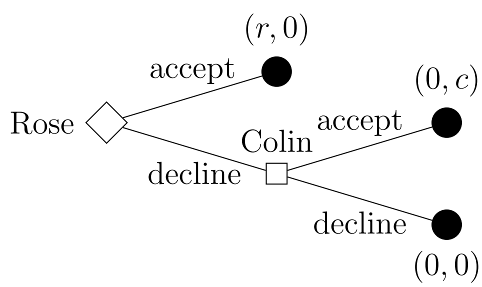

Figure 1: A simple extensive-form game.

In Figure 1, we have drawn the *tree* corresponding to the game we described. We play the game left-to-right by starting at the left-most point, or *node*. This node is labelled "Rose", representing that at the start of the game, it is Rose's turn to make a decision. There are two *edges* going to the right from Rose's node, indicating that she has two choices: to either accept or decline the offered item.

If she accepts the item, we follow the edge marked "accept" to the next node. In this case, we reach a *terminal node* (a black circle), one with no edges going right, so the game has ended. The node is labelled with the payoffs to each player in the form (Rose's payoff, Colin's payoff). In this case, Rose gets the payoff of r for accepting the item, and Colin is unaffected.

If she instead declines the item, we follow the edge marked "decline" to the next node. This is not a terminal node, and instead it is now Colin's turn, as the node indicates. From here, Colin can choose either "accept" or "decline". In either case, following the corresponding edges leads to a terminal node, and the game ends. The overall diagram is called the *game tree*, and this representation of the game is called its *extensive form*.[1]

Drawing a game in this form has a number of advantages. For one, this more clearly illustrates how we actually think about games with sequential decisions. This also allows for a more natural way to include *randomness* as part of our games. For example, when we roll the dice in Monopoly, we are letting a random event impact the game. In principle, this can be included in each player's strategy considerations when viewing the game in normal form (as in the first lesson), but it is easier to visualize in extensive form. To do so, we pretend that *nature* is another "player" in the game, but one who gets no payoff. We assign some nodes to nature, where the edges leading to the right from them represent choices made randomly by nature. We mark nature's nodes with a *grey* circle, and we label each of a nature node's edges with the probability p that it is the chosen strategy such that the sum of all probabilities heading out of a nature node is 1.

As an example, let us consider calling a coin flip. The game starts with Rose choosing heads or tails, and then Colin flips a coin. Even though Colin is the one flipping the coin, assuming he is not cheating the actual result of the coin flip is random, so this is a play by nature. Then either Rose or Colin wins depending on the result of the flip. This game is illustrated in Figure 2.

The game in 2 is actually a one-player game (since nature doesn't really count). What are its Nash equilibria?[2]

If you wish to practice more with understanding extensive games before continuing, you can try Exercises 1 and 2 now at the end of the lesson.

[1] So to write a game in extensive form, you write the steps left-to-right. Each node is either a terminal node (a black circle) labelled with the payoffs of the players, or a decision node (some other shape filled in white) labelled with the player whose turn it is. From a decision node, the edges (lines) going to the right are labelled with the possible choices that the player can make at that decision node.

[2] Again, since nature doesn't count, we only consider Rose's strategies.

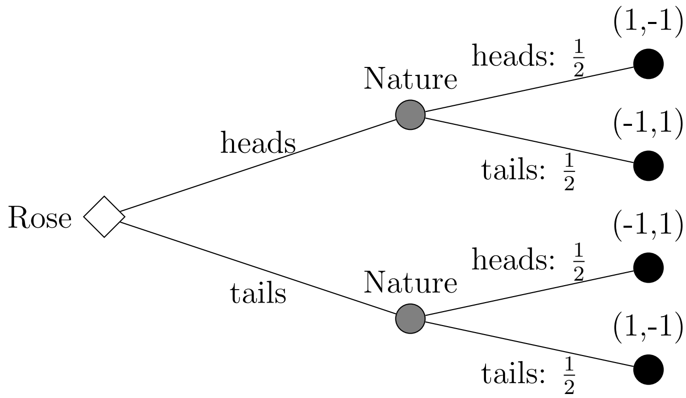

Figure 2: An extensive-form coin flip game.

