

Solutions to Problem Set 13-14

Problem 1. A gambler plays 120 hands of draw poker, 60 hands of black jack, and 20 hands of stud poker per day. He wins a hand of draw poker with probability $1/6$, a hand of black jack with probability $1/2$, and a hand of stud poker with probability $1/5$. Assume the outcomes of the card games are mutually independent.

(a) What is the expected number of hands the gambler wins in a day?

Solution. $120(1/6) + 60(1/2) + 20(1/5) = 54$. ■

(b) What is the variance in the number of hands won per day?

Solution. The variance can also be calculated using linearity of variance. For an individual hand the variance is $p(1 - p)$ where p is the probability of winning. Therefore the variance is $120(1/6)(5/6) + 60(1/2)(1/2) + 20(1/5)(4/5) = 50/3 + 15 + 16/5 = 250/15 + 225/15 + 48/15 = 523/15 = 34 \frac{13}{15}$. ■

(c) What would the Markov bound be on the probability that the gambler will win 108 hands on a given day?

Solution. The expected number of games won is 54, so by Markov, $\Pr\{R \geq 108\} \leq 54/108 = 1/2$. ■

(d) What would the Chebyshev bound be on the probability that the gambler will win 108 hands on a given day?

Solution.

$$\Pr\{R - 54 \geq 54\} \leq \Pr\{|R - 54| \geq 54\} \leq \frac{V}{54^2} = \frac{523/15}{54^2} \approx 1/85.$$

(e) Apply the Chernoff bound for the probability that the gambler will win the 108 hands to find the largest integer, n , such that the bound is less than e^{-n} . *Hint:* $\ln 2 \approx 0.7$.

Solution. Answer: $n = 21$.

$$\begin{aligned} \Pr\{R \geq 108\} &= \Pr\{R \geq 2 \cdot 54\} \\ &\leq \exp(-(2 \ln 2 - 2 + 1)54) \\ &= \exp(-108 \ln 2 + 54) \\ &\approx \exp(-108 \cdot 0.7 + 54) \\ &\approx \exp(-75.6 - 54) \\ &\approx \exp(-21.6) \\ &\approx \frac{1}{2,400,000,000} \end{aligned}$$

■

Problem 2. A man has a set of n keys, one of which fits the door to his apartment. He tries the keys until he finds the correct one. Give the expected number and variance for the number of trials until success if

(a) he tries the keys at random (possibly repeating a key tried earlier)

Solution. There are $n - 1$ wrong keys and 1 right key. Each time the man chooses a key, he tries it and then places it back to the key set. Let T be the waiting time for the man to pick the right key. $T = k$ means that on the k -th trial the man picks the right key. The probability of picking the right key on the first trial is $1/n$, and the probability of picking the wrong key is $(n - 1)/n$. The probability stays the same on all subsequent trials since the picked key is always returned to the key set. Thus

$$P(T = k) = \left(\frac{n-1}{n}\right)^{k-1} \frac{1}{n}.$$

Note: this is the same as the waiting time for the first head in the Bernoulli process with a biased coin. Let $p = 1/n$ and let $q = 1 - p$. Then, the expectation of T is

$$E(T) = \sum_{k=1}^{\infty} kq^{k-1}p = p \sum_{k=1}^{\infty} kq^{k-1} = p \frac{1}{(1-q)^2} = \frac{1}{p} = n$$

The variance of T can be computed directly from the formula,

$$\text{Var}[T] = E(T^2) - E^2(T).$$

Again, by definition we have

$$E(T^2) = \sum_{k=1}^{\infty} k^2 pq^{k-1}.$$

How do we evaluate this sum? We use the following trick. We know that

$$\sum_{k=1}^{\infty} kq^k = q \sum_{k=1}^{\infty} kq^{k-1} = \frac{q}{(1-q)^2}.$$

Moreover,

$$\begin{aligned} \sum_{k=1}^{\infty} k^2 q^{k-1} &= \sum_{k=1}^{\infty} \frac{d}{dq} kq^k \\ &= \frac{d}{dq} \sum_{k=1}^{\infty} kq^k \\ &= \frac{d}{dq} \frac{q}{(1-q)^2} \\ &= \frac{1}{(1-q)^2} + \frac{2q}{(1-q)^3} = \frac{1}{p^2} + \frac{2q}{p^3}. \end{aligned}$$

Finally, the expectation $E(T^2)$ is

$$E(T^2) = p \left(\frac{1}{p^2} + \frac{2q}{p^3} \right) = \frac{1}{p} + \frac{2q}{p^2},$$

and the variance is

$$\text{Var}[T] = \frac{1}{p} + \frac{2q}{p^2} - \frac{1}{p^2} = \frac{q}{p^2} = n(n-1).$$

■

(b) he chooses keys randomly from among those he has not yet tried.

