Date 11-09-2026

Full name (in CAPITAL LETTERS) _____ **SUNet ID** (the part before @stanford.edu) _____ **Section** (e.g., TUES 9:30) _____

Please write legibly so the teaching team can log your attendance: if we can't read your name, we can't give you credit.

PSYCH 10 · Fall 2026 · Practice exam

Question 1 · Gym visits and heart rate

A gym recorded, for some of its members, the number of times each visited in a month (from 2 to 14 visits) and each member's resting heart rate, in beats per minute. Each member chooses their own number of visits. Below is the R output for a regression model.

```
hr_model <- lm(HeartRate ~ Visits, data = gym)
hr_model
Coefficients:
(Intercept)       Visits
       80.0         -2.5
supernova(hr_model)
Analysis of Variance Table (Type III SS)
Model: HeartRate ~ Visits
             SS  df       MS      F    PRE      p
Model    992.250   1  992.250 27.000 0.3600  .0000
Error   1764.000  48   36.750
Total   2756.250  49   56.250
```

1.1 A bigger branch

The gym's other branch fits the same model to four times as many of its members and also finds PRE = .36. Would that branch's F be larger than 27, smaller than 27, or about the same? Explain your answer in terms of the numerator and the denominator of F.

The other branch's F is: ☐ larger than 27 ☐ smaller than 27 ☐ about the same Reason: _____

1.2 Reading the table

(i) How many members' data were used to fit this model? Name the table entry you used, and show how you got from that entry to the number of members.

(i) _____ members; table entry used: _____

(ii) The table shows df Model = 1. Which of the model's two coefficients does that one degree of freedom belong to? Why not the other coefficient?

(ii) _____

1.3 The AI assistant's summary

The gym asked an AI assistant to describe the output. Below is its summary, one claim per sentence. AI assistants can write summaries that sound convincing but contain mistakes.

[1] Members who visit more tend to have lower resting heart rates: each additional visit per month is associated with a decrease of about 2.5 beats per minute in resting heart rate. **[2]** So a member who adds four visits a month will lower their resting heart rate by 10 beats per minute. **[3]** A member who never visits would have a resting heart rate of 80. **[4]** Visits explain 36% of the variation in resting heart rate across members. **[5]** Because members chose how often to visit, these results tell us nothing about how visits and heart rate are related.

Rate each claim by how well the R output above justifies the claim: well, partly, or not at all. For each claim rated "partly" or "not at all", explain the reasoning behind your rating.

