Learning Goals

- Understand the difference between theoretical and empirical probability and how to calculate each using formulas and simulations.
- Describe and simulate common discrete probability distributions: Bernoulli, Binomial, Uniform Discrete, and Poisson.
- Explore continuous probability distributions such as Uniform Continuous, Gamma, Exponential, and Normal, including their properties and R simulations.
- Apply R functions like `matrix()` and `apply()` to simulate repeated trials and summarize outcomes in probability problems.
- Explain the Law of Large Numbers and how increasing sample size improves the accuracy of empirical probability estimates.

Key Functions

For each of the following R functions/distributions, write down a brief definition of what it does and one short example.

- **matrix():**
- **apply():**
- **Bernoulli:**
- **Binomial:**
- **Poisson:**
- **Uniform:**
- **Gamma:**
- **Exponential:**
- **Normal:**

Key Concepts

1. What is the difference between a population and a sample?

2. What is the difference between theoretical and empirical probability?

3. As the number of simulations increases, what happens to the difference between empirical and theoretical probabilities?

4. How does the Law of Large Numbers connect to the results of a simulation?

R Practice Problems

- Simulate flipping 10 coins for 100 trials using `matrix()` and `apply()`. How many heads appear in each trial?

- Using the results above, how could you determine the probability that the number of heads is less than or equal to 4?