Solution. $T = k$ means that the man picks the wrong key on the first trial, and he picks the wrong key on the second trial, etc, and he picks the right key on the k -th trial. Let K_i be the indicator random variable for the i -th trial, i.e. $K_i = 1$ if he picks the right key on the i -th trial, and 0 otherwise. Then

$$P(T = k) = P((K_1 = 0) \& (K_2 = 0) \& \cdots \& (K_{k-1} = 0) \& (K_k = 1)).$$

By the Multiplication Theorem we can compute

$$\begin{aligned} P(T = k) &= P(K_1 = 0)P(K_2 = 0|K_1 = 0)P(K_3 = 0|K_1 = 0, K_2 = 0) \cdots \\ &\quad \cdots P(K_k = 1|K_1 = 0, \dots, K_{k-1} = 0) \\ &= \frac{n-1}{n} \frac{n-2}{n-1} \frac{n-3}{n-2} \cdots \frac{n-k+1}{n-k+2} \frac{1}{n-k+1} \\ &= \frac{1}{n}. \end{aligned}$$

The expectation and variance are now easy to compute from the definitions.

$$\begin{aligned}
 E(T) &= \sum_{k=1}^n kP(T=k) = \frac{1}{n} \sum_{k=1}^n k = \frac{n+1}{2}. \\
 \text{Var}[T] &= E(T^2) - E^2(T) \\
 &= \sum_{k=1}^n k^2 P(T=k) - \left(\frac{n+1}{2}\right)^2 \\
 &= \frac{1}{n} \sum_{k=1}^n k^2 - \left(\frac{n+1}{2}\right)^2 \\
 &= \frac{1}{n} \frac{n(n+1)(2n+1)}{6} - \left(\frac{n+1}{2}\right)^2 \\
 &= \frac{n^2-1}{12}.
 \end{aligned}$$

■

Problem 3. We have two coins: one is a fair coin and the other is a coin that produces heads with probability $3/4$. One of the two coins is picked, and this coin is tossed n times. How many tosses suffice to make us 95% confident which coin was chosen? Explain.

Solution. To guess which coin was picked, set a threshold t between $1/2$ and $3/4$. If the proportion of heads is less than the threshold, guess it was the fair coin; otherwise, guess the biased coin. Let the random variable H_n be the number of heads in the first n flips. We need to flip the coin enough times so that $\Pr(H_n/n > t) \leq 0.05$ if the fair coin was picked, and $\Pr(H_n/n < t) \leq 0.05$ if the biased coin was picked. A natural threshold to choose is $5/8$, exactly in the middle of $1/2$ and $3/4$.

H_n is the sum of independent Bernoulli variables, which each have variance $1/4$ for the fair coin and $3/16$ for the biased coin. Using Chebyshev's Inequality for the fair coin,

$$\begin{aligned}
 \Pr\left(\frac{H_n}{n} > \frac{5}{8}\right) &= \Pr\left(\frac{H_n}{n} - \frac{1}{2} > \frac{5}{8} - \frac{1}{2}\right) = \Pr\left(H_n - \frac{n}{2} > \frac{n}{8}\right) \\
 &= \Pr\left(H_n - E[H_n] > \frac{n}{8}\right) \leq \Pr\left(|H_n - E[H_n]| > \frac{n}{8}\right) \\
 &\leq \frac{\text{Var}[H_n]}{(n/8)^2} = \frac{n/4}{n^2/64} = \frac{16}{n}
 \end{aligned}$$

For the biased coin, we have

$$\begin{aligned}
 \Pr\left(\frac{H_n}{n} < \frac{5}{8}\right) &= \Pr\left(\frac{3}{4} - \frac{H_n}{n} > \frac{3}{4} - \frac{5}{8}\right) = \Pr\left(\frac{3n}{4} - H_n > \frac{n}{8}\right) \\
 &= \Pr\left(E[H_n] - H_n > \frac{n}{8}\right) \leq \Pr\left(|H_n - E[H_n]| > \frac{n}{8}\right) \\
 &\leq \frac{\text{Var}[H_n]}{(n/8)^2} = \frac{3n/16}{n^2/64} = \frac{12}{n}
 \end{aligned}$$

We are 95% confident if these are at most 0.05, which is satisfied if $n \geq 320$.