☐ [1] ☐ well ☐ partly ☐ not at all Reason: _____

☐ [2] ☐ well ☐ partly ☐ not at all Reason: _____

☐ [3] ☐ well ☐ partly ☐ not at all Reason: _____

☐ [4] ☐ well ☐ partly ☐ not at all Reason: _____

☐ [5] ☐ well ☐ partly ☐ not at all Reason: _____

Question 2 · Background noise and reaction time

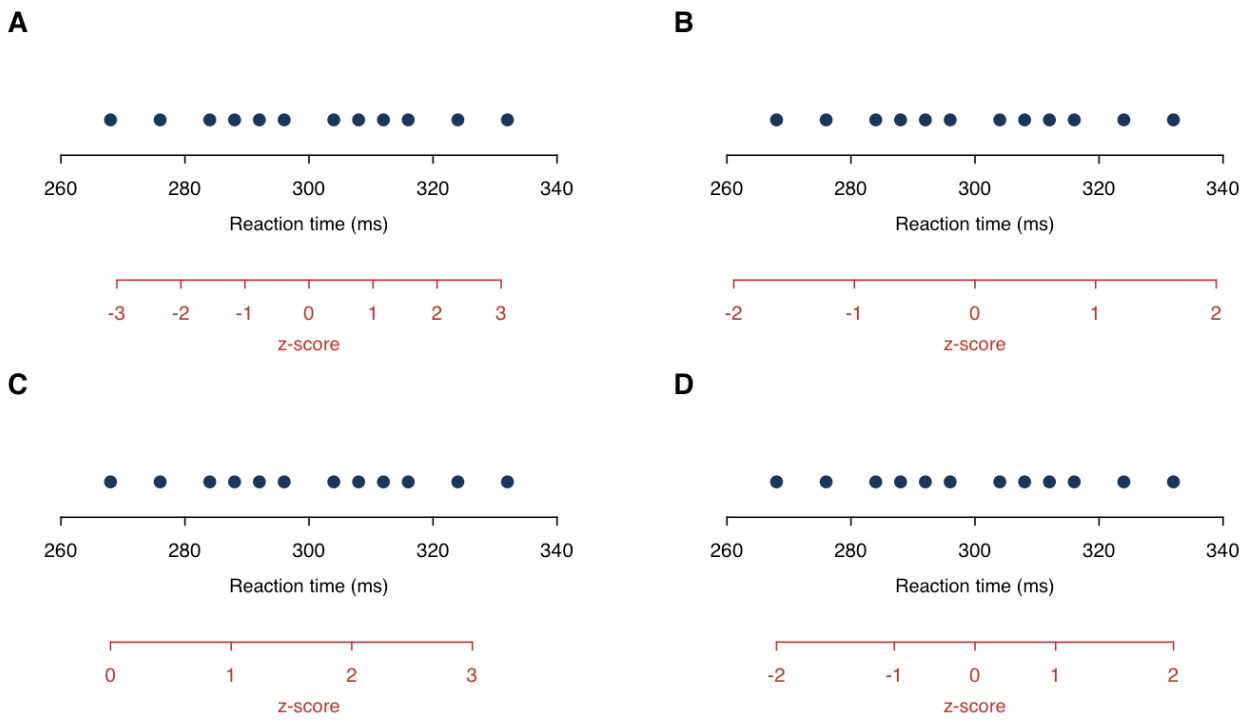
A class randomly assigned 12 volunteers to do a reaction-time task in one of three rooms, 4 volunteers per room: quiet, low background noise, or loud background noise. Below is the R output for reaction time, in milliseconds (ms), overall and by room.

```
favstats(~ RT, data = noise)
 min  Q1 median  Q3 max mean       sd  n missing
 268 287    300 313 332  300 19.44694 12       0
favstats(RT ~ Room, data = noise)
   Room min Q1 median  Q3 max mean       sd n missing
1  Loud 308 314    320 326 332  320 10.32796 4       0
2   Low 288 294    300 306 312  300 10.32796 4       0
3 Quiet 268 274    280 286 292  280 10.32796 4       0
```

The empty model predicts 300 ms, the overall mean, for every volunteer; its SS, the sum of squared errors, is about 4,160.

2.1 The slowest volunteer

(i) Each panel shows the 12 reaction times with a second axis labeled "z-score". Which panel's second axis correctly shows each reaction time's z-score, computed from the mean and standard deviation of all 12 volunteers? Choose one, and give your reason.



(i) ☐ A ☐ B ☐ C ☐ D Reason: _____

_____

(ii) The slowest volunteer, at 332 ms, was in the Loud room. About how many standard deviations above the mean is 332 ms, using the mean and standard deviation of all 12 volunteers? And using the mean and standard deviation of the 4 volunteers in the Loud room? A rough value is enough for each (for example, "about 0.5"). Which of the two values is larger?

(ii) Compared with all 12: about _____ SDs; compared with the Loud room: about _____ SDs

The larger value: ☐ all 12 ☐ the Loud room Reason: _____

(iii) Which of the two comparisons in (ii) is the right one for judging how unusual this time is for a volunteer in the Loud room? Explain the reason for the difference between the two values in (ii).

(iii) The right comparison: ☐ all 12 ☐ the Loud room Reason: _____

_____

_____

_____

2.2 One volunteer drops out

The class learns that one volunteer in the Low room misread the instructions. That volunteer's time was 296 ms, and the class removes it from the data. Compared with the standard deviation of all 12 times, about 19.4 ms, is the standard deviation of the 11 remaining times larger, smaller, or about the same? Choose one, and give your reason.

A. Smaller, because the total of squared errors now has one fewer term, so the total is smaller.
B. Larger, because 296 is near the mean: the total of squared errors barely changes, but it is now divided by 10 instead of 11.
C. About the same, because removing one volunteer out of 12 makes little difference to the spread of the times.
D. Smaller, because a smaller group of volunteers has less variation among its times, so the summary of that variation is smaller.

☐ A ☐ B ☐ C ☐ D Reason: _____

_____

_____

_____

2.3 What the rooms explain

The Room model, `lm(RT ~ Room)`, predicts each volunteer's reaction time with the mean for their room instead of 300 ms.

(i) What would the Room model's PRE be if the Loud, Low and Quiet rooms each had a mean of exactly 300 ms? What would it be if every volunteer's time equaled the mean of their room? Give each number and explain your reasoning.

