Key for Long Problem Set 3

1. The data file "BloodPressure.RData" contains results from a clinical study of the efficacy of calcium supplements as a treatment for blood pressure in males. In this study 21 subjects received either a calcium supplement ("calcium.yes") or a placebo that contained no calcium ("calcium.no") for 12 weeks. Each subject's blood pressure was measured before and after the treatment period and the difference recorded (a positive value represents a decrease in blood pressure). Determine at $\alpha = 0.05$ whether there is any evidence that calcium lowers blood pressure.

Answer. This problem calls for a one-tailed t-test of two mean values using unpaired data. Before we complete the t-test, we first use a two-tailed F-test to determine whether there is evidence to suggest that we cannot pool the standard deviations. The null hypothesis and the alternative hypothesis are

$H_0 : s^2_{no} = s^2_{yes}$

$H_A : s^2_{no} \neq s^2_{yes}$

```
var.test(calcium.no, calcium.yes, ratio = 1, conf.level = 0.95)
```

```
##
##  F test to compare two variances
##
## data:  calcium.no and calcium.yes
## F = 0.45547, num df = 10, denom df = 9, p-value = 0.2365
## alternative hypothesis: true ratio of variances is not equal to 1
## 95 percent confidence interval:
##  0.1149056 1.7212056
## sample estimates:
## ratio of variances
##          0.4554704
```

Both the p-value and the 95% confidence interval (which includes the expected ratio of 1) lead us to accept the null hypothesis; that is, we find no evidence at $\alpha = 0.05$ to suggest there is a difference in the precisions of the two data sets. We will, therefore, pool the standard deviations.

To compare the means we use a one-tailed t-test as we are interested only in whether calcium lowered the blood pressure The null hypothesis and the alternative hypothesis for the t-test are

$H_0 : \bar{X}_{yes} = \bar{X}_{no}$

$H_A : \bar{X}_{yes} > \bar{X}_{no}$

Note that because a decrease in blood pressure is reported as a positive value, the alternative hypothesis is greater than). The data are not paired as there is no overlap of patients between the two data sets.

```
t.test(calcium.yes, calcium.no, mu=0, alternative="greater",
       conf.level=0.95, var.equal=TRUE)
```

```
##
##  Two Sample t-test
##
## data:  calcium.yes and calcium.no
## t = 1.6341, df = 19, p-value = 0.05935
## alternative hypothesis: true difference in means is greater than 0
## 95 percent confidence interval:
##  -0.3066129        Inf
## sample estimates:
##  mean of x  mean of y
```

```
##  5.0000000 -0.2727273
```

Both the p-value and the 95% confidence interval lead us to retain the null hypothesis; that is, at $\alpha = 0.05$ the difference between the two mean values is not sufficiently large to suggest that it is not the result of uncertainty in the measurements. Nevertheless, with an $alpha$ of just less than 0.06, we might wish to explore further the potential affect of calcium on blood pressure, perhaps by using a larger sample of patients or by considering if there are differences between patients that need more careful control.

2. The file "SpeedLight.RData" contains two sets of results from Michelson's 1879 determination of the speed of light ("sol.one" and "sol.two"). To make it easier to enter the data into a file, all values are offset by subtracting 299,000 km/sec from the measured value. Determine at $\alpha = 0.05$ if there is a significant difference between the results of these two experiments.

Answer. This problem calls for a two-tailed t-test of two mean values with unpaired data. Before we complete the t-test, we first use a two-tailed F-test to determine whether there is evidence to suggest we cannot pool the standard deviations. The null hypothesis and the alternative hypothesis are

$H_0 : s^2_{one} = s^2_{two}$

$H_A : s^2_{one} \neq s^2_{two}$

```
var.test(sol.one, sol.two, alternative = "two.sided", conf.level = 0.95)
```

```
##
##  F test to compare two variances
##
## data:  sol.one and sol.two
## F = 0.93758, num df = 19, denom df = 19, p-value = 0.8897
## alternative hypothesis: true ratio of variances is not equal to 1
## 95 percent confidence interval:
##  0.3711061 2.3687531
## sample estimates:
## ratio of variances
##          0.9375813
```

Both the p-value and the 95% confidence interval lead us to accept the null hypothesis; that is, we find no evidence at $\alpha = 0.05$ to suggest there is a difference in the precisions of the two data sets. We will, therefore, pool the standard deviations. The null hypothesis and the alternative hypothesis for the t-test are