Because the variance of the biased coin is less than that of the fair coin, we can do slightly better if we make our threshold a bit bigger, to about 0.634, which gives 95% confidence with 279 coin flips.

Because H_n has a binomial distribution, we can get a much better bound using the estimates from Lecture 21, giving 95% confidence when $n > 42$. ■

Problem 4. An Unbiased Estimator

Suppose we are trying to estimate some physical parameter p . When we run our experiments and process the results, we obtain an estimator of p , call it p_e . But if our experiments are probabilistic, then p_e itself is a random variable which has a pdf over some range of values. We call the random variable p_e an *unbiased* estimator if $E[p_e] = p$.

For example, say we are trying to estimate the height, h , of Green Hall. However, each of our measurements has some noise that is, say, Gaussian with zero mean. So each measurement can be viewed as a sample from a random variable X . The expected value of each measurement is thus $E[X] = h$, since the probabilistic noise has zero mean. Then, given n independent trials, x_1, \dots, x_n , an unbiased estimator for the height of Green Hall would be

$$h_e = \frac{x_1 + \dots + x_n}{n},$$

since

$$E[h_e] = E\left[\frac{x_1 + \dots + x_n}{n}\right] = \frac{E[x_1] + \dots + E[x_n]}{n} = E[x_1] = h.$$

Now say we take n independent observations of a random variable Y . Let the true (but unknown) variance of Y be $\text{Var}[Y] = \sigma^2$. Then, using the alternative definition of variance:

$$\text{Var}[Y] = E[Y^2] - E^2[Y],$$

we could by the same reasoning use σ_e as an estimate of σ , where

$$\sigma_e^2 ::= \frac{y_1^2 + y_2^2 + \dots + y_n^2}{n} - \left(\frac{y_1 + y_2 + \dots + y_n}{n}\right)^2.$$

Now it turns out that $E[\sigma_e^2] \neq \sigma^2$, but a minor modification of the definition yields an unbiased estimator $(\sigma'_e)^2$.

(a) What is $E[\sigma_e^2]$?

Solution. Let $\sigma^2 = \text{Var}[X]$, $\mu = E[X]$. Then our estimator σ_e^2 is given by

$$\begin{aligned}\sigma_e^2 &= \frac{\sum y_i^2}{n} - \left(\frac{\sum y_i}{n}\right)^2 \\ E[\sigma_e^2] &= E\left[\frac{\sum y_i^2}{n} - \left(\frac{\sum y_i}{n}\right)^2\right] \\ &= \frac{\sum E[y_i^2]}{n} - \frac{E[(\sum y_i)^2]}{n^2} \\ &= \frac{\sum(\sigma^2 + \mu^2)}{n} - \frac{\text{Var}[\sum y_i] + E^2[\sum y_i]}{n^2} \\ &= \frac{n(\sigma^2 + \mu^2)}{n} - \frac{n\sigma^2 + n^2\mu^2}{n^2} \\ &= \sigma^2 \left(1 - \frac{1}{n}\right)\end{aligned}$$

So this gives a biased estimator. ■

(b) How should the unbiased estimator $(\sigma'_e)^2$ be defined?

Solution. We can make the previous estimate unbiased simply by defining

$$(\sigma'_e)^2 ::= \frac{n}{n-1} \sigma_e,$$

so

$$E[(\sigma'_e)^2] = \frac{n}{n-1} E[\sigma_e] = \frac{n}{n-1} \left(1 - \frac{1}{n}\right) \sigma^2 = \sigma^2.$$
■

Problem 5. The covariance, $\text{Cov}(X, Y)$, of two random variables, X and Y , is defined to be $E(XY) - E(X)E(Y)$. Note that if two random variables are independent, then their covariance is zero.

(a) Give an example to show that having $\text{Cov}(X, Y) = 0$ does not necessarily mean that X and Y are independent.

