Homework 7

PSTAT 5A: Spring 2023, with Ethan P. Marzban

Instructions

- Please submit your work to Gradescope by no later than **11:59pm on MONDAY, May 22**. As a reminder, late homework will not be accepted.
- Recall that you will be asked to upload a **single** PDF containing your work for *both* the programming and non-programming questions to Gradescope.
  - You can merge PDF files using either Adobe Acrobat, or using adobe's online PDF merger at this link.

Problem 1: Deriving the Lower-Tailed Hypothesis Test

Consider testing the set of hypothesis

$$\left[\begin{array}{rl} H_0: & p = p_0 \\ H_A: & p < p_0 \end{array}\right.$$

at an arbitrary α level of significance. Define the test statistic TS to be

$$\mathrm{TS} = \frac{\widehat{P} - p_0}{\sqrt{\frac{p_0(1-p_0)}{n}}}$$

a. Show that $\mathrm{TS} \stackrel{H_0}{\sim} \mathcal{N}(0,\ 1)$. If your answer depends on a set of conditions to be true, explicitly state those conditions.

Solution: So long as we are able to invoke the CLT for Proportions, we will be good. Hence, we need to first assure that both:

1) $np_0 \geq 10$

2) $n(1 - p_0) \geq 10$

Assume the above conditions are true. Then, under the null (i.e. assuming the true value of p is actually p_0), the CLT for proportions tells us

$$\widehat{P} \sim \mathcal{N}\left(p_0,\ \sqrt{\frac{p_0(1-p_0)}{n}}\right)$$

which means (by our familiar Standardization result)

$$\frac{\widehat{P} - p_0}{\sqrt{\frac{p_0(1-p_0)}{n}}} \stackrel{H_0}{\sim} \mathcal{N}(0,\ 1)$$

and we are done.

b. Argue, in words, that the test should be of the form

$$\texttt{decision}(\mathrm{TS}) = \begin{cases} \texttt{reject } H_0 & \text{if TS} < c \\ \texttt{fail to reject } H_0 & \text{otherwise} \end{cases}$$

for some constant c. As a hint, look up the logic we used in Lecture 13 to derive the two-tailed test, and think in terms of statements like "$\hat{p}$ is *far away* from p_0". **You do not have to find the value of c in this part**.

Solution: If the null hypothesis states that the true value of p is p_0, and if we observe an instance of $\hat{p}$ that is much less than p_0, we are more inclined to believe the alternative (i.e. that $p < p_0$) is true. In other words, we would reject the null for *small* values of TS; namely, our rejection region takes the form $(-\infty, c)$.

The key assertion, however, is that we would only really reject the null in favor of the alternative that $p < p_0$ if TS were small in *raw value*, **NOT** in absolute value. Said differently, observing a very large value of TS would **NOT** necessarily lead credence to the claim that $p < p_0$, and hence we would **NOT** reject the null in favor for the alternative if TS were large in the positive direction.

c. Now, argue that c must be the α^{th} percentile of the distribution of the standard normal distribution (**NOT** scaled by negative 1), thereby showing that the full test takes the form

$$\texttt{decision}(\mathrm{TS}) = \begin{cases} \texttt{reject } H_0 & \text{if TS} < z_\alpha \\ \texttt{fail to reject } H_0 & \text{otherwise} \end{cases}$$

where z_α denotes the $(\alpha) \times 100^{\text{th}}$ percentile of the standard normal distribution.

Solution: Recall that the level of significance α is precisely the probability of committing a Type I error; i.e. the probability of rejecting the null when the null were true:

$$\alpha = \mathbb{P}_{H_0}(\mathrm{TS} < c)$$

Since, under the null, $\mathrm{TS} \sim \mathcal{N}(0,\ 1)$ (as was shown in part (a) above), this means that c must satisfy

$$\mathbb{P}(Z < c) = \alpha$$

where $Z \sim \mathcal{N}(0,\ 1)$; i.e. c is the α^{th} percentile of the standard normal distribution.

