1 Balls in Bins Estimation

Note 20

We throw $n > 0$ balls into $m \geq 2$ bins. Let X and Y represent the number of balls that land in bin 1 and 2 respectively.

(a) Calculate $\mathbb{E}[Y \mid X]$. [*Hint*: Your intuition may be more useful than formal calculations.]

(b) What is $L[Y \mid X]$ (where $L[Y \mid X]$ is the best linear estimator of Y given X)? [*Hint*: Your justification should be no more than two or three sentences, no calculations necessary! Think carefully about the meaning of the conditional expectation.]

(c) Unfortunately, your friend is not convinced by your answer to the previous part. Compute $\mathbb{E}[X]$ and $\mathbb{E}[Y]$.

(d) Compute $\text{Var}(X)$.

(e) Compute $\text{cov}(X,Y)$.

(f) Compute $L[Y \mid X]$ using the formula. Ensure that your answer is the same as your answer to part (b).

Solution:

(a) $\mathbb{E}[Y \mid X = x] = (n-x)/(m-1)$, because once we condition on x balls landing in bin 1, the remaining $n-x$ balls are distributed uniformly among the other $m-1$ bins. Therefore,

$$\mathbb{E}[Y \mid X] = \frac{n-X}{m-1}.$$

(b) We showed that $\mathbb{E}[Y \mid X]$ is a linear function of X. Since $\mathbb{E}[Y \mid X]$ is the best *general* estimator of Y given X, it must also be the best *linear* estimator of Y given X, i.e. $\mathbb{E}[Y \mid X]$ and $L[Y \mid X]$ coincide.

(c) Let X_i be the indicator that the ith ball falls in bin 1. Then, $X = \sum_{i=1}^{n} X_i$, and by linearity of expectation, $\mathbb{E}[X] = \sum_{i=1}^{n} \mathbb{E}[X_i] = n/m$, since there are n indicators and each ball has a probability $1/m$ of landing in bin 1. By symmetry, $\mathbb{E}[Y] = n/m$ as well.

(d) The number of balls that falls into the first bin is binomially distributed with parameters n and $1/m$. Hence the variance is $n(1/m)(1-1/m)$.

(e) Let X_i be as before, and let Y_i be the indicator that the ith ball falls into bin 2.

$$\text{cov}(X,Y) = \sum_{i=1}^{n} \sum_{j=1}^{n} \text{cov}(X_i, Y_j)$$

We can compute $\text{cov}(X_i, Y_i) = \mathbb{E}[X_i Y_i] - \mathbb{E}[X_i]\,\mathbb{E}[Y_i] = 0 - (1/m)(1/m) = -1/m^2$ (note that $\mathbb{E}[X_i Y_i] = 0$ because it is impossible for a ball to land in both bins 1 and 2). Also, we have $\text{cov}(X_i, Y_j) = 0$ because the indicator for the ith ball is independent of the indicator for the jth ball when $i \neq j$. Hence, $\text{cov}(X, Y) = n(-1/m^2) = -n/m^2$.

(f)

$$\begin{aligned} L[Y \mid X] &= \mathbb{E}[Y] + \frac{\text{cov}(X,Y)}{\text{var}(X)}(X - \mathbb{E}[X]) \\ &= \frac{n}{m} + \frac{-n/m^2}{n(1/m)(1-1/m)}\left(X - \frac{n}{m}\right) \\ &= \frac{n}{m} - \frac{1}{m-1}\left(X - \frac{n}{m}\right) \\ &= \frac{mn - n - mX + n}{m(m-1)} = \frac{n - X}{m-1} \end{aligned}$$

2 Analyze a Markov Chain

Note 22 Consider a Markov chain with the state diagram shown below where $a, b \in (0, 1)$.

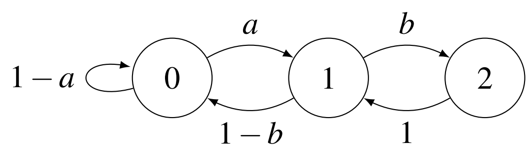

Here, we let $X(n)$ denote the state at time n.

