The F-distribution

Recall that the sum of squares of n iid standard normal variables follows a $\chi^2(n)$ distribution

If we were to wish to compare two sets of iid standard normal variables, sums are less appropriate: the more copies you sum up, the larger the variance $(2n)$ and the mean (n).

Instead, the quantities to compare would be the **means** of the sets of squared std normal variables: in other words, if $W_1 \sim \chi^2(n_1)$ and $W_2 \sim \chi^2(n_2)$, then we would want to compare W_1/n_1 to W_2/n_2.

Definition Suppose $W_1 \sim \chi^2(n_1)$ and $W_2 \sim \chi^2(n_2)$. Then the ratio

$$F = \frac{W_1/n_1}{W_2/n_2}$$

follows the F distribution with n_1 numerator degrees of freedom n_2 denominator degrees of freedom.

In R, this is handled by pf, qf, rf and df.

One-sample test for variance

Let $Y_1, \ldots, Y_n \sim N(\mu, \sigma^2)$ iid with unknown mean & variance.

We know $X = \dfrac{(n-1)S^2}{\sigma^2} \sim \chi^2(n-1)$ and we can use this to create statistical tests for population variance.

These tests we will want to follow a one-point null hypothesis $H_0 : \sigma^2 = \sigma_0^2$ and the usual upper/lower/twotailed alternative hypotheses.

If S^2 comes out remarkably large, this supports $\sigma^2 > \sigma_0^2$; if S^2 is remarkably small, then that supports $\sigma^2 < \sigma_0^2$. For a twosided RR, we run into the same issues as we had for the CIs: the χ^2 distribution is not symmetric, so picking equal probability weights in the two tails does not produce a largest possible rejection region. It is however one much easier to calculate with.

This gives us a test:

$H_0 : \sigma^2 = \sigma_0^2$ $\quad$ $H_A : \sigma^2 \overset{>}{\underset{<}{\neq}} \sigma_0^2$

Test statistic: $(n-1)S^2/\sigma_0^2 =: X$ $\quad$ RR: $\{X > X_{1-\alpha}\}$, $\{X > X_{1-\alpha/2}\} \cup \{X < X_{\alpha/2}\}$, $\{X < X_\alpha\}$

where $X_\alpha = F^{-1}_{\chi^2(n-1)}(\alpha)$.

This test is <u>not</u> robust to non-normal data.

Two-sample test of variances

Assume $X_1, \ldots, X_{n_X} \overset{iid}{\sim} N(\mu_X, \sigma_X^2)$ and $Y_1, \ldots, Y_{n_Y} \overset{iid}{\sim} N(\mu_Y, \sigma_Y^2)$

Recall that $W = \dfrac{(n-1)S^2}{\sigma^2} \sim \chi^2(n-1)$. It follows that

$W/(n-1) = S^2/\sigma^2$.

Applied to our Xs and Ys, we get that

$$F = \frac{W_X/(n_X-1)}{W_Y/(n_Y-1)} = \frac{S_X^2/\sigma_X^2}{S_Y^2/\sigma_Y^2} \sim F_{n_Y-1}^{n_X-1}$$

From this known distribution we can construct a test for $H_0: \sigma_X^2 = \sigma_Y^2$ vs. lower/upper/two-tailed alternatives.
The construction is familiar by now:

$$H_0: \sigma_X^2 = \sigma_Y^2 \qquad H_A: \sigma_X^2 \overset{>}{\underset{<}{\neq}} \sigma_Y^2$$

Under H_0, the factors σ_X^2 and σ_Y^2 in F cancel, leaving our test statistic: $F = S_X^2/S_Y^2$.

$$\text{Rejection Region}: \begin{array}{l} \{F > F_{n_Y-1,1-\alpha}^{n_X-1}\} \\ \{F > F_{n_Y-1,1-\alpha/2}^{n_X-1}\} \cup \{F < F_{n_Y-1,\alpha/2}^{n_X-1}\} \\ \{F < F_{n_Y-1,\alpha}^{n_X-1}\} \end{array}$$

where $F_{m,\alpha}^{n} = F_{F_m^n}^{-1}(\alpha)$.

cont...

The F distribution has some nice symmetries:
if $F \sim F^n_m$, then $\frac{1}{F} \sim F^m_n$.
With printed tables, this symmetry is important to cut down on the tables required and can both be used to get a lower threshold from an upper one, and to force $n > m$ in the lookup tables.

Nowadays, we have computers.

The F test is not robust to non-normality.