CS 70 Spring 2019	Discrete Mathematics and Probability Theory Babak Ayazifar and Satish Rao	HW 13

1 Markov's Inequality and Chebyshev's Inequality

A random variable X has variance $\text{var}(X) = 9$ and expectation $\mathbb{E}[X] = 2$. Furthermore, the value of X is never greater than 10. Given this information, provide either a proof or a counterexample for the following statements.

(a) $\mathbb{E}\left[X^2\right] = 13$.

(b) $\mathbb{P}[X = 2] > 0$.

(c) $\mathbb{P}[X \geq 2] = \mathbb{P}[X \leq 2]$.

(d) $\mathbb{P}[X \leq 1] \leq 8/9$.

(e) $\mathbb{P}[X \geq 6] \leq 9/16$.

(f) $\mathbb{P}[X \geq 6] \leq 9/32$.

Solution:

(a) TRUE. Since $9 = \text{var}(X) = \mathbb{E}[X^2] - \mathbb{E}[X]^2 = \mathbb{E}[X^2] - 2^2$, we have $\mathbb{E}[X^2] = 9 + 4 = 13$.

(b) FALSE. Construct a random variable X that satisfies the conditions in the question but does not take on the value 2. A simple example would be a random variable that takes on 2 values, where $\mathbb{P}[X = a] = \mathbb{P}[X = b] = 1/2$, and $a \neq b$. The expectation must be 2, so we have $a/2 + b/2 = 2$. The variance is 9, so $\mathbb{E}[X^2] = 13$ (from Part (a)) and $a^2/2 + b^2/2 = 13$. Solving for a and b, we get $\mathbb{P}[X = -1] = \mathbb{P}[X = 5] = 1/2$ as a counterexample.

(c) FALSE. Construct a random variable X that satisfies the conditions in the question but does not have an equal chance of being less than or greater than 2. A simple example would be a random variable that takes on 2 values, where $\mathbb{P}[X = a] = p, \mathbb{P}[X = b] = 1 - p$. Here, we use the same approach as part (b) except with a generic p, since we want $p \neq 1/2$. The expectation must be 2, so we have $pa + (1-p)b = 2$. The variance is 9, so $\mathbb{E}\left[X^2\right] = 13$ and $pa^2 + (1-p)b^2 = 13$. Solving for a and b, we find the relation $b = 2 \pm 3/\sqrt{x}$, where $x = (1-p)/p$. Then, we can find an example by plugging in values for x so that $a, b \leq 10$ and $p \neq 1/2$. One such counterexample is $\mathbb{P}[X = -7] = 1/10, \mathbb{P}[X = 3] = 9/10$.

(d) TRUE. Let $Y = 10 - X$. Since X is never exceeds 10, Y is a non-negative random variable. By Markov's inequality,

$$\mathbb{P}[10 - X \geq a] = \mathbb{P}[Y \geq a] \leq \frac{\mathbb{E}[Y]}{a} = \frac{\mathbb{E}[10 - X]}{a} = \frac{8}{a}.$$

Setting $a = 9$, we get $\mathbb{P}[X \leq 1] = \mathbb{P}[10 - X \geq 9] \leq 8/9$.

(e) TRUE. Chebyshev's inequality says $\mathbb{P}[|X - \mathbb{E}[X]| \geq a] \leq \text{var}(X)/a^2$. If we set $a = 4$, we have

$$\mathbb{P}[|X - 2| \geq 4] \leq \frac{9}{16}.$$

Now we observe that $\mathbb{P}[X \geq 6] \leq \mathbb{P}[|X - 2| \geq 4]$, because the event $X \geq 6$ is a subset of the event $|X - 2| \geq 4$.

(f) FALSE. Construct a random variable X that satisfies the conditions in the question but does not have an equal chance of being less than -2 or greater than 6. A simple example would be a random variable that takes on 2 values, where $\mathbb{P}[X = a] = p, \mathbb{P}[X = b] = 1 - p$. The expectation must be 2, so we have $pa + (1 - p)b = 2$. The variance is 9, so $\mathbb{E}\left[X^2\right] = 13$ and $pa^2 + (1 - p)b^2 = 13$. Solving for a and b. One example is $\mathbb{P}[X = 0] = 9/13, \mathbb{P}[X = 13/2] = 4/13$.

