Lesson 2 Exercise Solutions

Logan Crew and Sophie Spirkl

October 2021

1. We give the game tree below. All horizontal edges represent "Continue" and all diagonal edges represent "End".

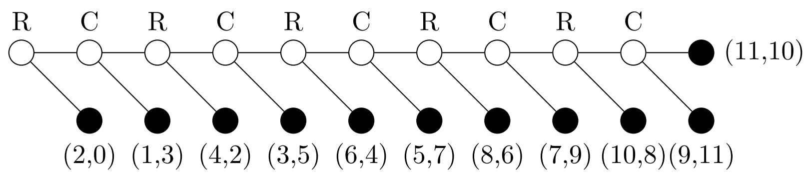

Figure 1: The two pots game (aka the Centipede Game). The initial R or C above a node indicates whose turn it is.

2. We give the game tree below. Each Nature node corresponds to either Rose or Colin flipping the coin as appropriate (with equal probability of heads (H) and tails (T)).

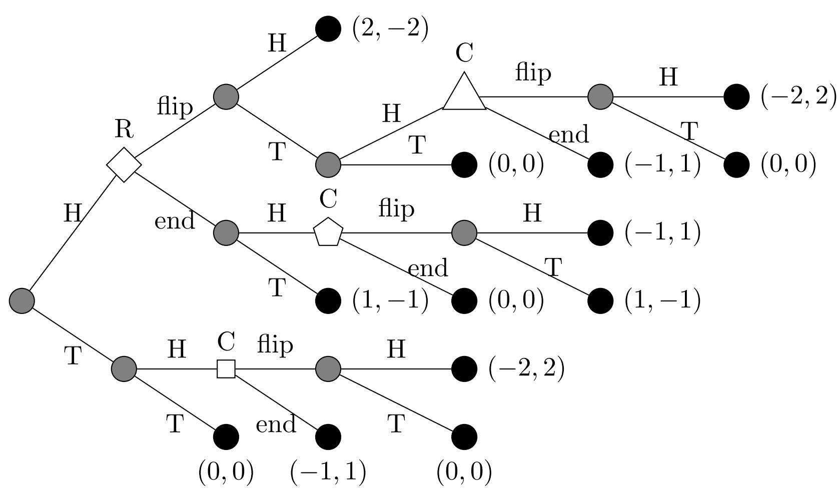

Figure 2: A game with Nature.

3. (a) If a player makes a guess for one of their information sets, then they are correct for one of the corresponding hat configurations, and incorrect for the other. Thus, if a player guesses for two information sets, they guess incorrectly on at least two of the eight hat configurations.

   (b) If each player guesses for at most one of their information sets, then each player guesses for at most two of the eight hat configurations. Even if there is no configuration for which multiple players guess, there will be at least $8 - 2 - 2 - 2 = 2$ configurations on which no player guesses.

4. The game tree for the Sleeping Beauty game is given below. Note that unlike in the previous games, which had perfect information, in this game *all* of Sleeping Beauty's nodes are in the same information set!

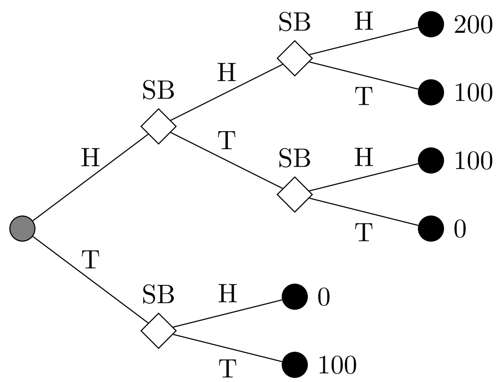

Figure 3: The Sleeping Beauty Game.

Now, suppose Sleeping Beauty guesses heads with probability x at each of her nodes. There is a 50% chance that heads was flipped, and in this case her payoff is $200x^2+100x(1-x)+100(1-x)x+0(1-x)^2$ based on whether she guesses heads or tails in her consecutive plays. There is also a 50% chance that tails was flipped, and in this case the expected payoff is $0x+100(1-x)$. Overall her expected winnings is

$$100x^2 + 100x(1-x) + 50(1-x) = 50 + 50x.$$

Thus, to maximize her expected winnings, Sleeping Beauty should maximize x, meaning that she should always guess heads!

5. Let us analyze the game by first looking at pure strategies. The simplest way to start is to look at the end of the game: what should Colin do on Turn 10? If he selects "continue", his payoff is 10, but if he selects "end", his payoff is 11! Thus, if the game reaches this point, Colin must choose "end" for a Nash equilibrium.

   But now consider Rose's strategy on Turn 9. If we are in a Nash equilibrium, then Colin will definitely play "end" if the game reaches Turn 10, and then Rose will get a payoff of 9. But if instead of continuing to Turn 10 Rose ends the game now, she gets a payoff of 10! Thus, if the game reaches this point, Rose must always choose "end".

   Working backwards in this way, it appears that the only possible Nash equilibrium for this game is if Rose always ends the game immediately on the first turn!

   And it is easy to verify that this is true even for mixed strategies. Suppose we have a Nash equilibrium. If there is a positive probability of reaching Turn 10, then Colin must always choose end there. But then by assumption Rose necessarily chooses "continue" in Turn 9 with positive probability, which cannot be a Nash equilibrium since she can improve her expected payoff by never choosing "continue" in Turn 9. Similarly, if turn i for $i > 1$ is the latest turn that can be reached with positive probability (so that in particular, whoever chooses that turn always chooses "end"), then there cannot be equilibrium because on the previous turn the other player is choosing "continue" with positive probability when they would do better never choosing "continue". Thus, in an equilibrium only Turn 1 can be reached with positive probability.

