STAT 0116: Introduction to Statistical Sciences

Practice Exam 1 **Name:** Key

Fall 2026

**Instructions**

- This exam consists of 5 questions. The point value for each question is provided at the start of the prompt.
- You may use **one hand-written sheet of notes** (front and back).
- No phones or other electronic devices are permitted.
- Show all work clearly. Unsupported answers may not receive full credit.

I have neither given nor received unauthorized aid on this examination, nor have I concealed any violation of the Honor Code. I have not used any outside resources, including the internet or generative AI, to complete this exam, nor have I spoken about this exam with anyone else.

**Signature:** _____

| Question | Points |
|---|---|
| 1 | 5 |
| 2 | 12 |
| 3 | 11 |
| 4 | 12 |
| 5 | 16 |
| **Total** | **57** |

1. [5 pt] In a coastal town, 36% of the residents fish regularly. Among those who fish regularly, 68% catch at least one fish on any given trip. For those who don't fish regularly, only 12% catch a fish if they happen to go out fishing. If a randomly selected person from the town is observed catching a fish, what is the probability that they are someone who fishes regularly?

   Let $R$ = fishes regularly and $C$ = catches a fish. Then $P(R) = 0.36$, $P(C \mid R) = 0.68$, and $P(C \mid R^c) = 0.12$.

   $$P(R \mid C) = \frac{P(C \mid R)\,P(R)}{P(C \mid R)\,P(R) + P(C \mid R^c)\,P(R^c)} = \frac{(0.68)(0.36)}{(0.68)(0.36) + (0.12)(0.64)} = \frac{0.2448}{0.3216} \approx 0.761$$

   About 76% of people observed catching a fish are regular fishers.

2. [12 pt] In youth hockey in Canada, players are grouped into age divisions by calendar year, so a child born in January can be almost a full year older (and bigger) than a teammate born in December. Researchers studying this "relative age effect" suspected that players born early in the year are more likely to reach the professional level. To investigate, they took a random sample of 300 National Hockey League (NHL) players and recorded each player's birth month, position (forward, defense, or goalie), jersey number, and height (in inches). Of the 300 players, 105 were born in January, February, or March.

   a. [1 pt] Who or what are the observational units in this study?

   > NHL players (each of the 300 sampled players).

   b. [2 pt] Classify each of the four variables the researchers recorded as categorical or quantitative.

   > - **Birth month:** categorical.
   > - **Position:** categorical.
   > - **Jersey number:** categorical. Even though it is a number, it is a label: averaging jersey numbers doesn't mean anything.
   > - **Height:** quantitative.

   c. [2 pt] If birth month had nothing to do with reaching the NHL, about one quarter of players would be born in January–March. Write hypotheses, in words and in notation, that address the researchers' suspicion.

   > $H_0 : \pi = 0.25$. NHL players are no more likely to be born in January–March than in any other quarter of the year.
   > $H_A : \pi > 0.25$. NHL players are more likely to be born in January–March.

   d. [1 pt] Give the value of the statistic, and describe what it measures.

   > $\hat{p} = 105/300 = 0.35$, the proportion of the 300 sampled NHL players who were born in January, February, or March.

3. [11 pt] The amount of time Middlebury students spend waiting in line at the dining hall during the lunch rush is strongly right-skewed: most students wait only a few minutes, but a few wait much longer. Across all students, the mean wait time is $\mu = 6$ minutes and the standard deviation is $\sigma = 4$ minutes.

   a. [2 pt] Would it be appropriate to use a Normal distribution to find the probability that a single randomly selected student waits more than 10 minutes? Explain.

   > No. The population of wait times is strongly right-skewed, not Normal. A single student's wait time follows the population distribution, so a Normal model would give a misleading probability.

   b. [2 pt] Interpret the standard deviation in the context of the problem.

   > The average distance from a student's wait time to the sample average wait time of 6 minutes is about 4 minutes.

