

HYPOTHESIS TESTS SUMMARY

We start with a **null hypothesis** H_0 , which we'll assume to be true until we have evidence to the contrary. Exactly what constitutes contrary evidence is determined by our choice of **alternative hypothesis** H_1 . We **reject** H_0 in favor of H_1 if an appropriate test statistic falls in a **critical region** (or **rejection region**). If the test statistic does not fall in the critical region, then we **fail to reject** H_0 . Alternatively, we reject H_0 if the p-value of our test statistic is **less than a set significance level**.

When testing a hypothesis, there are two ways to be wrong:

- **Type I error**: reject H_0 when H_0 is actually true;
- **Type II error**: fail to reject H_0 when H_0 is actually false.

The probability of a type I error is α ; the probability of a type II error is β . The critical region is chosen so that the test statistic lands in the critical region with probability α when H_0 is true. The **p-value** (or **observed significance level**) of your data is the smallest value for α that allows you to reject H_0 with your data.

Problems in which hypotheses aren't explicitly stated should be completed in three steps:

- (1) State your null and alternative hypotheses. (Does your null hypothesis make sense as a default assumption?)
- (2) Use the data to test H_0 against H_1 , testing either at a specified significance level or giving a P -value.
- (3) State your conclusion clearly (e.g. reject H_0 at significance level 0.05 or fail to reject H_0 at significance level 0.05).

1. TEST STATISTICS

For **tests about the mean** ($H_0 : \mu = \mu_0$) test statistics are:

- $z = \frac{\bar{x} - \mu_0}{\frac{\sigma}{\sqrt{n}}}$ (known variance σ^2 , all sample sizes if the pop. is normal, otherwise just large samples)
- $t = \frac{\bar{x} - \mu_0}{\frac{s}{\sqrt{n}}}$ (samples from approximately normally distributed populations, $n-1$ degrees of freedom)
R command: `t.test(x)`

For **tests about the difference of two means** ($H_0 : \mu_1 - \mu_2 = \delta_0$) some test statistics are:

- $z = \frac{\bar{x}_1 - \bar{x}_2 - \delta_0}{\sqrt{\frac{\sigma_1^2}{n_1} + \frac{\sigma_2^2}{n_2}}}$ (known variances, all sample sizes if pops are normal, otherwise just large samples)
- $t = \frac{\bar{x}_1 - \bar{x}_2 - \delta_0}{s_p \sqrt{\frac{1}{n_1} + \frac{1}{n_2}}}$ (normally distributed populations with the same variance, $n_1 + n_2 - 2$ d.f.).
 $s_p^2 = \frac{(n_1 - 1)s_1^2 + (n_2 - 1)s_2^2}{n_1 + n_2 - 2}$
R command: `t.test(x, y, var.equal = T)`
- $t = \frac{\bar{x}_1 - \bar{x}_2 - \delta_0}{\sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}}$ (normally distributed populations with different variances, ν degrees of freedom)
 $\nu \approx \frac{\left(\frac{s_X^2}{m} + \frac{s_Y^2}{n}\right)^2}{\frac{\left(\frac{s_X^2}{m}\right)^2}{m-1} + \frac{\left(\frac{s_Y^2}{n}\right)^2}{n-1}}$
R command: `t.test(x, y)`

For **tests about a population proportion** ($H_0 : \theta = \theta_0$) we can use the sample proportion $\hat{\Theta} = X/n$ or the sample total $X = n\hat{\Theta}$ and the test statistics are:

- x (X is binomial with parameters n and θ_0)
R command : `binom.test(x, n, p=theta_0)`
- $z = \frac{\hat{\theta} - \theta_0}{\sqrt{\frac{1}{n}\theta_0(1 - \theta_0)}} = \frac{x - n\theta_0}{\sqrt{n\theta_0(1 - \theta_0)}}$ (large samples, both $n\theta_0 \geq 10$ and $n(1 - \theta_0) \geq 10$)

For **tests about the variance** ($H_0\sigma^2 = \sigma_0^2$) the test statistic is :

- $\chi^2 = \frac{(n - 1)s^2}{\sigma_0^2}$ (chi-square distribution, $n - 1$ degrees of freedom)

For **tests about the ratio of two variances** ($H_0 : \frac{\sigma_1^2}{\sigma_2^2} = 1$) the test statistic is

- $\frac{s_1^2}{s_2^2}$ (F distribution with $n_1 - 1$ and $n_2 - 1$ degrees of freedom, order matters).
R command: `var.test(x, y)`