Result: Upper-Tailed Test

When testing the hypotheses

$$\left[\begin{array}{lll} H_0 & : & p = p_0 \\ H_A & : & p > p_0 \end{array}\right.$$

at an α level of significance, the test takes the form

$$\texttt{decision(TS)} = \begin{cases} \texttt{reject } H_0 & \text{if TS} > z_{1-\alpha} \\ \texttt{fail to reject } H_0 & \text{otherwise} \end{cases}$$

where:

- $\mathrm{TS} = \dfrac{\widehat{P} - p_0}{\sqrt{\frac{p_0(1-p_0)}{n}}}$
- $z_{1-\alpha}$ denotes the $(1-\alpha) \times 100^{\text{th}}$ percentile of the standard normal distribution.

provided that

- $np_0 \geq 10$
- $n(1-p_0) \geq 10$

Problem 2: Airplanes (not in the Night Sky)

According to *USAToday*, around 2.75% of flights in 2022 were cancelled. To test this claim, Jaime collects data on a representative sample of 500 flights from 2022 and finds that only 2.01% of these flights were cancelled. Assume that Jaime wishes to perform a two-sided test, at an $\alpha = 0.05$ level of significance.

a. What is the population?

Solution: The population is the set of all flights in 2022.

b. What is the sample?

Solution: The sample is the set of 500 sampled flights.

c. Write down the null and alternative hypotheses for this problem. Use mathematical notation.

Solution: Let p denote the true proportion of flights in 2022 that were delayed. Then

$$\left[\begin{array}{ll} H_0 & p = 0.0275 \\ H_A & p \neq 0.0275 \end{array}\right.$$

d. Compute the value of the test statistic.

Solution:

$$\text{TS} = \frac{\hat{p} - p_0}{\sqrt{\frac{p_0(1-p_0)}{n}}} = \frac{0.0201 - 0.0275}{\sqrt{\frac{0.0275 \cdot 0.9725}{500}}} = \boxed{-1.01}$$

e. Compute the critical value of the test.

Solution: Because we are performing a two-sided hypothesis test at an $\alpha = 0.05$ level of significance, we find the $(0.05/2) \times 100 = 2.5^{\text{th}}$ percentile of the standard normal distribution and scale by negative 1: 1.96.

f. Conduct the test, and phrase your conclusions in the context of the problem.

Solution: We reject only when $|\text{TS}| > 1.96$. In this case, $|\text{TS}| = |-1.01| = 1.01 < 1.96$ and so we fail reject the null; that is,

> At an $\alpha = 0.05$ level of significance, there was insufficient evidence to reject the claim that 2.75% of flights in 2022 were delayed in favor of the alternative that the true proportion was *not* 2.75%.

Problem 3: Airplanes (still not in the Night Sky)

Consider again the setup of Problem 2, except now suppose Jaime wishes to conduct an upper-tailed test (still at an $\alpha = 0.05$ level of significance).

a. Does the value of the test statistic change from what you found in Problem 2(d)? If so, provide the new value.

Solution: The value does not change.

b. Does the critical value change from what you found in Problem 2(e)? If so, provide the new value.

Solution: The critical value *does* change: now, because we are conducting an upper-tailed test the critical value becomes the $(1-0.05)\times100 = 95^{\text{th}}$ percentile of the standard normal distribution, which is 1.645.

c. Conduct the test, and phrase your conclusions in the context of the problem.

Solution: We now compare the *raw* value of the test statistic to the new critical value: $-1.01 < 1.645$ which is *not* in the rejection region of the test; i.e. we fail to reject:

> At an $\alpha = 0.05$ level of significance, there was insufficient evidence to reject the claim that 2.75% of flights in 2022 were delayed in favor of the alternative that

the true proportion was less than 2.75%.

Problem 4: Watch The Time (Review Problem)

In a 2015 article, *CBC News* predicted that in 2018 31% of people would wear a watch. Suppose a representative sample of 204 people, taken in 2018, contained 65 people that wore a watch.

a. Assuming *CBC*'s claim is correct, what is the probability that a representative sample (assume it was taken with replacement) contained 65 people that wore a watch? **State your logic clearly, and check all assumptions that may need to be checked.**

Solution: Let X denote the number of people, in a representative sample of size 204, that wear a watch. We check the Binomial Criteria:

1) **Independent Trials?** Yes, since the sample was taken with replacement.
2) **Fixed number of Trials?** Yes; $n = 204$
3) **Well-defined notion of success?** Yes; "success" = "finding a person that wears a watch"
4) **Fixed probability of success?** Yes; assumed to be $p = 0.31$

Therefore, we conclude that $X \sim \mathrm{Bin}(204,\ 0.31)$ and so

$$\mathbb{P}(X = 65) = \binom{204}{65}(0.31)^{65}(1 - 0.31)^{204-65} \approx \boxed{0.0578 = 5.87\%}$$

b. Assuming *CBC*'s prediction was correct, what is the expected number of people who would be wearing a watch in a sample of 204 people (again, assume the sample was taken with replacement)?