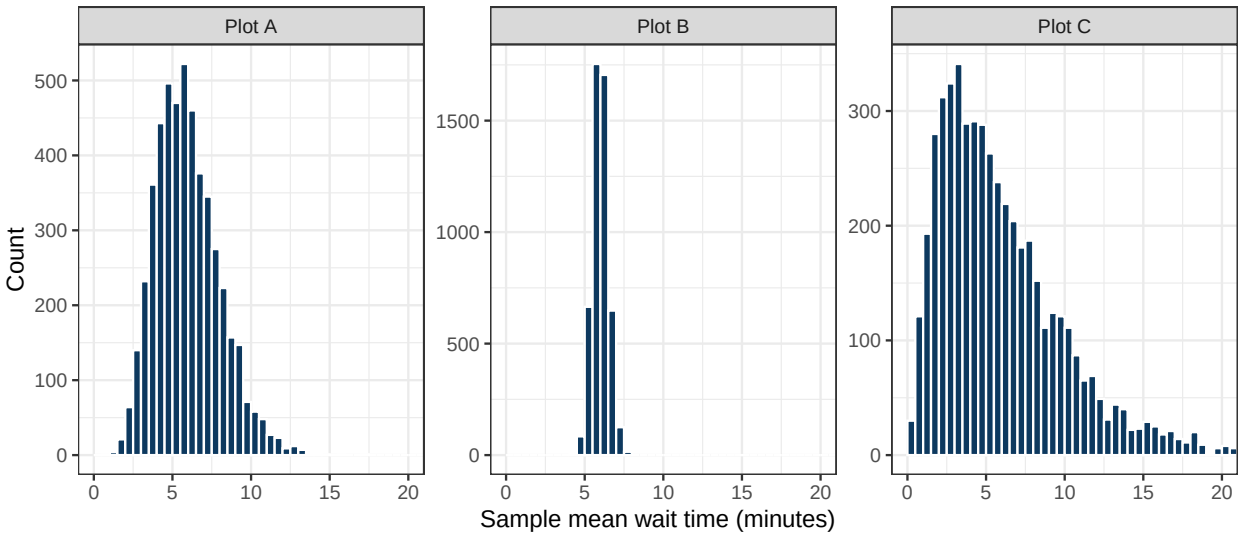
   c. [3 pt] Suppose we take a random sample of $n = 64$ students and compute their mean wait time, $\bar{x}$. Describe the sampling distribution of $\bar{x}$: give its shape, center, and spread, and justify your answer for the shape.

   > - **Shape:** approximately Normal. By the Central Limit Theorem, the sample size ($n = 64$) is large enough that the sampling distribution of $\bar{x}$ is approximately Normal even though the population is skewed.
   > - **Center:** $\mu = 6$ minutes.
   > - **Spread:** $SD = \sigma/\sqrt{n} = 4/\sqrt{64} = 0.5$ minutes.

   d. [2 pt] Between what two values would you expect the middle 95% of sample means (for samples of size 64) to fall?

   > $6 \pm 2(0.5)$, so between **5 and 7 minutes**.

   e. [3 pt] The histograms below each show 5,000 sample means computed from repeated random samples from this population. One plot used samples of size $n = 1$, one used $n = 4$, and one used $n = 64$. Match each plot to its sample size and explain your reasoning.

