CSTBC Exam 3 Solutions
Due: August 13, 2007, 11:59pm

August 14, 2007

This exam is open notes/open lecture and covers material from lectures 17-24. You are welcome to use any of the course material linked from the CSTBC website. You should not use other reference materials. Please send me your solutions to the exam via email. If you have any questions, please ask me.

1 Balls in Bins

Suppose that m balls are thrown into n bins uniformly and independently at random. In terms of m and n, what is the expected number of bins that contain at least one ball?

Solution Let X be the number of bins that have at least one ball, and for each $1 \leq j \leq n$, let X_j be an indicator random variable which has value 1 if the jth bin contains at least one ball, and zero otherwise. We have that $X = \sum_{j=1}^{n} X_j$, and by linearity of expectation, we have that

$$\mathbf{E}[X] = \sum_{j=1}^{n} \mathbf{E}[X_j].$$

By the definition of expectation,

$$\begin{aligned} \mathbf{E}[X_j] &= 0 \cdot \mathbf{Pr}(X_j = 0) + 1 \cdot \mathbf{Pr}(X_j = 1) \\ &= \mathbf{Pr}(X_j = 1). \end{aligned}$$

Because $X_j = 1$ and $X_j = 0$ are complementary events, $\mathbf{Pr}(X_j = 1) = 1 - \mathbf{Pr}(X_j = 0)$; the advantage is that $\mathbf{Pr}(X_j = 0)$ is easier to compute. By definition, $X_j = 0$ is the event that no balls land in the jth bin. For each $1 \leq r \leq m$, let A_r be the event that the rth ball does not land in the jth bin; note that $X_j = 0$ if and only if $\bigcap_r A_r$ occurs. Because the balls are thrown uniformly, $\mathbf{Pr}(A_r) = \frac{n-1}{n}$. Because the balls are thrown independently, the events $\{A_r : 1 \leq r \leq m\}$ are mutually independent and

$$\begin{aligned} \mathbf{Pr}(X_j = 0) &= \mathbf{Pr}\left(\bigcap_{r=1}^{m} A_r\right) \\ &= \prod_{r=1}^{m} \mathbf{Pr}(A_r) \\ &= \left(\frac{n-1}{n}\right)^m. \end{aligned}$$

Putting all the pieces together, we find that

$$\begin{aligned} \mathbf{E}[X] &= \sum_{j=1}^{n} \mathbf{E}[X_j] \\ &= \sum_{j=1}^{n} \mathbf{Pr}(X_j = 1) \\ &= \sum_{j=1}^{n} (1 - \mathbf{Pr}(X_j = 0)) \\ &= \sum_{j=1}^{n} \left(1 - \left(\frac{n-1}{n}\right)^m\right) \\ &= n\left(1 - \left(\frac{n-1}{n}\right)^m\right). \end{aligned}$$

2 A Hat Game

A group of three people play the following game. Each person is blindfolded and given a hat to wear. The hats are one of two colors (red or blue), and the colors are assigned independently and uniformly at random. Once each person is wearing a hat, the blindfolds are removed and each person can see the other two hats, but is unable to see his/her own hat. Without communicating in any way, each person simultaneously writes down one of three responses on a sheet of paper: "I guess my hat is red", "I guess my hat is blue", or "I choose not to guess".

The group wins if it is the case that no one guesses incorrectly and at least one person guesses correctly. Describe a way for the group to play the game so that they win with probability 3/4. (The group is allowed to discuss their strategy before the game begins, but once the game begins, no communication is allowed.)

Solution The group plays the game as follows. Each person plays according to the same rules. If a player sees that the other two people are wearing differently colored hats, the player writes down "I choose not to guess". If a player sees that the other two people are wearing hats of the same color, the player writes down a guess that his/her hat color is different from the common color of the two other players. Note that the players win the game if and only if the hats are not all the same color, which happens with probability 3/4.

3 Interviewing Candidates

A manager at a software company must hire a new programmer. The manager has the resumes of n people who applied for the job and schedules interviews as follows. First, the manager shuffles all the resumes so that they are in a random order. Next, the manager looks at each resume in order and schedules an interview whenever the current applicant's resume is superior to all previous resumes. What is the expected number of interviews that the manager schedules?

Solution Let X be the number of interviews scheduled, and for $1 \leq j \leq n$ let X_j be an indicator random variable which has value 1 if the jth best candidate is scheduled for an interview, and 0 otherwise. Note that $X = \sum_{j=1}^{n} X_j$, and therefore

$$\mathbf{E}[X] = \sum_{j=1}^{n} \mathbf{Pr}(X_j = 1)$$

(see problem 1 for the details we have skipped here). What is the probability that $X_j = 1$? By definition, $X_j = 1$ if and only if the jth best candidate is scheduled for an interview. For $j = 1$, the probability of this event is 1; for $j = n$, note that the worst candidate is interviewed if and only if the worst candidate appears first in the shuffled order of resumes, which happens with probability $1/n$.

In general, the jth best candidate is scheduled for an interview if and only if the jth candidate's resume appears in the shuffled order before the resumes of the $j-1$ superior candidates. That is, the jth best candidate is scheduled for an interview if and only if the jth candidate appears first in the shuffled order after removing the resumes from the $(j+1)$st, $(j+2)$nd, ..., nth best candidates. This event happens with probability $1/j$. Therefore

$$\mathbf{E}[X] = \sum_{j=1}^{n} \mathbf{Pr}(X_j = 1) = \sum_{j=1}^{n} \frac{1}{j} = H_n.$$

4 A Random Walk

In Lecture 23, we introduce the concept of a random walk on a line of n spaces. In that case, the token moves to the left with probability 1/2 and to the right with probability 1/2. Suppose instead that the token moves to the left with probability 1/3 and to the right with probability 2/3. Now what is the probability that we win (i.e. the token falls off of the right hand side)?

