

1. Let a biased coin have $Pr[Heads] = p$, where p is not necessarily 0.5. This coin can still be used to simulate fair coin tosses by tossing it twice, where if the actual sequence of tosses is HT, the output is “H”, while if the actual sequence of tosses is TH, the output is “T”. If the actual tosses are both the same, repeat the process. Let an “event” be a pair of tosses, and let the random variable X count the number of events until “H” or “T” is output, including the last, so $X \geq 1$.

- (a) What kind of distribution does X have (give a name)?

X has an exponential distribution with parameter $P = 2p(1 - p)$

- (b) What is $E[X]$ in terms of p ?

For an exponential r.v. $E[X] = 1/P$ and since $P = 2p(1 - p)$, we have

$$E[X] = 1/(2p - 2p^2)$$

- (c) How could you use the same coin to simulate a fair 3-sided dice, i.e. to output 1, 2, or 3 with equal probability?

Method 1: Toss the coin 3 times. The outcomes THH, HTH, HHT are all equally likely. Output 1, 2 or 3 respectively if they occur, repeat otherwise.

Method 2: Toss the coin 3 times. If all 3 results are the same, repeat. If not, one toss is different from the other two. Output 1, 2 or 3, depending on whether the odd outcome is on the first, second or third toss.

2. Let X be a random variable with $E[X] = 3$, $Var[X] = 4$, and let Y be a random variable with $E[Y] = 5$.

- (a) What is $E[3X]$?

By linearity of expected values, $E[3X] = 9$

- (b) What is $E[X + Y]$?

By linearity of expected values, $E[X + Y] = 8$

- (c) What is $E[XY]$?

Assuming X and Y independent, $E[XY] = 15$

- (d) What is $E[X^2]$?

Using the variance formula $VAR[X] = E[X^2] - (E[X])^2$ and rearranging, $E[X^2] = 4 + 3^2 = 13$

3. Pizza Palace has implemented a new electronic dispatching system but are having trouble with it. Because of that, m pizzas are dispatched independently and uniformly at random to n customers.

(a) How large should m be (in terms of n) to be confident that everyone gets a pizza?

This is an instance of the coupon collectors problem. Therefore we want $m > n \ln n + \Omega(n)$, say m at least $n \ln n + 5n$.

(b) How large should m be to be confident that no-one gets more than one pizza?

This is an instance of the birthday problem. The probability of no double deliveries is $\exp(-m(m-1)/2n)$. We want this probability close to one. Choosing $m = \sqrt{(n)}/4$ would work well.

(c) Assume $m = n$, that everyone ordered a pizza, and that one pizza gets delivered to each person but in a random permutation of the orders. What is the probability that no-one got the pizza that they ordered?

This is the probability that the permutation is a derrangement (a permutation with no fixed points). The probability of zero fixed points is $\approx 1/e$.

4. Let a fair coin be tossed 16 times. Let X be the number of heads, then $E[X] = 8$ and $Var[X] = 4$

(a) Give a Markov bound for $Pr[X \geq 12]$.

By Markov $Pr[X \geq 12] \leq E[X]/12 = 2/3$

(b) Give a Chebyshev bound for $Pr[X \geq 12]$.

By Chebyshev $Pr[X \geq 12 = E[X] + 2\sigma] \leq 1/2^2 = 1/4$

(c) Give a Chernoff bound for $Pr[X \geq 12]$.

By Chernoff (upper tail, $\delta < 2e - 1$) we have $Pr[X \geq 12 = \mu(1 + 0.5)] \leq \exp(-\mu\delta^2/4) = \exp(-8 \times (0.5^2)/4) = e^{-0.5}$