Quiz 1 - DSC 152, Spring 2026

Full Name:

PID:

Quiz Time: ◯ 3pm ◯ 4pm

Instructions:

- This quiz consists of **5** questions. You have **50 minutes** to complete it.
- Please write **clearly** in the provided answer boxes; we will not grade work that appears elsewhere. Completely fill in bubbles and square boxes; if we cannot tell which option(s) you selected, you may lose points.
  - ◯ A bubble means that you should only **select one choice**.
  - ☐ A square box means you should **select all that apply**.
- Show all work and R code where requested. Partial credit may be awarded.
- You may use one double-sided handwritten sheet of notes. No calculators, no computers.
- Assume we have already run all necessary `library()` calls in R.

By signing below, you are agreeing that you will behave honestly and fairly during and after this quiz.

Signature:

Version A

Please do not open your quiz until instructed to do so.

A food delivery company is interested in understanding customer tipping behavior across different orders. The company collects data on individual deliveries from a random sample of orders over a one-week period. Each row in the dataset represents a single completed delivery. The resulting dataset, called `deliveries`, contains the following columns:

- `order_id` (double): unique identifier for each order
- `restaurant_type` (character): type of restaurant (`"Fast Food"`, `"Casual"`, `"Premium"`)
- `distance` (double): distance of the delivery in miles
- `time` (double): total delivery time in minutes
- `order_total` (double): total cost of the order in dollars
- `tip_amount` (double): tip given by the customer in dollars
- `peak_hour` (logical): `TRUE` if the order was placed during peak hours, `FALSE` otherwise

The first five rows of `deliveries` are shown below.

order_id	restaurant_type	distance	time	order_total	tip_amount	peak_hour
20001	Fast Food	2.1	28.5	18.40	3.20	TRUE
20002	Premium	4.8	42.1	52.75	6.00	TRUE
20003	Casual	3.2	35.0	27.10	4.50	FALSE
20004	Fast Food	1.5	22.3	14.80	2.00	FALSE
20005	Casual	2.7	31.6	24.60	3.80	TRUE

The company claims that the **average** tip amount across all deliveries is **$4.00**. You have been asked to assess this claim using the data, along with other related questions.

Reminders: Write your **PID** on the top right of this page. Show all work where requested.

Question 1

You want to test whether the true average tip amount differs from \$4.00 using a two-sided one-sample t-test. For this question, we will use $\alpha = 0.10$ instead of the usual 0.05 level.

a) Write the null and alternative hypotheses in terms of the appropriate parameter.

H_0: [ ] H_A: [ ]

b) Write one line of R code to run this t-test on the tip amount in `deliveries` and output *only* the p-value.

[ ]

c) Suppose that this test outputs a p-value of approximately 0.14. Which of the following statements are correct? **Select all that apply**

- ☐ If H_0 is true, there is a 14% probability of observing a test statistic at least as extreme as the one we computed.
- ☐ There is a 14% probability that the true average tip amount is \$4.00
- ☐ At $\alpha = 0.05$, we would fail to reject H_0
- ☐ There is a 0.14 probability that H_A is true.
- ☐ None of the above

d) Suppose a manager instead wanted a one-sided test, with the alternative hypothesis being that the true mean tip is less than \$4.00. Based on the p-value stated in the previous part, and assuming that the sample mean tip is less than \$4.00, is it possible to obtain a one-sided p-value without performing any new calculations? If so, state what it is and your statistical decision at $\alpha = 0.05$. If not, explain why not.

[ ]

e) Briefly state any conditions that must be met for the one-sample t-test to be valid.

f) Suppose it is known that tip amounts are often highly right-skewed. If we collect a small sample of n = 6 deliveries from this skewed distribution and run a t-test at $\alpha = 0.10$ for the hypotheses stated in part (a), what will likely be true about the Type I Error rate?

- ◯ It will be approximately 0.10
- ◯ It is likely to be be inflated over 0.10
- ◯ It is likely to be deflated below 0.10
- ◯ You will always make a Type I Error

g) A colleague suggests estimating the Type I Error rate in the situation just discussed in part (f) using simulation. Fill in the three blanks in the code below to estimate the Type I Error rate when data come from a right-skewed Gamma distribution with a mean of 4.

```
count <- 0
for(i in 1:10000){
    skewed_data <- rgamma(n = 6, shape = 1.5, scale = (4 / 1.5))
    p.val <- t.test(skewed_data, __(a)__)$__(b)__
    if(___(c)___){
        count <- count + 1
    }
}
TypeI <- count / 10000
```

(a): ______ (b): ______ (c): ______

PID: ______________________________

Reminders: Write your **PID** on the top right of this page. Show all work where requested.

Question 2

Because tip amounts are right-skewed, a colleague suggests the use of the sign test instead of the t-test.

a) Briefly explain what **parameter** the sign test is testing, and how this differs from the one-sample t-test.

b) Suppose we collect the following sample of tip amounts (in dollars):

$$5.5, \quad 4.8, \quad 6.1, \quad 3.2, \quad 4.9, \quad 1.7, \quad 7.2$$

Using 4 as the null hypothesis value of the parameter of interest, write out the sign sequence (using $+$ and $-$) for the sign test.

c) Using the sign sequence from part (b), compute the p-value of a two-sided test using the Binomial distribution. You may leave your answer in terms of `pbinom()` or `dbinom()` R expressions.

d) Explain in 1-2 sentences how you would determine the significance level that is closest to 0.05. Your answer may include `pbinom` or `dbinom` expressions.

