1 Playing Blackjack

You are playing a game of Blackjack where you start with \$100. You are a particularly risk-loving player who does not believe in leaving the table until you either make \$400, or lose all your money. At each turn you either win \$100 with probability p, or you lose \$100 with probability $1-p$.

(a) Formulate this problem as a Markov chain i.e. define your state space, transition probabilities, and determine your starting state.

(b) Find the probability that you end the game with \$400.

Solution:

(a) Since it is only possible for us to either win or lose \$100, we define the following state space $\chi = \{0, 100, 200, 300, 400\}$. The following are the transition probabilities:

$$\mathbb{P}(0,0) = \mathbb{P}(400,400) = 1$$
$$\mathbb{P}(i,i+100) = p \text{ for } i \in \{100,200,300\}$$
$$\mathbb{P}(i,i-100) = 1-p \text{ for } i \in \{100,200,300\}$$

(b) We want to find the probability that we are "absorbed" by state 400 before we are absorbed by state 0. We can calculate this probability by leveraging the memoryless property of Markov Chains. Define a_i as the probability of reaching state 400 before 0 starting at state i.

We also know that for $i \in \{100,200,300\}$, we have the following relation:

$$a_i = (1-p)a_{i-100} + pa_{i+100} \text{ for } i \in \{100,200,300\}$$

We also know that $a_0 = 0$, since if you are at state 0, then there is no chance that you end up at state 400. We also have $a_{400} = 1$ since if we are at state 400, then we have already succeeded in our goal to reach 400.

We have three unknowns $(a_{100}, a_{200}, a_{300})$ and three equations, and we can now solve this system of equations for a_{100}.

$$
\begin{gathered}
a_0 = 0, a_{400} = 1 \\
\implies a_i = (1-p)a_{i-100} + pa_{i+100} \text{ for } i \in \{100, 200, 300\} \\
a_{100} = pa_{200} \\
a_{200} = (1-p)a_{100} + pa_{300} \implies a_{200}[1 - p(1-p)] = pa_{300} \\
\implies a_{200} = \frac{pa_{300}}{1-p(1-p)} \\
a_{300} = (1-p)a_{200} + p \implies a_{300} = \frac{(1-p)pa_{300}}{1-p(1-p)} + p \\
\implies a_{300} = \frac{p(1-p(1-p))}{1-2p(1-p)} \\
\implies a_{200} = \frac{p^2}{1-2p(1-p)} \\
\implies a_{100} = \frac{p^3}{1-2p(1-p)}
\end{gathered}
$$

This problem is called Gambler's Ruin, where it is used to show that even if p is decently large, after playing a large number of games without stopping, you will end up at 0 dollars with high probability. Let's look at a nicer way to solve the recurrence relation that gives a somewhat more insightful answer to the problem.

Suppose we have states 0 through N, and you start at state k. You go up a state with probability p and go down with probability $1-p$. You win if you end up at state N, and lose if you end up at state 0.

Again, we can write the recurrence relation as

$$a_i = pa_{i+1} + (1-p)a_{i-1}$$

for $1 \le i \le N-1$. We also know that $a_N = 1$ and $a_0 = 0$. We can rewrite the recurrence relation into the following form:

$$(1-p)(a_i - a_{i-1}) = p(a_{i+1} - a_i) \Rightarrow a_{i+1} - a_i = \frac{1-p}{p}(a_i - a_{i-1})$$

Define $w = \frac{1-p}{p}$, which is often called the odds ratio, and define $b_i = a_{i+1} - a_i$. Note that this tells us $a_i = b_0 + b_1 + \cdots + b_{i-1}$. So, the recurrence we have derived is $b_i = w \cdot b_{i-1}$. This tells us that $b_i = w^i b_0$, and

$$a_i = b_0 + \cdots + b_{i-1} = (1 + w + w^2 + \cdots + w^{i-1})b_0$$

What is b_0? We can now use our information that $a_N = 1$, to see that $b_0 = \frac{1}{1+w+w^2+\cdots+w^{N-1}}$. Thus, we finally see that

$$a_i = \frac{1 + w + w^2 + \cdots + w^{i-1}}{1 + w + w^2 + \cdots + w^{N-1}} = \frac{w^i - 1}{w^N - 1}$$

where we used the geometric series formula in the last step: $1+w+w^2+\cdots+w^{i-1}=\frac{w^i-1}{w-1}$. Note that the formula only works if $w\neq 1$.

