Unit 6: Inference for Proportions

Merrick Fanning

February 22, 2026

Unit 6 Outline: Inference for Proportions

Confidence Intervals as Probability Statements

Let θ be a fixed but unknown population parameter. A $100(1-\alpha)\%$ **confidence interval** for θ satisfies:

$$P\left(a < \theta < b\right) = 1 - \alpha$$

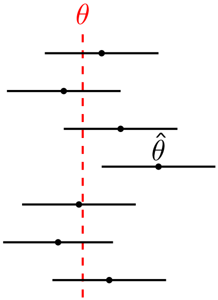

After the sample is taken, the interval is fixed and either contains θ or it does not - but the method captures θ in $100(1-\alpha)\%$ of repeated samples.

Confidence Interval for a Proportion

Suppose we take a random sample of size n with sample proportion $\hat{p}$. We want to construct a $100(1-\alpha)\%$ confidence interval for the population proportion p.
Assumptions:

- Random sampling
- Independence: $n < 10\%$ of the population N
- Normality: $np \geq 10$, $n(1-p) \geq 10$

1. We don't know p, how do we check the normality assumption?
2. What is the distribution of $\hat{p}$ under the assumptions?
3. Construct a $100(1-\alpha)\%$ confidence interval for p

How Do We Check the Normality Assumption?

The Problem

To use the normal model for proportions, we must check the normality assumption:

$$np \geq 10 \quad \text{and} \quad n(1-p) \geq 10$$

But p is unknown!

The Solution

We estimate p using the sample proportion $\hat{p}$. So we instead check:

$$n\hat{p} \geq 10 \quad \text{and} \quad n(1-\hat{p}) \geq 10$$

- This check uses the observed number of "successes" and "failures" in the sample.
- If both are greater than 10, we proceed with the normal model.

Constructing a Confidence Interval for p

We want a $100(1-\alpha)\%$ confidence interval for the population proportion p.
Under assumptions we have:

$$Z = \frac{\hat{p} - p}{\sqrt{\frac{p(1-p)}{n}}}$$

$$P\left(-z_{\alpha/2} \leq \frac{\hat{p} - p}{\sqrt{\frac{p(1-p)}{n}}} \leq z_{\alpha/2}\right) = 1 - \alpha$$

Rewriting:

$$P\left(\hat{p} - z_{\alpha/2}\sqrt{\frac{p(1-p)}{n}} \leq p \leq \hat{p} + z_{\alpha/2}\sqrt{\frac{p(1-p)}{n}}\right) = 1 - \alpha$$

$$\text{Standard Error: } SE(\hat{p}) = \sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$$

$$\Rightarrow \boxed{\hat{p} \pm z_{\alpha/2} \cdot \sqrt{\frac{\hat{p}(1-\hat{p})}{n}}}$$

Example: One-Sample z-Interval for p

A random sample of 250 people found that 172 support a new public transit initiative. Construct a 95% confidence interval for the true proportion p of supporters.

1. **State Information:**
   Procedure: One-Sample z-Interval for a Proportion
   Confidence Level: 95%, Goal: Estimate the true population proportion p
2. **Check Conditions:**
   - Random Sampling: Stated
   - Independence: $250 < 10\%$ of population of entire public
   - Normality: $n\hat{p} = 172 > 10, \quad n(1-\hat{p}) = 78 > 10$
3. **Calculate Interval:**
   $\hat{p} = \frac{172}{250} = 0.688,\ z^* = 1.96$
   $$\text{CI} = \hat{p} \pm z_{\alpha/2} \cdot \sqrt{\frac{\hat{p}(1-\hat{p})}{n}} = 0.688 \pm 1.96 \cdot \sqrt{\frac{0.688(1-0.688)}{250}}$$
   $$= (0.6317,\ 0.7443)$$
4. **Interpretation:** We are 95% confident that the true proportion of supporters lies between 63.2% and 74.4%.

How Wet Is the Earth?

Estimate the proportion of Earth's surface that is covered in water:

1. Using http://www.geomidpoint.com/random/, select 50 random points on Earth. What proportion fall in water?
   Sample Result: A sample of 50 points produced $\hat{p} = 0.75$
2. What is the parameter of interest in this investigation?
3. Are the conditions for inference met?
4. Construct and Interpret a 95% confidence interval estimate for p
5. How would the sample size you select effect the variability for the sampling distribution of $\hat{p}$.
6. Explain how each of the following effects the width of a confidence interval:
   - Increased confidence
   - Increased sample size

Determining Required Sample Size

Suppose you want a confidence interval for p that is within 0.1 of your estimate. What is the required sample size?

