Hypothesis Testing pt. 4

More on Strength of Evidence and Test-stats

Grinnell College

Review

Previously we introduced the idea of null distributions and test-statistics

If we were to collect many samples and compute $\overline{x}$, the null distribution refers to the distribution of statistics

$$t = \frac{\overline{x} - \mu_0}{\hat{\sigma}/\sqrt{n}}$$

when $H_0 : \mu = \mu_0$, i.e., the null hypothesis is true

Test-Statistic

Consider the pieces of a t-statistic

$$T = \frac{\overline{x} - \mu_0}{\hat{\sigma}/\sqrt{n}}$$

1. $\overline{x} - \mu_0$ indicates the distance between my observed data and my null hypothesis, though we cannot use this alone as it does not include a degree of "certainty" associated with $\overline{x}$
2. $\hat{\sigma}$ is my estimate of the population's standard deviation. When this is large, there will be more uncertainty in my estimate of $\overline{x}$
3. n represents the number of observations in my sample – the more observations I have, the more confidence I will have in my estimate

We should have a sense of how each of these components impact my *t*-statistic

Test-Statistic

$$T = \frac{\overline{x} - \mu_0}{\hat{\sigma}/\sqrt{n}}$$

We can think of the t-statistic as being a measure of evidence *against* the null hypothesis

If:

1. $\overline{x}$ is far from μ_0 and
2. Our certainty in $\overline{x}$ is high (i.e., low $\hat{\sigma}$ or large n)

then our statistic T will be larger.

A larger T statistic is less likely than a smaller one when H_0 is true

- What we consider "large" will depend on the t-distribution

t-distribution

When the null hypothesis is true,

$$t = \frac{\overline{x} - \mu_0}{\hat{\sigma}/\sqrt{n}}$$

follows a t-distribution with $n - 1$ degrees of freedom

The degrees of freedom tells us, relatively speaking, what values are considered "large"

$T = 1.69$ may be considered "large" when $df = 30$ but not when $df = 5$

t-distribution Comparison

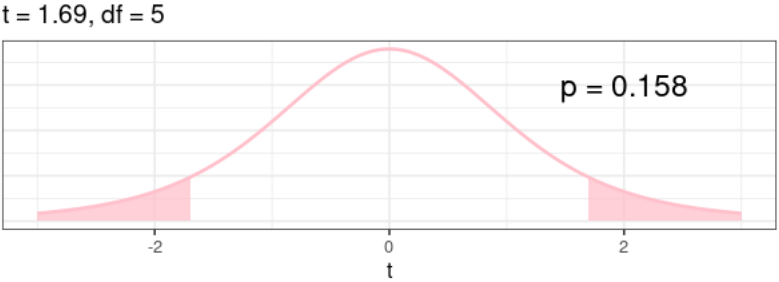

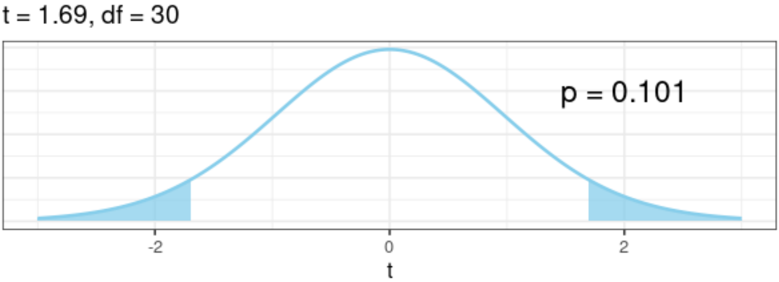

Hypothesis Testing Summary

The process goes like this:

1. Assume our null hypothesis $H_0 : \mu = \mu_0$ is true
2. Compute a T-statistic with our observed data
3. Ask: is this T-statistic "large"?
    - This will depend on the degrees of freedom
4. If we know our T statistic and we know our degrees of freedom, we can find the probability of observing our data if the null is true. This probability is our **p-value**

P-values

Hopefully by this point the utility of p-values is becoming clear

- quantify how unlikely the data is *if* H_0 is true
- take into account both the test-stat and the associated distribution

When we say 'there is (evidence) against the null hypothesis', what we are doing is quantifying how far off the data is from what we would expect.