

Continuous Random Variables, Joint Distributions

Before you leave lab, make sure you [click here](#) so that you're marked as having attended this week's section. The CA leading your discussion section can enter the password needed once you've submitted.

1 Warmups

1.1 Joint Distributions

- Given a Normal RV $X \sim N(\mu, \sigma^2)$, how can we compute $P(X \leq x)$ from the standard Normal distribution Z with CDF ϕ ?
- What is a continuity correction and when should we use it?
- If we have a joint PMF for discrete random variables $p_{X,Y}(x, y)$, how can we compute the marginal PMF $p_X(x)$?

1.2 Independent Random Variables

- What distribution does the sum of two independent binomial RVs $X + Y$ have, where $X \sim \text{Bin}(n_1, p)$ and $Y \sim \text{Bin}(n_2, p)$? Include any parameters with your answer.
- What distribution does the sum of two independent Poisson RVs $X + Y$ have, where $X \sim \text{Poi}(\lambda_1)$ and $Y \sim \text{Poi}(\lambda_2)$? Include any parameters with your answer.

2 Marguerite Gets Some Competition

In the late 1880s, Stanford began running a horse and twelve-person buggy service from the Stanford Quad to the train station just off campus. The name of this shuttling service was chosen to be **Marguerite**, which was the name of the favorite horse of some Stanford bigwig of the time. The horse-and-carriage operation was retired around 1910 and replaced with electric streetcars, which themselves were replaced with buses around 1930. The service has grown substantially since, and the buses have been upgraded several times. The service, however, has retained its name since the very beginning.

Several Stanford horse enthusiasts have recently revived the horse-and-buggy service to compete with Marguerite, and they've given it the name **Hildegard**. Now, when you need a ride from the Quad to the train station, you have two options!

You arrive at the Quad, headed to the train station, and you're equally happy to take either of the two independent services. You arrive precisely at 8:00am, which is the time that both services start for the day. The number of minutes you need to wait for a Marguerite bus is modeled by a **discrete** Uniform random variable $M \sim \text{Uni}(0, 20)$, whereas the number of minutes you need to wait for a Hildegard horse-and-buggy is modeled by a discrete Poisson random variable $H \sim \text{Poi}(10)$. (Yes, it's technically possible that Hildegard never arrives.)

- a. What is the probability that Marguerite and Hildegard both arrive at $t = 6$ minutes?
- b. What is the conditional probability that $H < M$, given $M = m$ —that is, what is $P(H < M | M = m)$? Express your answer as a sum.
- c. What is the unconditional probability that $H < M$, i.e., what is $P(H < M)$? Express your answer as a double sum that leverages your answer to part b.
- d. What is the CDF of your waiting time for the first of the two to arrive? You should leave your answer in summation form.

3 Burrow Smoke Detectors and Joint Probability Distributions

Burrow Labs has taken on other startups in the home safety and security space and has recently started marketing a new smoke detector. Burrow’s smoke detectors rely on CO_2 sensors that eventually fail, and that failure time dictates the average product lifetime of the smoke detector. Burrow manufactures three quarters of its smoke detectors in central Idaho, and the rest are manufactured in suburban Maine. Any single smoke detector’s product lifetime can be modeled as a Exponential random variable.

Each of the two locations sources its CO_2 sensors from different suppliers, so the smoke detectors manufactured in Maine have an average product lifetime of 7 years and the smoke detectors manufactured in Idaho have an average product lifetime of 6 years. All smoke detectors are sold online, so aside from the fact that a smoke detector is three times more likely to ship from the Idaho facility, you can’t tell by looking at a single smoke detector where it was manufactured.

Let T model the amount of time that passes until the CO_2 sensor (and therefore the smoke detector) fails, and let M be a discrete random variable that takes on the value of 1 for a smoke detector manufactured in Maine, and 0 otherwise.

- a. Present the cumulative distribution and probability density functions for the random variable T . Both your CDF and your PDF should be analytic functions on t .
- b. Compute the probability that a smoke detector was manufactured in Maine, given that it lasts more than 15 years. If needed, you can keep your answer in terms of $F_T(15)$ or $f_T(15)$ from part (a). However, any conditional expression of the form $P(\cdot|\cdot)$ or $f(\cdot|\cdot)$ must be evaluated.

4 Elections

We would like to see how we could predict an election between two candidates in France (A and B), given data from 10 polls. For each of the 10 polls, we report below their sample size, how many people said they would vote for candidate A, and how many people said they would vote for candidate B. Not all polls are created equal, so for each poll we also report a value “weight” which represents how accurate we believe the poll was. The data for this problem can be found on the class website in polls.csv:

Poll	N samples	A votes	B votes	Weight
1	862	548	314	0.93
2	813	542	271	0.85
3	984	682	302	0.82
4	443	236	207	0.87
5	863	497	366	0.89
6	648	331	317	0.81
7	891	552	339	0.98
8	661	479	182	0.79
9	765	609	156	0.63
10	523	405	118	0.68
Totals:	7453	4881	2572	

a. First, assume that each sample in each poll is an independent experiment of whether or not a random person in France would vote for candidate A (disregard weights).

- Calculate the probability that a random person in France votes for candidate A.
- Assume each person votes for candidate A with the probability you’ve calculated and otherwise votes for candidate B. If the population of France is 64,888,792, what is the probability that candidate A gets more than half of the votes?

b. Nate Silver at fivethirtyeight pioneered an approach called the “Poll of Polls” to predict elections. For each candidate A or B, we have a random variable S_A or S_B which represents their strength on election night (like ELO scores). The probability that A wins is $P(S_A > S_B)$.

- Identify the parameters for the random variables S_A and S_B . Both S_A and S_B are defined to be normal with the following parameters:

$$S_A \sim \mathcal{N}\left(\mu = \sum_i p_{A_i} \cdot \text{weight}_i, \sigma^2\right) \quad S_B \sim \mathcal{N}\left(\mu = \sum_i p_{B_i} \cdot \text{weight}_i, \sigma^2\right)$$

where p_{A_i} is the ratio of A votes to N samples in poll i , p_{B_i} is the ratio of B votes to N samples in poll i , weight_i is the weight of poll i , m_i is the N samples in poll i and:

$$\sigma = \frac{K}{\sqrt{\sum_i m_i}} \text{ s.t. } K = 350; \text{ thus } \sigma = 4.054.$$

- We will calculate $P(S_A > S_B)$ by simulating 100,000 fake elections. In each fake election, we draw a random sample for the strength of A from S_A and a random sample for the strength of B from S_B . If S_A is greater than S_B , candidate A wins. What do we expect to see if we simulate so many times? What do we actually see?

- c. Which model, the one from (a) or the model from (b) seems more appropriate? Why might that be the case? On election night candidate A wins. Was your prediction from part (b) "correct"?