- We want the margin of error to be at most 0.1:

$$z^* \cdot \sqrt{\frac{p(1-p)}{n}} \leq 0.1$$

- Solve for n:

$$n \geq \left(\frac{z^*}{0.1}\right)^2 p(1-p)$$

- What is the most conservative estimate for p?

What Is the Most Conservative Estimate for p?

The standard error for a sample proportion is:

$$SE(\hat{p}) = \sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$$

When planning a study (e.g., determining sample size), we want the **maximum** standard error-this gives the **most conservative estimate**.
Let:

$$f(p) = \sqrt{\frac{p(1-p)}{n}}$$

Maximizing $f(p)$ is equivalent to maximizing $p(1-p)$.
Take the derivative:

$$f(p) = p(1-p) \Rightarrow f'(p) = 1 - 2p$$

Set $f'(p) = 0 \Rightarrow 1 - 2p = 0 \Rightarrow p = \frac{1}{2}$
So, $SE(\hat{p})$ is maximized when:

$$\hat{p} = 0.5$$

Determining Required Sample Size

Suppose you want a confidence interval for some p that is within 0.1 of your estimate. What is the required sample size?

- Use the conservative estimate $p = 0.5$, $z^* = 1.96$ (for 95% confidence):

$$n \geq \left(\frac{1.96}{0.1}\right)^2 (0.5)(0.5) = 96.04$$

- **Answer:** Round up $\rightarrow$ $\boxed{n = 97}$

The Lady Tasting Tea Experiment

- **Participants:** Muriel Bristol (phycologist) and Ronald A. Fisher (statistician) at Rothamsted Experimental Station.
- **Claim:** Bristol asserted she could taste whether milk was poured before or after the tea.
- Fisher devised a blind test: 8 cups total—4 milk-first, 4 tea-first—served in random order.
- **Null hypothesis:** Bristol is merely guessing (no real ability).
- What is the probability Bristol correctly Identifies 8 cups in a row?
- What would make the experimental results statistically significant? Numberphile Video (17min)
- A more modern example: Speedrunning in Minecraft (40min)

(Ronald Fisher, 1913)

Introducing Hypothesis Testing - Paper Toss

Let's study a player's accuracy in a game of paper toss. We consider their true shooting percentage to be a fixed value p, the population proportion we wish to estimate.

Setup: A game of paper toss consists of 10 shots from 3 meters away.

Claim

Suppose a player's accuracy is claimed to be 90%, or $p_0 = 0.9$.

Do you believe this claim?

As statisticians, we test such claims using **hypothesis testing**.

Introducing Hypothesis Testing - Paper Toss

We test the claim about the population proportion p.

Hypothesis:

$$H_0 : p = 0.9$$
$$H_a : p < 0.9$$

We suspect the player may be overestimating their skill.

Play the Game:

Shot Number	Hit or Miss
1	
2	
$\vdots$	$\vdots$
10	

Calculate your sample proportion:

$$\hat{p} = \frac{\#\text{ hits}}{10}$$

This is our statistic - it estimates the parameter p.

Introducing Hypothesis Testing - Paper Toss

1. In this scenario, will we have $\hat{p} \sim \mathrm{Normal}\left(\mu_{\hat{p}} = p, \sigma_{\hat{p}} = \sqrt{\frac{p(1-p)}{n}}\right)$?
2. Let X be the number of shots that you make assuming that your shots are independent with a completion rate of 0.9 (assuming the null hypothesis is true). What is the distribution of X?
3. Let x_0 be the number of shots made in your 'experiment'. What is $P(X \leq x_0)$?
4. Interpret the probability (the p-value) calculated in the previous question. Is it convincing evidence against H_0?
5. What p-value is low enough to convince us to reject H_0?

Type I and Type II Errors and Power

What Could Go Wrong?

Even with good data collection and analysis, there are two types of incorrect conclusions we could make when testing a hypothesis.

	H_0 True	H_0 False
Reject H_0	Type I Error α	✓
Fail to reject H_0	✓	Type II Error β

The **power** of a test tells us the probability that we reject H_0 when it is indeed false. This is $1-\beta$.