(a) Is this Markov chain irreducible? Is this Markov chain aperiodic? Justify your answers.

(b) Calculate $\mathbb{P}[X(1) = 1, X(2) = 0, X(3) = 0, X(4) = 1 \mid X(0) = 0]$.

(c) Calculate the invariant distribution. Do all initial distributions converge to this invariant distribution? Justify your answer.

Solution:

(a) The Markov chain is irreducible because $a, b \in (0, 1)$. Also, $P(0,0) > 0$, so that

$$\gcd\{n > 0 \mid P^n(0,0) > 0\} = \gcd\{1, 2, 3, \ldots\} = 1,$$

which shows that the Markov chain is aperiodic.

We can also notice from the definition of aperiodicity that if a Markov chain has a self loop with nonzero probability, it is aperiodic. In particular, a self loop implies that the smallest number of steps we need to take to get from a state back to itself is 1. In this case, since $P(0,0) > 0$, we have a self loop with nonzero probability, which makes the Markov chain aperiodic.

(b) As a result of the Markov property, we know our state at timestep n depends only on timestep $n-1$. Looking at the transition probabilities, we see that the final expression is

$$P(0,1) \times P(1,0) \times P(0,0) \times P(0,1) = a(1-b)(1-a)a.$$

(c) The balance equations are

$$\begin{cases} \pi(0) = (1-a)\pi(0) + (1-b)\pi(1) \\ \pi(1) = a\pi(0) + \pi(2) \end{cases} \implies \begin{cases} a\pi(0) = (1-b)\pi(1) \\ \pi(1) = a\pi(0) + \pi(2) \end{cases}$$
$$\implies \begin{cases} a\pi(0) = (1-b)\pi(1) \\ \pi(1) = a\left(\frac{1-b}{a}\pi(1)\right) + \pi(2) \end{cases}$$
$$\implies \begin{cases} a\pi(0) = (1-b)\pi(1) \\ b\pi(1) = \pi(2) \end{cases}$$

As a side note, these last equations express the equality of the probability of a jump from i to $i+1$ and from $i+1$ to i, for $i=0$ and $i=1$, respectively. These relations are also called the "detailed balance equations".

From these equations we find successively that

$$\pi(1) = \frac{a}{1-b}\pi(0) \qquad \pi(2) = b\pi(1) = \frac{ab}{1-b}\pi(0).$$

The normalization equation is

$$1 = \pi(0) + \pi(1) + \pi(2) = \pi(0)\left(1 + \frac{a}{1-b} + \frac{ab}{1-b}\right)$$
$$1 = \pi(0)\left(\frac{1-b+a+ab}{1-b}\right)$$

so that

$$\pi(0) = \frac{1-b}{1-b+a+ab}.$$

Thus,

$$\pi(0) = \frac{1-b}{1-b+a+ab} \qquad \pi(1) = \frac{a}{1-b+a+ab} \qquad \pi(2) = \frac{ab}{1-b+a+ab}$$

Or in vector form,

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

Since the Markov chain is irreducible and aperiodic, all initial distributions converge to this invariant distribution by the fundamental theorem of Markov chains.

3 A Bit of Everything

Note 22

Suppose that $X_0, X_1, \ldots$ is a Markov chain with finite state space $S = \{1, 2, \ldots, n\}$, where $n > 2$, and transition matrix P. Suppose further that

$$P(1,i) = \frac{1}{n} \quad \text{for all states } i \text{ and}$$
$$P(j, j-1) = 1 \quad \text{for all states } j \neq 1,$$

with $P(i,j) = 0$ everywhere else.

(a) Prove that this Markov chain is irreducible and aperiodic.

(b) Suppose you start at state 1. What is the distribution of T, where T is the number of transitions until you leave state 1 for the first time?

