CS 70 Discrete Mathematics and Probability Theory
Spring 2021 HW 12

1 Gambling Woes

Forest proposes a gambling game to you (uh oh!). Every day, you flip two independent fair coins. If both of the coins come up heads, then your fortune triples on that day. If one coin comes up heads and the other coin comes up tails, then your fortune is cut in half. If both of the coins comes up tails, then game over: you lose all of your money! Forest claims that you can get rich quickly with this scheme, but you decide to calculate some probabilities first.

(a) Let M_0 denote your money at the start of the game, and let M_n denote the amount of money you have at the end of the nth day. Compute $\mathbb{E}[M_{n+1} \mid M_n]$.

(b) Use the law of iterated expectation to calculate $\mathbb{E}[M_{n+1}]$ in terms of $\mathbb{E}[M_n]$. Solve your recurrence to obtain an expression for $\mathbb{E}[M_{n+1}]$. Do you think this is a fair game?

(c) Calculate $\mathbb{P}(M_n > 0)$. What is the behavior as $n \to \infty$? Would you still play this game?

Solution:

(a) Suppose that you have $M_n = m$ dollars at the end of day n. At the end of day $n+1$: with probability $1/4$, your fortune is $3m$; with probability $1/2$, your fortune is $m/2$; with probability $1/4$, your fortune is 0. Therefore,

$$\mathbb{E}[M_{n+1} \mid M_n = m] = \frac{1}{4} \cdot 3m + \frac{1}{2} \cdot \frac{m}{2} = m$$

and $\mathbb{E}[M_{n+1} \mid M_n] = M_n$.

(b) $\mathbb{E}[M_{n+1}] = \mathbb{E}[\mathbb{E}[M_{n+1} \mid M_n]] = \mathbb{E}[M_n]$. Therefore, we can see that

$$\mathbb{E}[M_{n+1}] = \mathbb{E}[M_n] = \mathbb{E}[M_{n-1}] = \cdots = \mathbb{E}[M_0] = M_0.$$

Your expected fortune never changes! Technically, we would be justified in calling this a fair game.

(c) The probability that your fortune is non-zero on day n is the probability that you never flipped two tails on a day, which has probability $(3/4)^n$. As $n \to \infty$, $(3/4)^n \to 0$, so after many days of playing the game, it is highly probable that we are broke!

Remark: In part (b), we have shown that $\mathbb{P}(M_n > 0) \to 0$ as $n \to \infty$, which means that M_n converges to 0 in probability. In part (a), we have shown that $\mathbb{E}[M_n] = M_0$ for all n. That is, $M_n \to 0$ but $\mathbb{E}[M_n] \not\to 0$.

2 Iterated Expectation

In this question, we will try to achieve more familiarity with the law of iterated expectation.

1. You lost your phone charger! It will take D days for the new phone charger you ordered to arrive at your house (here, D is a random variable). Suppose that on day i, the amount of battery you lose is B_i, where $\mathbb{E}[B_i] = \beta$. Let $B = \sum_{i=1}^{D} B_i$ be the total amount of battery drained between now and when your new phone charger arrives. Apply the law of iterated expectation to show that $\mathbb{E}[B] = \beta\,\mathbb{E}[D]$. (Here, the law of iterated expectation has a very clear interpretation: the amount of battery you expect to drain is the average number of days it takes for your phone charger to arrive, multiplied by the average amount of battery drained per day.)

2. Consider now the setting of independent Bernoulli trials, each with probability of success p. Let S_i be the number of successes in the first i trials. Compute $\mathbb{E}[S_m \mid S_n]$. (You will need to consider three cases based on whether $m > n$, $m = n$, or $m < n$. Try using your intuition rather than proceeding by calculations.)

Solution:

1. This is simply Wald's Identity from lecture. Condition on $D = d$; then $B = \sum_{i=1}^{d} B_i$ and

$$\mathbb{E}[B \mid D = d] = \sum_{i=1}^{d} \mathbb{E}[B_i] = \beta d.$$

   Therefore, $\mathbb{E}[B \mid D] = \beta D$ and by the law of iterated expectation, $\mathbb{E}[B] = \mathbb{E}[\mathbb{E}[B \mid D]] = \beta\,\mathbb{E}[D]$.