Solution: Let X be defined as in part (a) above. We seek $\mathbb{E}[X]$, which we know can be computed using the formula for the expected value of the Binomial Distribution:

$$\mathbb{E}[X] = np = (204)(0.31) = \boxed{63.24}$$

c. Assuming *CBC*'s prediction was correct, what is the variance of the number of people who would be wearing a watch in a sample of 204 people (again, assume the sample was taken with replacement)?

Solution: We again let X be defined as in part (a); now we use the formula for the variance of the Binomial distribution:

$$\text{Var}(X) = np(1-p) = (204)(0.31)(1-0.31) = \boxed{43.636}$$

d. Assuming *CBC*'s prediction was correct, what is the probability that between 27.8% and 37.5% of people in a sample of size 204, taken with replacement, wear a watch?

Solution: If *CBC*'s claim is correct, then the true proportion of people that wear a watch in 2018 is 0.31. We therefore check the following success-failure conditions:

1) $np_0 = (204)(0.31) = 63.24 \geq 10$

2) $n(1-p_0) = (204)(1-0.31) = 140.76 \geq 10$

Since both conditions are satisfied, we can invoke the CLT for Proportions to conclude

$$\widehat{P} \sim \mathcal{N}\left(0.31,\ \sqrt{\frac{0.31 \cdot (1-0.31)}{204}}\right) \sim \mathcal{N}(0.31,\ 0.0324)$$

where $\widehat{P}$ denotes the proportion of people in a sample of 204 that wear a watch. We seek

$$\mathbb{P}(0.278 \leq \widehat{P} \leq 0.375)$$

which we compute as

$$\begin{aligned}
\mathbb{P}(0.278 \leq \widehat{P} \leq 0.375) &= \mathbb{P}(\widehat{P} \leq 0.375) - \mathbb{P}(\widehat{P} \leq 0.278) \\
&= \mathbb{P}\left(\frac{\widehat{P}-0.31}{0.0324} \leq \frac{0.375-0.31}{0.0324}\right) - \mathbb{P}\left(\frac{\widehat{P}-0.31}{0.0324} \leq \frac{0.278-0.31}{0.0324}\right) \\
&\approx \mathbb{P}\left(\frac{\widehat{P}-0.31}{0.0324} \leq 2.01\right) - \mathbb{P}\left(\frac{\widehat{P}-0.31}{0.0324} \leq -0.99\right) \\
&= 0.9778 - 0.1611 = \boxed{81.67\%}
\end{aligned}$$

where we obtained the final two values from the $z-$table.

e. Now, assume we wish to test *CBC*'s prediction against the two-sided alternative that the true proportion of people that wore a watch in 2018 was not equal to 31%. State the null and alternative hypotheses for this test in mathematical terms.

Solution: Letting p denote the true proportion of people that wore a watch in 2018, our hypotheses can be phrased as

$$\left[\begin{array}{l} H_0 \ :\ p = 0.31 \\ H_A \ :\ p \neq 0.31 \end{array}\right.$$

f. Conduct a test of the two hypotheses you formulated in part (e) above, using an $\alpha = 0.01$ level of significance.

Solution: Our first step is to compute the value of the test statistic.

$$\mathrm{TS} = \frac{\hat{p} - p_0}{\sqrt{\frac{p_0(1-p_0)}{n}}} = \frac{\left(\frac{65}{204}\right) - 0.31}{\sqrt{\frac{0.31\cdot(1-0.31)}{204}}} = 0.2664$$

Next, we compute the critical value. Since we are using an $\alpha = 0.01$ level of significance and a two-sided alternative, our critical value will be negative one times the

$$\left(\frac{0.01}{2}\right) \times 100 = 0.5^{\text{th}}$$

percentile of the standard normal distribution, which we see is around 2.575. Finally, we compare the absolute value of the test statistic to the critical value:

$$|\mathrm{TS}| = |0.2664| = 0.2664 < 2.575$$

which means we fail to reject the null:

> At an $\alpha = 0.01$ level of significance, there was insufficient evidence to reject the null hypothesis that the true proportion of people who wore a watch in 2018 was 31% in favor of the alternative that the true proportion was *not* 31%.

Problem 5: Random Variables (Review Problem)

Let X be a random variable with probability mass function

$\boldsymbol{k}$	-3.1	0	0.7	1.2
$\mathbb{P}(\boldsymbol{X} = \boldsymbol{k})$	$\boldsymbol{a}$	0.19	0.21	0.48

a. What is the value of $\boldsymbol{a}$?

Solution: We know that the probability values in a PMF must sum to 1; as such, we have

$$a + 0.19 + 0.21 + 0.48 = 1 \quad \Longrightarrow \quad a = 0.12$$

b. Compute $\mathbb{P}(\{X = -3\} \cup \{X = 0.7\})$.