(i) All means 300: PRE = _____ Every time at its room mean: PRE = _____

_____

_____

(ii) An AI assistant summarized the Room model in the three sentences below. Rate each sentence by how well the R output above justifies the sentence: well, partly, or not at all, and explain the reasoning behind each rating.

(ii) [1] is justified: ☐ well ☐ partly ☐ not at all Reason: _____

[2] is justified: ☐ well ☐ partly ☐ not at all Reason: _____

[3] is justified: ☐ well ☐ partly ☐ not at all Reason: _____

_____

_____

_____

[1] Of all models that make one prediction per room, the model that predicts each room's mean has the smallest total squared error. **[2]** The Room model's predictions are exactly right for every volunteer in the Low room. **[3]** Because the Room model's SS Error is smaller than the empty model's, the three rooms must differ in their effect on reaction time.

2.4 Checking a classmate's work

(i) A classmate finds the empty model's SS by squaring the overall standard deviation and multiplying by 12, the number of volunteers. Is the classmate's SS too large, too small, or correct? Explain your answer; you do not need to compute the SS.

(i) The classmate's SS is: ☐ too large ☐ too small ☐ correct Reason: _____

_____

_____

(ii) Another classmate says, "The standard deviation of about 19 ms means that each volunteer's time is about 19 ms from the mean of 300." How well justified is that statement: well, partly, or not at all? Explain the reasoning behind your rating.

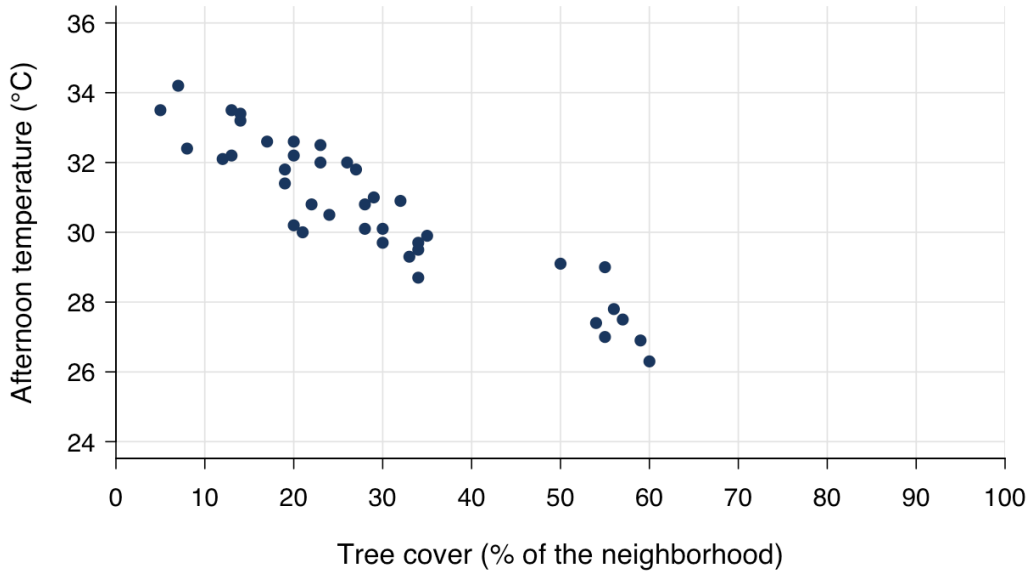
(ii) The statement is justified: ☐ well ☐ partly ☐ not at all Reason: _____

_____

_____

Question 3 · Trees and summer heat

A city measured tree cover (the percent of a neighborhood shaded by trees) and the mean summer afternoon temperature in 40 of its neighborhoods. Here are the data.



3.1 Which prediction?

A planner fits a straight line to the 40 points (the line is not drawn on the plot). At 21% tree cover the line gives about 31.5 °C, and the one neighborhood that was measured at 21% cover had a mean temperature of 30.0 °C. The planner considers four predictions. Which one is best supported by these data? Choose one, and give your reason.

A. A new neighborhood in this city with 21% cover will be about 30.0 °C, the same as the one measured neighborhood with 21% cover.
B. A new neighborhood in this city with 21% cover will be about 31.5 °C, the line's value at 21%.
C. A planned neighborhood in this city with 95% cover will be about 23 °C, the line's value extended to 95%.
D. A neighborhood with 21% cover in a distant desert city will be about 31.5 °C, the line's value at 21%.

☐ A ☐ B ☐ C ☐ D Reason: _____

3.2 The hills

In this city, the neighborhoods with the most trees are in the hills, which are higher than the rest of the city. A council member looks at the plot and says, "Planting trees cools neighborhoods."

