

Problem Set 13-14

Reading: [Week 13-14 Notes](#)

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?
- (b) What is the variance in the number of hands won per day?
- (c) What would the Markov bound be on the probability that the gambler will win 108 hands on a given day?
- (d) What would the Chebyshev bound be on the probability that the gambler will win 108 hands on a given day?
- (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:* $\ln 2 \approx 0.7$.

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

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.

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 $E[\sigma_e^2]$?
- (b) How should the unbiased estimator $(\sigma'_e)^2$ be defined?

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

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

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 ?

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 $r \leq p$.

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

(d) Explain why $r = pd$.

(e) Argue that

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

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

(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.

(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.