(c) Again starting from state 1, what is the expected number of transitions until you reach state n for the first time?

(d) Again starting from state 1, what is the probability you reach state n before you reach state 2?

(e) Compute the stationary distribution of this Markov chain.

Solution:

(a) For any two states i and j, we can consider the path $(i, i-1, \ldots, 2, 1, j)$, which has nonzero probability of occurring. Thus, this chain is irreducible. To see that it is aperiodic, observe that $d(1) = 1$, as we have self-loop from state 1 to itself.

(b) At any given transition, we leave state 1 with probability with probability $\frac{n-1}{n}$, independently of any previous transition. Thus, the distribution is Geometric, with parameter $\frac{n-1}{n}$.

(c) Suppose that $\beta(i)$ is the expected number of transitions necessary to reach state n for the first time, starting from state i. We have the following first step equations:

$$\beta(1) = 1 + \sum_{j=1}^{n} \frac{1}{n}\beta(j),$$
$$\beta(i) = 1 + \beta(i-1) \quad \text{for } 1 < i < n, \text{ and}$$
$$\beta(n) = 0.$$

We can simplify the second recurrence to

$$\beta(i) = i - 1 + \beta(1) \quad \text{for } 1 < i < n.$$

Substituting this simplified recurrence into the first equation, we get that

$$\beta(1) = 1 + \frac{1}{n}\sum_{i=1}^{n-1}(i - 1 + \beta(1)) = 1 + \frac{1}{n}\sum_{i=1}^{n-1}(i-1) + \frac{1}{n}\sum_{i=1}^{n-1}\beta(1) = 1 + \frac{(n-2)(n-1)}{2n} + \frac{n-1}{n}\beta(1),$$

which we can solve to get that

$$\beta(1) = \boxed{n + \frac{1}{2}(n-1)(n-2)}.$$

(d) Suppose that $\alpha(i)$ is the probability that we reach state n before we reach state 2, starting from state i. One immediate observation we can make is that from any state i in $\{2,\dots,n-1\}$, we are guaranteed to see state 2 before state n, as we can only take the path $(i, i-1, \dots, 2, 1)$. Hence, $\alpha(i) = 0$ if $i \in \{2,\dots,n-1\}$. Moreover, $\alpha(n) = 1$, so

$$\alpha(1) = \sum_{i=1}^{n} \frac{1}{n}\alpha(i) = \frac{1}{n}\alpha(1) + \frac{1}{n},$$

hence $\alpha(1) = \boxed{\frac{1}{n-1}}$.

(e) We have the balance equations

$$\pi(i) = \frac{1}{n}\pi(1) + \pi(i+1) \quad \text{if } i \neq n, \text{ and}$$
$$\pi(n) = \frac{1}{n}\pi(1).$$

We can collapse the first recurrence to

$$\pi(i) = \frac{n-i}{n}\pi(1) + \pi(n) = \frac{n-i+1}{n}\pi(1),$$

so we can express each stationary probability in terms of the stationary probability of state 1. We can finish by using the normalization equation:

$$\pi(1) + \pi(2) + \cdots + \pi(n) = 1 \implies \frac{1}{n}\pi(1)\sum_{i=1}^{n} n-i+1 = 1.$$

The last sum can be rearranged to be the sum of the integers from 1 up to n, so we get that

$$\pi(1) = \frac{2}{n+1} \implies \pi = \boxed{\frac{2}{n(n+1)}\begin{bmatrix} n & n-1 & \cdots & 1 \end{bmatrix}}.$$

4 Playing Blackjack

Note 22

Suppose you start with \$1, and at each turn, you win \$1 with probability p, or lose \$1 with probability $1-p$. You will continually play games of Blackjack until you either lose all your money, or you have a total of n dollars.

(a) Formulate this problem as a Markov chain.

(b) Let $\alpha(i)$ denote the probability that you end the game with n dollars, given that you started with i dollars.

