Lecture 13: Two-Sample t-tests
Notation and Equations

Notation:

n_A, M_A, μ_A, n_B, M_B, μ_B: Statistics and parameters for samples A and B and populations A and B in an independent-samples t-test

$\sigma_{M_A - M_B}$: Standard error of the difference between sample means; standard deviation of the sampling distribution of $M_A - M_B$, according to the null hypothesis

MS: Mean square. This is a generalization of the idea of sample variance that comes up in more complicated statistical tests, as a way of estimating the population variance, σ^2.

X_A, X_B: The two scores, or samples, in a paired-samples t-test

X_{diff}: Difference score in a paired-samples t-test.

Formula for difference score in paired-samples design. When each subject (or event, etc.) has two scores on the same variable, we can compute a difference score by subtracting one from the other.

$$X_{\text{diff}} = X_A - X_B \qquad (1)$$

t statistic for paired samples. A paired-samples t-test works just like a single-sample t-test applied to the difference scores. The null hypothesis states $\mu_A = \mu_B$, which implies that the mean difference score equals zero. Therefore, we use the formula for a single-sample t, with $\mu_0 = 0$.

$$t = \frac{M_{\text{diff}}}{s_{\text{diff}}/\sqrt{n}} \qquad (2)$$

It's important to keep in mind that everything in this formula is based on the difference scores, X_{diff}. That is, M_{diff} is the mean of X_{diff}, s_{diff} is the sample standard deviation of X_{diff}, and n is the number of difference scores (*not* the total number of observations; that's $2n$).

Mean Square for independent-samples t-test. Like a single-sample t-test, the independent-samples t-test requires an estimate of the population variance, σ^2. We could estimate σ^2 using either sample by itself, i.e. using s_A or s_B, but we get a better estimate by using all of the data. The formula for sample variance from a single sample is the average of $(X - M)^2$, and when we have two samples we do essentially the same thing.

$$MS = \frac{\sum_A \left(X - M_A\right)^2 + \sum_B \left(X - M_B\right)^2}{df} \tag{3}$$

In this formula, the $\sum_A$ symbol means sum over all of Sample A, and $\sum_B$ means sum over all of Sample B. Notice that $(X - M_A)^2$ or $(X - M_B)^2$ for each individual score is an estimate of the population variance, σ^2, and all Equation 3 does is average these estimates to get one single best estimate for σ^2. The degrees of freedom for MS for the independent-samples t-test is discussed below.

Standard error of $M_A - M_B$. The standard error of any statistic is the standard deviation of its sampling distribution, which in turn is the square root of its variance. So, let's think about the variance of $M_A - M_B$. It turns out this is just the variance of M_A plus (not minus!) the variance of M_B. If you think of variance as uncertainty, then the uncertainty in $M_A - M_B$ is the sum of the uncertainties coming from each of M_A and M_B. We know the variance of each sample mean alone from the Central Limit Theorem, so it's straightforward to add them together.

$$\begin{aligned} \sigma^2_{M_A - M_B} &= \sigma^2_{M_A} + \sigma^2_{M_B} \\ &= \tfrac{\sigma^2}{n_A} + \tfrac{\sigma^2}{n_B} \\ &= \sigma^2\left(\tfrac{1}{n_A} + \tfrac{1}{n_B}\right) \end{aligned}$$

Since the standard error is the standard deviation of $M_A - M_B$, we want the square root of what we just computed.

$$\sigma_{M_A - M_B} = \sigma \cdot \sqrt{\frac{1}{n_A} + \frac{1}{n_B}}$$

Compare this to the formula for the standard error of a single sample mean. Notice that if you took away one of the $1/n$ terms, you'd be left with $\sigma/\sqrt{n}$.

Since we don't know σ, we estimate it using $\sqrt{MS}$. This gives the final formula for the (estimated) standard error of $M_A - M_B$. (Eq. 4 is the only formula you have to remember here; the rest is just to help it make sense.)

$$\sigma_{M_A - M_B} = \sqrt{MS \cdot \left(\frac{1}{n_A} + \frac{1}{n_B}\right)} \tag{4}$$

Independent-samples t statistic. The t statistic for an independent-samples t-test measures how far apart the two sample means are, relative to how far apart they could be expected to be by chance. The first part of this, the distance between the sample means, is just $M_A - M_B$. The second part, the distance that could be expected by chance, is the "typical" difference expected by the null hypothesis, i.e. under the assumption that the true difference in the population is zero. In other words, it's the standard deviation of the sampling distribution of $M_A - M_B$, which is the standard error of $M_A - M_B$.

$$t = \frac{M_A - M_B}{\sigma_{M_A - M_B}} \tag{5}$$

Equations 4 and 5 can be combined into a single formula for t.

$$t = \frac{M_A - M_B}{\sqrt{MS \cdot \left(\frac{1}{n_A} + \frac{1}{n_B}\right)}} \tag{6}$$

Degrees of freedom for independent-samples t-test. The degrees of freedom for a t statistic equals the degrees of freedom from the estimate of σ^2 in its denominator. This estimate, the MS in Equation 3, is the average of n_A numbers from Sample A and n_B numbers from Sample B. However, because the formula depends on M_A and M_B, which are both computed from the data, we "lose" two degrees of freedom.

$$\text{Independent-samples t-test: } df = n_A + n_B - 2 \tag{7}$$

Degrees of freedom in general. Any Mean Square formula is an average of numbers of the form $\left(X - \hat{X}\right)^2$, where $\hat{X}$ is some statistic like M. The question is, how many numbers are we *really* averaging? The answer isn't what it naively seems to be, and this is why we need the concept of degrees of freedom.

Take the sample variance as an example. Sample variance is the simplest version of a mean square, where we have a single sample and we compare each score to M, the sample mean.

$$s = \frac{\sum (X - M)^2}{n - 1}$$

When we first introduced s, you learned that we have to divide by $n - 1$ instead of n because M depends on X. If we substitute the formula for the mean, $M = \sum X/n$, into the formula for s, we can rewrite the formula for s so that it's a sum of only $n - 1$ squares. In other words, even though $\sum(X - M)^2$ looks like it's a sum of n squares (because there are n numbers in X), it's effectively a sum of only $n - 1$ squares.

In general, for every sample statistic that $\hat{X}$ is based on, we can use algebra to rearrange the Mean Square formula to make that statistic disappear, and the result is always that we end up with one less summand than we started with. This is why we say that we "lose" one degree of freedom for every statistic that's used to define $\hat{X}$.

$$df = (\text{number of items being added}) - (\text{number of statistics used to define } \hat{X}) \quad (8)$$

For the independent-samples t-test, $\hat{X}$ is based on M_A and M_B, so we lose 2 degrees of freedom and the final df equals $n_A + n_B - 2$.