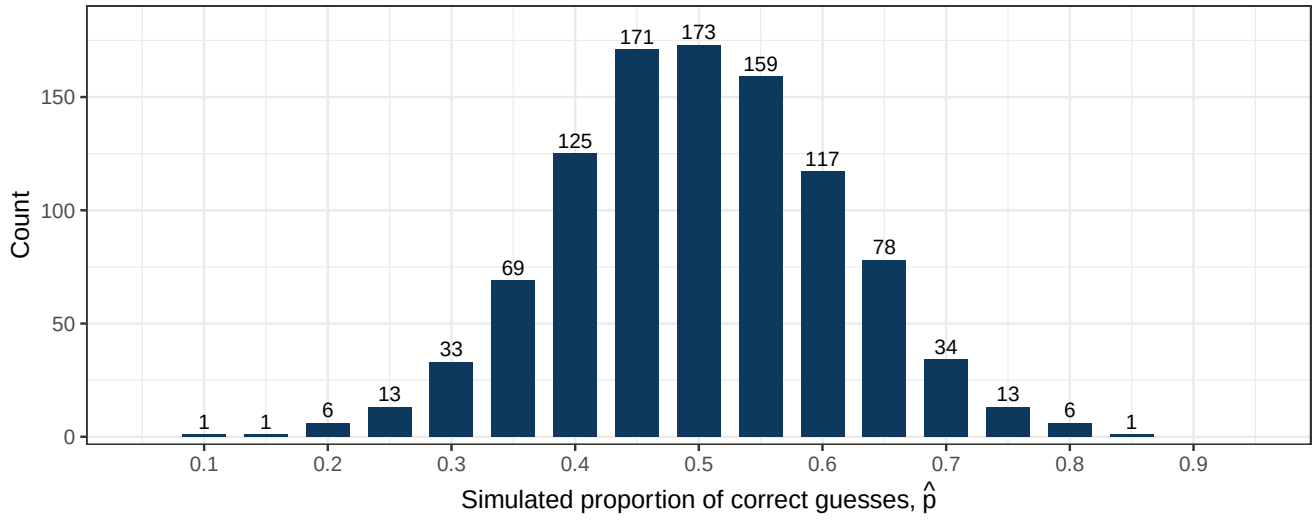



- **Plot C is $n = 1$:** it looks like the population itself, strongly right-skewed with a spread of about 4 ($\sigma/\sqrt{1} = 4$).
- **Plot A is $n = 4$:** it is still slightly right-skewed and moderately spread out ($SE = 4/\sqrt{4} = 2$).
- **Plot B is $n = 64$:** it is symmetric, bell-shaped, and very narrow ($SE = 4/\sqrt{64} = 0.5$).

As $n$ increases, the sampling distribution of $\bar{x}$ gets narrower (its SE is $\sigma/\sqrt{n}$) and more Normal (CLT). All three plots are centered at 6.

4. [12 pt] A Middlebury student claims she can tell real Vermont maple syrup from imitation "pancake syrup" by taste alone. To test her claim, her roommate runs 20 blind trials. In each trial, the student is blindfolded and tastes two cups, one with each type of syrup, in an order decided by a coin flip. She then says which cup has the real maple syrup. She identifies the real maple syrup correctly in 15 of the 20 trials.

   The null distribution for this problem is provided below.




   a. [2 pt] Describe the parameter of interest in words, and calculate the observed statistic, $\hat{p}$.

   > The parameter $\pi$ is the student's long-run proportion (probability) of correctly identifying the real maple syrup on a trial.
   > $\hat{p} = 15/20 = 0.75$

   b. [2 pt] State the null and alternative hypotheses in words and in symbols.

   > $H_0 : \pi = 0.5$. The student is just guessing; the probability that she is correct is 0.5.
   > $H_a : \pi > 0.5$. The probability that she is correct is greater than 0.5.

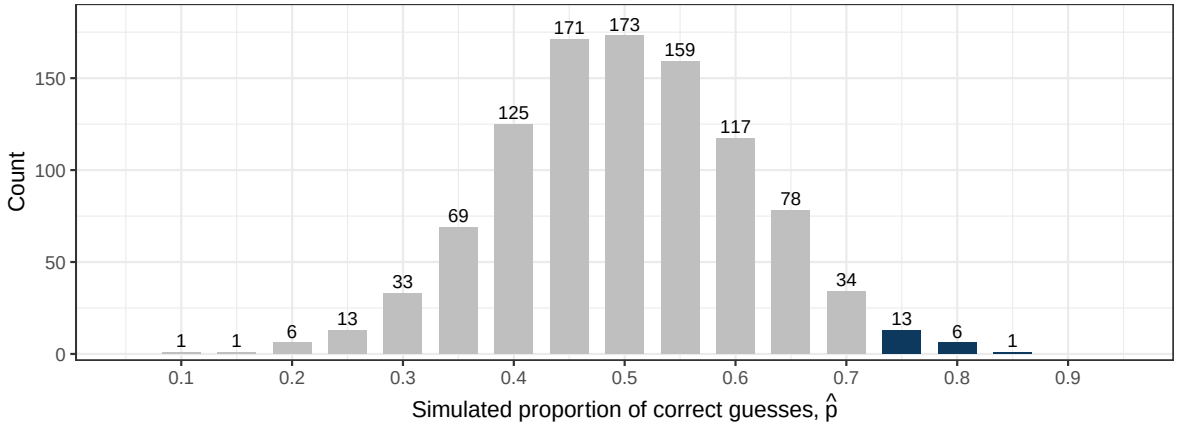
   c. [2 pt] Why is the simulated distribution centered at 0.5?

   > The null distribution is simulated assuming $H_0$ is true. If she is guessing, each trial is like a coin flip with a 0.5 chance of being correct, so the simulated proportions of correct guesses center at $\pi_0 = 0.5$.

   d. [2 pt] The null distribution is based on 1000 simulated samples. Find the p-value.

   > We count the simulations with $\hat{p} \geq 0.75$ (shaded below): $13 + 6 + 1 = 20$.

$$\text{p-value} \approx \frac{20}{1000} = 0.02$$




e. [2 pt] Interpret the p-value from part (d) in the context of this study.