$H_0 : \bar{X}_{one} = \bar{X}_{two}$

$H_A : \bar{X}_{one} \neq \bar{X}_{two}$

```
t.test(sol.one, sol.two, mu = 0, alternative = "two.sided",
       conf.level = 0.95, var.equal = TRUE)
```

```
##
##  Two Sample t-test
##
## data:  sol.one and sol.two
## t = -0.829, df = 38, p-value = 0.4123
## alternative hypothesis: true difference in means is not equal to 0
## 95 percent confidence interval:
##  -55.07133  23.07133
## sample estimates:
## mean of x mean of y
##     820.5     836.5
```

Both the p-value and the 95% confidence interval lead us to retain the null hypothesis; that is, at $\alpha = 0.05$ the difference between the two mean values is not sufficiently large to suggest that the difference between the means is not the result of uncertainty in the measurements.

3. The data at this link (http://www.rsc.org/images/CO2_methods_tcm18-57755.txt) reports results for the determination of CO_2 by six different methods. The data itself uses Na_2CO_3 as a reference standard; presumably, a portion of the standard was treated to release the CO_2, which subsequently was determined and reported as % w/w CO_2 in the sample. Using the data for the gravimetric method, determine at $\alpha = 0.05$ if there is any evidence for a determinate error in the analysis. Note that the authors report a known value of 41.518% w/w CO_2 for the reference standard.

Answer. This problem calls for a two-way t-test of the mean value for the experimental results to a theoretical value of 41.518%. The null hypothesis and the alternative hypothesis for the t-test are

$H_0 : \bar{X} = \mu$

$H_A : \bar{X} \neq \mu$

```
grav = c(41.41, 41.62, 41.48, 41.44, 41.50, 41.51, 41.43, 41.51, 41.59)
mean(grav)
```

```
## [1] 41.49889
```

```
sd(grav)
```

```
## [1] 0.07043516
```

```
t.test(grav, mu = 41.518, alternative = "two.sided", conf.level = 0.95)
```

```
##
##  One Sample t-test
##
## data:  grav
## t = -0.81399, df = 8, p-value = 0.4392
## alternative hypothesis: true mean is not equal to 41.518
## 95 percent confidence interval:
##  41.44475 41.55303
## sample estimates:
## mean of x
##  41.49889
```

Both the p-value and the 95% confidence interval lead us to retain the null hypothesis; that is, at $\alpha = 0.05$ the difference between the experimental mean of 41.499 and the theoretical mean of 41.518 is not sufficiently large to suggest that it cannot be explained by uncertainty in the measurements.

4. The file "Clouds.RData" contains results for the amount of rainfall recorded from 26 clouds, half of which were randomly seeded with AgI (the units are in acre-feet, or a volume equivalent to the feet of rain covering one acre of ground). For each data set, "seeded" and "unseeded", make a convincing argument that the data are not normally distributed and then, using an appropriate statistical test, determine at $\alpha = 0.05$ if seeding clouds has any effect. By the way, the writer Kurt Vonnegut's brother, Bernard, was an atmospheric scientest at General Electric who discovered that AgI could be used to seed clouds.

Answer. There are several ways to evaluate the two data sets, but perhaps the simplest is to exam qqnorm plots, which, as shown in Figure 1, suggest that the data are strongly skewed to the right; that is, the data in both cases tails strongly toward larger values.

```
old.par = par(mfrow = c(1, 2))
qqnorm(unseeded, main = "unseeded", pch = 19, cex = 0.5)
qqline(unseeded)
qqnorm(seeded, main = "seeded", pch = 19, cex = 0.5)
qqline(seeded)
```