2. Suppose $m > n$. Then we already know that the first n trials resulted in S_n successes, and there are $m - n$ trials for which we do not know the outcome. Each of these $m - n$ trials has probability of success p, so we expect $(m-n)p$ additional successes. Hence, $\mathbb{E}[S_m \mid S_n] = S_n + (m-n)p$.

   Next, consider when $m = n$. Here, $\mathbb{E}[S_m \mid S_n] = S_n$.

   Finally, suppose that $m < n$. In n trials, we have S_n successes, and due to symmetry, we expect the S_n successes to be distributed uniformly among the n trials. In particular, if we look at the first m trials only, then we expect a proportion m/n of the total successes to be distributed among the first m successes. Therefore, $\mathbb{E}[S_m \mid S_n] = mS_n/n$.

3 Strange Dilution

You have a jar of red marbles and blue marbles. At each time step, you draw a marble, and you note the color of the marble. Then, you dilute the proportion of the opposite-colored marbles by a factor of γ, where $0 < \gamma < 1$. (For example: if you pick a red marble, then the proportion of blue marbles is reduced by a factor of γ.) If p is the fraction of marbles that started off as red, what is the expected proportion of red marbles at time n?

Solution:

Let X_n be the fraction of marbles that are red at time n. With probability X_n, you choose a red marble, so the fraction of blue marbles is now $\gamma(1-X_n)$. Therefore, the new fraction of red marbles is $1-\gamma(1-X_n) = 1-\gamma+\gamma X_n$. With probability $1-X_n$, you choose a blue marble, so the fraction of red marbles is now γX_n. We have

$$\mathbb{E}[X_{n+1} \mid X_n] = X_n(1-\gamma+\gamma X_n) + (1-X_n)\gamma X_n = X_n - \gamma X_n + \gamma X_n = X_n$$

By the law of iterated expectation, $\mathbb{E}[X_{n+1}] = \mathbb{E}[\mathbb{E}[X_{n+1} \mid X_n]] = \mathbb{E}[X_n]$, so $\mathbb{E}[X_n] = p$.

4 Oski's Markov Chain

When Oski Bear is studying for CS70, he splits up his time between reading notes and working on practice problems. To do this, every so often he will make a decision about what kind of work to do next.

When Oski is already reading the notes, with probability a he will decide to switch gears and work on a practice problem, and otherwise, he will decide to keep reading more notes. Conversely, when Oski is already working on a practice problem, with probability b he will think of a topic he needs to review, and will decide to switch back over to the notes; otherwise, he will keep working on practice problems.

Assume that (unlike real life, we hope!) Oski never runs out of work to do.

(a) Draw a 2-state Markov chain to model this situation.

(b) In the remainder of this problem, we will learn to work with the definitions of some important terms relating to Markov Chains. These definitions are as follows:

  (a) (Irreducibility) A Markov chain is irreducible if, starting from any state i, the chain can transition to any other state j, possibly in multiple steps.

  (b) (Periodicity) $d(i) := \gcd\{n > 0 \mid P^n(i,i) = \mathbb{P}[X_n = i \mid X_0 = i] > 0\}$, $i \in \mathscr{X}$. If $d(i) = 1\ \forall i \in \mathscr{X}$, then the Markov chain is aperiodic; otherwise it is periodic.

  (c) (Matrix Representation) Define the transition probability matrix P by filling entry (i,j) with probability $P(i,j)$.

  (d) (Invariance) A distribution π is invariant for the transition probability matrix P if it satisfies the following balance equations: $\pi = \pi P$.

  For what values of a and b is the Markov chain irreducible?

(c) For $a = 1$, $b = 1$, prove that the Markov chain is periodic.

(d) For $0 < a < 1$, $0 < b < 1$, prove that the Markov chain is aperiodic.

(e) Construct a transition probability matrix using the Markov chain.

