The Upgrade: Calculating the Evidence

1. From Theory to Script

In the previous block, we saw the "Safety Net" visually. Now, we apply it to a real experimental audit. You are testing whether a control group of plants meets the industry standard height of **5.0 units**.

Coding Corner: Prepping the Audit

Step 1: Prep the Audit Data
In your R console, create the dataset by filtering the built-in `PlantGrowth` data:
```
auditdata <- PlantGrowth$weight[ PlantGrowth$group == "ctrl" ]
```

Step 2: Get the Template
Navigate to Section 10.2.3: One-Sample Hypothesis Test Code on the Statypus site.

Step 3: Run the Test
Copy the code block from the website into a new R script. Adapt the variables to test if the mean of your `auditdata` is significantly different from $\mu = 5.0$.

Calculation: The Manual vs. The Machine

Record the following results from your script output:

- Sample Mean ($\bar{x}$):
- Calculated t-statistic:
- The p-value:

Reflection: The Safety Net Width

Question: With only $n = 10$ plants ($df = 9$), your "Safety Net" is quite wide. Look at your t-score. Does it need to be *larger* or *smaller* than 1.96 to be considered significant? Why?

2. Enter the Hero: The `t.test()` Shortcut

Statypus Insight: The Hero Function

After years of Lead Engineers sweating over manual calculations, R introduced a single function designed to do everything you just did—the t-calculation, the degrees of freedom check, and the p-value lookup—in one heroic stroke.

Coding Corner: Running the Shortcut

Task: In your console, type and run the following:

```
t.test( auditdata, mu = 5.0 )
```

Reflection: Comparison Audit

Compare the output of this one-line “Hero” to the manual script you built on Page 1.

Did the results match exactly? ☐ Yes ☐ No

Question: If the hero `t.test()` is so powerful, why did we spend the first half of this section learning about the “Safety Net” and calculating things by hand? What happens to a Lead Engineer who trusts the shortcut without understanding the distribution underneath?

3. Interpreting the Evidence

Statypus Insight: The Verdict Rules

Our Null Hypothesis was $\mu = 5.0$. To make a decision, we compare our p-value against our chosen Significance Level (α):

- If your p-value is less than α, you **Reject** the Null.
- If your p-value is greater than α, you **Fail to Reject**.

Reflection: The "Common" Event

The "Common" Event: Usually, we hunt for $p < \alpha$ (the rare event). But here, your p-value is significantly high (> 0.50).
In your own words, explain why a p-value this high means the variation we are seeing in the sample "happens more often than not" if the Null is actually true. If a result is this ordinary, can we justify claiming these plants are different from the 5.0 standard?

Reflection: The Cost of Certainty

As a Lead Engineer, you often have to choose between a small, fast audit ($n = 10$) or a large, expensive audit ($n = 500$). How does the t-distribution physically "punish" you for choosing the smaller sample size? (Hint: Think about the distance your t-score has to travel to reach the tail).

Please provide your final written reflections for the tasks above on the blank back of this page.