The Carnival Calibration: Intuition

Bill the Statypus says: Welcome to the Statypus Carnival! To win the big stuffed platypus, you have to shoot basketballs into carnival sized basketball hoops. But before we play, let's look at the setup. A standard NBA basketball rim is 9 inches in radius. The carnival rim is 8 inches.

1. **The Baseline:** Imagine taking a single free throw on a standard 9-inch NBA rim. Based on your own basketball skill, estimate your personal probability (p) of making the shot.

2. **The Squeeze:** Now imagine taking a single free throw on the 8-inch carnival rim. What is your estimated probability (p) of making this shot?

3. **The Core Relationship:** In statistics, the "target size" allowed for a success is a **Margin of Error (MOE)**. A smaller rim requires higher precision, meaning it has a smaller MOE. For a fixed sample (one single shot), describe the relationship between the MOE (rim size) and your **confidence** (your probability of success).

Reflection: The Trade-Off

Notice that a large value of MOE means you have low precision. Based on this, what is the relationship between confidence and precision?

Sally the Statypus says: Notice that my friend MOE walks in the same direction as your confidence. A bigger MOE offers more confidence, but that precision wants to go in the opposite direction! Estimation is a game of tug-of-war between confidence and precision.

The Carnival's "Hidden Math"

Bill the Statypus says: You noticed your confidence dropped when we shrunk the target. If we treat a basketball shot as a 2D physics problem mapping to a normal distribution, shrinking the rim by 1 inch reduces the probability of a make.

Sally the Statypus says: You are so right, Bill! I looked into it and a 60% shooter on a 9 inch rim becomes around a 47% on a carnival 8 inch rim. However, they don't stop there! They often raise the rim up to a height of 12 feet instead of the regulation 10 feet. When stepping up to the smaller, slightly taller carnival rim, the tighter margin of error drops the success rate of a typical 60% free throw shooter to around only **40% ($p = 0.40$)**.

Bill the Statypus says: Exactly! Carnivals, like casinos, always make sure the games are in their favor! I went to a carnival where they had the smaller and higher rims AND they made me make 2 out of 3 shots!

4. **The Standard Game:** To win the prize, the player must make **at least 2 out of 3** shots ($X \geq 2$). Using the techniques from Chapter 8 (look up `qbinom()`), calculate the probability of winning on the standard NBA rim ($p = 0.60$).

5. **The Carnival Reality:** Now calculate the probability of winning the game on the carnival rim ($p = 0.40$).

Sally the Statypus says: Wait, Bill! Shrinking the rim just a little bit dropped the player from winning the majority of the time to losing most of the time!

Reflection: The Business of Precision

Compare your two win probabilities above. How does a carnival entice players to play this game?

Chapter 9

The Candy Net: Transition to Inference

Bill the Statypus says: Let's return to the candy factory from 9.1. You open a single sample bag of $n = 52$ candies and find $x = 31$ white candies. Your point estimate is $\hat{p} = 31/52 \approx 0.596$. We need to build a "net" around this estimate.

Sally the Statypus says: To measure the spread of our sampling distribution, we want the Standard Deviation: $\sigma_{\hat{p}} = \sqrt{\frac{p(1-p)}{n}}$. But we face a Catch-22: we are sampling precisely because we do not know p! Because we are forced to approximate p by using our sample $\hat{p}$, we can no longer call it the Standard Deviation. We call this approximation the **Standard Error** **(SE)**.

1. **The Standard Error:** Write down the formula for the Standard Error and calculate the Standard Error for your candy bag.

2. **The Confidence Multiplier:** To build a 95% Confidence Interval, we need to know how many SEs to reach out in both directions. Using the Empirical Rule as a rough guide, we need roughly 2 SEs as SE is out best estimate for $\sigma_{\hat{p}}$.

   To find the exact boundaries that trap the middle 95% of a standard normal curve, find the code snippet just below Definition 9.7 in the textbook. Calculate the value of z^* for 95% confidence.

   Bill the Statypus says: The term z^* is known as the **Critical z-Value** for the given level of confidence.