Notice that for $0 < i < n$, we can write $\alpha(i+1) - \alpha(i) = k(\alpha(i) - \alpha(i-1))$. Find k.

(c) Using part (b), find $\alpha(i)$, where $0 \leq i \leq n$. (You will need to split into two cases: $p = \frac{1}{2}$ or $p \neq \frac{1}{2}$.)

Hint: Try to apply part (b) iteratively, and look at a telescoping sum to write $\alpha(i)$ in terms of $\alpha(1)$. The formula for the sum of a finite geometric series may be helpful when looking at the case where $p \neq \frac{1}{2}$:

$$\sum_{k=0}^{m} a^k = \frac{1 - a^{m+1}}{1 - a}.$$

Lastly, it may help to use the value of $\alpha(n)$ to find $\alpha(1)$ for the last few steps of the calculation.

(d) As $n \to \infty$, what happens to the probability of ending the game with n dollars, given that you start with i dollars, with the following values of p?

(i) $p > \frac{1}{2}$

(ii) $p = \frac{1}{2}$

(iii) $p < \frac{1}{2}$

Solution:

(a) We have the following state transition diagram:

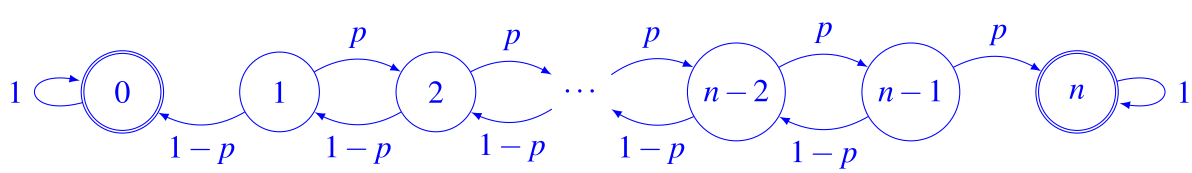

In particular, we have $n+1$ states, $\{0, 1, 2, \ldots, n\}$, where the transition probability from i to $i+1$ is p, and the transition probability from i to $i-1$ is $1-p$. The transition probabilities for $i = 0$ and $i = n$ are edge cases, where we stay in place with probability 1.

(b) If we start with i dollars, this means that we start at state i. The next transition can either be to state $i+1$ with probability p, or to state $i-1$ with probability $1-p$. This means that we have

$$\alpha(i) = p\alpha(i+1) + (1-p)\alpha(i-1).$$

Here, a trick is to expand $\alpha(i) = p\alpha(i) + (1-p)\alpha(i)$. Substituting this in, we can rewrite

$$\begin{aligned} p\alpha(i) + (1-p)\alpha(i) &= p\alpha(i+1) + (1-p)\alpha(i-1) \\ (1-p)(\alpha(i) - \alpha(i-1)) &= p(\alpha(i+1) - \alpha(i)) \\ \alpha(i+1) - \alpha(i) &= \frac{1-p}{p}(\alpha(i) - \alpha(i-1)) \end{aligned}$$

(c) Now that we have a relationship between $\alpha(i+1)-\alpha(i)$ and $\alpha(i)-\alpha(i-1)$, notice that we can iteratively apply the recurrence to get

$$\begin{aligned}
\alpha(i+1)-\alpha(i) &= \frac{1-p}{p}(\alpha(i)-\alpha(i-1)) \\
&= \left(\frac{1-p}{p}\right)^2(\alpha(i-1)-\alpha(i-2)) \\
&\vdots \\
&= \left(\frac{1-p}{p}\right)^i(\alpha(1)-\alpha(0)) \\
&= \left(\frac{1-p}{p}\right)^i\alpha(1)
\end{aligned}$$

since $\alpha(0)=0$ (once we lose all our money, we stop and can never reach n).