2 An old hat game

Let us consider[3] the following game: Rose, Colin, and Larry are sitting in a circle. The game master Mario puts either a black hat or a white hat on each of their heads, and each player has exactly a 50% of chance of getting either colour chosen randomly. Each player can see the colour of the the hats on the other two players' heads, *but not the colour of the hat on their own head*. There are three players[4] (Rose, Colin, Larry) and eight possible configurations of hats.

Simultaneously, all three of them either try to guess the colour of their own hat ("White" or "Black") or say "No Guess". The game is cooperative: all three players win if

- At least one player correctly says the colour of their own hat, **and**
- No player guesses incorrectly.

Conversely, all three players lose if

- At least one player guesses incorrectly, **or**
- All three players say "No Guess."

The players would like to work together to find a strategy that allows them to win as often as possible. Clearly they can win with probability 50%: just have Rose guess "White" and each of the others say "No Guess". Can they do better?

[3] There are several versions of this game, often given as a riddle. Feel free to come up with your own variants; you could add more hat colors, more players, more hats per person, and more restrictions on who can see what.

[4] Mario, as the game master, is considered as a force of nature in this game.

The answer, perhaps surprisingly, is yes! Let the players all agree in advance that they will each do the following: Say "No Guess" unless two hats of the same colour are on the other players' heads, in which case guess the *other* colour. Let us briefly explain how this works. For brevity, let us express hat configurations with a three-letter sequence, where each letter is W or B corresponding to a white hat or a black hat, and the first letter represents Rose's hat colour, the second Colin's, and the third Larry's.

First, suppose the configuration is WWB.[5] Then Rose and Colin both see two hats of different colours, so they say "No Guess". Larry sees two white hats, and so guesses "Black". Since Larry is correct, this results in a win for the three players. You can easily check that the three players also win for the configurations WBW, BWW, BBW, BWB, WBB.

In the remaining cases, WWW and BBB, every player sees two hats of the same colour, and every player then guesses the other colour, which is wrong. So in these two cases, every player guesses wrong, and the players lose.

Overall, with this strategy the players win on 6 out of 8 configurations, or 75% of the time![6] Exercise 3 will ask you to show that this is the best possible rate of success for this game.

3 The hat game as an extensive game: imperfect information

The hat game is very interesting as a game and as an exercise in creative thinking! But we bring it up because it is also illustrates an interesting wrinkle in writing extensive-form games. The first move may be considered to be by Nature, assigning hats to the players randomly. Now, we need the consider the players as guessing sequentially to write them as separate entities in extensive form. So to simulate this, let us picture that the three players look at each other, and then Rose writes her guess on a piece of paper and puts it in a leftover hat,[7] followed by Colin, and then Larry. This is effectively the same as if they all shout their guesses simultaneously but is better suited for extensive form.

The beginning of the game tree is shown in Figure 3. The game's first move is made by nature, which picks one of the eight possible configurations of hats, each with probability $\frac{1}{8}$. Each of the edges leads to a node in which it is now Rose's turn to guess (or not). We label these nodes with the configuration leading to it in the order Rose's hat, Colin's hat, Larry's hat as before.

[5]So Rose and Colin have white hats, and Larry has a black hat.

[6]If you'd like to think about this game with more players, the seven-player version has a nice answer where players can win in all but 16 out of the 128 configurations.

[7]We assume that Mario owns enough hats that the colour of the leftover hat doesn't tell the players anything.

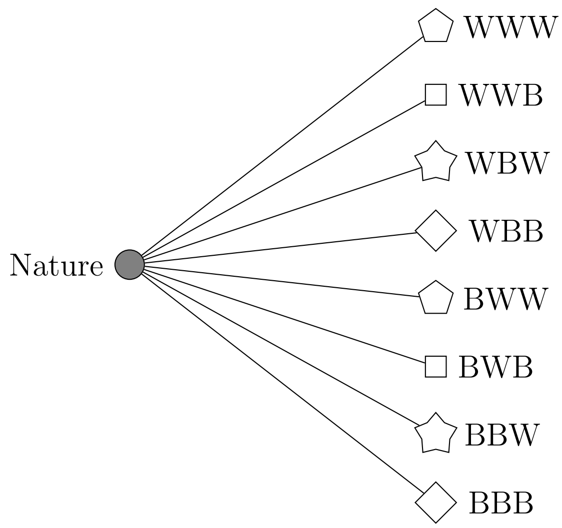

Figure 3: The beginning of the hat game.