2 Subset Card Game

Jonathan and Yiming are playing a card game. Jonathan has $k > 2$ cards, and each card has a real number written on it. Jonathan tells Yiming (truthfully), that the sum of the card values is 0, and that the sum of squares of the values on the cards is 1. Specifically, if the card values are $c_1, c_2, \ldots, c_k$, then we have $\sum_{i=1}^{k} c_i = 0$ and $\sum_{i=1}^{k} c_i^2 = 1$. Jonathan and Yiming also agree on a positive target value of α.

The cards are then going to be dealt randomly in the following fashion: for each card in the deck, a fair coin is flipped. If the coin lands heads, then the card goes to Yiming, and if the coin lands tails, the card goes to Jonathan. Note that it is possible for either player to end up with no cards/all the cards.

A player wins the game if the sum of the card values in their hand is at least α, otherwise it is a tie.

(a) Prove that the probability that Yiming wins is at most $\frac{1}{8\alpha^2}$.

(b) Find a deck of k cards and target value α where the probability that Yiming wins is exactly $\frac{1}{8\alpha^2}$.

Solution:

(a) Let I_i be the indicator random variable indicating whether or not card i goes to Yiming. Define $S = \sum_{i=1}^{k} c_i I_i$ as the value of Yiming's hand. Then, we see that $\mathbb{E}[S] = \sum_{i=1}^{k} c_i \cdot \frac{1}{2} = 0$ and

$$\begin{aligned}
\text{var}(S) &= \sum_{i=1}^{k} \text{var}(c_i I_i) \qquad \text{(due to independence) of } I_i \\
&= \sum_{i=1}^{k} c_i^2 \,\text{var}(I_i)
\end{aligned}$$

We know that I_i is a Bernoulli random variable, so its variance is $\frac{1}{4}$. Thus, we see that $\mathrm{var}(S) = \frac{1}{4}$.

By Chebyshev, we see that $\mathbb{P}(|S| \geq \alpha) \leq \frac{1}{4\alpha^2}$. Now we need to make a symmetry argument, specifically that for each value of x, $\mathbb{P}(S = x) = \mathbb{P}(S = -x)$. This is true because for each outcome where Yiming gets x, Jonathan gets $-x$, since the sum of the card values is 0. However, we also know that the reverse outcome, where Jonathan gets Yiming's cards and vice versa, has the same probability of happening.

Since the distribution of S is symmetric around 0, we see that $\mathbb{P}(|S| \geq \alpha) = 2\mathbb{P}(S \geq \alpha)$, and plugging into our bound yields $\mathbb{P}(S \geq \alpha) \leq \frac{1}{8\alpha^2}$.

(b) We now need to appeal to the equality case of Chebyshev's inequality. Recall that the derivation of Chebyshev's inequality uses Markov's inequality on the quantity $(S - \mathbb{E}[S])^2$. Let's walk through the proof that $\mathbb{P}(S^2 \geq \alpha^2) \leq \frac{\mathbb{E}[S^2]}{\alpha^2}$ again:

$$\begin{aligned}
\mathbb{E}[S^2] &= \sum_v \mathbb{P}(S^2 = v) \cdot v \\
&= \sum_{0 \leq v < \alpha^2} \mathbb{P}(S^2 = v) \cdot v + \sum_{v \geq \alpha^2} \mathbb{P}(S^2 = v) \cdot v \\
&\geq \sum_{v \geq \alpha^2} \mathbb{P}(S^2 = v) \cdot v \\
&\geq \mathbb{P}(S^2 = \alpha^2) \cdot \alpha^2
\end{aligned}$$

In order for equality to hold, then equality must hold in both the third and fourth steps. We got the third step by saying that v is always at least 0, so we can drop them from the sum. Equality holds if it is not possible for S^2 to be anything strictly between 0 and α^2. We get the fourth line by observing that since $\mathbb{P}(S^2 = v) \cdot v \geq 0$ for all $v > \alpha^2$, we can also drop them from the sum. If we want equality to hold, then these values must also be 0, meaning that S^2 cannot take on values beyond α^2. This means that S^2 is either 0 or α.