3. **The Margin of Error (MOE):** Multiply your exact Z^* by your SE to find the width of your "net".

$$MOE = Z^* \cdot SE$$

4. **The Confidence Interval:** Our **Confidence Interval** is now the interval between our point estimate plus or minus MOE. Find the bounds of that interval now.

Chapter 9

The Estimation Engine: DIY

Bill the Statypus says: We've been using step-by-step techniques to find the components of a confidence interval, but what happens when you're out in the field and you don't have me or Sally around? You need to become an operator of your own tools. Your textbook is your engineering manual.

Coding Corner: Tool Retrieval & Execution

Open your `CarnivalInference.r` script. Part 1 is currently empty except for instructions. Follow the instructions there to find the 95% confidence interval for 31 white candies out of 52 total candies.

95% CI: [________________ , ________________]

Sally the Statypus says: Bill, you did it! You made this to work as basically just a function. If we input the values of `x`, `n`, and `C.Level`, the rest of the code just does all of the work for us!

Bill the Statypus says: That's the power of repeatable code!

5. **The Cost of Certainty:** Now, change your `C.Level` in your script to **0.99** (99% confidence) and run the script again. Record your new Margin of Error and Confidence Interval.

   99% MOE: ________________ 99% CI: [________________ , ________________]

Reflection: Balancing the Scale

Compare the 95% interval to the 99% interval.

- Which interval is wider?
- Explain, in your own words, why the 99% interval requires a wider "net." How does this mathematically align with your intuition from the basketball carnival on Page 1?

Applying the Estimation Engine

Now that you have your functional Estimation Engine, let's put it to work on some real data scenarios.

1. **The Union Poll:** A local manufacturing plant has around 500 workers who are considering unionizing. Management randomly samples $n = 50$ workers to gauge support for unionizing. Of those sampled, $x = 32$ state they would vote "yes".

   - Use the Code Template we used on Page 4 to calculate the **90% Confidence Interval** for the true proportion of workers supporting the union.

   - Write a sentence interpreting this interval in the context of the problem. See Example 9.7 in your book if you need a model.

2. **Melody's Shoes (Example 9.13):** In Section 8.4 of the textbook, we were introduced to Melody's struggle regarding putting on her shoes. (Reread Section 8.4 if you are unfamiliar with her story.) Navigate to **Example 9.13** in your textbook to see how confidence intervals apply to her high-stakes situation.

   - From the textbook example, what are Melody's values for x and n?

   - Use the Code Template we used on Page 4 to calculate a **98% Confidence Interval** using Melody's data. (Note: The value of 98% is not one that is done for you in Example 9.13!)

   - Based on this 98% interval, does it seem plausible that her true success rate is actually at least as good as a Blind Ape in a Dark Room? Explain.

What is "Confidence" Anyway?

Bill the Statypus says: We've built a 95% confidence net for a single candy bag. But what happens if we simulate 100 different people each pulling their own bag of $n = 52$ candies and building their own nets?

Coding Corner: The Confidence Ripple

1. Navigate to **Part 2** of your `CarnivalInference.r` script.
2. Download the `candy_data.csv` file using the direct link provided in the script comments and save it to your working directory.
3. Run the code to automatically build 100 different 95% confidence intervals and count exactly how many successfully capture the true factory proportion ($p = 0.50$).

Sally the Statypus says: Wait, Bill! The R output says exactly 93 out of the 100 intervals successfully trapped 0.50. But we used a 95% confidence level! Did our math break?

Bill the Statypus says: Not at all, Sally! We only took 100 samples. The 95% represents the success rate of our net-building method. Getting 93 captures out of a single batch of 100 is just natural variation—it's totally expected when we only run the process a few times!

Statypus Insight: Confidence is in the Process, Not the Interval.

Once a specific confidence interval is calculated and written down, the true population parameter p is either inside that specific interval or it is not. We use the term **95% Confidence** to describe the reliability of our *method*—if we construct thousands of intervals using this exact process, 95% of those intervals will successfully capture the true parameter.

Reflection: The Final Word

Explain in your own words why a statistician will say "We are 95% confident," but will refuse to say "There is a 95% probability that p is in our interval." Refer to Definition 9.6 and the following Big Idea from the textbook for assistance.