In the paper toss scenario:

1. What is a Type I error?
2. What is a Type II error?
3. What is the power?

One-Sample z-Test for a Proportion

Goal: Test a claim about a population proportion p based on a sample.

Assumptions:

- The sample is a **simple random sample**.
- Individual observations are **independent** ($n < 10\%$ of population).
- **Normality:** $np_0 \geq 10$, $n(1 - p_0) \geq 10$

Hypotheses:

$$H_0 : p = p_0 \quad \text{vs.} \quad H_a : p < p_0, \; p > p_0, \; \text{or } p \neq p_0$$

Test Statistic:

$$z = \frac{\hat{p} - p_0}{\sqrt{\frac{p_0(1-p_0)}{n}}}$$

Decision Rule:

- Find the p-value based on the direction of H_a.
- Compare to significance level α.
- **Reject** H_0 if p-value $< \alpha$.

1. State Hypothesis, significance level, statistics, parameter
2. Check assumptions for inference
3. Calculate Test-Statistic and P-Value (DRAW PICTURE)
4. Interpret p-Value in plain language

Example: One-Sample z-Test for a Proportion

A company claims that 60% of customers are satisfied. In a random sample of 100 customers, 52 report being satisfied. Test the claim at the 5% significance level.

Step 1: State Hypotheses and Information

- $H_0 : p = 0.60, \quad H_a : p < 0.60$
- $\alpha = 0.05, \quad \hat{p} = \frac{52}{100} = 0.52$

Example: One-Sample z-Test for a Proportion

A company claims that 60% of customers are satisfied. In a random sample of 100 customers, 52 report being satisfied. Test the claim at the 5% significance level.

Step 2: Check Conditions

- Random sample given.
- Independence: $100 < 10\%$ of population.
- Normality: $np_0 = 60 > 10$, $n(1 - p_0) = 40 > 10$

Example: One-Sample z-Test for a Proportion

A company claims that 60% of customers are satisfied. In a random sample of 100 customers, 52 report being satisfied. Test the claim at the 5% significance level.

Step 3: Compute z-Statistic and p-Value

$$z\text{-stat} = \frac{0.52 - 0.60}{\sqrt{\frac{0.60(0.40)}{100}}} = \frac{-0.08}{0.049} \approx -1.633$$

$$\text{p-value} = P(Z < -1.633) \approx 0.0512$$

Example: One-Sample z-Test for a Proportion

A company claims that 60% of customers are satisfied. In a random sample of 100 customers, 52 report being satisfied. Test the claim at the 5% significance level.

Step 4: Conclusion

- There is a 5.12% chance of observing a sample proportion of 0.52 or lower, assuming the true proportion is 0.60.
- Since p-value $= 0.0512 > \alpha = 0.05$, we **fail to reject** H_0. There is not enough evidence to conclude satisfaction is below 60%.

Two-Sample Z-Intervals for a Difference in Proportions

- From unit 5, what conditions are required for $p_1 - p_2 \sim \text{Normal}\left(p_1 - p_2, \sqrt{\frac{p_1(1-p_1)}{n_1} - \frac{p_2(1-p_2)}{n_2}}\right)$?
- Assuming all conditions are met construct a 95% confidence interval for $p_1 - p_2$:

$$(\hat{p}_1 - \hat{p}_2) \pm z^* \sqrt{\frac{\hat{p}_1(1-\hat{p}_1)}{n_1} + \frac{\hat{p}_2(1-\hat{p}_2)}{n_2}}$$

Two-Sample Z-Intervals for a Difference in Proportions

Each year, researchers test different flu vaccines. Suppose two vaccines were given to different groups:

- Vaccine A: Given to 1500 people, 1350 did not get the flu.
- Vaccine B: Given to 1200 people, 1020 did not get the flu.

Can we estimate the difference in effectiveness between the two vaccines?

Two-Sample Z-Intervals for a Difference in Proportions

We want to estimate the difference in the proportions of people who **did not get the flu**:

$$p_1 - p_2$$

- p_1: true proportion for Vaccine A
- p_2: true proportion for Vaccine B

We will construct a **two-sample Z-interval** to estimate $p_1 - p_2$.

