Practice Quiz 1 - DSC 152, Spring 2026

Full Name: SOLUTIONS

PID:

Quiz Time: ◯ 3pm ◯ 4pm

Instructions:

- This quiz consists of 8 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, no phones or any other devices.
- Assume we have already run all necessary `library()` calls in R.

Additional Practice Quiz notes:

- This practice quiz is meant to reflect the style, difficulty, and possible content that may appear on your actual quiz.
- However, please note that it is NOT meant to reflect comprehensive coverage of all concepts that may appear on your actual quiz (as there is no way to put all of that on one quiz). The content for Quiz 1 is everything that is in Lectures 1 through 5, Labs 1-3 (only the first part of Lab 3), and HW1.
- While an answer key will eventually be provided, it is recommended that you do not simply read through the key. It will be much better preparation if you: (1) Actually do the practice quiz as if it were the real thing; (2) Check your answers with the key, but make sure that you actually understand WHY each answer is true.

In light of the recent partial government shutdown, operations at San Diego International Airport have experienced intermittent staffing shortages, including among Transportation Security Administration (TSA) personnel. As a result, travelers have reported increased wait times, particularly during peak travel periods. To better understand these patterns, we will investigate the `tsa` data frame (first few rows shown below), which contains one row per passenger observation, collected randomly from all passengers over the time period of the partial shutdown. The columns are:

- `passenger_id` (double): unique identifier for each passenger
- `checkpoint` (character): which TSA checkpoint was used (`"Terminal 1"` or `"Terminal 2"`)
- `day_type` (character): `"Weekday"` or `"Weekend"`
- `hour` (double): hour of day the passenger entered the line (0–23)
- `wait_min` (double): wait time in minutes
- `tsa_precheck` (logical): `TRUE` if passenger used TSA PreCheck, `FALSE` otherwise

The first five rows of `tsa` are shown below. Not all columns will be needed on this quiz.

	passenger_id	checkpoint	day_type	hour	wait_min	tsa_precheck
0	10001	Terminal 1	Weekday	7	12.4	FALSE
1	10002	Terminal 2	Weekend	14	27.1	FALSE
2	10003	Terminal 1	Weekday	9	8.3	TRUE
3	10004	Terminal 2	Weekday	6	5.7	TRUE
4	10005	Terminal 1	Weekend	11	31.6	FALSE

The airport claims that the **average** TSA wait time across all passengers is **20 minutes**. You have been asked to assess this claim using the data, along with other related questions.

Question 1

You want to test whether the true average wait time differs from 20 minutes.

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

H_0: $\mu = 20$ H_A: $\mu \neq 20$

b) Write one line of R code to run this t-test on the wait times in `tsa` and output the p-value.

```
t.test(tsa$wait_min, mu=20)$p.value
```

c) Suppose that the p-value you obtain is 0.7043. Your colleague says: "This means there is a 70% chance that the true average wait time is 20 minutes." Is this interpretation correct? Select the best answer.

○ Yes, the p-value is the probability that H_0 is true.

● No. The p-value is the probability of observing a test statistic at least as extreme as ours if H_0 is true.

○ No. The p-value is the probability that H_A is true.

○ No. The p-value gives the probability of making a Type I Error.

d) What are the two key conditions discussed in class that must be met for the one-sample t-test to be valid (i.e., to guarantee that its Type I Error rate equals the nominal α level)? Briefly explain why each matters.

Solution: (1) **The data are a random sample** — so that the observations are independent of one another and the sample is representative of the population of interest. (2) **The data come from a normal distribution** (or n is large enough for the CLT to apply) — so that the test statistic $t_s = (\bar{x} - \mu_0)/(s/\sqrt{n})$ truly follows a t-distribution under H_0. If this is not satisfied, the actual Type I Error rate may be inflated above α.

Question 2

Suppose the true average wait time at SAN is actually $\mu = 50$ minutes (not 20), and that wait times follow a **normal distribution** with standard deviation $\sigma = 10$ minutes.

a) What is the name of the quantity that measures the probability of correctly rejecting H_0 in this scenario?

Statistical power

b) In a few words, describe what must happen every time the statistical test is performed, for this quantity to equal 1.0.

Solution: The test would need to reject H_0 in every possible sample drawn under H_A, i.e. the test statistic always falls in the rejection region.

c) A colleague runs the following R code to compute the statistical power of the one-sample t-test with $n = 10$, $\delta = 30$, $\sigma = 10$, $\alpha = 0.05$.

```
power.t.test(n = 10, delta = 30, sd = 10,
             sig.level = 0.05,
             type = "one.sample",
             alternative = "two.sided")
## power = 0.9999
```

Briefly explain in plain English what this power value of approximately 1.0 means in the context of the SAN wait time study.