If that's the case then the values of the cards can only be $-\alpha, 0$, or α, since it is possible for Yiming to get exactly one card. There also cannot exist two cards with value α, since otherwise Yiming could potentially end up with a hand value of $2\alpha \neq \alpha$. Thus, the deck must be of the form $(\alpha, -\alpha, 0, 0, \ldots, 0)$, and we pick $\alpha = \frac{1}{\sqrt{2}}$ to ensure that the sum of squares must be 1.

3 Sampling a Gaussian With Uniform

In this question, we will see one way to generate a normal random variable if we have access to a random number generator that outputs numbers between 0 and 1 uniformly at random.

As a general comment, remember that showing two random variables have the same CDF or PDF is sufficient for showing that they have the same distribution.

(a) First, let us see how to generate an exponential random variable with a uniform random variable. Let $U_1 \sim Uniform(0,1)$. Prove that $-\ln U_1 \sim Expo(1)$.

(b) Let $N_1, N_2 \sim \mathcal{N}(0,1)$, where N_1 and N_2 are independent. Prove that $N_1^2 + N_2^2 \sim Expo(1/2)$.

Hint: You may use the fact that over a region R, if we convert to polar coordinates $(x,y) \to (r,\theta)$, then the double integral over the region R will be

$$\iint_R f(x,y)\,dx\,dy = \iint_R f(r\cos\theta, r\sin\theta) \cdot r\,dr\,d\theta.$$

(c) Suppose we have two uniform random variables, U_1 and U_2. How would you transform these two random variables into a normal random variable with mean 0 and variance 1?

Hint: What part (b) tells us is that the point (N_1, N_2) will have a distance from the origin that is distributed as the square root of an exponential distribution. Try to use U_1 to sample the radius, and then use U_2 to sample the angle.

Solution:

(a) The CDF of an exponential $Expo(\lambda)$ distribution is $1 - e^{-\lambda t}$. Let us prove that the $-\ln(U_1)$ also has the same CDF.

We see that

$$\begin{aligned}
\mathbb{P}(-\ln(U_1) \leq t) &= \mathbb{P}(\ln(U_1) \geq -t) \\
&= \mathbb{P}(U_1 \geq e^{-t}) \\
&= 1 - e^{-t}
\end{aligned}$$

This shows that $-\ln(U_1)$ has an exponential distribution with $\lambda = 1$.

(b) We compute the CDF of $N_1^2 + N_2^2$. We want the probability that $N_1^2 + N_2^2 \leq t$ for some t. This means that we are integrating the joint distribution over a circle of radius $\sqrt{t}$, centered at the origin. We therefore compute the following integral

$$\iint_{(x,y):x^2+y^2 \leq t} \frac{1}{2\pi} e^{-(x^2+y^2)/2}\,dx\,dy = \int_0^{2\pi} \int_0^{\sqrt{t}} \frac{1}{2\pi} r e^{-r^2/2}\,dr\,d\theta$$

Evaluating this integral yields

$$\int_0^{2\pi} -\frac{e^{-r^2/2}}{2\pi}\bigg|_0^{\sqrt{t}}\,d\theta = \int_0^{2\pi} \frac{1 - e^{-t/2}}{2\pi}\,d\theta = 1 - e^{-t/2}.$$

This proves that $N_1^2 + N_2^2 \sim Expo(1/2)$.

(c) We will sample the point (N_1, N_2) using uniform random variables U_1 and U_2. We first sample the radius, which we know is an exponential distribution. Therefore, we know that $-2\ln(U_1)$ is an exponential $1/2$ distribution, so $\sqrt{-2\ln(U_1)}$ can be our radius. Since the (N_1, N_2) joint distribution is rotationally symmetric, we know that we can pick our angle uniformly at random once the radius is determined. Therefore, we let $\theta = 2\pi U_2$.