If the student were only guessing, there would be about a 2% chance of her correctly identifying the real maple syrup in 15 or more of the 20 trials.

f. [2 pt] State your conclusion in the context of the problem. Which type of error (Type I or Type II) could you have made? Describe what that error would mean here in the context of the problem..

A p-value of 0.02 provides moderately strong evidence against the null hypothesis that the student is just guessing. Therefore, there is fairly strong evidence that the student can identify real Vermont maple syrup more often than she would by guessing.

Because we rejected $H_0$, we could have made a Type I error. That would mean concluding she can tell the syrups apart when she was really just guessing and got lucky.

5. [16 pt] In 2022, Middlebury's sustainability office found that 40% of students bought single-use bottled water at least once a week. Since then, the College has installed water-bottle refill stations in every academic building. This fall, the office surveyed a random sample of 150 Middlebury students, and 45 of them said they buy bottled water at least once a week. The office wants to know whether weekly bottled-water buying has become less common. It designed its test to have a significance level of $\alpha = 0.025$.

   a. [2 pt] What population does the sustainability office want to learn about? What is the parameter, and what statistic do we use to estimate it? Use notation, and give a value for anything you can compute.

   - **Population:** all current Middlebury students.
   - **Parameter:** $\pi$, the proportion of *all* current Middlebury students who buy bottled water at least once a week (unknown).
   - **Statistic:** $\hat{p} = 45/150 = 0.30$, the proportion of the 150 surveyed students who buy bottled water at least once a week.

   b. [2 pt] Write the hypotheses in symbols, and explain in one sentence each what they claim.

   $H_0 : \pi = 0.40$. Nothing has changed: the true proportion of students who buy bottled water weekly is still 0.40.
   $H_A : \pi < 0.40$. The true proportion of students who buy bottled water weekly has decreased since 2022.

   c. [2 pt] Can the Central Limit Theorem be used to approximate the null distribution of $\hat{p}$? Check the necessary conditions.

   Yes.

   - **Independence:** the students were randomly sampled, so there is no reason to think one surveyed student's answer affects another's.
   - **Success–failure:** under $H_0$, we expect $n\pi_0 = 150(0.40) = 60 \geq 10$ successes and $n(1-\pi_0) = 150(0.60) = 90 \geq 10$ failures.

   d. [3 pt] How many standard errors below the null value is the observed sample proportion? Show how you calculate the standard error.

$$SE_0 = \sqrt{\frac{\pi_0(1-\pi_0)}{n}} = \sqrt{\frac{0.40(0.60)}{150}} = \sqrt{0.0016} = 0.04$$

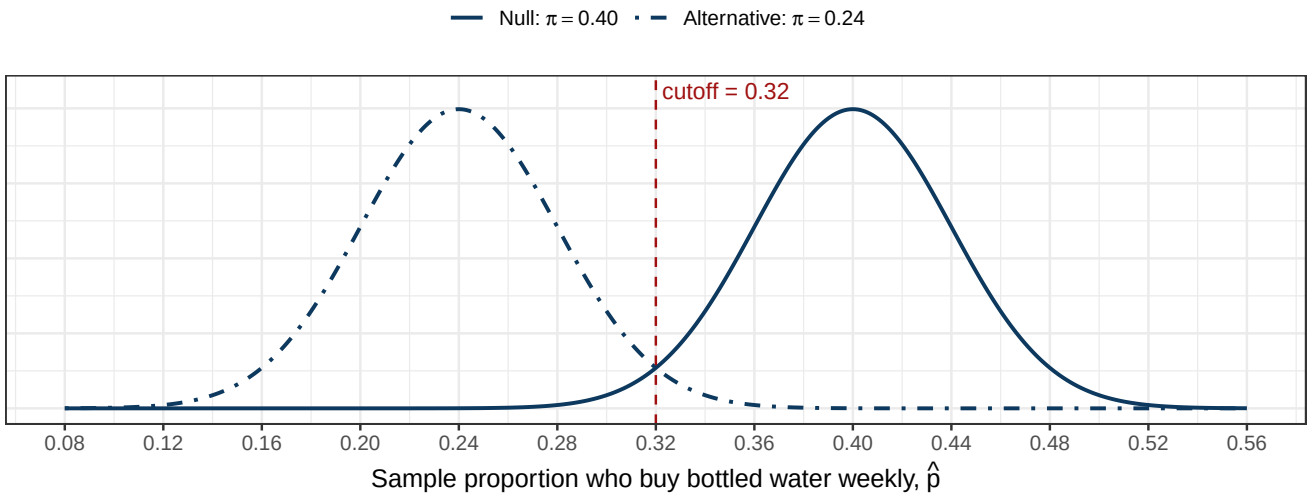
$$z = \frac{\hat{p} - \pi_0}{SE_0} = \frac{0.30 - 0.40}{0.04} = -2.5$$