By the way, if you are interested in how to derive the geometric series, first write it like this:

$$S=1+w+w^2+\cdots+w^{i-1}$$

multiply both sides by w, to get

$$wS=w+w^2+\cdots+w^i$$

and subtracting these two equations will cancel most of the terms! We get:

$$(w-1)S=w^i-1$$

solving for S yields $\frac{w^i-1}{w-1}$.

2 Markov's Coupon Collecting

Courtney is home for Thanksgiving and needs to make some trips to the Traitor Goes grocery store to prepare for the big turkey feast. Each time she goes to the store before the holiday, she receives one of the n different coupons that are being given away. You may recall that we studied how to calculate the expected number of trips to the store needed to collect at least one of each coupon. Using geometric distributions and indicator variables, we determined that expected number of trips to be $n(\ln n+\gamma)$.

Let's re-derive that, this time with a Markov chain model and first-step equations.

(a) Define the states and transition probabilities for each state (explain what states can be transitioned to, and what probabilities those transitions occur with).

(b) Now set up first-step equations and solve for the expected number of grocery store trips Courtney needs to make before Thanksgiving so that she can have at least one of each of the n distinct coupons.

Solution:

(a) We model the coupon collector's problem as a Markov chain with states $X_1,X_2,\ldots,X_n,X_{n+1}$ where X_i represents the state we are at if we have collected $i-1$ of the unique coupons and are seeking the i^{th} coupon. State X_{n+1} represents the terminal state, after we successfully collected all n coupons and don't need to make any more grocery store trips.

If we are at state X_i, we either transition back to X_i with probability $(i-1)/n$, or we collect a new coupon and transition to state X_{i+1} with probability $(n-i+1)/n$. Transitioning to any other state is not possible.

(b) For each state X_i:

$$\begin{aligned}
\beta(X_i) &= 1 + \frac{i-1}{n} \cdot \beta(X_i) + \frac{n-i+1}{n} \cdot \beta(X_{i+1}) \\
\frac{n-i+1}{n} \cdot \beta(X_i) &= 1 + \frac{n-i+1}{n} \cdot \beta(X_{i+1}) \\
\beta(X_i) &= \frac{n}{n-i+1} + \beta(X_{i+1})
\end{aligned}$$

We know that for the terminal state, $\beta(X_{n+1}) = 0$. Then:

$$\begin{aligned}
\beta(X_n) &= n \\
\beta(X_{n-1}) &= n + \frac{n}{2} \\
&\vdots \\
\beta X_1 &= n + \frac{n}{2} + \cdots + \frac{n}{n} \\
&= n \sum_{i=1}^{n} \frac{1}{i} \\
&= n(\ln n + \gamma).
\end{aligned}$$

3 Reflecting Random Walk

Alice starts at vertex 0 and wishes to get to vertex n. When she is at vertex 0 she has a probability of 1 of transitioning to vertex 1. For any other vertex i, there is a probability of $1/2$ of transitioning to $i+1$ and a probability of $1/2$ of transitioning to $i-1$.

(a) What is the expected number of steps Alice takes to reach vertex n? Write down the hitting-time equations, but do not solve them yet.

(b) Solve the hitting-time equations. [*Hint*: Let R_i denote the expected number of steps to reach vertex n starting from vertex i. As a suggestion, try writing R_0 in terms of R_1; then, use this to express R_1 in terms of R_2; and then use this to express R_2 in terms of R_3, and so on. See if you can notice a pattern.]

Solution:

Formulate hitting time equations; the hard part is solving them. R_i represents the expected number of steps to get to vertex n starting from vertex i. In particular, $R_n = 0$ and we are interested in calculating R_0. We have the equations:

$$\begin{aligned}
R_0 &= 1 + R_1, \\
R_1 &= 1 + \frac{1}{2}R_0 + \frac{1}{2}R_2, \\
&\vdots \\
R_i &= 1 + \frac{1}{2}R_{i-1} + \frac{1}{2}R_{i+1}, \\
&\vdots \\
R_{n-1} &= 1 + \frac{1}{2}R_{n-2} + \frac{1}{2}R_n.
\end{aligned}$$