We will actually arrive at two Gaussians, so we can just take N_1, which will be

$$\boxed{\sqrt{-2\ln(U_1)}\cos(2\pi U_2)}$$

4 Optimal Gambling

Jonathan has a coin that may be biased, but he doesn't think so. You disagree with him though, and he challenges you to a bet. You start off with X_0 dollars. You and Jonathan then play multiple rounds, and each round, you bet an amount of money of your choosing, and then coin is tossed. Jonathan will match your bet, no matter what, and if the coin comes up heads, you win and you take both yours and Jonathan's bet, and if it comes up tails, then you lose your bet.

(a) Now suppose you actually secretly know that the bias of the coin is $\frac{1}{2} < p < 1$! You use the following strategy: on each round, you will bet a fraction q of the money you have at the start of the round. Let X_n denote the amount of money you have on round n. X_0 represents your initial assets and is a constant value. What is $\mathbb{E}[X_n]$ in terms of X_0, p, and q?

(b) What value of q will maximize $\mathbb{E}[X_n]$? For this value of q, what is the distribution of X_n? Can you predict what will happen as $n \to \infty$? [*Hint*: Under this betting strategy, what happens if you ever lose a round?]

(c) The problem with the previous approach is that we were too concerned about expected value, so our gambling strategy was too extreme. Let's start over: again we will use a gambling strategy in which we bet a fraction q of our money at each round. Express X_n in terms of n, q, X_0, and W_n, where W_n is the number of rounds you have won up until round n. [*Hint*: Does the order in which you win the games affect your profit?]

(d) By the law of large numbers, $W_n/n \to p$ as $n \to \infty$. Using this fact, what does $(\ln X_n)/n$ converge to as $n \to \infty$?

(e) The rationale behind $(\ln X_n)/n$ is that if $(\ln X_n)/n \to c$, where c is a constant, then that means for large n, X_n is roughly e^{cn}. Therefore, c is the asymptotic growth rate of your fortune! Find the value of q that maximizes your asymptotic growth rate.

(f) Using the value of q you found in the previous part, compute $\mathbb{E}[X_n]$.

(g) Now, Jonathan is not going to always believe that the coin is fair. What he will do is play 100 rounds with you, and observe how many times the coin comes up heads. Jonathan will then construct a 95% confidence interval for the true bias of the coin, p. True or false, the probability that Jonathan's interval captures p is at least 95%.

(h) Let's say after playing 100 rounds, Jonathan observes that 74 heads have appeared. Help Jonathan construct a 95% confidence interval using the CLT. Also, answer true or false: the probability that this interval contains the true bias p of the coin is 95%. You may assume that $\Phi(2) - \Phi(-2) \approx 0.95$, where Φ is the CDF of the standard Gaussian.

Solution:

(a) At the start of round n, the amount of money you have is X_n. With probability $1-p$, you will lose the round and $X_{n+1} = (1-q)X_n$. With probability p, you will win the round and

$X_{n+1} = (1-q)X_n + 2qX_n = (1+q)X_n$. Therefore,

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

By the law of iterated expectation,

$$\mathbb{E}[X_{n+1}] = \mathbb{E}[\mathbb{E}[X_{n+1} \mid X_n]] = ((1-p)(1-q) + p(1+q))\,\mathbb{E}[X_n].$$

Therefore,

$$\mathbb{E}[X_n] = ((1-p)(1-q) + p(1+q))^n X_0.$$

(b) We want $(1-p)(1-q)+p(1+q)$ to be as large as possible. Note that this is linear in q, and the coefficient for q is $p-(1-p) > 0$. Hence, we should take q to be as large as possible, which is 1 (you cannot bet more money than you actually have).