(f) Write down the balance equations for the Markov chain and solve them. Assume that the Markov chain is irreducible.

Solution:

(a) Let state 0 represent Oski working on a practice problem, and let state 1 represent reading notes.

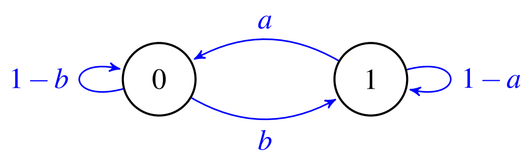

(b) The Markov chain is irreducible if both a and b are non-zero. It is reducible if at least one is 0.

(c) We compute $d(0)$ to find that:

$$d(0) = \gcd\{2,4,6,\ldots\} = 2.$$

Thus, the chain is periodic.

(d) We compute $d(0)$ to find that:

$$d(0) = \gcd\{1,2,3,\ldots\} = 1.$$

Thus, the chain is aperiodic.

(e)

$$\begin{bmatrix} 1-b & b \\ a & 1-a \end{bmatrix}$$

(f)

$$\pi(0) = (1-b)\pi(0) + a\pi(1),$$
$$\pi(1) = b\pi(0) + (1-a)\pi(1).$$

One of the equations is redundant. We throw out the second equation and replace it with $\pi(0)+\pi(1) = 1$. This gives the solution

$$\pi = \frac{1}{a+b}\begin{bmatrix} a & b \end{bmatrix}.$$

5 Markov Chains: Prove/Disprove

Prove or disprove the following statements, using the definitions from the previous question.

(a) There exists an irreducible, finite Markov chain for which there exist initial distributions that converge to different distributions.

(b) There exists an irreducible, aperiodic, finite Markov chain for which $\mathbb{P}(X_{n+1} = j \mid X_n = i) = 1$ or 0 for all i, j.

(c) There exists an irreducible, non-aperiodic Markov chain for which $\mathbb{P}(X_{n+1} = j \mid X_n = i) \neq 1$ for all i, j.

(d) For an irreducible, non-aperiodic Markov chain, any initial distribution not equal to the invariant distribution does not converge to any distribution.

Solution:

(a) False. Every finite irreducible Markov chain has a unique stationary distribution. If it's possible for the Markov chain to converge to two different distributions given different starting distributions, it implies there are two stationary distributions. To elaborate further, we know in the long run the fraction of time spent in each state converges to the stationary distribution. So if the distribution converges, the long-run fraction of time will be whatever distribution it converges to, which we see must be the stationary distribution.

(b) True, you can have one state pointing to itself. However for number of states > 1 it is false. Consider the initial distribution of having a probability of 1 of being in an arbitrary state. After a transition, the resulting distribution must be a probability 1 of being in a different state (if it were the same state, this would immediately imply that the Markov chain is reducible). Further transitions have the same effect. Therefore this initial distribution does not converge. Therefore this Markov chain cannot be aperiodic and irreducible (since it would converge in that case).

(c) True. Consider the states $\{0,1,2,3\}$. Set $P(i,j) = 1/2$ if $i \equiv j \pm 1 \mod 4$ and 0 otherwise. In other words, the Markov chain is a square with each side replaced with two links pointing in opposite directions with probabilities of $1/2$. Consider the period of state 0. Any path from 0 back to itself, such as $0-1-2-1-0$, alternates in parity of each consecutive state since each state only points to the state above or below it mod 4. Therefore state 0 has period 2. Therefore this Markov chain is not aperiodic (and all states have period 2).

(d) False. Take the initial distribution $\begin{bmatrix} 0.25 & 0.30 & 0.25 & 0.20 \end{bmatrix}$ for the above Markov chain. After one transition it goes to the invariant distribution, $\begin{bmatrix} 0.25 & 0.25 & 0.25 & 0.25 \end{bmatrix}$.