The observed proportion is 2.5 standard errors below the null value.

e. [3 pt] Give a range of plausible values for the p-value. Then describe the strength of evidence against the null hypothesis, and state your conclusion in the context of the problem.

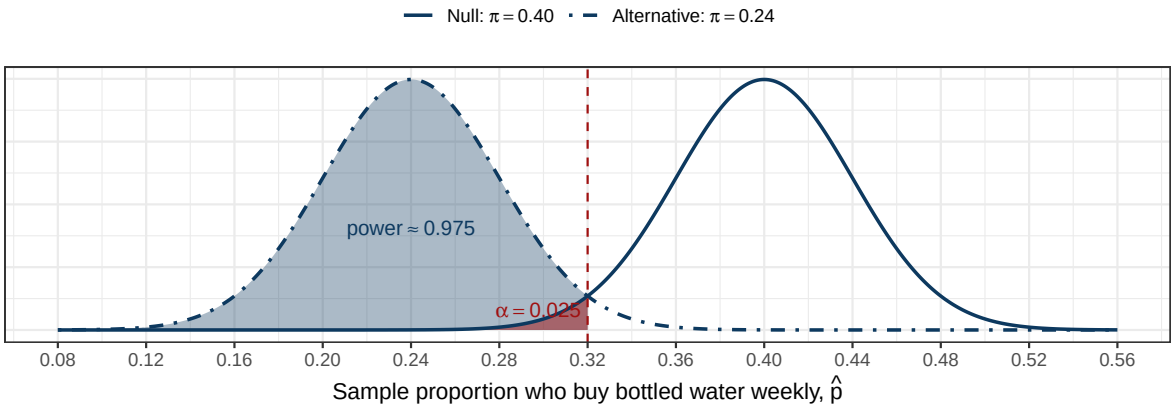
> This is a lower-tailed test, so the p-value is the area below $z = -2.5$. Since $-2.5$ falls between $-3$ and $-2$, the empirical rule gives an area below $-2$ of about 0.025 and below $-3$ of about 0.0015:
>
> $$0.0015 < \text{p-value} < 0.025.$$
>
> (The exact p-value is about 0.006.)
> A p-value this small (less than 0.025, and likely well below it) provides strong evidence against the null hypothesis that nothing has changed. Therefore, there is strong evidence that the true proportion of Middlebury students who buy bottled water at least once a week has decreased since 2022.

The plot below shows two possible sampling distributions of $\hat{p}$ for random samples of 150 students: one if $H_0$ is true ($\pi = 0.40$), and one if the true proportion has dropped to $\pi = 0.24$. Treat the standard error as approximately 0.04 for both. At $\alpha = 0.025$, the office rejects $H_0$ whenever $\hat{p}$ falls to the left of the dashed line.




f. [2 pt] On the plot, shade the area that represents the power of the test if $\pi = 0.24$. Then use the empirical rule to approximate the power.

> Power is the probability of rejecting $H_0$ when $\pi = 0.24$: the area under the **alternative** curve to the *left* of the cutoff (shaded blue). The cutoff $0.32 = 0.24 + 2(0.04)$ is 2 SE above the center of the alternative distribution:
>
> $$\text{Power} \approx 0.50 + 0.475 = 0.975$$




g. [2 pt] Suppose the office had used $\alpha = 0.05$ instead of $\alpha = 0.025$. What would happen to the location of the cutoff, the power of the test, and the probability of a Type I error? Then name one change, other than changing $\alpha$, that would increase the power of the test.

With $\alpha = 0.05$, the red area under the null curve must be larger, so the cutoff moves to the right (closer to 0.40). More of the alternative curve then falls to the left of the cutoff, so power increases. The probability of a Type I error increases from 0.025 to 0.05. Another way to increase power: survey more students. A larger $n$ shrinks the standard error, so both curves get narrower and overlap less. (Power would also be higher if the true proportion were farther below 0.40, but the office can't control that.)