Solution:

$$\begin{aligned}\mathbb{P}(\{X=-3\}\cup\{X=0.7\}) &= \mathbb{P}(X=-3)+\mathbb{P}(X=0.7)-\mathbb{P}(\{X=-3\}\cap\{X=0.7\})\\ &= \mathbb{P}(X=-3)+\mathbb{P}(X=0.7)\\ &= 0+0.21 = \boxed{0.21}\end{aligned}$$

c. Compute $\mathbb{P}(X \le 1)$.

Solution:

$$\mathbb{P}(X\le 1)=\mathbb{P}(X=-3.1)+\mathbb{P}(X=0)+\mathbb{P}(X=0.7)=0.12+0.19+0.21=\boxed{0.52}$$

Alternatively, using the complement rule,

$$\mathbb{P}(X\le 1)=1-\mathbb{P}(X>1)=1-\mathbb{P}(X=1.2)=1-0.48=\boxed{0.52}$$

d. Compute $\mathbb{E}[X]$, the expected value of X.

Solution:

$$\begin{aligned}\mathbb{E}[X] &= \sum_{\text{all } k} k\cdot\mathbb{P}(X=k)\\ &= (-3.1)\cdot\mathbb{P}(X=3.1)+(0)\cdot\mathbb{P}(X=0)+(0.7)\cdot\mathbb{P}(X=0.7)+(1.2)\cdot\mathbb{P}(X=1.2)\\ &= (-3.1)\cdot(0.12)+(0)\cdot(0.19)+(0.7)\cdot(0.21)+(1.2)\cdot(0.48)=\boxed{0.351}\end{aligned}$$

e. Compute $\mathrm{SD}(X)$, the standard deviation of X.

Solution: Using the second formula for variance, we would first find

$$\begin{aligned}\sum_{\text{all } k} k^2\cdot\mathbb{P}(X=k) &= (-3.1)^2\cdot\mathbb{P}(X=3.1)+(0)^2\cdot\mathbb{P}(X=0)+(0.7)^2\cdot\mathbb{P}(X=0.7)+(1.2)^2\cdot\mathbb{P}(X=1.2)\\ &= (-3.1)^2\cdot(0.12)+(0)^2\cdot(0.19)+(0.7)^2\cdot(0.21)+(1.2)^2\cdot(0.48)=1.9473\end{aligned}$$

and so

$$\mathrm{Var}(X)=\left(\sum_{\text{all } k} k^2\cdot\mathbb{P}(X=k)\right)-(\mathbb{E}[X])^2=1.9473-(0.351)^2=1.824099$$

Alternatively, we could have used the first formula for variance:

$$\begin{aligned}\mathrm{Var}(X) &= \sum_{\text{all } k} (k - \mathbb{E}[X])^2 \cdot \mathbb{P}(X = k) \\ &= (-3.1 - 0.351)^2 \cdot \mathbb{P}(X = 3.1) + (0 - 0.351)^2 \cdot \mathbb{P}(X = 0) + (0.7 - 0.351)^2 \cdot \mathbb{P}(X = 0.7) \\ &\qquad + (1.2 - 0.351)^2 \cdot \mathbb{P}(X = 1.2) \\ &= (-3.1 - 0.351)^2 \cdot (0.12) + (0 - 0.351)^2 \cdot (0.19) + (0.7 - 0.351)^2 \cdot (0.21) \\ &\qquad + (1.2 - 0.351)^2 \cdot (0.48) = 1.824099\end{aligned}$$

Either way, we find

$$\mathrm{SD}(X) = \sqrt{\mathrm{Var}(X)} = \sqrt{1.824099} \approx 1.35$$

Problem 6: Programming: The Exponential Distribution

Another continuous distribution that we haven't discussed thus far is the so-called **Exponential distribution**. It takes a single parameter, called the *rate* parameter (denoted λ) and has probability density function (p.d.f.):

$$f_X(x) = \begin{cases} \lambda e^{-\lambda x} & \text{if } x \geq 0 \\ 0 & \text{otherwise} \end{cases}$$

We use the notation $X \sim \text{Exp}(\lambda)$ to denote the fact that a random variable X follows the Exponential distribution with parameter λ. The density curves of the $\text{Exp}(\lambda)$ distribution look like:

