Stat 390, Spring, Test 3, June 8, 2021; Marzban

60 minutes

READ the following paragraph!!

$12+2+2+2+3+3=24$

Points
0

BY SUBMITTING THIS TEST, I HEREBY PLEDGE ON MY HONOR THAT I HAVE TAKEN IT IN PERSON AND WITHOUT ASSISTANCE FROM ANY OTHER PERSON. I ACKNOWLEDGE THAT THE PENALTY FOR VIOLATING ACADEMIC INTEGRITY IS MOST SEVERE.

- This is an online test given during the Corona Virus outbreak.
- For questions that have only 1 correct response, the options will appear as circles.
- For questions that may have multiple response, the options will appear as squares.
- It is open book/web/hw/solutions/past_tests/calculator/etc, but closed collaboration.
- Googling will will simply waste your time. Instead, just make sure you are organized.
- The questions are presented to students in random order.
- You have the option of changing your answer to an already answered question.
- The list of questions and the remaining time for the whole test appear at the bottom of each page.
- Questions 1-12 are worth 1 point and do not require much calculation or writing.
- Questions 13-15 are worth 2 points each, and require a bit more work.
- Questions 16-17 are worth about 2 or 3 points and require varying levels of calculation, all of which are to be done on paper (or Tablet), and saved/scanned/photographed, and uploaded to canvas before 3:31. For these non-multiple-choice question, SHOW WORK. NO CREDIT FOR CORRECT ANSWER WITHOUT EXPLANATION/WORK. - With file-upload questions, I am giving you several choices: You may 1) upload individual files for each question, or 2) submit a single file containing answers to all of your file-upload questions, or 3) email all your files to me (marzban@uw.edu), but ONLY IF all else has failed.

lect18

1 **1.** We have learned that the sampling distribution of $t = \frac{\overline{x}-\mu_x}{s_x/\sqrt{n}}$ is given by the t-distribution with $df = (n-1)$. Now, if the n itself also changes from sample to sample, the sampling distribution of t will be ____ the t-distribution with $df = (n-1)$.

a) wider than ← 3 sources of variability: $\overline{x}, s_x, \sqrt{n}$.

b) comparable to

c) narrower than

1 **2.** On a given data set, if a 1-sided p-value is near 1, then the 2-sided p-value is also near 1. Hint: draw pictures for yourself. True or False?

=> 2-sided = sum of the near-0 areas

hw-lect16-2

1 **3.** Here is a CI (NOT p-value) question. At a given confidence level, if σ_x is known, we have seen that the probability of a random sample mean $\overline{x}$ falling in the observed CI for μ_x is given by $pr(z_{obs} - z^* < z < z_{obs} + z^*)$. So, it **looks like** one can actually compute that prob once observed data are available. But, in fact, it is NOT possible to compute that probability, because (select the correct **reason**)

a) z_{obs} doesn't have a prob.

b) z^* doesn't have a prob.

c) z_{obs} cannot be computed. $z_{obs} = \frac{\overline{x}_{obs} - \mu_x}{\sigma_x/\sqrt{n}}$ ← not known

d) z^* cannot be computed.

hw-lect22-2, hw-lect22-1

1 **4.** Suppose we observe the following 95% CI for $\mu_2 - \mu_1$: (1.1, 2.3), and then realize that all we really care about is whether or not μ_2 exceeds μ_1 by at least 1.5. Then, the appropriate H_0/H_1 is

a) $H_0: \mu_2 - \mu_1 = 1.1$ $H_1: \mu_2 - \mu_1 > 1.1$

b) $H_0: \mu_2 - \mu_1 = 1.1$ $H_1: \mu_2 - \mu_1 > 1.5$

c) $H_0: \mu_2 - \mu_1 = 1.5$ $H_1: \mu_2 - \mu_1 > 1.5$

d) $H_0: \mu_2 - \mu_1 = 1.5$ $H_1: \mu_2 - \mu_1 \quad 1.5$

Rosetta Stone

1 **5.** Suppose in a test of $H_0: \pi > \pi_0$, $H_1: \pi < \pi_0$, we find a near-1 p-value, suggesting that there is no evidence for $\pi < \pi_0$. On the same data, testing $H_0: \pi < \pi_0$, $H_1: \pi > \pi_0$, we will conclude

a) there is no evidence for $\pi < \pi_0$

b) there is no evidence for $\pi > \pi_0$

p-value = near-zero

reject H_0 in favor of H_1

c) there is evidence for $\pi < \pi_0$

d) there is evidence for $\pi > \pi_0$

lect23

1 **6.** We know that we can use the chi-squared test to test whether k proportions corresponding to k categories in one population have specific values. Can we use that test to test whether k proportions are equal? Yes/No? $H_0: \pi_1 = \pi_{01}, \pi_2 = \pi_{02}, \ldots$ set $\pi_{01} = \pi_{02} = \cdots = \frac{1}{k}$.

