

# Stat 1300: The Clocktower Census

## Baseline Inference & The Statypus Audit

### Statypus Insight: The Task & The Method

**The Task:** How “Billiken” is Saint Louis University? Later today, you will spend exactly **5 minutes** at the Clocktower 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:** Since we know your prior guess ( $p_0$ ) cannot be *exactly* perfect, write a one-sentence claim in plain English about what you suspect the true proportion actually is compared to your  $p_0$ . Then, provide the formal notation for  $H_0$  and  $H_a$ . It should be a one-tailed test.

5. **The Field Prediction ( $n_{\text{pred}}$ ):** 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:

### Coding Corner: Simulating the Null

```
# 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.

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

**Manual Documentation Check:** Review **Section 9.5.1** to find the syntax for the `binom.test` function. Construct the R command yourself using your field values ( $x_0$ ,  $n_0$ ,  $p_0$ ) and the alternative hypothesis ( $H_a$ ) you established in Part 1.

**In your Google Doc, address the following:**

- **The  $p$ -value:** Explain the  $p$ -value by describing how *rare* or *common* your observation is in a world where  $p_0$  is actually true.
- **Strength of Evidence:** Based on your  $p$ -value, how strong is your evidence against the null? Do not provide a simple yes/no; comment on the qualitative weight of the evidence.
- **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.”

### Statypus Insight: AI Prompting Guardrail

To ensure the output matches the tools we use in this course, explicitly instruct Copilot to use **base R**.

1. **The Prompting Challenge:** Open a new R Script and use comments (#) to prompt Copilot to generate a simulation. Your generated code must do the following:
  - Simulate **10,000 trials** using your observed number of students  $n_0$  and your proportion guess  $p_0$ .
  - Generate a **histogram of counts of students who were wearing SLU gear**.
  - Include a **vertical line** marking your observed field observation  $x_0$ .
2. **Documentation:** In your Google Doc, provide the final prompt(s) used, the final R code, and an image of the resulting histogram.
3. **Code Intuition:** Take a “guess” at what the specific functions are doing. What part of the code is responsible for the “10,000 trials”? What part is drawing the vertical line?
4. **Visual Analysis:**
  - Where does your line ( $x_0$ ) 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.*

### Reflection: The Census Synthesis

1. **The Mathematical Goal:** In Part 1, you calculated a  $p_0$  that you likely didn't believe was true. Why do we bother testing against a Null Hypothesis if we already suspect it is false? What is the actual goal of this statistical process?
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. **Temporal Context:** How did the specific *timing* of your 5-minute window (e.g., passing time between classes vs. middle of a lecture block) impact your data? How does this short snapshot differ from a 24-hour observation?
5. **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 human error impact your confidence in the audit?

---

### FINAL SUBMISSION CHECKLIST:

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