Further, notice that we have the telescoping sum

$$[\alpha(i)-\alpha(i-1)]+[\alpha(i-1)-\alpha(i-2)]+\cdots+[\alpha(1)-\alpha(0)] = \alpha(i)-\alpha(0) = \alpha(i).$$

This means that we have the summation

$$\begin{aligned}
\alpha(i) &= \sum_{k=0}^{i-1}(\alpha(k+1)-\alpha(k)) \\
&= \sum_{k=0}^{i-1}\left(\frac{1-p}{p}\right)^k\alpha(1) \\
&= \alpha(1)\sum_{k=0}^{i-1}\left(\frac{1-p}{p}\right)^k \\
&= \alpha(1)\cdot\frac{1-\left(\frac{1-p}{p}\right)^i}{1-\frac{1-p}{p}}
\end{aligned}$$

[Note that if $p=\frac{1}{2}$, the last step is not valid; in fact, since $\frac{1-p}{p}=1$, this means that $\alpha(i)=i\alpha(1)$. We'll come back to this case later.]

The previous formula applies for all $0<i\le n$, so we can let $i=n$ and simplify to find $\alpha(1)$:

$$\begin{aligned}
1 = \alpha(n) &= \alpha(1)\cdot\frac{1-\left(\frac{1-p}{p}\right)^n}{1-\frac{1-p}{p}} \\
\frac{1-\frac{1-p}{p}}{1-\left(\frac{1-p}{p}\right)^n} &= \alpha(1)
\end{aligned}$$

Plugging this back in for $\alpha(i)$, we have

$$\alpha(i) = \frac{1-\frac{1-p}{p}}{1-\left(\frac{1-p}{p}\right)^n} \cdot \frac{1-\left(\frac{1-p}{p}\right)^i}{1-\frac{1-p}{p}} = \frac{1-\left(\frac{1-p}{p}\right)^i}{1-\left(\frac{1-p}{p}\right)^n}.$$

Going back to the case where $p = \frac{1}{2}$, we saw that the summation simplifies to $\alpha(i) = i\alpha(1)$. Since $\alpha(n) = 1$, this means that $1 = n\alpha(1)$, or $\alpha(1) = \frac{1}{n}$. This means that we have

$$\alpha(i) = i\alpha(1) = \frac{i}{n}.$$

Together, we have the following formula for any $0 \le i \le n$:

$$\alpha(i) = \begin{cases} \frac{1-\left(\frac{1-p}{p}\right)^i}{1-\left(\frac{1-p}{p}\right)^n} & p \neq \frac{1}{2} \\ \frac{i}{n} & p = \frac{1}{2} \end{cases}.$$

(d) (i) If $p > \frac{1}{2}$, then $\frac{1-p}{p} < 1$, and as $n \to \infty$, the $\left(\frac{1-p}{p}\right)^n$ term in the denominator vanishes. This means that all we're left with is the numerator, and as such

$$\lim_{n\to\infty} \alpha(i) = 1 - \left(\frac{1-p}{p}\right)^i.$$

(ii) If $p = \frac{1}{2}$, then we know that $\alpha(i) = \frac{i}{n}$. As $n \to \infty$, this fraction goes to 0, and we have

$$\lim_{n\to\infty} \alpha(i) = 0.$$

(iii) If $p < \frac{1}{2}$, then $\frac{1-p}{p} > 1$, and as $n \to \infty$, the $\left(\frac{1-p}{p}\right)^n$ term in the denominator blows up. This means that the denominator tends to $-\infty$, while the numerator remains bounded for any fixed i. This means that the entire fraction tends to 0, i.e,

$$\lim_{n\to\infty} \alpha(i) = 0.$$

Note that this problem shows that, even in the case of a fair game (i.e., $p = \frac{1}{2}$), the probability that a gambler wins $n before going broke tends to zero as $n \to \infty$. This is one version of the so-called "Gambler's Ruin" problem. Only in the case where $p > \frac{1}{2}$, i.e., when the game is strictly in the gambler's favor, does the gambler come out on top with positive probability.