Reminders: Write your **PID** on the top right of this page. Show all work where requested.

Question 3

A colleague proposes using a bootstrap confidence interval (instead of the t-test) to decide whether the mean tip amount differs from \$4.00. The following R code constructs a 95% bootstrap confidence interval for the mean:

```
set.seed(152)
tip_sample <- c(2.50, 5.80, 3.10, 6.20, 1.90, 4.70, 3.60, 5.40)
boot_means <- NA
for(i in 1:10000){
    boot_means[i] <- __(a)__
}
ci <- __(b)__
ci
##  2.5% 97.5%
## 3.10  5.19
```

a) Fill in the blanks above to compute the 95% bootstrap confidence interval for the mean.

`(a):` [ ]

`(b):` [ ]

b) Based on this 95% bootstrap CI, what is the decision regarding $H_0 : \mu = 4$ at $\alpha = 0.05$?

◯ Reject H_0 because 4 is inside the interval

◯ Reject H_0, because 4 is outside the interval

◯ Fail to reject H_0, because 4 is outside the interval

◯ Fail to reject H_0, because 4 is inside the interval

c) A colleague says: "Bootstrap confidence intervals should have a Type I Error rate of approximately 0.05 regardless of the shape of the data distribution, because the bootstrap is a non-parametric procedure." Is this claim correct? Select all that apply.

☐ Yes, everything the colleague said is correct

☐ No. The bootstrap can have an inflated Type I error rate

☐ No. The colleague is incorrect because the bootstrap is actually a parametric procedure

☐ Yes. However, since the bootstrap is a non-parametric procedure, it only works for testing the median, not the mean

☐ No. The accuracy of bootstrap confidence intervals relies on the sample being an accurate representation of the population, which is not always the case.

☐ None of the above

Question 4

Researchers want to estimate the statistical power of the one-sample t-test for detecting whether the average tip amount is greater than \$4.00. They use simulation to estimate power when the true mean tip amount is \$6.00, at $\alpha = 0.05$.

a) Consider the following code:

```
count <- 0
for(i in 1:10000){
    samp <- rnorm(n = 30, mean = 6, sd = 3)
    pval <- ___(a)___
    if(pval < 0.05){
        count <- count + 1
    }
}
count / 10000
```

Fill in the blank to make this code correctly calculate the power in this situation.

b) Using the same code and setup from part (a), can a theory-based equivalent of the power estimated by this simulation be calculated? If so, write appropriate R code to calculate it. If not, explain why not.

Reminders: Write your **PID** on the top right of this page. Show all work where requested.

c) Now consider this code:

```
count <- 0
for(i in 1:10000){
    samp <- rgamma(n = 30, shape = 1.2, scale = (6 / 1.2))
    pval <- (reference your answer to 5(a))
    if(pval < 0.05){
        count <- count + 1
    }
}
count / 10000
```

Can a theory-based equivalent of the power estimated by this simulation be calculated? If so, write appropriate R code to calculate it. If not, explain why not.

d) Which of the following changes, applied individually, would **increase** the statistical power of the test in part (a)? **Select all that apply.**

- ☐ Increase the sample size n
- ☐ Increase the effect size δ
- ☐ Increase the population standard deviation σ
- ☐ Decrease the significance level α from 0.05 to 0.01
- ☐ None of the above

Question 5

The food delivery company collects a very large random sample of **40,000** deliveries and tests:

$$H_0 : \mu = 4 \qquad \textbf{vs} \qquad H_A : \mu \neq 4$$

where μ is the true average tip amount. The sample mean tip is **\$4.03**, the sample standard deviation is **\$1.6**, and the resulting p-value is **0.0002**. The 95% confidence interval for the mean is approximately (\$4.01, \$4.05).

a) At $\alpha = 0.05$, select **all** that are true:

- ☐ Since the p-value is extremely small, we reject H_0 at alpha=0.05
- ☐ Since the sample mean is only \$0.03 away from the null hypothesis value, we actually do not reject H_0 at alpha=0.05 even though the p-value is small
- ☐ Since the 95% confidence interval is very close to the null hypothesis value, the difference from \$4 may not be enough to act on even though we should reject H_0
- ☐ Since the p-value is extremely small, this indicates that there is a meaningful difference from \$4 even though the estimated mean is not very different from \$4
- ☐ None of the above

b) Now consider a different scenario. A small pilot study of $n = 12$ deliveries from Premium restaurants finds a sample mean tip of \$6.80 with a p-value of 0.11 for the test $H_0 : \mu = 4$ vs $H_A : \mu \neq 4$. Select **all** correct responses.

- ☐ We reject H_0 at alpha=0.05 since the sample mean is so different from \$4
- ☐ If the true mean tip is \$4 and the conditions of the test were met, then the probability of observing a tip that is at least \$6.80 or at most \$1.20 is 11%
- ☐ We fail to reject H_0 at alpha=0.05 and should end the investigation here, concluding that there is no meaningful difference from \$4
- ☐ The result indicates that a larger sample size may be needed to determine whether there is a difference from \$4.
- ☐ None of the above