Now, here we run into an issue we have never previously had. Before, when choosing how a player acts in an extensive-form game, we would treat each node separately, and analyze the choices there. In particular, if the same player has multiple nodes in the tree, they can consider each of them separately, and may make different decisions at each of them.

But consider at the start of the hat game in Figure 3 the two nodes marked with a pentagon. From Rose's perspective, in both cases she sees a white hat on each of the other two players' heads. Can she act separately at each pentagon node? Intuitively she cannot, because if she could, **then she could distinguish the colour of hat on her own head**! The issue is that Rose has *imperfect information* about the game: she does not know all of the choices made previously. In this case, she does not know the exact path that Nature chose at the start. This is in contrast to other extensive-form games we have studied so far, which all have had *perfect information*, meaning that every player knows exactly where they are in the game tree when they make a decision.[8]

How we have chosen to draw the start of the game tree in Figure 3 hints at a way to model imperfect information: we can group a player's nodes based on whether they can be distinguished by that player or not. In the hat game, Rose has eight nodes, but they come in pairs where Rose cannot distinguish between the two nodes in a pair (cases where the other players have the same configuration of hats). These pairs are marked in Figure 3 as nodes that have the same shape. A set of nodes for a player that are indistinguishable to that player is called an *information set*. So, in playing the game from Figure 3, Rose has to make her move knowing the shape of the node at which we are (so pentagon, square, star, or diamond) but not the precise node.

[8]Many real-life games have imperfect information: for example, if there are cards involved, we usually know our own cards but not the ones that other players are holding. On the other hand, Chess is an example of a game with perfect information.

We can model extensive-form games with imperfect information by requiring that a player's strategy is identical on every node in each information set. In the example of the hat game, we require that Rose always uses the same strategy to make a decision at either of the two pentagon nodes, whether that is always choosing a fixed option or choosing a mixed strategy among her set of options from these nodes. In other words, Rose's strategy should read, "if at a pentagon, do X; if at a square, do Y" and so on. Try drawing more of the game tree: at each of our eight nodes for Rose's turn, Rose has three choices ("White," "Black," "No guess") and therefore, the next level should have $8 \times 3 = 24$ nodes. Since Colin doesn't learn what Rose has guessed, these 24 nodes should be grouped into four information sets (each representing one of the four configurations of hats that Colin can see). For Larry's turn, there will be $24 \times 3 = 72$ nodes, each with three choices he can make. Still, Larry's nodes consists of four information sets as well.[9] In total, we would have $1+8+24+72+216 = 321$ nodes.

4 The Monty Hall game

A discussion of extensive-form games with imperfect information would be incomplete without talking about the famous Monty Hall game. Monty Hall was was the original host of the TV game show *Let's Make a Deal*, and at the end of each episode he would invite an audience member as a contestant to play the following game:

1. Monty shows the contestant three enormous closed doors, and tells them that behind one of the doors is a prize (usually a vacation, car, or something of the sort), and behind the other two doors are goats, representing a loss. The doors have the numbers 1, 2, and 3 on them. The door containing the prize has been randomly chosen. Monty knows which door has the prize; the contestant of course does not.

2. Monty instructs the contestant to pick one of the doors.

3. Then, Monty opens one of the two doors the contestant did not choose to reveal a goat.

4. He then asks the contestant if they would like to stick with the door they initially chose, or switch to the remaining unopened door. In either case, the contestant wins the prize if it is behind the final door they choose.

The Monty Hall problem asks if the contestant will be better off switching doors, and in particular what the probability is that they win if they switch.

Most people presented with this problem reach the conclusion that the probability of winning is $\frac{1}{2}$ regardless of your final choice, because there are two unopened doors at the end and it seems equally likely that the prize is in either of them.