(i) Suppose trees had no effect on temperature. How would elevation have to be related to temperature for the plot to show this same downward pattern? Explain your answer.

(i) _____

(ii) How well do these data support the council member's claim: well, partly, or not at all? Explain the reasoning behind your rating, using your answer to (i).

(ii) The data support the claim: ☐ well ☐ partly ☐ not at all Reason: _____

3.3 The cool-pavement lottery

Trees take years to grow, so the city wants to test a quicker alternative: a reflective coating for streets. It can afford to coat 30 streets this summer, and residents of 60 streets have asked for the coating. A planner proposes a lottery: 30 of the 60 are drawn at random to be coated this summer, and the other 30 are coated next year. Over the summer, the city will compare afternoon temperatures on the two sets of streets.

(i) Could this comparison show whether the coating cools streets? Explain your reasoning.

(i) _____

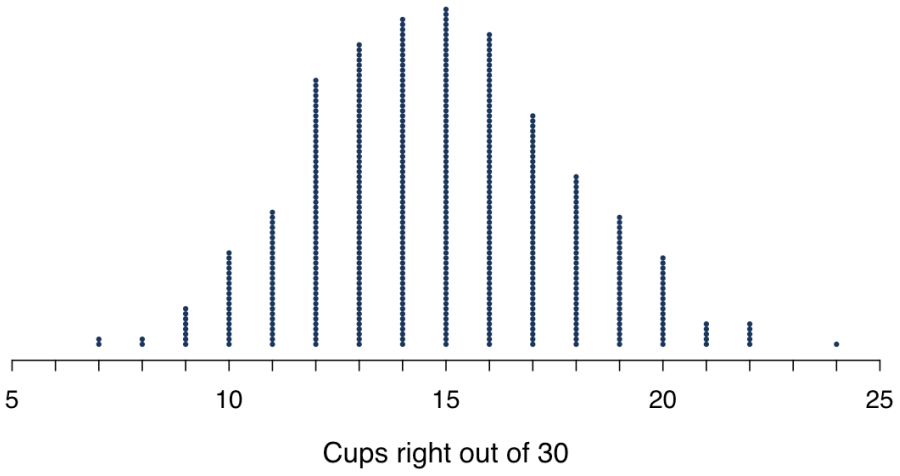
(ii) A resident objects: "A lottery is unfair; you're experimenting on people." How well justified is this objection: well, partly, or not at all? Judge whether the lottery is fair to these residents, and explain the reasoning behind your rating.

(ii) The objection is justified: ☐ well ☐ partly ☐ not at all Reason: _____

Question 4 · Telling fresh coffee from stale

A barista says she can taste whether a cup of coffee was made from beans roasted this week or last month. To test this claim, a friend prepares 30 cups. For each cup the friend flips a coin to decide which beans to use, and the barista tastes the coffee, without seeing the beans, and says whether the beans were fresh or old. She is right for 19 of the 30 cups.

To judge this result, the friend ran a computer simulation of a taster who only guesses: for each of the 30 cups, the taster is right with probability one half. The friend repeated this simulation 500 times and made the plot below. Each dot is one simulated guesser, placed at the number of cups out of 30 that guesser got right.



4.1 Nineteen of thirty

(i) From the plot, estimate the number of the 500 simulated guessers who got 19 or more cups right (to the nearest 10 is fine), and say about how many that is out of every 100 guessers.

(i) Roughly _____ of the 500 guessers, about _____ in every 100

(ii) The friend says, "So this shows she can't really tell the difference." How well justified is the friend's conclusion: well, partly, or not at all? Explain the reasoning behind your rating.

(ii) The conclusion is justified: ☐ well ☐ partly ☐ not at all Reason: _____

4.2 A longer test

The barista takes a longer test: 120 cups, chosen by coin flips in the same way. She gets 76 right, again about 63%. How does this result compare with her 19 of 30, as evidence that she can taste the difference? Choose one, and give your reason.

A. About the same evidence, because she got the same share of cups right both times.
B. Weaker evidence, because with 120 cups she had more chances to make mistakes.
C. Stronger evidence, because a guesser's share stays nearer one half on a longer test.
D. Stronger evidence, because she got 57 more cups right than she did on the first test.

☐ A ☐ B ☐ C ☐ D Reason: _____

4.3 The tasting contest

A chain of coffee shops holds a contest. Each of its 40 baristas tastes 30 cups prepared in the same way. The chain will call a barista "gifted" if the barista gets 20 or more cups right. A guesser gets 20 or more right about 6 times in 100 (29 of the 500 simulated guessers in the plot).

