Sampling, Uncertainty, and Hypothesis Testing

Mathematical Intuition for Seminar & Worksheet

James Rice

Monday, February 2, 2026

Motivation: Why Do We Need Probability?

- In practice, we observe **samples**, not populations
- Samples vary due to **random chance**
- Statistics gives us tools to reason about this uncertainty

Key question: How confident should we be that an observed pattern reflects the population?

Sampling and Central Tendency

Let $Y_1, \ldots, Y_n$ be a random sample from a population with mean μ.

$$\bar{Y} = \frac{1}{n}\sum_{i=1}^{n} Y_i$$

- $\bar{Y}$ is our **best guess** of μ
- Different samples $\Rightarrow$ different sample means

This variability across samples is called **sampling variation**.

The Sampling Distribution

If we repeatedly draw samples and compute $\bar{Y}$:

- The collection of sample means forms a distribution
- This is the **sampling distribution of the mean**

$$\bar{Y} \sim \mathcal{N}\left(\mu, \frac{\sigma^2}{n}\right)$$

- Centered at the true mean μ
- Spread shrinks as n increases

Central Tendency + CLT

Idea: Individual observations can be skewed, but **averages** across many random samples cluster near the true mean.

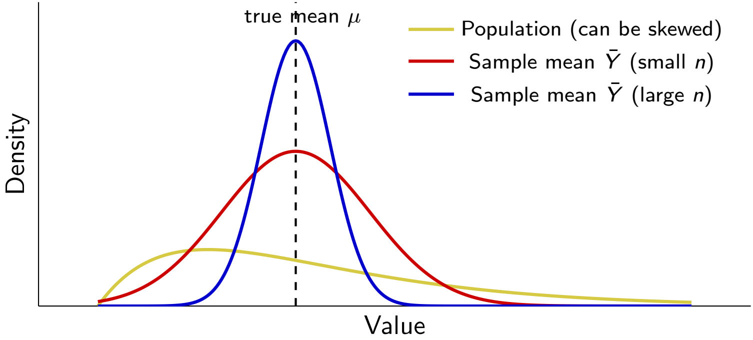

CLT (one line): As sample size n increases, the distribution of $\bar{Y}$ becomes more normal and more concentrated around μ (since $SE = \sigma/\sqrt{n}$ shrinks).

Standard Error: Measuring Uncertainty

The standard deviation of the sampling distribution is the **standard error**:

$$SE(\bar{Y}) = \frac{\sigma}{\sqrt{n}}$$

- Larger $n \Rightarrow$ smaller uncertainty
- Explains why large samples give more precise estimates

The same principle applies to:

- Differences in means
- Regression coefficients

Hypothesis Testing: Formal Setup

We test claims about the population using two hypotheses:

$$H_0 : \theta = \theta_0 \quad \text{(Null)}$$

$$H_A : \theta \neq \theta_0 \quad \text{(Alternative)}$$

- H_0 represents **no effect / no difference**
- H_A represents the presence of an effect

We assume H_0 is true and ask: *How surprising is our data?*

Sampling Distributions: When H_0 is True vs Not True

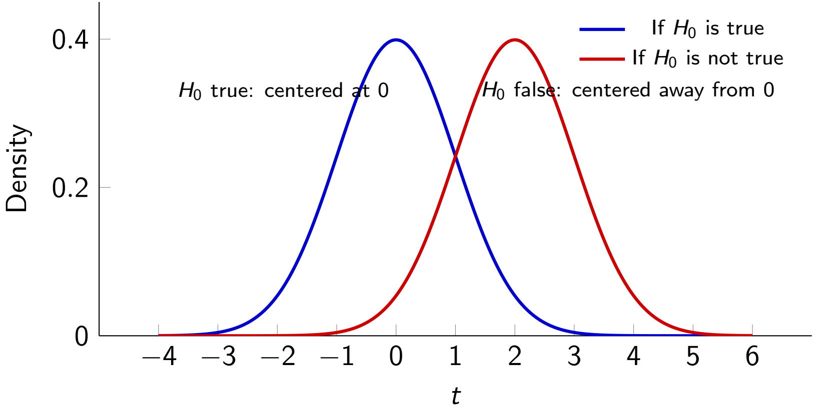

- If H_0 is true, repeated samples give test statistics clustered around 0.
- If H_0 is not true (a real effect exists), the distribution shifts away from 0.