> **Solution:** If the true mean wait time is 50 minutes (30 minutes above the null value of 20), then with a sample size of 10 passengers, we would almost certainly detect this difference and correctly reject H_0.

Question 3

One important condition for the t-test to be valid is that the data come from a **normal distribution** (especially for small samples). Suppose instead that wait times follow a **highly right-skewed** distribution with a true mean of 20 minutes.

a) If we run the t-test at $\alpha = 0.05$ with $n = 5$ observations drawn from this skewed distribution, what happens to the actual Type I Error rate compared to the nominal level of 0.05?

- ○ It will be exactly 0.05.
- ● It may be inflated above 0.05.
- ○ It will be deflated below 0.05.
- ○ It will always be exactly 0.

b) The following R code estimates the Type I Error rate via simulation when data are drawn from a right-skewed Gamma distribution with mean 20. Fill in the three blanks.

```
count <- 0
for(i in 1:10000){
    skewed_data <- rgamma(n = 5, shape = 1.2,
                          scale = (20 / 1.2))
    p.val <- t.test(skewed_data,
                    mu = __(a)__)$__(b)__
    if(p.val < __(c)__){
        count <- count + 1
    }
}
TypeI <- count / 10000
```

(a): `20` (b): `p.value` (c): `0.05`

Question 4

A colleage suggests that since the distribution of wait times may be skewed, we should use the **sign test**.

a) While the sign test would be a valid approach for this situation, it tests a different thing than the t-Test. Explain.

> **Solution:** The sign test works by checking whether 50% of observations fall above and below the null value, which is a property of the median (not the mean). For skewed distributions, the mean does not necessarily have this 50/50 property.

b) Suppose a sample of $n = 7$ passengers has the following wait times (in minutes):

$$13.2, \quad 25.0, \quad 8.7, \quad 31.4, \quad 17.9, \quad 22.3, \quad 9.5$$

Using $\tilde{\mu}_0 = 20$, write out the sign sequence (using $+$ and $-$) for the sign test.

> $-, +, -, +, -, +, -$

c) Using the sign sequence from part (b), compute the p-value using the Binomial distribution with $p = 0.5$ and $n = 7$. Let X be the number of $+$ signs. Then:

$$\text{p-value} = 2 \times P(X \leq \min(n_+, n_-))$$

You may leave your answer in terms of `pbinom()` or `dbinom()` R expressions.

> **Solution:** We have 3 positive signs and 4 negative signs. $\min(3, 4) = 3$.
> p-value $= 2 \times P(X \leq 3) =$ `pbinom(3, size=7, prob=0.5) * 2` ≈ 1.0
> (fail to reject; the data are quite consistent with a median of 20)

Question 5

The following R code estimates the Type I Error rate of the **sign test** when data are normally distributed with mean 20 and standard deviation 3, using $n = 7$ and $\alpha = 0.015625 + \epsilon$.

```
sign_test <- function(x, mu0 = 20){
    signs <- x - mu0
    n_pos <- sum(signs > 0)
    n_neg <- sum(signs < 0)
    n <- n_pos + n_neg
    min_count <- min(n_pos, n_neg)
    p_val <- pbinom(min_count, size = n, prob = 0.5) * 2
    return(p_val)
```

```
}

count <- 0
for(i in 1:10000){
    norm_data <- rnorm(n = 7, mean = 20, sd = 3)
    p.val <- sign_test(norm_data)
    if(p.val < 0.015626) count <- count + 1
}
TypeI <- count / 10000
```

a) Approximately what will `TypeI` be equal to, and why?

> **Solution:** Approximately 0.015625, because we are simulating under H_0 being true, and if H_0 is actually true, then there is a 0.015625 probability of observing a result that is in the rejection region.

b) Notice that $0.015625 + \epsilon$ is quite far away from the typical α level of 0.05. Yet, it is in fact the closest that we can get to 0.05 in this situation. Explain.

> **Solution:** The key is the discreteness of the possible rejection regions; any answer that gives a reasonable explanation of this would be accepted. Here, with $n = 7$ and $p = 0.5$, the Binomial distribution dictates that the achievable significance levels jump from one discrete value to the next (in this case, from about 0.015625 to about 0.125), so we cannot get any closer to 0.05 than 0.015625.