(i) If all 40 baristas were only guessing, about how many would you expect the chain to call gifted? Give the number and show the calculation.

(i) About _____ baristas

(ii) Three baristas score 20 or more. Is that more than the number expected from 40 guessers? What do the three scores show, if anything, about whether any of the 40 baristas is gifted?

(ii)

4.4 The retest

The chain retests the three high scorers on 30 new cups each.

(i) If the three were only guessing, about what would you expect each to score on the retest? Give a number and explain your reasoning.

(i) About _____ cups right

(ii) The chain sets a rule: "Anyone who scores below 20 on the retest has no gift." Consider a taster who is genuinely right on 65 of every 100 cups. In a simulation of such a taster on 30 cups, the score was below 20 about half the time. How well justified is the chain's rule: well, partly, or not at all? Explain the reasoning behind your rating.

(ii) The rule is justified: ☐ well ☐ partly ☐ not at all Reason:

Question 5 · A reading app

A company makes a reading app used by 40,000 students in 120 school districts.

5.1 Two surveys

The company emailed a satisfaction survey to all 40,000 users. Of the 4,000 who replied, 85% said they liked the app. A researcher instead chose 200 users at random and kept asking until 190 of the 200 had answered the same survey; 62% of those 190 said they liked the app.

(i) As an estimate of the percentage of all 40,000 users who like the app, is the emailed survey's 85% likely too high, too low, or about right? Explain your reasoning, based on which users chose to answer the emailed survey.

(i) The 85% is likely: ☐ too high ☐ too low ☐ about right Reason: _____

(ii) Which estimate would you trust more, the 85% or the 62%? Explain your reasoning, taking into account both the way each sample was chosen and the number of people who answered.

(ii) Trust more: ☐ the 85% ☐ the 62% Reason: _____

5.2 The district rankings

The company ranks the 120 districts by their students' mean gain on a reading test over the year. The 10 districts at the top of the ranking and the 10 at the bottom all have fewer than 30 students using the app. A sales manager says, "Small districts use the app best. Let's market it to small districts."

(i) Explain how the smallest districts could end up at both the top and the bottom of the ranking even if district size had no relation to the app's effectiveness.

(i) _____

(ii) How well justified is the sales manager's conclusion: well, partly, or not at all? Explain the reasoning behind your rating.

(ii) The conclusion is justified: ☐ well ☐ partly ☐ not at all Reason: _____

5.3 The big trial

The company then ran a trial. It randomly assigned 20,000 students to two groups: one group used the app for a year, and the other group did not. On a 100-point reading test, the app group's mean was higher by 0.4 points. The researchers then shuffled the group labels 1,000 times, each time reassigning "app" and "no app" to the students at random and recomputing the difference between the two groups' means. No shuffled difference was as large as 0.4 points. Teachers consider a gain of about 5 points meaningful. The company's ad says, "Proven to dramatically improve reading scores!"

(i) What does the shuffle result show about whether the app has any effect on reading scores?

(i) _____

_____

(ii) How well justified is the ad's claim: well, partly, or not at all? Explain the reasoning behind your rating, and say whether each of the two words "proven" and "dramatically" is justified.

(ii) The ad's claim is justified: ☐ well ☐ partly ☐ not at all Reason: _____

_____

_____

_____

Question 6 · A city bike-share

A city's bike-share system recorded every trip taken in May, 2,000 trips in all. For each trip, the system recorded the start station's number (an ID from 1 to 80 printed on each station), the neighborhood where the trip started, the trip's length in minutes, and the temperature when the trip started.

6.1 The analyst's two models

An analyst fits two regression models to the 2,000 trips and writes a finding from each:

- **(a)** From `lm(Minutes ~ StationNumber)`, with slope 0.05: "Trips that start at a station numbered one higher are, on average, 0.05 minutes longer."
- **(b)** From `lm(Minutes ~ Temperature)`, with slope 0.3: "Trips that start when it is one degree warmer are, on average, 0.3 minutes longer."

(i) Rate each finding by how well its slope supports it: well, partly, or not at all. Give a one-sentence reason for each rating.