For this value of q, note that on each round you either double your money or go broke. Hence, the distribution is:

$$X_n = \begin{cases} 2^n X_0, & \text{with probability } p^n \\ 0, & \text{with probability } 1-p^n \end{cases}$$

Uh-oh. As $n \to \infty$, the probability that you are broke approaches 1. The issue here is that your expected fortune grows large, but the probability that you are rich grows vanishingly small. In general, $X_n \to 0$ as $n \to \infty$ does not necessarily imply that $\mathbb{E}[X_n] \to 0$, which is what we see here.

(c) You win W_n times and each time you win, your fortune is multiplied by $1+q$; you lose $n-W_n$ times, and each time you lose, your fortune is multiplied by $1-q$. Therefore,

$$X_n = X_0(1-q)^{n-W_n}(1+q)^{W_n}.$$

(d) One has

$$\begin{aligned} \frac{\ln X_n}{n} &= \frac{\ln X_0}{n} + \left(1 - \frac{W_n}{n}\right)\ln(1-q) + \frac{W_n}{n}\ln(1+q) \\ &\xrightarrow[n\to\infty]{} (1-p)\ln(1-q) + p\ln(1+q). \end{aligned}$$

(e) We can use calculus to optimize c:

$$\frac{\mathrm{d}}{\mathrm{d}q}((1-p)\ln(1-q) + p\ln(1+q)) = -\frac{1-p}{1-q} + \frac{p}{1+q}.$$

Set the derivative to 0:

$$\frac{p}{1+q} = \frac{1-p}{1-q} \implies p - pq = 1 + q - p - pq \implies q = 2p-1.$$

This is known as the Kelly betting criterion. Notice the farther from $\frac{1}{2}$ that p is, the more that you bet. If $p = 1$, you would all-in every time. The CS 70 course staff is not responsible for any losses you incur with this betting strategy, but we do think it's pretty cool that you can analyze optimal gambling with the methods in this course.

(f) We can plug in $q = 2p - 1$ from our previous result:

$$\mathbb{E}[X_n] = ((1-p)(1-q) + p(1+q))^n X_0 = 2^n (p^2 + (1-p)^2)^n X_0.$$

(g) True. This is the definition of confidence interval. The important part is that this is the probability that the confidence interval captures the probability p *before* any observations/calculations are done.

(h) We will use W_n/n as our estimator for the bias of the coin. We see that

$$\operatorname{var}(W_n/n) = \frac{1}{n^2}\operatorname{var}(W_n) = \frac{p(1-p)}{n} \leq \frac{1}{4n}$$

Then, for $n = 100$, we see that the variance of our estimator is at most $\frac{1}{400}$, so the standard deviation is at most $\frac{1}{20}$.

By the central limit theorem, we therefore know that $W_{100}/100$ is roughly a normal distribution with mean p and standard deviation of $\frac{1}{20}$. Therefore, we have

$$\mathbb{P}\left(\left|\frac{W_{100}}{100} - p\right| \leq 2 \cdot \frac{1}{20}\right) \leq 0.95$$

Thus, our resulting confidence interval should be $0.74 \pm 2 \cdot \frac{1}{20} = (0.64, 0.84)$.

The answer to the true/false is FALSE. The bias of the coin p is not chosen randomly, so it does not make sense to ask what the probability of it landing between 0.64 and 0.84 is. We say that we are 95% confident because *prior to constructing the interval*, the probability that the resulting interval captures the true p is 95%.

5 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).

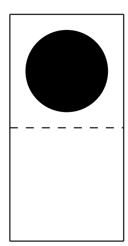

Figure 1: A straw with one boba currently inside. The straw only has enough room to fit two boba.

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.

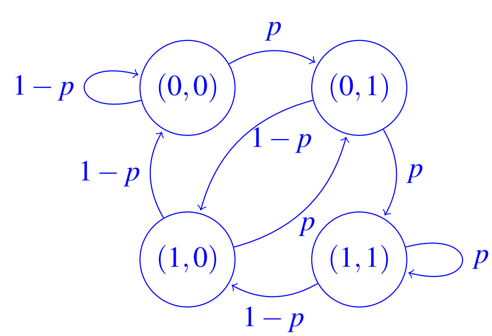

Figure 2: Transition diagram for the Markov chain.

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

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

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

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

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 $2p$.