6 Markov Property Practice

Let $X_0, X_1, \ldots$ be a Markov chain with state space S, such that i_j is the value that X_j takes in the j^{th} state. One of the properties that it satisfies is the Markov property:

$$\mathbb{P}(X_n = i_n | X_{n-1} = i_{n-1}, \ldots, X_0 = i_0) = \mathbb{P}(X_n = i_n | X_{n-1} = i_{n-1}), \text{ for all } i_0, i_1, \ldots, i_n \in S, \ n \in \mathbb{Z}_{>0}.$$

Use the Markov property and the total probability theorem to prove the following.

(a) $\mathbb{P}(X_3 = i_3 | X_2 = i_2, X_1 = i_1) = \mathbb{P}(X_3 = i_3 | X_2 = i_2)$, for all $i_1, i_2, i_3 \in S$.

Note: This is not exactly the Markov property because it does not condition on X_0.

(b) $\mathbb{P}(X_3 = i_3 | X_1 = i_1, X_0 = i_0) = \mathbb{P}(X_3 = i_3 | X_1 = i_1)$, for all $i_0, i_1, i_3 \in S$.

(c) $\mathbb{P}(X_1 = i_1 | X_2 = i_2, X_3 = i_3) = \mathbb{P}(X_1 = i_1 | X_2 = i_2)$, for all $i_1, i_2, i_3 \in S$.

Solution:

(a) By the total probability rule, we have

$$\begin{aligned}\mathbb{P}(X_3 = i_3 | X_2 = i_2, X_1 = i_1) &= \sum_{i_0 \in S} \mathbb{P}(X_3 = i_3, X_0 = i_0 | X_2 = i_2, X_1 = i_1) \\ &= \sum_{i_0 \in S} \mathbb{P}(X_3 = i_3 | X_2 = i_2, X_1 = i_1, X_0 = i_0) \mathbb{P}(X_0 = i_0 | X_2 = i_2, X_1 = i_1)\end{aligned}$$

The first probability is now in a form where we can apply the Markov property, leaving us with

$$\begin{aligned}\mathbb{P}(X_3 = i_3 | X_2 = i_2, X_1 = i_1) &= \sum_{i_0 \in S} \mathbb{P}(X_3 = i_3 | X_2 = i_2) \mathbb{P}(X_0 = i_0 | X_2 = i_2, X_1 = i_1) \\ &= \mathbb{P}(X_3 = i_3 | X_2 = i_2) \sum_{i_0 \in S} \mathbb{P}(X_0 = i_0 | X_2 = i_2, X_1 = i_1)\end{aligned}$$

where in the second step we used that $\mathbb{P}(X_3 = i_3 | X_2 = i_2)$ doesn't depend on i_0 in order to factor it out of the summation. We now note that X_0 must take on some value, so the summation in our previous equation must equal 1. Hence, we have

$$\mathbb{P}(X_3 = i_3 | X_2 = i_2, X_1 = i_1) = \mathbb{P}(X_3 = i_3 | X_2 = i_2)$$

which is the equality we wanted.

Intuition: The Markov property says that in order to predict the n^{th} state that will be visited, knowing only the $(n-1)^{st}$ state is just as good as knowing the entire history of the 0^{th} through $(n-1)^{st}$ states. We see that on both the left and right hand sides of the equation, we know state X_2 and want to predict X_3, so we can expect that any states before X_2 won't give us any extra information about X_3.

(b) Again applying the total probability rule, we have that

$$\begin{aligned}\mathbb{P}(X_3 = i_3 | X_1 = i_1, X_0 = i_0) &= \sum_{i_2 \in S} \mathbb{P}(X_3 = i_3, X_2 = i_2 | X_1 = i_1, X_0 = i_0) \\ &= \sum_{i_2 \in S} \mathbb{P}(X_3 = i_3 | X_2 = i_2, X_1 = i_1, X_0 = i_0)\mathbb{P}(X_2 = i_2 | X_1 = i_1, X_0 = i_0)\end{aligned}$$

We then note that the Markov property allows us to drop the conditioning on X_1 and X_0 in the first probability and the conditioning on X_0 in the second. This means that

$$\mathbb{P}(X_3 = i_3 | X_1 = i_1, X_0 = i_0) = \sum_{i_2 \in S} \mathbb{P}(X_3 = i_3 | X_2 = i_2)\mathbb{P}(X_2 = i_2 | X_1 = i_1)$$