We can write this in terms of the differences $D_i := R_{i+1} - R_i$. If we take the recurrence relation $R_i = 1 + \frac{1}{2}R_{i-1} + \frac{1}{2}R_{i+1}$, we can rearrange the equation:

$$\begin{aligned}
R_i &= 1 + \frac{1}{2}R_{i-1} + \frac{1}{2}R_{i+1} \\
2R_i &= 2 + R_{i-1} + R_{i+1} \\
R_i - R_{i-1} - 2 &= R_{i+1} - R_i \\
D_{i-1} - 2 &= D_i
\end{aligned}$$

Furthermore, we know that $D_0 := R_1 - R_0 = -1$ from the very first hitting time equation. Since we have shown that D_i decreases by 2 every time, we know that $D_i = -2i - 1$. How do we get back R_i from knowing D_i? Well, we see that

$$R_i = (R_i - R_{i-1}) + (R_{i-1} - R_{i-2}) + \cdots + (R_1 - R_0) + R_0 = D_{i-1} + D_{i-2} + \cdots + D_0 + R_0$$

Therefore, we have $R_i = -1 - 3 - 5 - \cdots - (2i - 1) + R_0$. What is the sum of the first i odd integers? Here is how you would derive it. Let $S = 1 + 3 + 5 + \cdots + (2i - 1)$. Then, we can also write S backwards, as $S = (2i - 1) + (2i - 3) + \cdots + 5 + 3 + 1$. Lining up the terms, we see:

$$\begin{aligned}
S &= 1 \quad\quad\;\; + 3 \quad\quad\;\; + \cdots + (2i-3) + (2i-1) \\
S &= (2i-1) + (2i-3) + \cdots + 3 \quad\quad\;\; + 1
\end{aligned}$$

Adding these together gives us $2S = 2i + 2i + 2i + \cdots + 2i = 2i^2$. Solving for S yields $S = i^2$.

Now that we know this fact, we see that $R_i = R_0 - i^2$. Since we know that $R_n = 0$, we see that $R_0 - n^2 = 0$, and thus $R_0 = n^2$.

4 Boba in a Straw

Imagine that Jonathan is drinking milk tea and he has a very short straw: it has enough room to fit two boba (see figure).

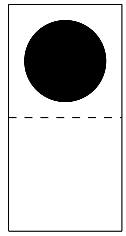

Figure 1: A straw with one boba currently inside. The straw only has enough room to fit two boba.

Here is a formal description of the drinking process: We model the straw as having two "components" (the top component and the bottom component). At any given time, a component can contain nothing, or one boba. As Jonathan drinks from the straw, the following happens every second:

1. The contents of the top component enter Jonathan's mouth.
2. The contents of the bottom component move to the top component.
3. With probability p, a new boba enters the bottom component; otherwise the bottom component is now empty.

Help Jonathan evaluate the consequences of his incessant drinking!

(a) At the very start, the straw starts off completely empty. What is the expected number of seconds that elapse before the straw is completely filled with boba for the first time? [Write down the equations; you do not have to solve them.]

(b) Consider a slight variant of the previous part: now the straw is narrower at the bottom than at the top. This affects the drinking speed: if either (i) a new boba is about to enter the bottom component or (ii) a boba from the bottom component is about to move to the top component, then the action takes two seconds. If both (i) and (ii) are about to happen, then the action takes three seconds. Otherwise, the action takes one second. Under these conditions, answer the previous part again. [Write down the equations; you do not have to solve them.]

(c) Jonathan was annoyed by the straw so he bought a fresh new straw (the straw is no longer narrow at the bottom). What is the long-run average rate of Jonathan's calorie consumption? (Each boba is roughly 10 calories.)

(d) What is the long-run average number of boba which can be found inside the straw? [Maybe you should first think about the long-run distribution of the number of boba.]

Solution:

(a) We model the straw as a four-state Markov chain. The states are $\{(0,0),(0,1),(1,0),(1,1)\}$, where the first component of a state represents whether the top component is empty (0) or full

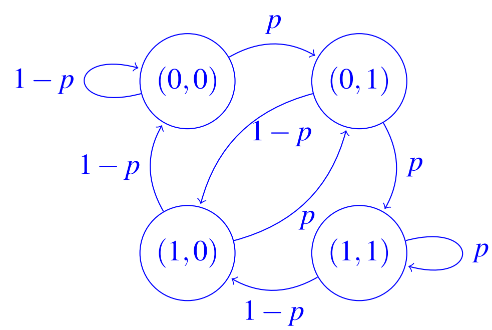

Figure 2: Transition diagram for the Markov chain.