(i) (a): ☐ well ☐ partly ☐ not at all Reason: _____

(b): ☐ well ☐ partly ☐ not at all Reason: _____

_____

_____

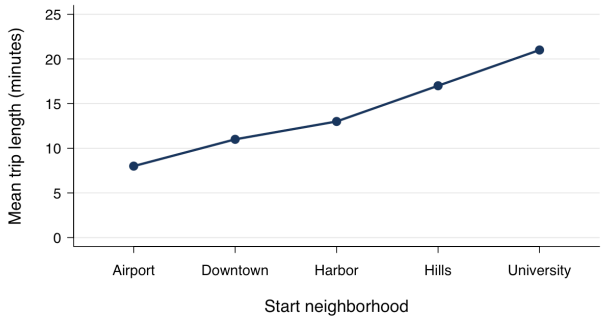
(ii) For each finding rated "partly" or "not at all", state the error in it, and describe a better way to use that variable in a model of trip length.

(ii) _____

_____

6.2 The neighborhood chart

The analyst made this chart of the mean trip length for trips starting in each of five neighborhoods, and a reader concludes, "Trips get longer as you move along the chart from Airport to University."



(i) Is the variable on the horizontal axis categorical or quantitative? Do its values have a natural order?

(i) _____

(ii) How well do these data support the reader's conclusion: well, partly, or not at all? Explain the reasoning behind your rating.

(ii) The conclusion is supported: ☐ well ☐ partly ☐ not at all Reason: _____

(iii) Describe the chart you would make to compare the five neighborhoods' mean trip lengths, and explain the reasoning behind your choice.

(iii) _____

6.3 Rain and riding

Over the past 90 days, the analyst compares the days with rain in the forecast with the days without rain in the forecast. The mean number of rides per day was 5% lower on rain-forecast days. When the analyst shuffled the "rain forecast" labels among the days many times, a difference of at least 5% in either direction came up in about 40 of every 100 shuffles. Which is the most accurate summary for the operations team? Choose one, and give your reason.

A. The forecast does not affect ridership, because a difference of this size is nothing more than chance.
B. The forecast lowers ridership by about 5%, since that is the difference observed in the data from these past 90 days.
C. Rides were a little lower, but chance often produces a difference of this size, so the data are not yet decisive.
D. There is about a 40-in-100 chance that the forecast has no effect at all on the number of rides.

☐ A ☐ B ☐ C ☐ D Reason: _____

Formula sheet

You are not expected to memorize how to compute these quantities. Each is given in words and in symbols. Do not assume that any given question will need one. When one does, deciding which to use, what to put into it, and what to do with the result is up to you.

One variable	How it is computed	In symbols
Mean	Add up all the values and divide by the number of values.	$\overline{Y} = \frac{\sum_{i=1}^{n} Y_i}{n}$
Median	Sort the values and take the middle one (with an even count, average the two middle ones).	middle value of the sorted Y_i
Range	The maximum minus the minimum.	$\max(Y) - \min(Y)$
Interquartile range (IQR)	Q3 minus Q1.	$\text{IQR} = Q3 - Q1$
Outlier cutoffs	Lower: Q1 minus 1.5 times the IQR. Upper: Q3 plus 1.5 times the IQR.	lower: $Q1 - 1.5 \times \text{IQR}$ upper: $Q3 + 1.5 \times \text{IQR}$
Relative frequency	The count for a level divided by the total number of values.	$\frac{\text{count}}{n}$
Deviation	A value minus the mean.	$Y_i - \overline{Y}$
Mean absolute deviation (MAD)	Add up the absolute value of each value minus the mean (ignore the sign), then divide by the number of values. (SAD is the sum before dividing.)	$\text{MAD} = \frac{\sum_{i=1}^{n} \lvert Y_i - \overline{Y} \rvert}{n}$ $\text{SAD} = \sum_{i=1}^{n} \lvert Y_i - \overline{Y} \rvert$
SS Total	Add up the squared difference between each value and the mean.	$SS_{\text{Total}} = \sum_{i=1}^{n} (Y_i - \overline{Y})^2$
Variance	SS Total divided by (the number of values minus 1).	$s^2 = \frac{SS_{\text{Total}}}{n-1}$
Standard deviation	The square root of the variance.	$s = \sqrt{s^2} = \sqrt{\frac{\sum_{i=1}^{n} (Y_i - \overline{Y})^2}{n-1}}$
SS Total from a variance	The variance times (the number of values minus 1).	$SS_{\text{Total}} = s^2(n-1)$

z-scores	How it is computed	In symbols
z-score	(A value minus the mean) divided by the standard deviation.	$z = \frac{Y - \overline{Y}}{s}$
Empirical rule	For a normal (bell-shaped) distribution: about 68%, 95% and 99.7% of values are within 1, 2 and 3 standard deviations of the mean, and about 16% are in each tail beyond 1, and 2.5% beyond 2. For other shapes these percentages don't hold.	$\approx 68\%$ within $\overline{Y} \pm 1s$ $\approx 95\%$ within $\overline{Y} \pm 2s$ $\approx 99.7\%$ within $\overline{Y} \pm 3s$