Conditions for Two-Proportion Z-Interval

Before using a two-sample Z-interval, check conditions:

1. **Random:** Each sample was randomly selected.
2. **Independence:** Each group is less than 10% of the population of all people, so we may assume observations are independent.
3. **Normalality:** At least 10 successes and 10 failures in each group, so we may assume the sampling distribution for $\hat{p}_1 - \hat{p}_2$ is normal.

In our case:

- Vaccine A: 1350 success, 150 failure
- Vaccine B: 1020 success, 180 failure

All conditions are met.

Z-Interval Formula

The two-sample Z-interval for a difference in proportions is:

$$(\hat{p}_1 - \hat{p}_2) \pm z^* \sqrt{\frac{\hat{p}_1(1-\hat{p}_1)}{n_1} + \frac{\hat{p}_2(1-\hat{p}_2)}{n_2}}$$

Where:

- $\hat{p}_1 = 1350 \ / \ 1500 = 0.9$
- $\hat{p}_2 = 1020 \ / \ 1200 = 0.85$
- $z^* = 1.96$ for a 95% confidence level

Calculations

- Point estimate: $0.9 - 0.85 = 0.05$
- Standard error:
$$\sqrt{\frac{0.9(0.1)}{1500} + \frac{0.85(0.15)}{1200}}$$
- Margin of error: $1.96 \times \sqrt{\frac{0.9(0.1)}{1500} + \frac{0.85(0.15)}{1200}}$

Confidence interval:

$$0.05 \pm 1.96 \times \sqrt{\frac{0.9(0.1)}{1500} + \frac{0.85(0.15)}{1200}} \Rightarrow (0.02473,\ 0.07527)$$

Interpretation

We are 95% confident that the true difference in the proportions of people who did not get the flu is between 2.47% and 7.53%, with Vaccine A performing better.

Context: This means Vaccine A likely prevents the flu in about 2 to 8 more people per 100 compared to Vaccine B.

1. Press `STAT` $\rightarrow$ `TESTS`
2. Select `2-PropZInt`
3. Enter:
   - $x_1 = 1350$, $n_1 = 1500$
   - $x_2 = 1020$, $n_2 = 1200$
   - C-Level $= 0.95$
4. Choose `Calculate`

You should get: $(0.025, \ 0.075)$

Example: Do Email Campaigns A and B Perform Differently?

Scenario: A marketing team tests two email designs.

- Campaign A: 300 clicks out of 1500 emails
- Campaign B: 360 clicks out of 1600 emails

Step 1: State the Hypotheses and Parameters

- Let p_1: true click rate for Campaign A
- Let p_2: true click rate for Campaign B
- $H_0 : p_1 = p_2 \quad$ (or $p_1 - p_2 = 0$)
- $H_a : p_1 \neq p_2$
- Significance level: $\alpha = 0.05$

Step 2: Check Conditions for Inference

Random: The emails were randomly sent to customers.
Independence: Less than 10% of the entire customer base was sampled for both samples.
Normality (Large Counts):

- Campaign A: successes = 300, failures = 1200
- Campaign B: successes = 360, failures = 1240

All conditions are met — proceed with the two-proportion z-test.

Step 3: Compute Test Statistic and P-Value

Sample Proportions:

$$\hat{p}_1 = \frac{300}{1500} = 0.20, \quad \hat{p}_2 = \frac{360}{1600} = 0.225$$

Pooled Proportion:

$$\hat{p}_{\text{pooled}} = \frac{300 + 360}{1500 + 1600} = \frac{660}{3100}$$

Standard Error:

$$SE_{(\hat{p}_1 - \hat{p}_2)} = \sqrt{\frac{660}{3100}\left(1 - \frac{660}{3100}\right)\left(\frac{1}{1500} + \frac{1}{1600}\right)}$$

Z-Statistic:

$$z = \frac{0.20 - 0.225}{SE_{(\hat{p}_1 - \hat{p}_2)}} \approx -1.69926$$

P-Value: Two-tailed $\rightarrow 2 \times P(Z < -1.69926) \approx 0.08927021$

Step 4: Interpret the p-value

Interpretation:

The p-value is approximately 0.09. This means:

> *If there really is no difference in click rates between Campaign A and Campaign B, there is a 9% chance of observing a difference as extreme (or more) than what we saw in the sample.*

Decision at $\alpha = 0.05$:

Since $p = 0.09 > 0.05$, we **fail to reject** the null hypothesis.

Conclusion:

We do not have strong enough evidence to say the click-through rates are different between Campaign A and B.