Lecture

1. Notice that for a function g and random variable X, the function $g(X)$ is also a random variable. Therefore, we can compute the expected value of $g(X)$, in the discrete case,

$$\mathbf{E}[g(X)] = \sum_y y\mathbf{Pr}[g(X) = y] = \sum_x g(x)\mathbf{Pr}[X = x].$$

and in the continuous case,

$$\mathbf{E}[g(X)] = \int_{-\infty}^{\infty} y f_{g(X)}(y)\, dy = \int_{-\infty}^{\infty} g(x) f_X(x)\, dx$$

The first equality is the definition of $\mathbf{E}$ and the second equality is proved in the book.

2. The probability density function may seem like a strange way to determine continuous probabilities, but it's actually quite natural. For a continuous random variable X with p.d.f. f_X we have, by definition,

$$\mathbf{Pr}[a \leq X \leq b] = \int_a^b f_X(x)\, dx.$$

Observe the pleasant similarity to the discrete case. If Y has p.m.f. p_Y then we also have

$$\mathbf{Pr}[a \leq Y \leq b] = \sum_{y=a}^{b} p_Y(y).$$

3. Consider breaking a stick uniformly at random. The stick spans $[0, 1]$ and we will model the breaking spot as a random variable U uniform over $[0, 1]$. To make the problem more interesting, let's say there is a bug on the stick at a point $p \in [0, 1]$. Let $L_p(U)$ be the random variable representing the length of the part of the stick with the bug after the break. So,

$$L_p(U) = \begin{cases} 1 - U & \text{for } U < p \\ U & \text{for } U > p \\ \text{dead bug} & \text{for } U = p \end{cases}$$

Let's calculate the expected value of $L_p(U)$. Using the formula above,

$$\mathbf{E}[L_p(U)] = \int_{-\infty}^{\infty} L_p(x) f_U(x)\, dx.$$

Because U is uniform over $[0, 1]$, that means $f_U(x) = 1$ only for $0 \leq x \leq 1$. Thus,

$$\begin{aligned} \mathbf{E}[L_p(U)] &= \int_0^1 L_p(x)\, dx \\ &= \int_0^p 1 - x\, dx + \int_p^1 x\, dx \\ &= x\Big|_0^p - \frac{x^2}{2}\Big|_0^p + \frac{x^2}{2}\Big|_p^1 \\ &= p - \frac{p^2}{2} + \frac{1}{2} - \frac{p^2}{2} \\ &= \frac{1}{2} + p(1-p). \end{aligned}$$

Exercises

1. For any real numbers a and b, show that $\mathbf{E}[aX + b] = a\mathbf{E}[X] + b$.

2. For any real numbers a and b, prove that

$$var(aX + b) = var(aX) = a^2 var(X).$$

3. Consider modeling the time (in hours) that a device is active before it breaks down as a random variable X with probability density function

$$f(x) = \begin{cases} \lambda e^{-x/100} & \text{for } x \geq 0\,, \\ 0 & \text{for } x < 0. \end{cases}$$

   (a) Determine the value of λ so that the probability density function f satisfies the normalization property. That is, so that $\int_{-\infty}^{\infty} f(x)\, dx = 1$.

   (b) Calculate the probability that the device breaks down within the first 100 hours of use.

   (c) Calculate the probability the device breaks down after 50 but before 150 hours of use.

4. Recall that the Poisson random variable with parameter λ has probability mass function defined as

$$\mathbf{Pr}[X = k] = p(x) = e^{-\lambda}\frac{\lambda^k}{k!}.$$

   There's a lot to say about a Poisson random variable and the best way to learn about it is to read the book and google around. For now, let's observe (without any real observation...) that you can use the nicely tractable form of the Poisson to approximate the binomial. Specifically, for large number n of samples and for small probability p of success, it turns out that setting $\lambda = np$ gives a good approximation to a binomial random variable Y with parameters n and p. What I really mean is that

$$\mathbf{Pr}[X = k] \approx \mathbf{Pr}[Y = k].$$

   See this for yourself by calculating both for $k = 1, 2, 3, \ldots$ with $n = 10$ and $p = 0.1$.
   Quoted from wikipedia,

   > "According to two rules of thumb, this approximation is good
   > if $n \geq 20$ and $p \leq 0.05$, or if $n \geq 100$ and $np \leq 10$."

Puzzle 1: Three Dice

You have an opportunity to bet \$1 on a number between 1 and 6. Three dice are then rolled. If your number fails to appear, you lose \$1. If it appears once, you win \$1; if twice, \$2; if three times, \$3.

Is this bet in your favor, fair, or against the odds? How can you determine this without doing many calculations?

Puzzle 2: Rolling all the numbers

On average, how many times must you roll a die before all six different numbers have turned up?
Comment: there are more elegant and less elegant analyses.