However, we have a problem here: the second probability is conditioned on the value of X_1 while the first is not, so we can't immediately combine them into a single probability. However, we can use the previous part to replace the $\mathbb{P}(X_3 = i_3 | X_2 = i_2)$ with $\mathbb{P}(X_3 = i_3 | X_2 = i_2, X_1 = i_1)$, allowing us to write

$$\begin{aligned}\mathbb{P}(X_3 = i_3 | X_1 = i_1, X_0 = i_0) &= \sum_{i_2 \in S} \mathbb{P}(X_3 = i_3 | X_2 = i_2, X_1 = i_1)\mathbb{P}(X_2 = i_2 | X_1 = i_1) \\ &= \sum_{i_2 \in S} \mathbb{P}(X_3 = i_3, X_2 = i_2 | X_1 = i_1) \\ &= \mathbb{P}(X_3 = i_3 | X_1 = i_1)\end{aligned}$$

where in the last step we applied the total probability rule in reverse.

Intuition: This says that knowing the value of X_0 does not give us any new information about state X_3, assuming that we already knew the value of X_1. This makes sense from the memoryless property of Markov Chains: to determine the probability distribution for X_3, we only need to know the value of X_2, and to determine the probability distribution for X_2, we only need to know X_1. So once we know about X_1, everything that comes before it is irrelevant in determining what happens with X_3.

(c) Using the definition of conditional probability, we can write

$$\begin{aligned}\mathbb{P}(X_1 = i_1 | X_2 = i_2, X_3 = i_3) &= \frac{\mathbb{P}(X_1 = i_1, X_3 = i_3 | X_2 = i_2)}{\mathbb{P}(X_3 = i_3 | X_2 = i_2)} \\ &= \frac{\mathbb{P}(X_3 = i_3 | X_2 = i_2, X_1 = i_1)\mathbb{P}(X_1 = i_1 | X_2 = i_2)}{\mathbb{P}(X_3 = i_3 | X_2 = i_2)}\end{aligned}$$

Note that this is effectively applying Bayes' Rule to switch the positions of X_1 and X_3 in the conditional probability. Applying our result from part (a), this becomes

$$\begin{aligned}\mathbb{P}(X_1 = i_1 | X_2 = i_2, X_3 = i_3) &= \frac{\mathbb{P}(X_3 = i_3 | X_2 = i_2)\mathbb{P}(X_1 = i_1 | X_2 = i_2)}{\mathbb{P}(X_3 = i_3 | X_2 = i_2)} \\ &= \mathbb{P}(X_1 = i_1 | X_2 = i_2)\end{aligned}$$

Intuition: We know that Markov Chains are "memoryless", so the probability distribution of X_3 depends only on the value of X_2 and not any earlier states (such as X_1). So, if we already know the value of X_2, finding out the value of X_3 doesn't give us any new information about what happened in state X_1.

7 Knight on a Chessboard (Optional)

This problem is optional and will not be required as part of your submission for this homework assignment. You will receive no additional credit for completing this problem.

(a) An irreducible Markov chain is said to be reversible if there exists a probability distribution π such that

$$\pi(i)P(i,j) = \pi(j)P(j,i) \qquad \forall i,j \in \mathscr{X}. \tag{1}$$

Show that if the chain is reversible, then the distribution π which satisfies (1) is the stationary distribution.

(b) Consider a random walk on a finite undirected connected graph: starting from any vertex, the walk transitions to any of the neighboring vertices with equal probability. The state space $\mathscr{X}$ is the set of vertices of the graph. Show that $\pi(v) = \deg(v)/\sum_{v' \in \mathscr{X}} \deg(v')$, for $v \in \mathscr{X}$, is the stationary distribution.