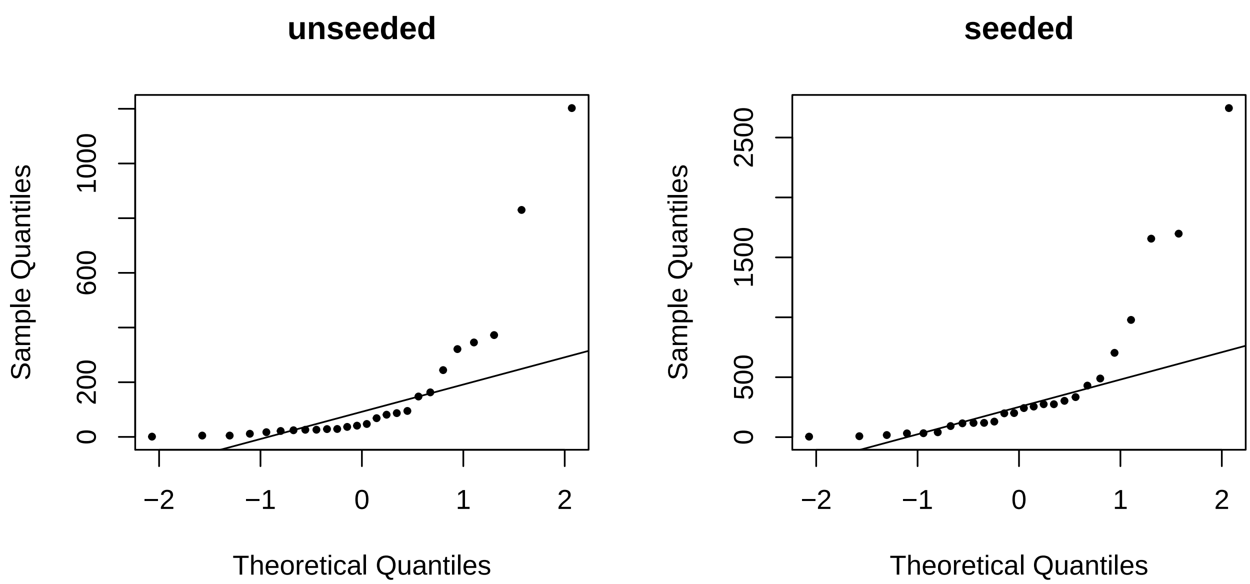

Figure 1: qqnorm plots for cloud data

```r
par(old.par)
```

Because the data sets clearly are not normally distributed, we cannot use the t-test as it assumes a normal distribution; instead, we will use the Wilcoxon Rank Sum test to compare the mean values for the two data sets using the following null hypothesis and alternative hypothesis

$H_0 : \bar{X}_{seeded} = \bar{X}_{unseeded}$

$H_A : \bar{X}_{seeded} \neq \bar{X}_{unseeded}$

```r
wilcox.test(seeded, unseeded, alternative = "two.sided", conf.level = 0.95)
```

```
## Warning in wilcox.test.default(seeded, unseeded, alternative =
## "two.sided", : cannot compute exact p-value with ties
```

```
##
##  Wilcoxon rank sum test with continuity correction
##
## data:  seeded and unseeded
## W = 473, p-value = 0.01383
## alternative hypothesis: true location shift is not equal to 0
```

With a p-value less than $\alpha = 0.05$, we have evidence that the difference between the two mean values is not explained by random error alone; thus, we have reason to believe there is a systematic difference between the two mean values. The warning message simply indicates that R's version of this test cannot give an exact p-value when two or more data points have the same value; this is a minor concern, but not a major concern. Note: a one-sided test to see if seeding the clouds increases rainfall is an acceptable choice here as well.

5. Sometimes it is possible to transform a strongly right-skewed distribution by using a log function. In R you can take the log of an object using the command log10(object). Transform the data from Problem 4 and show that the data are now normally distributed. Repeat your significance test using a test appropriate for normally distributed data and compare this result to your results from Problem 4.

Answer. The qqnorm plots for log10(seeded) and for log10(unseeded) in Figure 2 suggest that the transformed data sets closely approximate a normal distribution, although the transformation is more convincing for the unseeded clouds than for the seeded clouds.