Test Statistics and p-values

A generic test statistic:

$$z = \frac{\hat{\theta} - \theta_0}{SE(\hat{\theta})}$$

- Measures distance from the null in standard errors
- The **p-value** is:

$$P(|Z| \geq |z_{\text{obs}}| \mid H_0 \text{ true})$$

Small p-value $\Rightarrow$ data unlikely under the null.

Note: Often software will use a t-distribution for the test statistic. For our purposes, this is immaterial.

p-values vs Significance Level α (Two-Tailed Test)

Assume the null is true (e.g. $H_0 : \mu = \mu_0$). The **p-value** is the probability of results at least this extreme under H_0: $p = P(|Z| \geq |z_{\text{obs}}| \mid H_0)$.
Decision rule: reject H_0 if $p < \alpha$

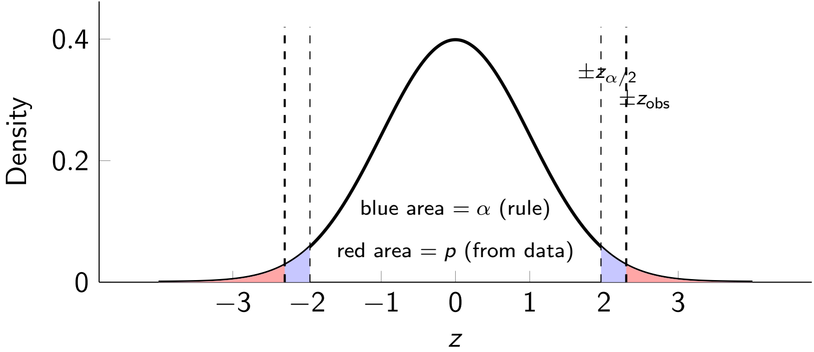

Two-tailed intuition: we count outcomes *too high or too low* relative to the null (both tails).

What Is a p-value?

The p-value answers:

If the null were true, what fraction of the normal distribution lies at values at least this far from μ_0?

- A p-value of 0.05 means **5% of the total area** of the distribution is more extreme than what we observed
- In a two-tailed test, this probability is split across **both tails**

Key intuition: small p-value $\Rightarrow$ the observed sample mean lies in a low-probability region of the null distribution.

Significance Levels and Decisions

Choose a significance level α (e.g. 0.05 or 0.01).

- If $p < \alpha$: reject H_0
- If $p \geq \alpha$: fail to reject H_0

Important:

- Rejecting H_0 does *not* prove H_A
- It means the data are inconsistent with H_0

Application: OLS Regression

Simple regression model:

$$Y_i = \beta_0 + \beta_1 X_i + \varepsilon_i$$

- β_1 measures the average change in Y for a 1-unit increase in X
- We test:

$$H_0 : \beta_1 = 0 \quad \text{vs} \quad H_A : \beta_1 \neq 0$$

Sampling Distribution of the OLS Slope

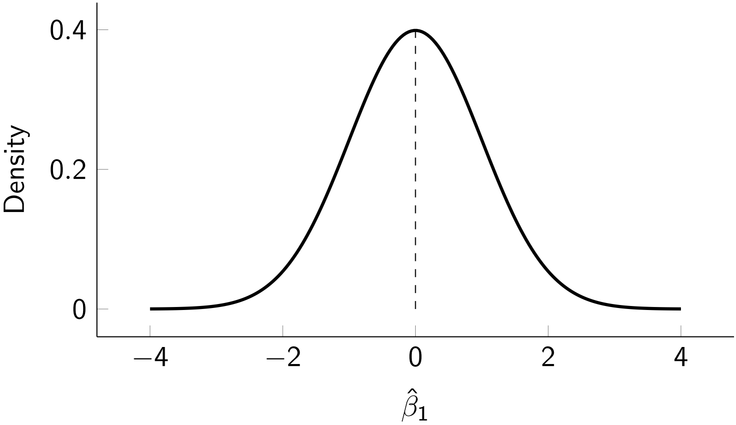

- Centered at 0 if H_0 is true
- Observed slope far from 0 ⇒ small p-value

Link to the Worksheet

In the worksheet you will:

- State null and alternative hypotheses
- Interpret t-statistics and p-values
- Compare results at 5% and 1% significance levels

Always ask:

- What is the parameter?
- What value does the null specify?
- How extreme is the observed estimate?