(1); similarly, the second component represents whether the bottom component is empty or full. See Figure **??**.

Now, we set up the hitting time equations. Let T denote the time it takes to reach state $(1,1)$, i.e. $T = \min\{n > 0 : X_n = (1,1)\}$. Let $\mathbb{E}_i[\cdot] = \mathbb{E}[\cdot \mid X_0 = i]$ denote the expectation starting from state i (for convenience of notation). The hitting-time equations are

$$\begin{aligned}
\mathbb{E}_{(0,0)}[T] &= 1 + (1-p)\,\mathbb{E}_{(0,0)}[T] + p\,\mathbb{E}_{(0,1)}[T], \\
\mathbb{E}_{(0,1)}[T] &= 1 + (1-p)\,\mathbb{E}_{(1,0)}[T] + p\,\mathbb{E}_{(1,1)}[T], \\
\mathbb{E}_{(1,0)}[T] &= 1 + (1-p)\,\mathbb{E}_{(0,0)}[T] + p\,\mathbb{E}_{(0,1)}[T], \\
\mathbb{E}_{(1,1)}[T] &= 0.
\end{aligned}$$

The question did not ask you to solve the equations. If you solved the equations anyway and would like to check your work, the hitting time is $\mathbb{E}_{(0,0)}[T] = (1+p)/p^2$.

(b) The new hitting-time equations are

$$\begin{aligned}
\mathbb{E}_{(0,0)}[T] &= (1-p)(1 + \mathbb{E}_{(0,0)}[T]) + p(2 + \mathbb{E}_{(0,1)}[T]), \\
\mathbb{E}_{(0,1)}[T] &= (1-p)(2 + \mathbb{E}_{(1,0)}[T]) + p(3 + \mathbb{E}_{(1,1)}[T]), \\
\mathbb{E}_{(1,0)}[T] &= (1-p)(1 + \mathbb{E}_{(0,0)}[T]) + p(2 + \mathbb{E}_{(0,1)}[T]), \\
\mathbb{E}_{(1,1)}[T] &= 0.
\end{aligned}$$

You did not have to solve the equations, but to get a sense for what the solution is like, solving the equations and plugging in $p = 1/2$ yields (after some tedious algebra) $\mathbb{E}_{(0,0)}[T] = 11$.

(c) This part is actually more straightforward than it might initially seem: the average rate at which Jonathan consumes boba must equal the average rate at which boba enters the straw, which is p per second. Hence, his long-run average calorie consumption rate is $10p$ per second.

(d) We compute the stationary distribution. The balance equations are

$$\begin{aligned}
\pi(0,0) &= (1-p)\pi(0,0) + (1-p)\pi(1,0), \\
\pi(0,1) &= p\pi(0,0) + p\pi(1,0), \\
\pi(1,0) &= (1-p)\pi(0,1) + (1-p)\pi(1,1), \\
\pi(1,1) &= p\pi(0,1) + p\pi(1,1).
\end{aligned}$$

Expressing everything in terms of $\pi(0,0)$, we find

$$\pi(0,1) = \pi(1,0) = \frac{p}{1-p}\pi(0,0),$$
$$\pi(1,1) = \frac{p^2}{(1-p)^2}\pi(0,0).$$

From the normalization condition we have

$$\pi(0,0)\left(1 + \frac{2p}{1-p} + \frac{p^2}{(1-p)^2}\right) = 1,$$

so $\pi(0,0) = (1-p)^2$. Hence, the stationary distribution is

$$\pi(0,0) = (1-p)^2,$$
$$\pi(0,1) = \pi(1,0) = p(1-p),$$
$$\pi(1,1) = p^2.$$

In states $(0,1)$ and $(1,0)$, there is one boba in the straw; in state $(1,1)$, there are two boba in the straw. Therefore, the long-run average number of boba in the straw is

$$\pi(0,1) + \pi(1,0) + 2\pi(1,1) = 2p(1-p) + 2p^2 = 2p.$$

Alternate Solution: The goal of the question was to have you work through the balance equations, but there is a simple solution. Observe that at any given time after at least two seconds have passed, each component has probability p of being filled with boba. Therefore, the number of boba in the straw is like a binomial distribution with 2 independent trials and success probability p, so the average number of boba in the straw is $2p$.