Andrew J. Spieker, PhD
BIOS 6342 - Contemporary Statistical Inference
Lab #7: The likelihood ratio test

Two-sample likelihood ratio tests: For most (but certainly not all) of this course, we've focused on one-sample, one-parameter problems. Part of the rationale is that such problems are much easier to see in pictures. They therefore do a marvelous job at building intuition. In class, we did (or will do) a two-sample example based on the Beta family. For this lab, we will focus on two-sample normal problems. We will also build on a theme from last week (and the subject of a future unit): model misspecification.

Setup: Suppose $X_1, \ldots, X_n \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mu, \sigma^2)$ and $Y_1, \ldots, Y_m \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\nu, \tau^2)$.

<u>Exercise</u>: Suppose first that it is known that $\mu = \nu = 0$. Derive a likelihood ratio test statistic to test $H_0 : \sigma^2 = \tau^2$ vs. $H_1 : \sigma^2 \neq \tau^2$. Provide the two useful asymptotic statements.

<u>Exercise</u>: Characterize the behavior of this test by simulation when the null is true.

<u>Exercise</u>: Let $n = m = 5$ and consider data distributed as $X_1, \ldots, X_n \overset{\text{i.i.d.}}{\sim}$ Exponential($\lambda = 1$) and $Y_1, \ldots, Y_m \overset{\text{i.i.d.}}{\sim}$ Exponential($\lambda = 1$). Evaluate (by simulation) how the likelihood ratio test statistic fares under the null.

<u>Exercise</u>: Does your answer to the above exercise change if you consider $n = m = 500$?

<u>Exercise</u>: Now consider testing the *strong* null hypothesis:

$$H_0 : \mu = \nu \text{ and } \sigma^2 = \tau^2 \text{ vs. } H_1 : \text{“not } H_0\text{.”}$$

Derive the likelihood ratio test statistic for this problem and provide the two useful asymptotic statements.

<u>Exercise</u>: Characterize the behavior (by simulation) of the test you just developed when the null is true.

<u>Exercise</u>: The two-sample Kolmogorov–Smirnov test is a nonparametric test of equality of two distributions based on the following statistic:

$$D_{n,m} = \sup_t |\widehat{F}_{X,n}(t) - \widehat{F}_{Y,m}(t)|,$$

where, for instance, $\widehat{F}_{X,n}(\cdot)$ is the empirical CDF for $X_1, \ldots, X_n$. Getting too deep into the behavior of this test statistic would require some background on Brownian bridges; instead, I will just tell you that it can be implemented in R using the function `ks.test()`. Conduct a simulation to evaluate the power of this test relative to the likelihood ratio test of the previous problem when the model is correct.