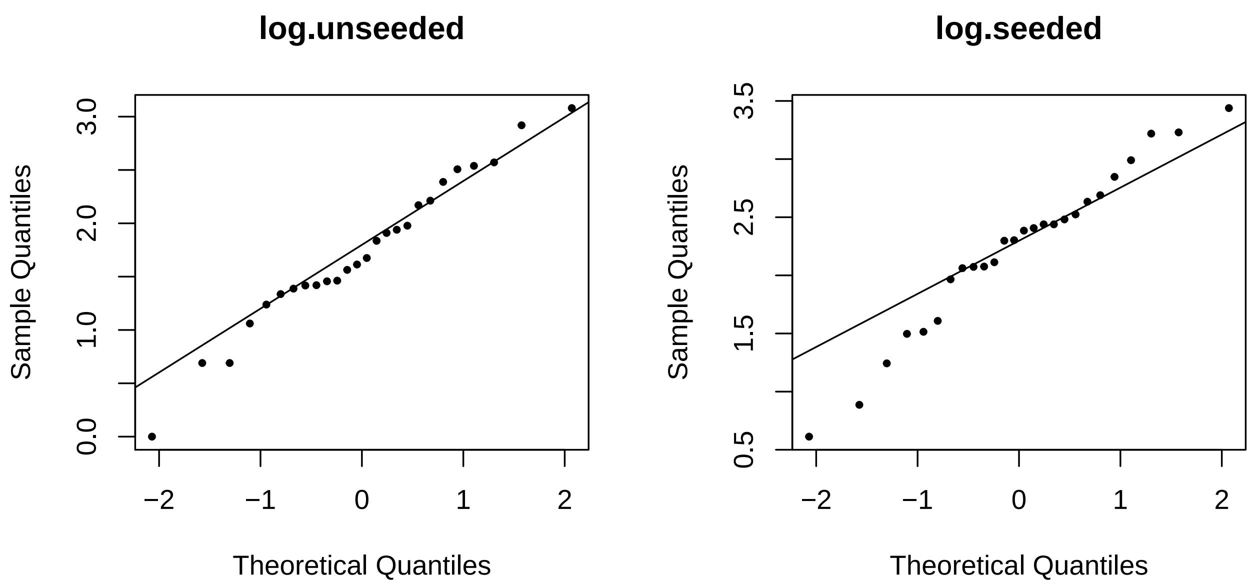

Figure 2: qqnorm plots for transformed cloud data

```
log.seeded = log10(seeded)
log.unseeded = log10(unseeded)
old.par = par(mfrow = c(1, 2))
qqnorm(log.unseeded, main = "log.unseeded", pch = 19, cex = 0.5)
qqline(log.unseeded)
qqnorm(log.seeded, main = "log.seeded", pch = 19, cex = 0.5)
qqline(log.seeded)
```

```
par(old.par)
```

If we take the transformed data as being normally distributed, then we can compare their respective means using a two-tailed t-test as unpaired data. Before we complete the t-test, we first use a two-tailed F-test to determine whether there is evidence to suggest that we cannot pool the standard deviations. The null hypothesis and the alternative hypothesis are

$$H_0 : s^2_{log.seeded} = s^2_{log.unseeded}$$

$$H_A : s^2_{log.seeded} \neq s^2_{log.unseeded}$$

```
var.test(log.seeded, log.unseeded, alternative = "two.sided", conf.level = 0.95)
```

```
##
##  F test to compare two variances
##
## data:  log.seeded and log.unseeded
## F = 0.9491, num df = 25, denom df = 25, p-value = 0.8971
## alternative hypothesis: true ratio of variances is not equal to 1
## 95 percent confidence interval:
##  0.4255461 2.1167714
## sample estimates:
## ratio of variances
##          0.9490963
```

Both the p-value and the 95% confidence interval lead us to accept the null hypothesis; that is, there is no evidence at $\alpha = 0.05$ to suggest there is a difference in the precisions of the two data sets. We will, therefore, pool the standard deviations. The null hypothesis and the alternative hypothesis for the t-test are

$$H_0 : \bar{X}_{log.seeded} = \bar{X}_{log.unseeded}$$

$$H_A : \bar{X}_{log.seeded} \neq \bar{X}_{log.unseeded}$$

```
t.test(log.seeded, log.unseeded, alternative="two.sided",
       var.equal=TRUE, conf.level = 0.95)
```

```
##
##  Two Sample t-test
##
## data:  log.seeded and log.unseeded
## t = 2.5444, df = 50, p-value = 0.01408
## alternative hypothesis: true difference in means is not equal to 0
## 95 percent confidence interval:
##  0.1046064 0.8888693
## sample estimates:
## mean of x mean of y
##  2.229749  1.733011
```

Both the p-value and the 95% confidence interval lead us to reject the null hypothesis; that is, there is evidence at $\alpha = 0.05$ to suggest that the difference in the mean values is not explained by random error only; thus, we have reason to believe there is a systematic difference between the two mean values. Note that our analysis of the transformed data gives results consistent with the analysis of the raw data in the previous problem; this is comforting as we should not expect that transforming the data will lead to starkly different conclusions.