(c) Let $T_i = \min\{n > 0 : X_n = i\}$ be the first time that the chain reaches state i. Use the result $\mathbb{E}[T_i \mid X_0 = i] = 1/\pi(i)$ to answer the following question:

> A knight begins at a corner of an 8×8 chessboard. At each step, it chooses one of its legal moves uniformly at random and moves to a new square. What is the expected number of steps before the knight returns to the same corner from which it started?

[*Hint*: Think about how this question can be formulated as a random walk on a graph.]

Solution:

(a) One has

$$\sum_{i \in \mathscr{X}} \pi(i)P(i,j) = \sum_{i \in \mathscr{X}} \pi(j)P(j,i) = \pi(j) \sum_{i \in \mathscr{X}} P(j,i) = \pi(j).$$

Hence, π satisfies the balance equations.

(b) π is a valid probability distribution: the entries are non-negative and

$$\sum_{v \in \mathscr{X}} \pi(v) = \sum_{v \in \mathscr{X}} \frac{\deg(v)}{\sum_{v' \in \mathscr{X}} \deg(v')} = \frac{\sum_{v \in \mathscr{X}} \deg(v)}{\sum_{v' \in \mathscr{X}} \deg(v')} = 1.$$

To check that the chain is reversible, note that if u and v are neighbors, then

$$\pi(u)P(u,v) = \frac{\deg(u)}{\sum_{v' \in \mathscr{X}} \deg(v')} \cdot \frac{1}{\deg(u)} = \frac{1}{\sum_{v' \in \mathscr{X}} \deg(v')}.$$

Also,

$$\pi(v)P(v,u) = \frac{\deg(v)}{\sum_{v' \in \mathscr{X}} \deg(v')} \cdot \frac{1}{\deg(v)} = \frac{1}{\sum_{v' \in \mathscr{X}} \deg(v')}.$$

So, $\pi(u)P(u,v) = \pi(v)P(v,u)$ if u and v are neighbors. If u and v are not neighbors, then $P(u,v) = P(v,u) = 0$, so the equation holds in this case as well. Therefore, the chain is reversible and π is stationary.

(c) Each square of the chessboard is a vertex on the graph, and there is an edge between two vertices on the graph if the knight can travel between the two vertices in one move. The degree of the corner vertex is 2, since the knight only has two legal moves starting from the corner. To compute $\pi(v)$ (where v is the corner vertex), we must compute the sum of the degrees in the graph.

Here is a table which represents the degree of each node:

2	3	4	4	4	4	3	2
3	4	6	6	6	6	4	3
4	6	8	8	8	8	6	4
4	6	8	8	8	8	6	4
4	6	8	8	8	8	6	4
4	6	8	8	8	8	6	4
3	4	6	6	6	6	4	3
2	3	4	4	4	4	3	2

Summing up all of the degrees in the graph is a little tedious, but it comes out to be 336. Hence,

$$\pi(v) = \frac{2}{336} = \frac{1}{168}$$

and the expected number of moves before the knight returns to the corner is 168.

8 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).

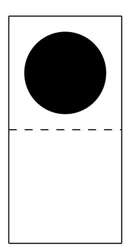

Figure 1: A straw with one boba currently inside. The straw only has enough room to fit two boba.

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 (1); similarly, the second component represents whether the bottom component is empty or full. See Figure 2.

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

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

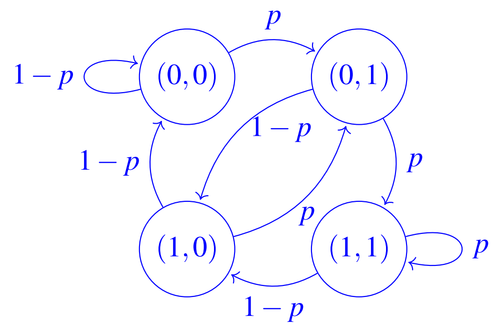

Figure 2: Transition diagram for the Markov chain.

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

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

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

$$\begin{aligned}
\pi(0,0) &= (1-p)^2, \\
\pi(0,1) &= \pi(1,0) = p(1-p), \\
\pi(1,1) &= p^2.
\end{aligned}$$

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 $2p$.