Stat 1300: The Clocktower Census

Baseline Inference & The Statypus Audit

The Task: How "Billiken" is Saint Louis University? Later today, you will spend exactly **5 minutes** at the Clocktower (or another location your instructor designates) tallying students to estimate the true proportion (p) of the campus population wearing SLU-branded apparel.
The Method: We begin with a **Null Hypothesis (H_0)**—a baseline assumption of how the world works (p_0). You will gather real-world data (x_0, n_0) and then return to the classroom to perform a formal statistical test.

Complete Part 1 and Part 2 before heading outside.

Part 1: The Formal Framework

1. **Defining the Population:** Who exactly are we studying? How might the day of the week or time of day influence your sample?

2. **The Parameter (p):** Define in words the specific population proportion you are investigating.

3. **The Null Value (p_0):** What is your "Prior Guess" for the proportion?

4. **The Claim and Hypotheses:** Write a one-sentence claim in plain English about what you suspect the true proportion is compared to your p_0, then provide the formal notation for H_0 and H_a.

5. **The Field Prediction (n_{pred}):** How many total students do you think will walk past in 5 minutes?

Part 2: The Mental Model (Pre-Game)

Run the following code in your R console using your specific values from Part 1:

```
# Fill in your predicted n_pred for 'size' and your p_0 for 'prob':
# x_sim <- rbinom(n = 1, size = _______, prob = _______)
# x_sim
```

1. **The Simulation Result (x_{sim}):** How many SLU-clad students did R generate?

2. **The Simulated Proportion ($\hat{p}_{sim}$):** Use n_{pred} and x_{sim} to calculate your proportion.

3. **Explain the Code:** Why is there a `#` symbol? What happens if you run it without removing it?

Part 3: Field Observations

Note: *You only need these first two pages for the field. You will receive the analysis section (Pages 3-4) back in class.*

Head to your designated location now. Observe for exactly 5 minutes.

Total Students (n_0)	**SLU Gear (x_0)**	**Observed Proportion ($\hat{p}_0$)**

The Intuition Check: Consider your gathered proportion ($\hat{p}_0$) relative to your assumptions.

a. How close was your $\hat{p}_0$ to your single simulation $\hat{p}_{sim}$?

b. How much evidence does $\hat{p}_0$ provide against p_0? Is it enough to suggest your guess was wrong, or is this just a lucky/unlucky 5-minute window?

STOP: Once Part 3 is complete, head back to the classroom. Keep this paper until you have finished the digital analysis in Parts 4–6.

DIGITAL SUBMISSION: Complete Parts 4, 5, and 6 in a Google Doc. Use RStudio to run your analysis and GitHub Copilot to assist with your simulation. Submit a PDF of your final work to Canvas.

Part 4: Formal Analysis

Use RStudio and your Google Doc to see if your intuition holds up.

The Exact Binomial Test: This test calculates the exact probability of observing your result (or something more extreme) assuming that your null hypothesis p_0 is the objective truth.
The R Command:

```
binom.test(x = [your x_0], n = [your n_0], p = [your p_0])
```

In your Google Doc, address the following:

- **The P-Value:** Explain the p-value by describing how *rare* or *common* your 5-minute observation is in a world where p_0 is actually true.
- **The Confidence Interval:** Look at the "95 percent confidence interval." Does your original p_0 fall inside this range? What does that tell you about your initial guess?

Part 5: The Statypus Audit (GitHub Copilot)

You will now use GitHub Copilot within RStudio to visualize the "Null Distribution"—the world where your p_0 is true.

1. **The Prompting Challenge:** Open a new R Script. Use comments (`#`) to prompt Copilot to generate a script. You may need to refine your request until the code works. Your code must meet these constraints:
   - It must use **Base R** (no Tidyverse or ggplot2).
   - It must simulate **10,000 trials** using your field n_0 and your guess p_0.
   - It must generate a **histogram of proportions**.
   - It must include a **vertical line** marking your observed field proportion $\hat{p}_0$.
2. **Documentation:** In your Google Doc, provide the following:
   - A copy of the **final prompt(s)** used to generate the code.
   - A copy of the **final R code** itself.
   - An image of the resulting **histogram**.
3. **Code Intuition:** Look at the script. Even if you didn't write it, take a "guess" at what the specific functions are doing. What part of the code do you think is responsible for the "10,000 trials"? What part is drawing the vertical line?
4. **Visual Analysis:**
   - Where does your line ($\hat{p}_0$) fall relative to the center of the simulations?
   - How does this visualization explain the p-value you calculated in Part 4?

Part 6: Full Reflection

Conclusion of the Audit.

1. **The Mathematical Goal:** In Part 1, you calculated a p_0 that you likely didn't believe was true. Since we know the probability of any guess being perfectly correct is zero, why do we bother testing against a Null Hypothesis at all? What is the actual goal of this statistical process if we already suspect the Null is false?

2. **Sensitivity to n:** If you had stayed for 60 minutes and observed the *exact same proportion* ($\hat{p}_0$), would your p-value get larger or smaller? Why?

3. **The Frequentist Perspective:** If you repeated this 5-minute census 100 times, how often would you expect to see a result as extreme as yours if p_0 were actually true?

4. **Data Fidelity:** Be honest—how much do you trust your n_0 and x_0? Did you miss anyone? Did you get distracted? How does the "shortness" of a 5-minute window impact the reliability of your data compared to a 60-minute or even a 24-hour observation?

FINAL SUBMISSION CHECKLIST:

1. Turn in the physical **Pages 1–2** to the instructor after finishing Parts 4–6.
2. Upload the **PDF of your digital work** (Parts 4–6) to Canvas.