   Note that in this case Rose gets 2 and Colin gets 0, even though they can both do much better with other terminal nodes! So this is something of an extensive-form Prisoners' Dilemma.

6. We could evaluate the expected payoff of switching with probability x by summing over each branch of the game tree, but there is also an easier way to do this using symmetry.

Note that our information sets come in pairs. But in each pair, the probability of reaching the node for which staying is a win is half the probability of reaching the node for which switching is a win![1] Thus, we see that we are twice as likely to win if we switch than if we stay.

7. (a) We will call the boxes WW, BB, and BW based on whether they contain two white, two black, or one black and one white marble respectively. The first nature node represents the box you choose at random, each with probability $\frac{1}{3}$. If you choose either the WW or BB box, the colour of marble you draw is determined, so we skip straight to the guessing step. If you choose the BW box, there is a 50% chance you draw the white marble, and a 50% chance you draw the black one, and then you make a guess.

   The player has two different information sets represented by the different node shapes, based on the color of the marble they initially draw. The strategies represent the colour guessed, and the payoffs are the winnings in dollars.

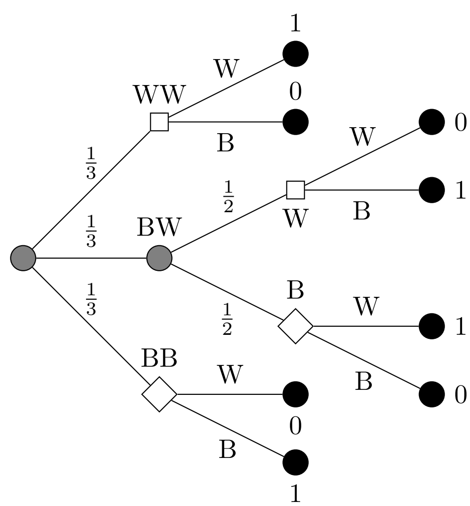

Figure 4: The marbles game.

   (b) The diamonds correspond to the information set in which you have drawn a black marble. There is a $\frac{1}{3}$ chance from the beginning that you went to the BB node, and a $\frac{1}{6}$ chance that you went to the BW node and then drew the black marble. Thus, it is twice as likely that the other marble in your box is black as white, so the other marble in the same box is black with probability $\frac{2}{3}$. Thus, your expected payoff is maximized by guessing black in this case.

   (c) If you have drawn a black marble, then by the previous part, if you guess white with probability x, your expected payoff is

$$\frac{1}{3}xw + \frac{2}{3}(1-x)b = \frac{x}{3}(w-2b) + \frac{2b}{3}.$$

   This value is independent of x if $w - 2b = 0$, or $w = 2b$.

8. (a) We can model this game by taking the original game tree, and assigning probabilities 1 and 0 instead of $\frac{1}{2}$ and $\frac{1}{2}$ when Monty has a choice.

   This changes the information sets! Some information set pairs now have a node that has 0 probability of being reached, so you know for certain what to do if you are in this set! This corresponds to the cases where Monty opens the higher-numbered of the two remaining doors. In these cases, you know for sure that switching is correct, because Monty would have opened the other door if there wasn't the prize behind it!

[1] Formally, these are conditional probabilities of which node we are at given that we know what information set we are in.

On the other hand, some information set pairs are now such that each of the two nodes is equally likely to be reached, so you have a 50% chance to win regardless of whether you stay or switch. This corresponds to when Monty does open the lower-numbered door. In these cases, it is equally likely that the prize is behind your door or the remaining unopened one.

(b) We can model this game by adding more nodes, so that Monty always has a 50% chance of opening either other door regardless of the door you pick. We can visualize these new nodes as just being new copies of the original game restarting if Monty opens the prize door[2]. Nodes of Monty's that previously only had one strategy now have two that are each chosen 50% of the time.

But now given an information set, the probability is equal that you are at either of its two nodes, so you have a 50% chance of winning regardless of whether you stay or switch.

If Monty is always choosing the lower-numbered door to open in this scenario, then each of Monty's nodes now only has one strategy, and again each of our information sets contains two nodes that are equally likely to be reached, so we win 50% of the time no matter our strategy. Intuitively, this is because Monty's door opening now gives us no information at all, since we always know which door he will open.

9. We claim that the maximum possible probability of winning is 87.5%, or $\frac{7}{8}$. It is easy to see that we cannot do better than this: the first time at least one player guesses, that player cannot distinguish between the two configurations of their information set, so there is at least one hat configuration of the eight such that the first player to guess is wrong.

It remains to show that there is a strategy that attains a win probability of $\frac{7}{8}$. There are many possibilities; we describe one below:

- In the first round, Colin and Larry never guess, and Rose guesses only if both Colin and Larry have white hats. If so, she guesses white. This results in a win if all hats are white, and a loss if Rose's hat is black and the other two are white.
- Rose never guesses after the first round. If we reach the second round, both Colin and Larry know that Rose does not see two white hats, so at least one of them has a black hat. Thus, they look at each other. If either Colin or Larry sees that the other has a white hat, they may safely guess Black for their own hat and win (and they do not guess if they see a black hat on the other player).
- If both Colin and Larry have black hats, neither will guess in the second round, so we advance to a third round. Now Colin and Larry know they both have black hats, so one or both of them may guess Black and win.

This strategy wins unless Rose gets a black hat and Colin and Larry get white hats, so the probability of winning is indeed $\frac{7}{8}$.

[2] Technically we should limit the game to a finite number of iterations, but that is not too important.