lect24

1 **7.** We know that we can use the 1-way ANOVA F test to test whether k means are equal. Can we use that test to test whether the k means take specific values? Yes/No? $H_0: \mu_1 = \mu_2 = \cdots = \mu_k$

lect23

1 **8.** There are currently three COVID-19 vaccine types in use in the USA. For each of the three types, there exists data on the number of vaccinated people who show no symptoms even after they are infected. Which of the tests we have studied can be used to test whether or not the three vaccine types are equally protective? Note that the three vaccine types are not equally common across the US. Counts in 3 categories with unequal probs. See Tornado example.

a) z-test b) chi-squared test c) 1-way ANOVA F-test d) F-test of model utility e) None of the above.

winter21/#4

1 **9.** If data from two populations are unpaired, then $V[\overline{x}_1 - \overline{x}_2] = V[\overline{x}_1] + V[\overline{x}_2]$, and this is the combination that shows-up in our formula for t_{obs}, used for computing the p-value. But if the data are (appropriately) paired, then it turns out that in general $V[\overline{x}_1 - \overline{x}_2] = V[\overline{x}_1] + V[\overline{x}_2] -$ a positive term. Therefore, t_{obs} in a paired design will be generally ____ that in an unpaired design. Hint: Look at the formula for t_{obs}. $t_{obs} = (\overline{x}_1 - \overline{x}_2 - (\mu_1 - \mu_2)) / \sqrt{V[\overline{x}_1] + V[\overline{x}_2] - (\ldots)}$ ← denom = smaller

a) larger than b) comparable to c) less than

∴ t_{obs} = larger

hw-lect24-1, hw-lect18-1

1 **10.** Suppose we have a data set on one continuous variable y and one categorical variable x with 3 levels. We want to know if x has an effect on y. For example, does the mean of y vary across the levels of x? Which test is most appropriate?

a) t-test b) chi-squared test c) 1-way ANOVA F-test d) F-test of model utility.

lect26

1 **11.** The true mean (expected value) of estimation error is $E[\hat{y}(x) - y(x)] = E[\hat{y}(x)] - y(x) = y(x) - y(x)$

a) 0 b) $y(x)$ c) $\sigma^2_{\hat{y}}$ d) None of the above.

many

1 **12.** A 95% PI for an individual y value covers $y(x)$ (mean of y, given x), _____ than 95% of the time.

a) less often b) equally often c) more often

PI = wider than CI

winter21/#11, hw-lect22-1

2 **13.** I believe the true proportion of people who carry the COVID antibody is less than 10%. But according to the latest sample of size 100, the proportion is 12%.

- The p-value for the appropriate test is a) 0.7486 b) 0.2514 c) 0.5028
- At $\alpha = 0.05$, does this data contradict my belief? a) Yes b) No

$H_0: \pi < 0.1$ ← Belief

$H_1: \pi > 0.1$, $\pi_0 = 0.1$

$z_{obs} = \frac{p_{obs} - \pi_0}{\sqrt{\frac{\pi_0(1-\pi_0)}{n}}} = \frac{0.12 - 0.1}{\sqrt{\frac{.1(.9)}{100}}} = \frac{2}{3} = 0.67$

p-value $= pr(p > p_{obs} \mid \pi = \pi_0) = Pr(z > z_{obs}) = pr(z > 0.67)$

$= 1 - 0.7486 = 0.2514 > \alpha$

q210

2 **14.** In the following R output, at $\alpha = 0.05$, is there an inconsistency between the p-value from the F-test of model utility, and the p-values shown in the remainder of the output? Hint: this is related to the moral of one of our quizzes.

a) Yes b) No

```
            Estimate Std. Error t value Pr(>|t|)
(Intercept) -0.079492   0.031664  -2.511   0.0122
x1          -0.065505   0.032332  -2.026   0.0430
x2           0.022564   0.030760   0.734   0.4634
```

x3	-0.001813	0.031155	-0.058	0.9536
x4	-0.004118	0.031949	-0.129	0.8975
x5	-0.008429	0.031927	-0.264	0.7918
x6	-0.019267	0.031581	-0.610	0.5419
x7	0.009918	0.033173	0.299	0.7650
x8	-0.016619	0.032835	-0.506	0.6129
x9	-0.067775	0.031213	-2.171	0.0301
x10	0.062442	0.031745	1.967	0.0495

```
Residual standard error: 0.9938 on 989 degrees of freedom
Multiple R-squared:  0.01375,Adjusted R-squared:  0.003775
F-statistic: 1.379 on 10 and 989 DF,  p-value: 0.1849
```

Explanation: The ANOVA p-value $> \alpha$, ie. There is no evidence against $H_0: \beta_1 = \beta_2 = \cdots \beta_{100} = 0$. So, it may seem inconsistent that several of these p-values suggest that some β's $\neq 0$. In fact, as explained in q#10, there is no inconsistency at all.

2 Test25_example, winter21 #13

15. The following output was obtained from R, using the model $y = \alpha + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_1 x_2$.