PID: ____________________

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

Question 6

Researchers want to estimate statistical **power** for the one-sample t-test of $H_0 : \mu = 20$ vs. $H_A : \mu \neq 20$ using simulation.

a) Complete the function below so that it returns a simulated estimate of the power of the one-sample t-test when data are normally distributed with mean $\mu_0 + \delta$ and standard deviation `sd`, using sample size `n` and significance level `alpha`. The null hypothesis value is $\mu_0 = 20$.

```
sim_power <- function(n, delta, sd, alpha = 0.05, reps = 10000){
    count <- 0
    for(i in 1:reps){
        samp <- rnorm(n = __(a)__,
                      mean = __(b)__,
                      sd   = __(c)__)
        pval <- t.test(samp, mu = 20)$p.value
        if(__(d)__){
            count <- count + 1
        }
    }
    return(__(e)__)
}
```

(a):	`n`	(b):	`20 + delta`
(c):	`sd`	(d):	`pval < alpha`
(e):	`count / reps`		

b) Suppose we call `sim_power(n=5, delta=2, sd=1)` and get a value near 0.91. What does this number represent in the context of TSA wait times at SAN?

> **Solution:** If the true average wait time is 22 minutes (2 minutes above the null value of 20), and the standard deviation is 1 minute, then with a sample of 5 passengers we would correctly reject H_0 about 91% of the time.

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

Question 7

The airport collects a large sample of $n = 50{,}000$ passengers and runs a t-test of $H_0 : \mu = 20$ vs. $H_A : \mu \neq 20$. The sample mean wait time is $\bar{x} = 20.3$ minutes, with sample standard deviation $s = 15$ minutes. The resulting p-value is 0.0001.

a) At $\alpha = 0.05$, what is the decision?

● Reject H_0. ○ Fail to reject H_0.

b) Despite the highly significant p-value, a manager argues that a 0.3-minute difference in average wait time is operationally meaningless for the airport. In 2–3 sentences, explain what concept this scenario illustrates and why the p-value alone can be misleading in large samples.

Solution: This illustrates the distinction between **statistical significance** and **practical (effect-size) significance**. With a very large sample size, even a tiny true difference from the null value will produce a very small p-value, causing us to reject H_0. However, a 0.3-minute difference is negligible from an operational standpoint — it provides no meaningful guidance for improving passenger experience. The p-value tells us whether a difference exists, but not whether it is large enough to matter.

Question 8

A senior analyst proposes using a **bootstrap confidence interval** (instead of the t-test) to decide whether the mean wait time differs from 20 minutes. The following R code constructs a 95% bootstrap CI for the mean.

```
set.seed(152)
wait_sample <- c(18.2, 24.5, 15.9, 30.1, 22.7, 11.3, 28.4, 19.6)
boot_means <- NA
for(i in 1:10000){
    boot_means[i] <- mean(sample(wait_sample, replace = TRUE))
}
ci <- quantile(boot_means, probs = c(0.025, 0.975))
ci
##  2.5% 97.5%
## 15.65 26.94
```

a) Based on this 95% bootstrap CI, what is the decision regarding $H_0 : \mu = 20$ at $\alpha = 0.05$?

○ Reject H_0, because 20 is inside the interval.

● Fail to reject H_0, because 20 is inside the interval.

○ Reject H_0, because 20 is outside the interval.

○ Fail to reject H_0, because 20 is outside the interval.

b) A colleague says: "Bootstrap confidence intervals are guaranteed to have a Type I Error rate of exactly 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 perfectly correct.

■ No. The bootstrap can have an inflated or deflated Type I Error rate, especially for small samples.

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

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

c) Suppose we wanted to estimate the Type I Error rate of the bootstrap hypothesis test procedure (reject H_0 when the 95% CI excludes 20) via simulation. Describe in 2–3 sentences what you would simulate and how you would estimate the Type I Error rate. You do **not** need to write R code.

Solution: We would repeatedly (e.g., 1000 times) draw a sample of size n from a distribution whose true mean is 20 (i.e., H_0 is true). For each sample, we would construct a 95% bootstrap confidence interval, by sampling with replacement from that dataset to get a bootstrap sample of size n. We do this repeatedly (e.g. 1000 times) to get an estimate of the sampling distribution of the mean, from these bootstrap samples. The 95% bootstrap confidence interval is then the 2.5th and 97.5th percentiles of estimated sampling distribution of the mean. Then we check whether 20 falls outside the interval. The estimated Type I Error rate is the proportion of simulations in which 20 fell outside the CI (i.e., in which H_0 was incorrectly rejected).