Models and fit	How it is computed	In symbols
Residual	The actual value minus the value the model predicts for it.	$e_i = Y_i - \widehat{Y}_i$
SS Error	Add up the squared difference between each value and the model's predicted value.	$SS_{\text{Error}} = \sum_{i=1}^{n} (Y_i - \widehat{Y}_i)^2$
SS Model	SS Total minus SS Error.	$SS_{\text{Model}} = SS_{\text{Total}} - SS_{\text{Error}}$
PRE (proportional reduction in error)	SS Model divided by SS Total. Equivalent: (SS Total minus SS Error) divided by SS Total.	$PRE = \frac{SS_{\text{Model}}}{SS_{\text{Total}}} = \frac{SS_{\text{Total}} - SS_{\text{Error}}}{SS_{\text{Total}}}$

Models and fit (continued)	How it is computed	In symbols
Degrees of freedom (df)	df Total: number of values minus 1. df Model: number of parameters minus 1 (groups minus 1 for a group model; 1 for a regression with one explanatory variable). df Error: df Total minus df Model.	$\text{df}_{\text{Total}} = n - 1$ $\text{df}_{\text{Model}} = k - 1$ $\text{df}_{\text{Error}} = \text{df}_{\text{Total}} - \text{df}_{\text{Model}} = n - k$ (k = number of parameters)
MS (error per degree of freedom)	A row's SS divided by its df. Do this for Model, Error and Total.	$MS = \frac{SS}{\text{df}}$
F ratio	MS Model divided by MS Error. Equivalent: (PRE divided by df Model) divided by ((1 minus PRE) divided by df Error).	$F = \frac{MS_{\text{Model}}}{MS_{\text{Error}}} = \frac{PRE/\text{df}_{\text{Model}}}{(1-PRE)/\text{df}_{\text{Error}}}$
Typical residual	The square root of MS Error.	$\sqrt{MS_{\text{Error}}}$
Correlation (r)	The slope of the regression line when both variables are z-scores. Equivalent: add up the products of each pair of z-scores and divide by (the number of pairs minus 1).	r = slope of z_y predicted from z_x $r = \frac{\sum z_x z_y}{n-1}$
r squared	The correlation times itself. With one explanatory variable, this equals PRE.	r^2 ($= PRE$ with one explanatory variable)
Cohen's d	(Mean of group 1 minus mean of group 2) divided by the pooled standard deviation. Pooled standard deviation: the square root of the average of the two groups' variances, each weighted by its df (number of values minus 1); for a group model, the square root of MS Error.	$d = \frac{\overline{Y}_1 - \overline{Y}_2}{s_{\text{pooled}}}$ $s_{\text{pooled}} = \sqrt{\frac{\text{df}_1 s_1^2 + \text{df}_2 s_2^2}{\text{df}_1 + \text{df}_2}} = \sqrt{MS_{\text{Error}}}$ for a group model

ANOVA table layout

	SS	df	MS	F	PRE
Model	SS_{Model}	df_{Model}	MS_{Model}	$\frac{MS_{\text{Model}}}{MS_{\text{Error}}}$	$\frac{SS_{\text{Model}}}{SS_{\text{Total}}}$
Error	SS_{Error}	df_{Error}	MS_{Error}		
Total	SS_{Total}	df_{Total}			

Blank cells are not shown. In R, supernova() adds a p-value column.

Inference	How it is computed	In symbols
F and t (two groups)	F equals the t value squared.	$F = t^2$
95% confidence interval	The estimate plus and minus 2 times its standard error. (The 2 applies to 95% only.)	$\text{estimate} \pm 2 \times SE$
Bootstrap 95% bounds	Sort the 1000 bootstrapped b_1s. The bounds are the 26th and 975th values.	26th and 975th of the 1000 sorted bootstrapped b_1s

R commands

You will not be asked to write R code, and reading R code or output is not guaranteed to appear in any question. These commands are here in case it does. Y is the outcome variable, X is the explanatory variable, and data_set is the data frame.