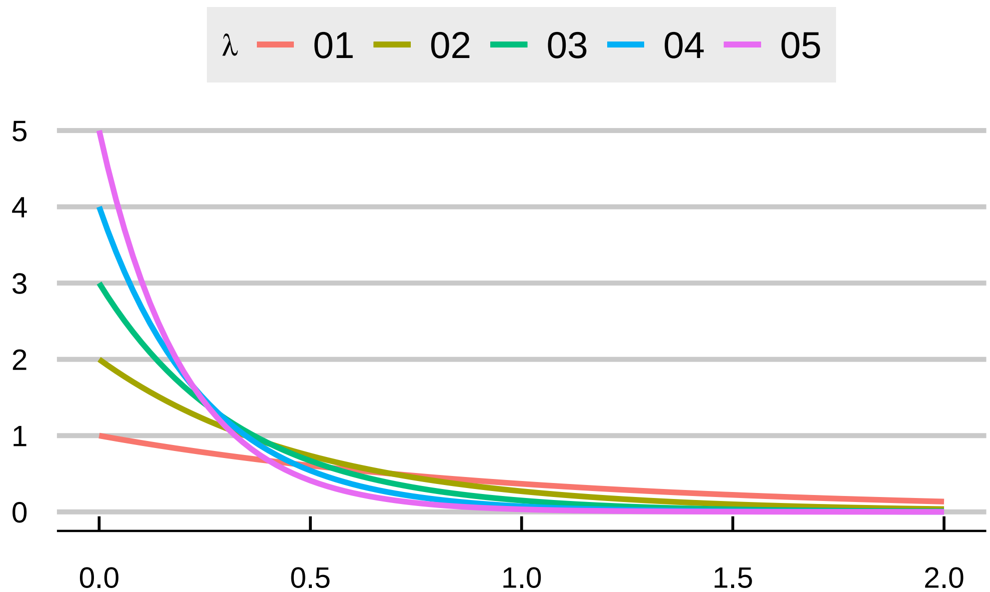

The Exponential distribution is often used for modeling lifetimes; e.g. the lifetime of a lightbulb, etc. It turns out that there is a nice closed-form expression for the area underneath a portion of an Exponential density curve: if $X \sim \text{Exp}(\lambda)$, then

$$\mathbb{P}(a \leq X \leq b) = e^{-a \cdot \lambda} - e^{-b \cdot \lambda}$$

assuming $0 < a < b < \infty$. For example, if $X \sim \text{Exp}(1)$, then $\mathbb{P}(1 \leq X \leq 2) = e^{-1 \cdot 1} - e^{-2 \cdot 1} = e^{-1} - e^{-2} \approx 0.2325$.

❗ Task 1

Write a function called `d_exp()` that takes in two arguments, `x` and `lam`, and returns the value of the p.d.f. of the Exp(`lam`) distribution at the point `x`. Your function should:

- have a default `lam` value of 1
- return zero for any negative values of `x`

Check that your function behaves as follows:

```
d_exp(3.5, 2.31) # specify both arguments
```

```
0.00071177231822478
```

```
d_exp(3.5) # use default lam value
```

```
0.03019738342231850
```

```
d_exp(-2, 4) # return, due to negative input
```

```
0
```

Solutions

```
import numpy as np

def d_exp(x, lam = 1) :
  """returns the Exp(lam) p.d.f. at x"""
  if x >= 0:
    return lam * np.exp(-lam * x)
  else:
    return 0
```

Task 2

Write a function called `p_exp()` that takes in three arguments: `a`, `b`, and `lam`, and returns the probability that an Exp(`lam`)-distributed random variable lies between `a` and `b`. Set `lam` to have a default value of `1`. **Think very carefully about any cases you might need to consider!** (You may assume that `a` is always less than `b`.)

Check that your function behaves as follows:

```
p_exp(1, 2, 1) # specify all three arguments
```

```
0.23254415793482963
```

```
p_exp(1, 2) # use default lam value
```

```
0.23254415793482963
```

```
p_exp(-1, 2) # specify negative `a` value
```

```
0.8646647167633873
```

NOTE: One quirk of python is that, when defining a function with multiple arguments, only *some* of which have default values, you must place the arguments with default values *after* those that do not. I think you will see what I mean when you try to define your `p_exp()` function above!

Solutions

```
def p_exp(a, b, lam = 1):
  if a < 0:
    return 1 - np.exp(-lam * b)
  else:
    return np.exp(-lam * a) - np.exp(-lam * b)
```

The key is to note that the p.d.f. $f_X(x)$ drops to zero for negative values of x. What this means is that the area underneath the density curve from a negative number a to a positive number b is equivalent to the area from 0 to b (the picture below shows $\lambda = 1$, but the result holds for general values of λ):