Solution We present the modifications to the argument seen in Lecture 23; see this lecture for more details. Let p_n be the probability that we win. We have

$$\begin{aligned} p_n = \mathbf{Pr}(\text{win}) &= \mathbf{Pr}(\text{win} \mid \text{1st move left}) \cdot \mathbf{Pr}(\text{1st move left}) + \\ &\quad \mathbf{Pr}(\text{win} \mid \text{1st move right}) \cdot \mathbf{Pr}(\text{1st move right}) \\ &= 0 \cdot \frac{1}{3} + \mathbf{Pr}(\text{win} \mid \text{1st move right}) \cdot \frac{2}{3} \\ &= \frac{2}{3} \cdot \mathbf{Pr}(\text{win} \mid \text{1st move right}). \end{aligned}$$

Let A be the event that the token falls off on the right side when the token starts from the second square, so that $\mathbf{Pr}(A) = \mathbf{Pr}(\text{win} \mid \text{1st move right})$, and let B be the event that the token falls off the smaller board consisting of the rightmost $n-1$ squares on the right side. We have

$$\begin{aligned} \mathbf{Pr}(A) &= \mathbf{Pr}(A|B) \cdot \mathbf{Pr}(B) + \mathbf{Pr}\left(A|\overline{B}\right) \cdot \mathbf{Pr}\left(\overline{B}\right) \\ &= 1 \cdot \mathbf{Pr}(B) + \mathbf{Pr}\left(A|\overline{B}\right) \cdot \mathbf{Pr}\left(\overline{B}\right) \\ &= 1 \cdot p_{n-1} + p_n (1 - p_{n-1}). \end{aligned}$$

Hence $p_n = \frac{2}{3}(p_{n-1} + p_n(1-p_{n-1}))$ which solves to the recurrence

$$p_n = \frac{2p_{n-1}}{1+2p_{n-1}}$$

with base case $p_1 = 2/3$. (Alternatively, we could set $p_0 = 1$ which yields to $p_1 = 2/3$.)

We can solve this recurrence with the guess and check method; the first few values of p_n are $p_1 = \frac{2}{3}$, $p_2 = \frac{4}{7}$, $p_3 = \frac{8}{15}$ which prompts us to guess

$$p_n = \frac{2^n}{2^{n+1}-1}$$

which we may verify by induction (left as an exercise). Note that as n grows, the probability that we win now approaches the constant $1/2$.

5 More Pirates

In Exam 2, we had the following problem: "After distributing their treasure, our n pirates from Exam 1 have worked up an appetite. The pirate ship's cafeteria offers three different, non-overlapping dinner times. As the strongest pirate, you are in charge of assigning each pirate to one of the three dinner slots. Unfortunately, not all pirates get along with each other. You have a list of k pairs of pirates that have fought each other in the past. Prove that you can assign pirates to dinners so that at most $k/3$ of the troublesome pairs eat dinner at the same time."

Suppose that n is divisible by three and that the ship's cafeteria has limited seating, so that each dinner can accommodate only $n/3$ pirates. Prove that even with this added restriction, you can assign the pirates to dinner slots so that at most $k/3$ troublesome pairs eat dinner at the same time. [Hint: although an inductive proof is possible, the probabilistic method (see Lecture 24) is easier.]

Solution We assign the pirates to dinner times with the following probabilistic experiment. First, we choose a permutation of the pirates uniformly at random. The first $n/3$ pirates in the permutation eat at the first dinner, the second $n/3$ pirates eat at the second dinner, and the third $n/3$ pirates eat at the third dinner. Let X be the number of pairs of troublesome pirates that eat together. For each $1 \le j \le k$, let X_j be the indicator random variable which has value 1 if the jth bad pair eats dinner together and 0 otherwise. Note that $X = \sum X_j$. Also, we can compute $\mathbf{Pr}(X_j = 1)$ noting that $X_j = 1$ if and only if weaker pirate in the jth bad pair eats at the same time as the stronger pirate in the jth bad pair. Regardless of when the stronger pirate in the jth bad pair eats, the probability that the weaker pirate eats at the same time is $\frac{(n/3)-1}{n}$, so that $\mathbf{Pr}(X_j = 1) = \frac{(n/3)-1}{n}$. (Formally, what we have done here is conditionalize on the three events that the stronger pirate in the jth bad pair eats during the first, second, and third time slots.) Therefore, we have

$$\begin{aligned} \mathbf{E}[X] &= \sum_{j=1}^{k} \mathbf{E}[X_j] \\ &= \sum_{j=1}^{k} \mathbf{Pr}(X_j = 1) \\ &= \sum_{j=1}^{k} \frac{\frac{n}{3}-1}{n-1} \\ &= \frac{1}{3}\sum_{j=1}^{k} \frac{n-3}{n-1} \\ &= \frac{k}{3}\left(\frac{n-3}{n-1}\right) \\ &\le \frac{k}{3} \end{aligned}$$

It follows that at least one of the permutations of pirates leads to a dinner schedule in which at most $\frac{k}{3}$ of the bad pairs eat at the same time.