Hypothesis Testing

More on Null Distributions and P-values

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)
- choose the sign to match the research question
- both hypotheses use the same value

Review

This is the general outline we will follow for Hypothesis Testing

1. Define hypotheses
2. Simulate what the parameter looks like under 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: Points 2 and 3 can be combined with the creation of a something called a 'test-statistic'

Null Distribution

The distribution of the statistics if the null hypothesis is true

- simulates what the null hypothesis looks like
- use this to compute p-values

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

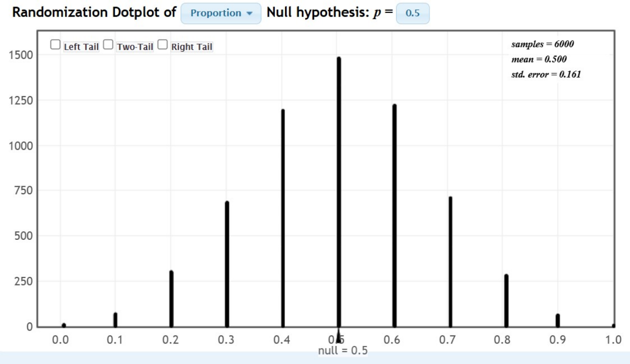

Null Distribution

What if we flipped 100 coins instead of 10

- Would getting $\hat{p} = .8$ be more common or less common?

Null Distribution

This represents 'H_0: p = 0.5', when the sample size (# of flips) is 100

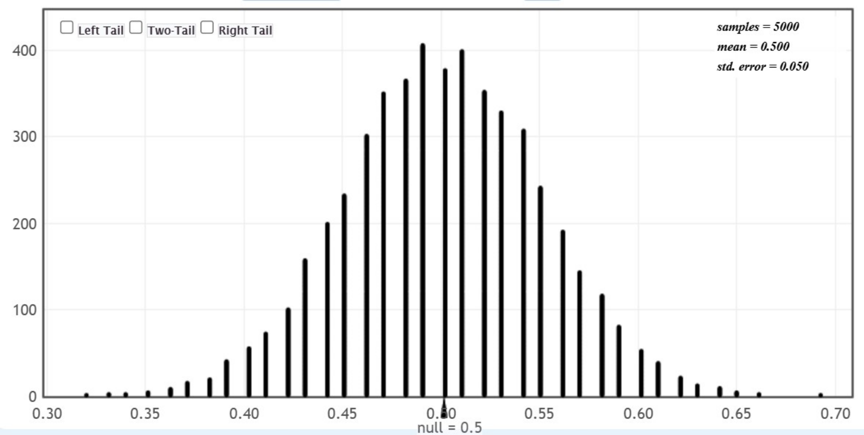

- What do we see?
- Are large or small values of $\hat{p}$ more/less common?

Null Distribution

The **null distribution** will look very similar to the sampling distribution stuff we saw before

- this time it is simulating the null hypothesis
- for means and proportions this looks like a Normal curve

These distributions looks Normal when certain conditions are met. Very similar to what we had when we were using confidence intervals.

Null Distribution – Proportion

Conditions:

- Random Sample
  - this doesn't affect our null distr. but makes sure answers are accurate
- $np_0 > 10$
- $n(1 - p_0) > 10$

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

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

- use pnorm() to get p-values with these values for mean & std. dev.

Note: we have p_0's in the distribution because we are simulating what the null hypothesis looks like

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...

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

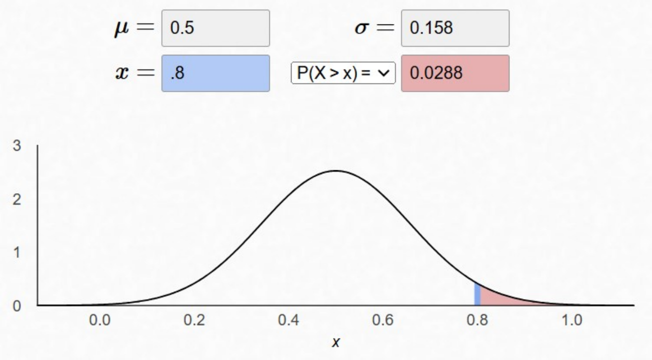

- can use Null distr. and $\hat{p}$ to get p-value

Coin Flip Example

For our scenario of 10 flips...

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

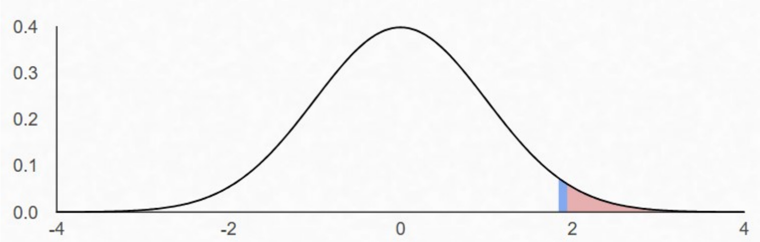

- can use N(0,1) and Test-statistic Z to get p-value

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 $Z = 1.9$

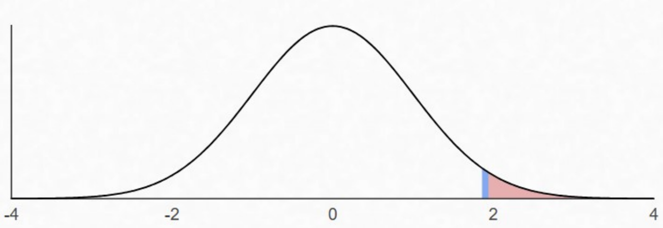


```
> pnorm(1.9, lower.tail=F)
[1] 0.02871656
```

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

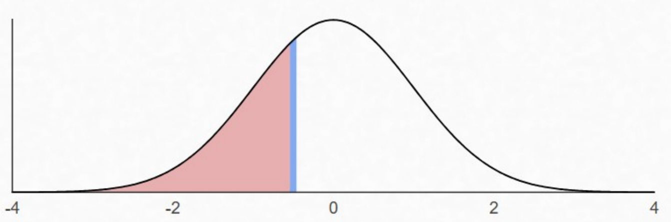


```
> pnorm(-.5)
[1] 0.3085375
```

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

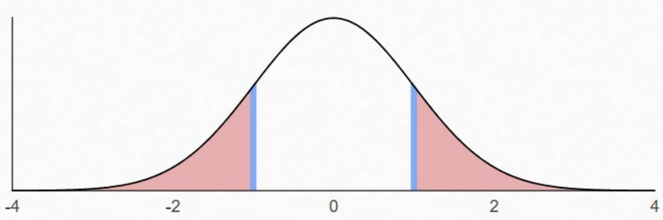


```
> 2*pnorm(1, lower.tail=F)
[1] 0.3173105
> 2*pnorm(-1)
[1] 0.3173105
```