[9] One thing to note here is that we don't care what Colin knows when it's Rose's turn. So only the nodes where it's Rose's turn to act are partitioned into information sets for Rose.

But this is incorrect! In fact, it is better to switch, and if you switch your chance of winning is $\frac{2}{3}$. We first explain why intuitively, and then represent the problem as an extensive-form game.

The reason has to do with the host's extra information. Regardless of whether you pick the correct door initially, the host can always pick a door not containing the prize to open in Step 3 of the game, so we always proceed to Step 4 no matter what. Thus, whether you have the correct door at the beginning of Step 4 depends only on whether you chose it correctly in Step 2. But clearly the probability you choose the correct door at the start is $\frac{1}{3}$, so since there is now only one other door, the probability of winning if you switch to this door must be $\frac{2}{3}$.

There is one slight caveat to this argument: we implicitly assume that if you do happen to choose the correct door at the beginning that Monty is equally likely to open either of the other two doors in Step 3. If this is not the case, then the problem changes, as we will see in the exercises!

We give the extensive form of the Monty Hall game in Figure 4. The grey circle nodes are where Nature (here played by Monty) acts, the white nodes are where the contestant acts, and the terminal nodes are black for a win and red for a loss (we could assign them explicit payoffs, but it doesn't change the optimal strategy as long as black nodes give higher payoff than red ones). Each terminal node of the game is marked by a sequence of four numbers (one for each step of the game) representing in order Nature's choice of door for the prize, the contestant's initial choice of door, Monty's choice of a door to open, and the contestant's final choice of door. Internal nodes are labelled with a partial sequence containing the information determined prior to that node. The contestant's nodes of the same shape are in an information set together, and are indistinguishable from each other (from the point of view of the contestant).

The left-most node, r, is the beginning of the game. Nature chooses the the winning door, 1, 2, or 3, each with probability 1/3 or $33.\overline{3}\%$. Let's say Nature chooses the second door; then the game proceeds to the middle square node, r_2. Now the contestant chooses a door; let's say the contestant picks the first door. This means that we're moving along the edge labelled "1" from r_2 to r_{21}. Here, Nature acts, but since the contestant chose the first door, and the second holds the prize, there is only one choice for Nature (Monty): open the third door! This moves us to r_{213}, one of the nodes marked with a five-pointed star. The contestant only knows that we're at one of r_{213} and r_{113} (in the upper branch), because the contestant does not know if the winning door is 1 or 2. However, the contestant has learned that switching is a better strategy, so from pentagon, the contestant picks door 2. In our play-through, this moves us from r_{213} to r_{2132}, a winning node!

The game tree gives us a more rigorous way of analyzing the Monty Hall game, which we will do in Exercise 6.