Solution. Let (X, Y) have joint probability given by the table below:

X	Y	P
-1	1	$\frac{1}{3}$
0	0	$\frac{1}{3}$
1	1	$\frac{1}{3}$

Note that X and Y are not independent. $P(X = 1 \& Y = 1) = \frac{1}{3} \neq \frac{2}{9} = P(X = 1)P(Y = 1)$. But $E(X) = 0$ and $E(XY) = 0$ (since $XY = X$). Thus $\text{Cov}(X, Y) = 0$. ■

(b) Let X_1, \dots, X_n be random variables. Prove that

$$\text{Var}(X_1 + \dots + X_n) = \sum_{i=1}^n \text{Var}(X_i) + 2 \sum_{i<j} \text{Cov}(X_i, X_j)$$

Solution.

$$\begin{aligned} \text{Var}(X_1 + \dots + X_n) &= E[X_1 + \dots + X_n]^2 - E[X_1 + \dots + X_n]^2 \\ &= E\left[\sum_i X_i^2 + \left(\sum_{i<j} 2X_i X_j\right)\right] - \left[\sum_i E(X_i)^2 + \sum_{i<j} 2E(X_i)E(X_j)\right] \\ &= \sum_i E(X_i^2) + \sum_{i<j} 2E(X_i X_j) - \sum_i E(X_i)^2 - \sum_{i<j} 2E(X_i)E(X_j) \\ &= \sum_i E(X_i^2) - E(X_i)^2 + \sum_{i<j} 2(E(X_i X_j) - E(X_i)E(X_j)) \\ &= \sum_i \text{Var}(X_i) + 2 \sum_{i<j} \text{Cov}(X_i, X_j) \end{aligned}$$

■

Problem 6. A Gambler plays a game in casino. Let p be some fixed probability he will win an individual \$1 bet and get back \$2. Let $q = 1 - p$ be the probability he will lose the \$1 bet. The Gambler has some fixed goal of T dollars.

(a) Let $e(n)$ be the expected number of bets the Gambler must make until the game ends (the Gambler either gets \$ T or is bankrupt), where $n \leq T$ is his initial capital, *i.e.*, the number of dollars he starts with. Then the function e satisfies a linear recurrence of the form

$$e(n) = a \cdot e(n+1) + b \cdot e(n-1) + c.$$

for constants $a, b, c \in \mathbb{R}$. What are a, b and c ?

Solution. $a = p, b = 1 - p, c = 1$.

From Notes 13: Let Q be the number of bets till the Gambler's Ruin game ends. So

$$e(n) = E[Q] = p E[Q \mid \text{gambler wins first bet}] + q E[Q \mid \text{gambler loses first bet}].$$

But after the gambler wins the first bet, his capital is $n + 1$, so he can expect to make another $e(n + 1)$ bets. That is,

$$E[Q \mid \text{gambler wins first bet}] = 1 + e(n + 1),$$

and similarly,

$$E[Q \mid \text{gambler loses first bet}] = 1 + e(n - 1).$$

So we have

$$e(n) = p(1 + e(n + 1)) + q(1 + e(n - 1)) = pe(n + 1) + qe(n - 1) + 1. \quad \blacksquare$$

We consider $T = \infty$. Define R to be the event that the the Gambler wins his first bet, but that his capital eventually gets back to where it started at n , and let $r ::= \Pr\{R\}$. Define D to be the event the Gambler's capital ever goes down to $n - 1$, and let $d ::= \Pr\{D\}$.

(b) Explain why $r \leq p$.

Solution. R can only hold if the Gambler wins the first bet, so by monotonicity

$$\Pr\{R\} \leq \Pr\{\text{Gambler wins first bet}\} = p. \quad \blacksquare$$

(c) Explain why r and d do not depend on n . Why isn't this true in a bounded game?

Solution. Take any walk starting with \$1 and ending when it gets back to \$1 for the first time after one or more bets. Note that the first step of such a walk has to be up, or the Gambler would have gone broke on the first bet,

Now add $n - 1$ to the y -coordinate of every point on this walk. Because the game is unbounded, the result is a path starting and with capital n and ending when it gets back to n for the first time. Also, the shifted path has the same probability as the original path. This correspondence is clearly a bijection, since it has an inverse—namely subtract $n - 1$. Hence the sum of the probabilities of the set of all paths starting with \$1 and ending with a first return to \$1 is the same as the sum for all paths starting at n and ending at the first step where the capital is again n . So the probability, r , of return to n is the same as the probability of a return to 1, and therefore does not depend on n . The same correspondence holds between walks which start at \$1 and end at \$0, and walks which start at n and end the first time they reach $n - 1$, so d similarly does not depend on n .