	Estimate	Std. Error	t value	Pr(>|t|)
(Intercept)	-0.3513	1.0399	-0.338	0.7382
x1	1.0957	2.0565	0.533	0.5987
x2	4.6774	1.7721	2.639	0.0138
x1*x2	1.7865	0.8108	.....	......

```
Residual standard error: 4.344 on 26 degrees of freedom
Multiple R-squared:  0.7092,Adjusted R-squared:  0.6756
F-statistic: 21.13 on 3 and 26 DF,  p-value: 3.805e-07
```

$n-(k+1)$

- If we have concerns about possible overfitting, would it be appropriate to consider a model that has no intercept at all? Yes/No.

p-value $= 0.7382 > \alpha$, i.e. intercept may be zero

- What is the p-value that must appear on the "x1*x2" line of the output.

a) 2.203 b) 0.1257 c) 0.036 d) None of the above.

$H_0: \beta_3 = 0$

$H_1: \beta_3 \neq 0$

$$t_{obs} = \frac{\hat{\beta}_3 - 0}{S_{\hat{\beta}_3}} = \frac{1.7865}{0.8108} = 2.2$$

$$\text{p-value} = 2\,pr(t > 2.2) = 2(.018) = 0.036.$$

$$df = n-(k+1) = 26$$

hw-lect 22-2, paired version

~ 3 **16.** As one more example to expose the connections between a CI and a p-value, consider this question: Is there evidence from data that $\mu_2 - \mu_1$ exceeds the upper side of the observed 95% CI for $\mu_2 - \mu_1$, in a **paired design.**

- Write down the appropriate H0/H1,
- Write an expression for the p-value that does NOT involve t,
- Standardize the expression you have written for p-value so that it DOES involve t, and
- Compute the p-value; the answer will depend on α.

Hint: $\overline{d}_{obs} = (\overline{x_2} - \overline{x_1})_{obs}$. Also, you may use $pr(t > -t^*) = 1 - \frac{\alpha}{2}$ without derivation.

1.5
$H_0: \mu_2 - \mu_1 = \Delta$ 0.25
$H_1: \mu_2 - \mu_1 > \Delta$ 0.25

$\Delta = \bar{d}_{obs} + t^* \frac{s_d}{\sqrt{n}}$ 0.5 = 0.25 + 0.25

1.5
$$\text{p-value} = pr\left((\bar{x}_2 - \bar{x}_1) > (\bar{x}_2 - \bar{x}_1)_{obs} \mid \mu_2 - \mu_1 = \Delta \right)$$ 0.5 0.5

$$= pr\left(\frac{(\bar{x}_2 - \bar{x}_1) - (\mu_2 - \mu_1)}{s_d/\sqrt{n}} > \frac{(\bar{x}_2 - \bar{x}_1)_{obs} - \Delta}{s_d/\sqrt{n}} \right) = pr\left(t > \frac{\cancel{(\bar{x}_2 - \bar{x}_1)_{obs}} - \cancel{\bar{d}_{obs}} - t^* s_d/\sqrt{n}}{s_d/\sqrt{n}} \right)$$ 0.25 0.25

$$= pr(t > -t^*) = 1 - \alpha/2$$

hw-lect 26-1

~ 3 **17.** In a simple regression model, let y_0 denote how far above and below the true fit one must go such that the probability of a random sample fit being in that interval is 0.98.

a) Write **one equation** that is the mathematical translation of the above information. (This may seem like a strange and new question, but it's not; it's what you've been doing all along. I'm separating it here only for grading purposes.)

1.5
$$pr\left(y(x) - y_0 < \hat{y}(x) < y(x) + y_0 \right) = 0.98.$$ 0.5 0.5 0.5

b) Now, suppose from a fit to a sample of size 16, we have found $s_e = 4$, and that at $x = 11$ we have $s_{est\ err} = 3$, $s_{pred\ err} = 5$. Find the numerical value of y_0.

Standardize!

standardization = 1 = 0.5 + 0.5 (No double penalty for prev. part)

1.5
$$pr\left(\frac{\cancel{y(x)} - y_0 - \cancel{y(x)}}{s_{est.err.}} < \frac{\hat{y}(x) - y(x)}{s_{est.err}} < \frac{\cancel{y(x)} + y_0 - \cancel{y(x)}}{s_{est.err.}} \right) = 0.98$$

$$pr\left(\frac{-y_0}{s_{est.err}} < t < \underbrace{\frac{y_0}{s_{est.err}}}_{= 2.6} \right) = 0.98$$

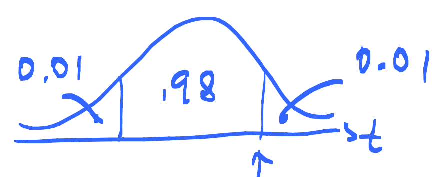

2.6 df = n − 2 = 14
0.5

$$\therefore y_0 = 2.6\, s_{est.err}$$
$$y_0 = 2.6\,(3)$$

If $y_0 = t^* s_{\hat{y}} \Rightarrow$ 0.5 only, even though the answer is same.