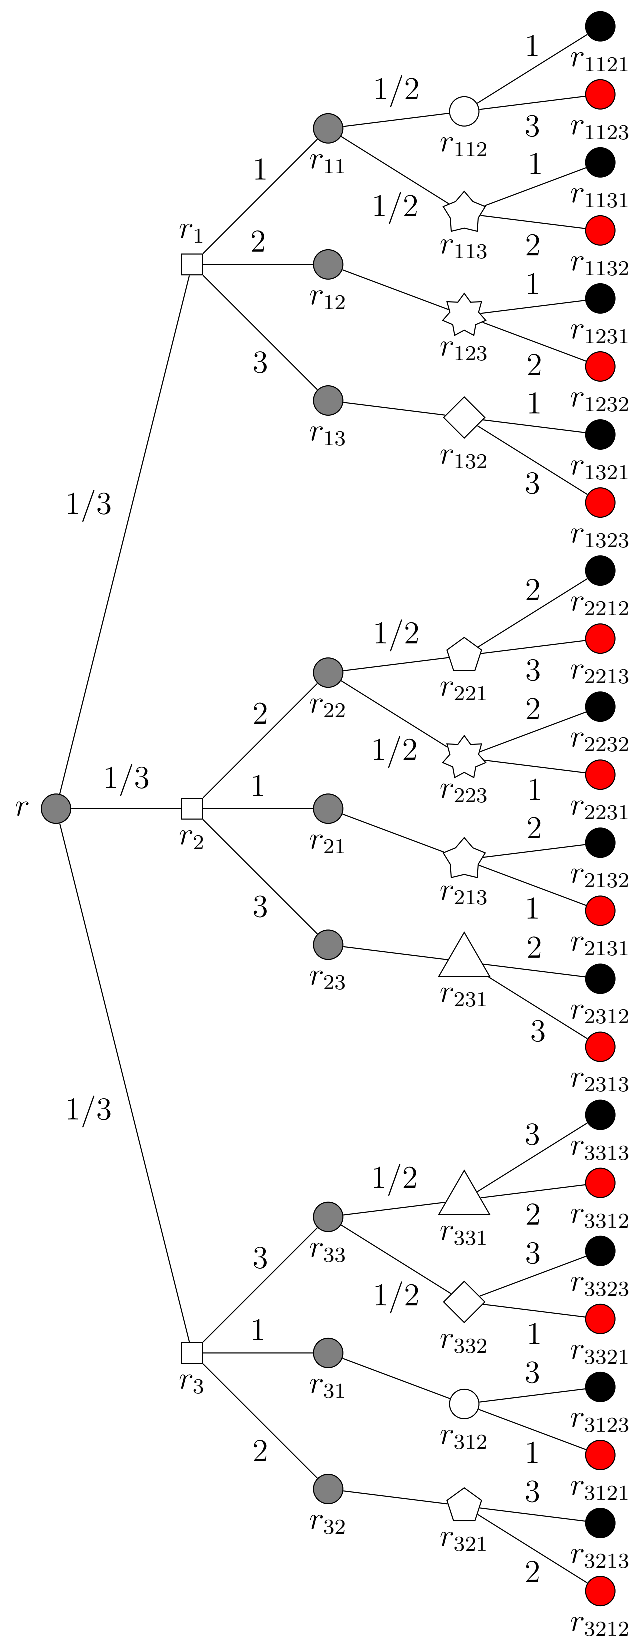

Figure 4: The game tree for the Monty Hall game.

5 Lesson 2 Exercises

1. Write the following as an extensive-form game:

   Rose and Colin have two pots in front of them, one red and one green. Initially the green pot has \$2 in it and the red pot has \$0. Rose takes the first turn. On a player's turn, they can choose between "continue" and "end". If they select "continue", \$1 is placed into each pot, and it is now the other person's turn. If they select "end", the game ends, and the player selecting "end" gets the money in the green pot, and the player selecting "continue" gets the money in the red pot.

   However, the game has finite length, so there is one exception. If it is the 10^{th} turn, and Colin (whose turn it must be) chooses continue, the game still ends, but Rose receives the money in the green pot, while Colin receives the money in the red pot plus a \$1 bonus.

2. Write the following as an extensive-form game involving nature:

   Rose and Colin both start with zero points, and Rose plays first. The game works as follows:

   - Rose flips a fair coin. If she flips tails, her turn ends. If she flips heads instead, then she scores a point. She now has the option to either flip again, or end her turn with 1 point. If she flips again and gets heads again, she scores another point, and the game ends immediately. If she flips tails on this second flip, her turn ends and she loses the first point she won (so she ends her turn with 0 points).
   - Once Rose's turn ends, if the game did not end it is now Colin's turn. Repeat the above process for Colin until the game ends or his turn ends (at which point the game also ends).

   If Rose scored r points and Colin scored c points, Rose gets a payoff of $r - c$ and Colin gets a payoff of $c - r$.

3. For this exercise, we show that in the hat game there is no pure strategy that succeeds more than 75% of the time.

   (a) Show that if any player makes a guess in at least two of their information sets, the probability that this player guesses wrong is at least 25%. Since the game is a loss if one player is wrong, such a strategy loses at least 25% of the time.

   (b) Suppose that instead each player makes a guess on at most one of their information sets. Show that in this case there are at least 2 configurations of hats in which no player guesses (so the probability that the players lose because none of them made a guess is at least $\frac{2}{8} = 25\%$).

4. Write the following as an extensive-form game with imperfect information:

   Sleeping Beauty takes part in an experiment. On Sunday, she goes into a deep sleep. The experimenters then flip a fair coin.

