Hypothesis Testing 3

More on Null Distributions and P-values, Test-stats

Grinnell College

Review

Hypothesis Testing: formal technique for answering a question with two competing possibilities

Null Hypothesis: represents a skeptical view or a perspective of no difference

- H_0: 'parameter' $=$ (some value)

Alternate Hypothesis: what the researchers actually want to show with the study

- H_A: 'parameter' $[< / > / \neq]$ (some value)

Review

This is the general outline we will follow for Hypothesis Testing

1. Define hypotheses
2. Create a null distribution (represents H_0)
3. See how our statistic compares to this
4. Compute a p-value
   - do the results look unlikely if H_0 is true?
5. Interpretations / Conclusions

Note: We will see alternatives to simulation for step 2 today

Null Distribution

The distribution of the statistics if the null hypothesis is true

- simulates what the null hypothesis looks like

We looked at the coin-flip scenario. Null distribution of fair coin, 10 flips

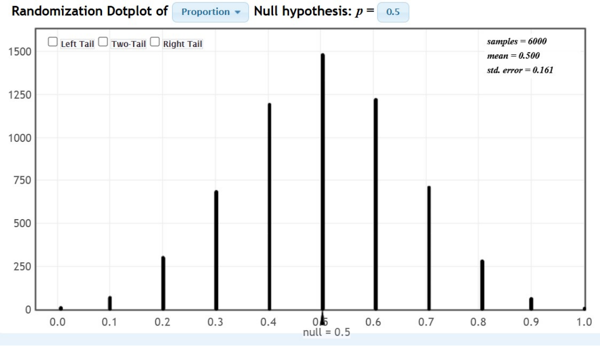

- Notice the shape of this (Normal). We could use Normal distribution methods to compute a p-value instead.

Null Distribution – Proportion

Conditions

- Representative Sample
- $np_0 \geq 10$ (Success condition)
- $n(1 - p_0) \geq 10$ (Failure condition)

With the conditions met, the null distribution for $\widehat{p}$ looks like this:

$$\widehat{p} \sim N(p_0, \sqrt{\frac{p_0(1-p_0)}{n}})$$

- distribution of $\widehat{p}$'s is Normal
- mean $= p_0$, sd $= \sqrt{\frac{p_0(1-p_0)}{n}}$

Note: We are going to do something so we can use Standard Normal distribution instead.

Test Statistics

$$\widehat{p} \sim N(p_0, \sqrt{\frac{p_0(1-p_0)}{n}})$$

The test-statistic saves us from having to plot this Normal distribution every time. We will *standardize* $\hat{p}$ to make the distribution simpler.

$$\underbrace{\frac{\widehat{p} - p_0}{\sqrt{\frac{p_0(1-p_0)}{n}}}}_{\text{Test Statistic}} \sim N(0,1)$$

- the whole term on the left is the Test-statistic
- we can compute this value and we know it follows N(0,1)

Coin Flip Example

For our scenario of 10 flips...

- $\hat{p} = 8/10 = 0.8$
- $p_0 = 0.5$
- $\sqrt{\frac{p_0(1-p_0)}{n}} = \sqrt{\frac{0.5(1-0.5)}{10}} = 0.158$

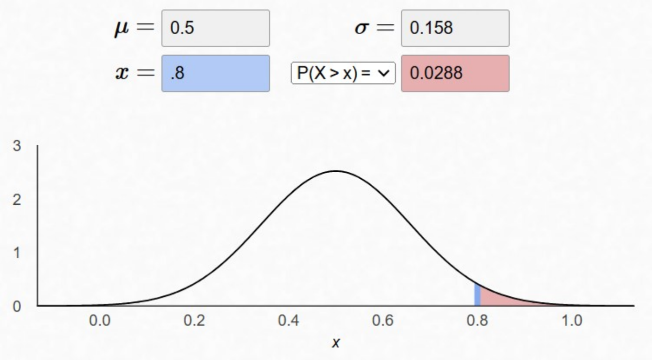

- can use Null distr. and $\hat{p}$ to get p-value (pnorm in R)

Coin Flip Example

For our scenario of 10 flips...

- $\hat{p} = 8/10 = 0.8, \ p_0 = 0.5$
- $Z = \frac{\widehat{p} - p_0}{\sqrt{\frac{p_0(1-p_0)}{n}}} = \frac{0.8-0.5}{\sqrt{\frac{0.5(1-0.5)}{10}}} = 1.897$

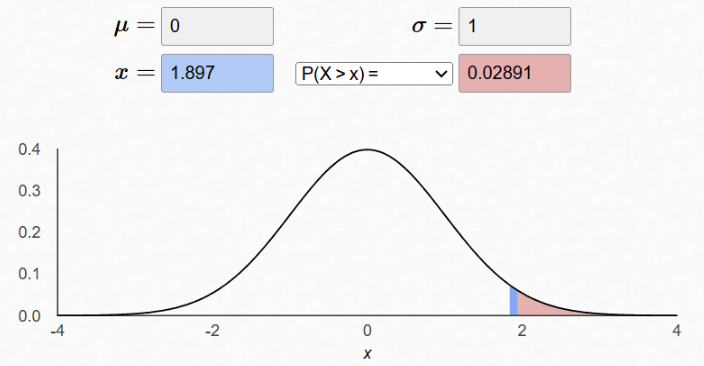

- std. Normal distribution and Z value to get p-value (‘pnorm()‘)
- Z-values are comparable across different studies

Test-statistics

From here on out, we will place almost all of our emphasis on Test-statistics instead of elaborately detailing what the null distribution should look like.