This correspondence fails if there is a bound on the game, *e.g.*, if the Gambler starts with $n = T$, then the game ends immediately and there is no chance that the Gambler will win his first bet, because there won't be any first bet. \blacksquare

(d) Explain why $r = pd$.

Solution. The only way R can occur is if the Gambler wins his first bet, and then his new stake of $n + 1$ eventually goes down to n . That is

$$r = \Pr\{R\} = \Pr\{\text{win first bet}\} \Pr\{\text{go from } n + 1 \text{ to } n\} = pd. \quad \blacksquare$$

(e) Argue that

$$r = pq + prq + pr^2q + \cdots + pr^kq + \cdots .$$

Solution. The stake can get back to n for the first time after going up to $n + 1$, by returning to $n + 1$ repeatedly k times for some $k \geq 0$ and then going down to n . The probability of this is the probability, p , of going up on the first bet, times the probability, r^k , of repeating k times going above and returning to $n + 1$, times the probability of going from $n + 1$ to n on the final bet. Hence the probability of going up and eventually getting back to n is the sum of the probabilities of these disjoint events. ■

(f) Conclude that if $p \leq 1/2$, then $r = p$.

Solution. From the previous part, we have

$$r = pq \sum_{k=0}^{\infty} r^k = \frac{pq}{1-r}$$

So we have a quadratic

$$r^2 - r + pq = 0$$

so

$$\begin{aligned} r &= 1 \pm \frac{\sqrt{1-4pq}}{2} \\ &= 1 \pm \frac{\sqrt{1-4p(1-p)}}{2} \\ &= 1 \pm \frac{\sqrt{1-4p+4p^2}}{2} \\ &= 1 \pm \frac{\sqrt{(2p-1)^2}}{2} \\ &= 1 \pm \frac{2p-1}{2} \\ &= \begin{cases} 1-p, & \text{or} \\ p. \end{cases} \end{aligned}$$

If $p = 1/2$, then $p = 1 - p$, so in any case $r = 1/2$. If $p < 1/2$, then $p < 1 - p$, and since $r \leq p$ by the part (a), we cannot have $r = 1 - p$, so it must be that $r = p$. ■

(g) Conclude that if $p \leq 1/2$, then the Gambler goes broke with probability 1, no matter what his initial stake. *Hint:* First assume his initial capital is \$1.

Solution. We know that if $p \leq 1/2$, then $r = p$, but we also know $r = pd$. So $p = pd$, and since $p \neq 0$, it follows that $d = 1$. In particular, the Gambler goes broke with probability 1 if his initial

capital in \$1. Now by induction, assume he goes broke with probability 1 when his initial capital is n . Then,

$$\begin{aligned}
 & \Pr \{ \text{Gambler goes broke with initial capital } n + 1 \} \\
 &= \Pr \{ \text{capital goes from } n + 1 \text{ to } n \} \\
 & \quad \cdot \Pr \{ \text{go broke} \mid \text{initial capital } n \} \\
 &= d \cdot 1 \qquad \qquad \qquad (\text{by induction hypothesis}) \\
 &= 1 \cdot 1 = 1.
 \end{aligned}$$

■

(h) Now let $t ::= E$ [number of bets till the gambler's stake first goes down by \$1]. Prove that if t is finite, then

$$E[\text{number of bets till the gambler goes bankrupt}] = \frac{n}{1 - 2p}.$$

(It turns out that t is finite iff $p < 1/2$, but we won't prove this today.)

Hint: Prove that $t = 1/(1 - 2p)$ by conditioning on the first bet.

Solution. To go down 1, a walk must go down immediately, or else go up 1 and then eventually down 1, and then eventually go down 1 again. So,

$$t = 1 \cdot \Pr \{ \text{lose first bet} \} + \Pr \{ \text{win first bet} \} (1 + 2t) = (1 - p) + p(2t + 1).$$

That is,

$$t = 1 + 2pt,$$

so if $t \neq \infty$, we conclude that

$$t = \frac{1}{1 - 2p}.$$

Hence, if the gambler starts off with \$ n , then the expected time to go bankrupt is $nt = n/(1 - 2p)$.

■