- If the result is heads, they wake her up on Monday and ask her whether the coin flip result was heads or tails. They then wake her up again on Tuesday and ask her the same question.
- If the result is tails, they wake her up and ask her the question only on Monday, and let her sleep through Tuesday.

Sleeping Beauty sleeps so deeply that when she is awakened, she cannot tell what day it is, or whether she has been questioned before.

In any case, the experiment ends on Wednesday, and she is woken up and asked no questions. Instead she is told what day it is, and she receives \$100 for each time she was questioned and answered correctly.

After writing the game tree, compute Sleeping Beauty's expected payoff if she guesses "heads" with probability x and "tails" with probability $(1-x)$, and recommend an optimal strategy for her.

5. In extensive-form games with perfect information, each player can specify a mixed strategy by specifying a mixed strategy for each of their nodes. For such games, the concepts of dominant strategies and Nash equilibria apply as well as they did for games written in table form. Find a Nash equilibrium for the game described in Problem 1. (Hint: if it is turn 10, what is Colin's optimal strategy? Knowing this, what is Rose's optimal strategy on turn 9?)

6. Use the game tree for the Monty Hall game to show that switching is a better strategy. If you're not sure how to do this, try computing, for each node, the probability that we arrive there given a strategy for the contestant (which consists of an initial guess 1, 2, or 3; and a decision to stay or switch doors after Monty has opened one of the doors).

7. Consider the following game:

   You are presented with three boxes. You are told that one of them contains two white marbles, one contains two black marbles, and one contains one white marble and one black marble, but you do not know which box is which. You reach into one of the boxes at random, and pull out one of its two marbles at random. You then have to guess whether the other marble in the box you drew from is white or black, and you win \$1 if you guess correctly.

   (a) The game tree for this game is given in Figure 5. Fill in the details about what the nodes are, what the strategies represent, and what the payoffs are.

   (b) Suppose you play the game and have drawn a black marble. What is the probability that the other marble in the box you drew from is also black? With this in mind, what is an optimal strategy for this game?

   (c) Suppose that instead of a straight \$1 reward for a correct guess, you instead receive \$$w$ for a correct guess of white and \$$b$ for a correct guess of black. Assuming that the first marble you draw is black, give a relationship between w and b such that your expected payoff is independent of the probability you guess white (or black) for the other marble in the same box.

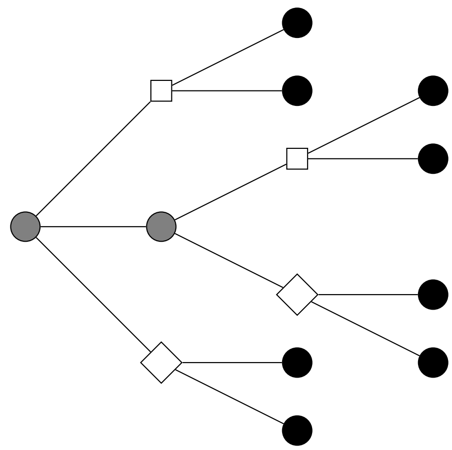

Figure 5: The marbles game.

8. (a) In the Monty Hall game, suppose that when you pick the correct door in Step 2, instead of opening a random door in Step 3 Monty always opens the lower numbered of the two doors. What is the optimal strategy now?

   (b) Suppose that we play a variant where now Monty does not know which door the prize is behind. However, in Step 3 he still opens (at random) one of the two doors you did not pick, but in the event he accidentally opens the prize door, we cancel the whole game and restart from the beginning. What is the optimal strategy now? What if Monty instead always chooses to open the lower numbered of the two doors you do not choose?

9. Challenge Problem: Consider a variant of the hat game in which the players do not lose if all of them say "No Guess". Instead, they are all allowed to guess (simultaneously) again. The usual win and loss conditions apply to the first round in which at least one player guesses.[10]

   Find a strategy for the players with the highest possible probability of winning, and prove that no better strategy exists.

[10] This game is infinite in theory, so let's also say that the players lose if they keep playing forever. We can also skip rounds in which the probability of a guess being made is zero, and so in practice, with reasonable strategies, the game will end fairly quickly.