Homework 6 R Exercise

R has simple functions to conduct most of the tests we have seen in the text, though you would need to install the package **EnvStats** to be able to test a population variance (or population standard deviation) from a single sample. To do this, you select **Tools**, then **Install Packages** from the menu bar. Type **EnvStats** in the Packages window and then ignore all the output to the Console as the package is installed.

For our R session, we will load the **EnvStats** package just for practice, then step through a t-test on a mean μ and a z-test on a binomial proportion p. We also discussed the functions needed to conduct a chi-squared test on a population variance σ^2, and an F test of a set of population variances σ_1^2 and σ_2^2.

For the t-test, we generate a random sample of size 10 from a normal distribution with mean $\mu = 10$, and then test $H_0 : \mu = 9$ against a one-sided alternative $H_A : \mu > 9$. Note that results will differ from sample to sample; what was your p-value? Is there strong evidence against H_0?

```
#We will not actually use EnvStats in this session, but
#it is good practice to install and then load a library
library(EnvStats)
T_Sample=rnorm(10,10,1)
t.test(T_Sample,mu=9,alternative="greater")
```

If we were running a two-sided test, we could either specify **alternative="two-sided"**, or leave the argument out altogether, since a two-sided alternative is the default choice. Note the unusual confidence interval, which is generated because we ran a one-sided test; if you would like to recover our usual two-sided confidence interval, you could run a two-sided test.

To run a test on a binomial proportion, we can use **prop.test**, though the test statistic is a chi-squared statistic that is the square of a Z-test statistic using a different denominator from our Z-test statistic $\left(\sqrt{\hat{p}(1-\hat{p})/n}\right.$ rather than $\left.\sqrt{p_0(1-p_0)/n}\right)$, so the p-value will be a little different from the p-value in the text. We actually have a couple different choices for the way to represent the outcome, but will study only one, in which we provide x as the first argument and n as the second argument. From Example 8.10 in the text (pp. 402-403), we have $x = 10$, $n = 300$, and are testing $H_0 : p = 0.05$ against $H_A : p < 0.05$. Other than the test statistic, do the results from R roughly match what you find in the book?

```
prop.test(x=10,n=300,p=0.05,alternative="less")
```

Tests on a population variance, using **varTest** in **EnvStats**, and a comparison of two population variances, using **var.test**, are conducted similarly; you can consult help for details, and perhaps use them to check your homework.