More on P-values

Forms of the Alternate Hypothesis

- H_A: parameter $>$ 'hypothesized value' (right-tailed test) **OR**
- H_A: parameter $<$ 'hypothesized value' (left-tailed test) **OR**
- H_A: parameter $\neq$ 'hypothesized value' (two-tailed test)

Up until now we have only seen how to calculate p-values Alternate Hypotheses that use the $>$ symbol.

- coin flip biased in favor of heads (H_A : p $>$ 0.5)
- Monday breakups (H_A: p $> \frac{1}{7}$)

More on P-values

H_A: parameter $>$ 'hypothesized value' (right-tailed test)

- to find a p-value, you find where the Test-statistic is on a N(0,1) distribution, then find the area to the right of it under the curve

Suppose test-stat $Z = 1.9$

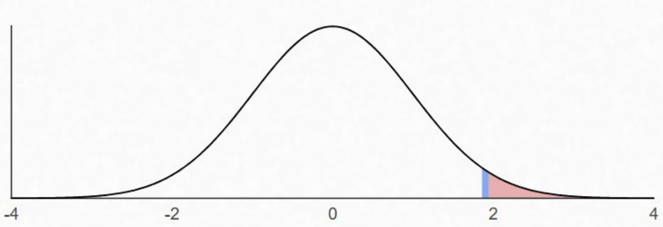


```
> pnorm(1.9, lower.tail=F)
[1] 0.02871656
```

- This is how we would get p-value for coinflip example

More on P-values

H$_A$: parameter $<$ 'hypothesized value' (left-tailed test)

- to find a p-value, you find where the Test-statistic is on a N(0,1) distribution, then find the area to the left of it under the curve

Suppose test-stat Z $=$ -.5

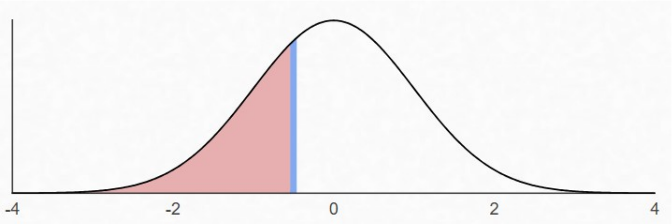


```
> pnorm(-.5)
[1] 0.3085375
```

More on P-values

H_A: parameter $\neq$ 'hypothesized value' (two-tailed test)

- to find a p-value, you find where the Test-statistic is on a N(0,1) distribution
- area to the right (Test-stat > 0) or area to the left (Test-stat < 0)
- multiply this area by 2

Suppose test-stat Z $= 1$

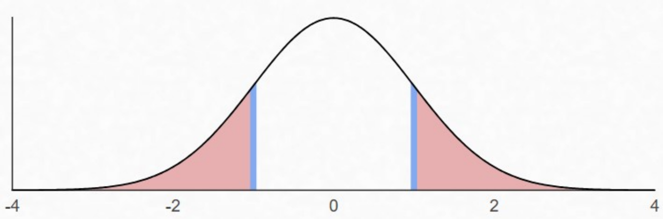


```
> 2*pnorm(1, lower.tail=F)
[1] 0.3173105
> 2*pnorm(-1)
[1] 0.3173105
```

Example: Iowa Smoking Rates

A public health researcher wants to know if adult rate of smoking cigarettes in Iowa is lower than the national rate of 11%. A random sample of 500 Iowans adults is surveyed and finds a sample rate of 10%. Do these data provide convincing evidence that the smoking rate in Iowa is lower than the national average?

H_0: p $= 0.11$ H_A: p < 0.11

$$Z = \frac{\widehat{p} - p_0}{\sqrt{\frac{p_0(1-p_0)}{n}}} = \frac{0.10 - 0.11}{\sqrt{\frac{.11 \times .89}{500}}} = \frac{-.01}{\sqrt{0.0001958}} = -0.71$$

```
> pnorm(-0.71, lower.tail=T)
[1] 0.2388521
```

Conclusion According to this study there is little to no evidence the adult rate of smoking in Iowa is less than the national rate.

Note: the real rate of smoking in Iowa is MUCH higher. Study from 2019: https://hhs.iowa.gov/media/11502/download?inline

Summary

We can use Normal distribution stuff we've seen previously to get p-values.

- Check conditions to make sure it works!
- Compute Z-score then use pnorm to get p-value

Which method to use?

- Conditions met $\rightarrow$ Normal distribution
- Conditions not-met $\rightarrow$ Simulation (Statkey)

Bonus: My plug for probability theory

Note: You will not be tested on this slide.
Where does the null distribution come from?

Suppose we have a sample of n observations from a population with a binary outcome (1 or 0) and p = P(success) = P(1). We can treat these as Bernoulli events, then using the Bernoulli random variable stuff from the RV lecture slides. Bernoulli: E(X) = p, var(X) = p(1-p)

n independent Bernoulli events, look at the average of them all as $\hat{p} = \bar{X}$

- $E(\widehat{p}) = E(\frac{1}{n}\sum X_i) = \frac{1}{n}(np) = p$
- $V(\widehat{p}) = V(\frac{1}{n}\sum X_i) = \frac{1}{n^2}np(1-p) = \frac{p(1-p)}{n}$
- large n $\rightarrow$ $\hat{p}$ will be approx. Normal (CLT)

And so $\widehat{p} \sim N(p, \sqrt{\frac{p(1-p)}{n}})$