Commands	What they do
Base R	
`x <- 5` `c(1, 2, 3)` `x[3]` `sum()` `print()`	Assign; combine values; third element; add; print
`==` `!=` `>` `<` `>=` `<=` `NA`	Comparisons; a missing value
Data	
`str(data_set)` `head(data_set)` `tail(data_set)` `data_set$Y`	Structure; first six rows; last six rows; one column
`select(data_set, Y1, Y2)` `filter(data_set, Y > 300)` `arrange(data_set, Y)`	Keep variables; keep rows that meet a condition; sort rows
`na.omit(data_set)` `is.na(Y)`	Drop rows with missing values; test for missing
`mutate()` `recode()` `factor(data_set$Y)` `as.numeric(data_set$Y)`	Create a variable; recode values; make categorical; make numeric
Summaries	
`favstats(~ Y, data = data_set)` `favstats(Y ~ X, data = data_set)`	Min, Q1, median, Q3, max, mean, sd, n, missing; the same by group
`tally(~ Y, data = data_set)` `tally(Y ~ X, data = data_set, margins = TRUE, format = "proportion")`	Frequency table; two-way frequency table
`mean()` `median()` `var()` `sd()` `sqrt()`	Mean, median, variance, standard deviation, square root
`cor(Y ~ X, data = data_set)` `cohensD(Y ~ X, data = data_set)`	Correlation; Cohen's d
`b0()` `b1()` `pre()` `f()`, as in `b1(Y ~ X, data = data_set)`	Just b_0, b_1, PRE or F (each takes a formula or a fitted model)
Graphs	
`gf_histogram(~ Y, data = data_set, bins = , binwidth = )` `gf_dhistogram()`	Histogram (counts); histogram (density)
`gf_boxplot(~ Y)` `gf_bar(~ Y)` `gf_props()` `gf_percents()`	Box plot; bar graph; proportions; percents
`gf_point(Y ~ X)` `gf_jitter(Y ~ X, width = , height = 0)` `gf_boxplot(Y ~ X)`	Scatterplot; jittered scatterplot; box plots by group
`gf_facet_grid(X ~ .)` `gf_model(model)` `gf_lm()`	Panels by group; add model predictions; add a regression line
`gf_vline(xintercept = )` `gf_hline(yintercept = )`	Add a vertical or horizontal line
Models	
`empty_model <- lm(Y ~ NULL, data = data_set)`	Fit the empty model (the intercept is the mean)
`my_model <- lm(Y ~ X, data = data_set)`	Fit a model with one explanatory variable
`predict(my_model)` `resid(my_model)` `sum(resid(my_model)^2)`	Predictions; residuals; SS Error
`supernova(my_model)`	ANOVA table (columns SS, df, MS, F, PRE, p)
Probability	
`zscore(data_set$Y)` `xpnorm(65.1, mean = , sd = )`	Z-scores; area under a normal curve
`xpf(sample_f, df1 = 2, df2 = 10)`	p-value from the F-distribution
Simulation	
`shuffle(data_set$Y)` `resample(data_set)` `sample(x, n, replace = TRUE)` `set.seed()`	Mix up values; sample rows with replacement; sample with replacement; fix the random seed
`sdob1 <- do(1000) * b1(shuffle(Y) ~ X, data = data_set)`	1000 b_1s from shuffled data (centered on 0)
`sdob1_boot <- do(1000) * b1(Y ~ X, data = resample(data_set))`	1000 bootstrapped b_1s (centered on the sample b_1)
`gf_histogram(~ b1, data = sdob1, fill = ~middle(b1, .95))` `lower()` `upper()`	Shade the middle 95% of a sampling distribution; the lower and upper values
`tally(~ b1 > sample_b1, data = sdob1, format = "proportion")` `tally(sdob1$b1 > sample_b1 \| sdob1$b1 < -sample_b1)`	Proportion of b_1s above the sample b_1 (one tail); count of b_1s at either extreme
`sd(~ b1, data = sdob1)`	Standard deviation of the b_1s (the standard error)
Intervals and tests	
`confint(my_model)` `confint(my_model, level = .99)`	Confidence intervals (columns 2.5% and 97.5%; the intercept row is β_0, the X row is β_1)
`summary(my_model)$coef`	Estimates, standard errors, t values and p-values
`t.test(Y ~ X, data = data_set, var.equal = TRUE)`	t-test using pooled variance
`pairwise(my_model, correction = "Bonferroni")` `plot = TRUE`	Pairwise comparisons (corrections: "Tukey", "Bonferroni", "none")
`xchisq.test(Y ~ X, data = data_set)`	Chi-square test