

1. Will I Get My Package?

Sneaky delivery guy of some company is out delivering n packages to n customers. Not only does he hand a random package to each customer, he tends to open a package before delivering with probability $\frac{1}{2}$. Let X be the number of customers who receive their own packages unopened.

- (a) Compute the expectation $E(X)$.
- (b) Compute the variance $\text{Var}(X)$.

2. Variance

This problem will give you practice using the “standard method” to compute the variance of a sum of random variables that are not pairwise independent (so you cannot use “linearity” of variance).

- (a) A building has n floors numbered $1, 2, \dots, n$, plus a ground floor G. At the ground floor, m people get on the elevator together, and each gets off at a uniformly random one of the n floors (independently of everybody else). What is the *variance* of the number of floors the elevator *does not* stop at? (In fact, the variance of the number of floors the elevator *does* stop at must be the same, but the former is a little easier to compute.)
- (b) A group of three friends has n books they would all like to read. Each friend (independently of the other two) picks a random permutation of the books and reads them in that order, one book per week (for n consecutive weeks). Let X be the number of weeks in which all three friends are reading the same book. Compute $\text{Var}(X)$.

3. Markov’s Inequality and Chebyshev’s Inequality

A random variable X has variance $\text{Var}(X) = 9$ and expectation $\mathbb{E}(X) = 2$. Furthermore, the value of X is never greater than 10. Given this information, provide either a proof or a counterexample for the following statements.

- (a) $\mathbb{E}(X^2) = 13$.
- (b) $\Pr[X = 2] > 0$.
- (c) $\Pr[X \geq 2] = \Pr[X \leq 2]$.
- (d) $\Pr[X \leq 1] \leq 8/9$.
- (e) $\Pr[X \geq 6] \leq 9/16$.
- (f) $\Pr[X \geq 6] \leq 9/32$.

4. Easy A's

A friend tells you about a course called “Laziness in Modern Society” that requires almost no work. You hope to take this course next semester to give yourself a well-deserved break after mastering CS70. At the first lecture, the professor announces that grades will depend only a midterm and a final. The midterm will consist of three questions, each worth 10 points, and the final will consist of four questions, also each worth 10 points. He will give an A to any student who gets at least 60 of the possible 70 points.

However, speaking with the professor in office hours you hear some very disturbing news. He tells you that, in the spirit of the class, the GSIs are very lazy, and to save time the grading will be done as follows. For each student's midterm, the GSIs will choose a real number randomly from a normal distribution with mean $\mu = 5$ and variance $\sigma^2 = 1$. They'll mark each of the three questions with that score. To grade the final, they'll again choose a random number from the same distribution, independent of the first number, and will mark all four questions with that score.

If you take the class, what will the mean and variance of your total class score be? Use Chebyshev's inequality to conclude that you have less than a 5% chance of getting an A.