Example 1: Rust inhibitors. Four brands of rust inhibitor were tested (A, B, C, D). Ten experimental units were assigned at random to each brand. Sample size is 40, number of replicates is 10. Response Y is a coded value for which higher values indicate less rust.

This is an experimental study with one categorical predictor variable. The correct model to use is the one-way ANOVA model; the single factor has four levels.

Read the dataset into R, rename variables, and make the group variable a `factor` object:

Omit from project write-up

```
> rust <- read.table("CH17TA02.txt", header=FALSE)
> str(rust)
'data.frame':        40 obs. of  3 variables:
 $ V1: num  43.9 39 46.7 43.8 44.2 47.7 43.6 38.9 43.6 40
 $ V2: int  1 1 1 1 1 1 1 1 1 1 ...     group variable
 $ V3: int  1 2 3 4 5 6 7 8 9 10 ...    exp. unit id
> names(rust) <- c("Y","brand","obs")
> rust$brand <- factor(rust$brand,
+ labels=c("A","B","C","D"))
```

This is Draft 1 of Q#1 for Project 1 (Example)

Let's improve the previous page, to make a good ~~model~~ example for you to imitate in Q#1 of Project 1.

The first step is to read the dataset into R, rename variables, and make the group variable a `factor` object in R. We then print out the first observation in each of the four groups as a partial check that we have read the data in correctly.

[Here, show the *clean R code* as you ran it, with the greater-than signs.]
[Let's edit the R code on p. 95, to complete and clean it up. Your paper should contain only *clean code*.]

```
> rust[c(1, 11, 21, 31), ] # prints rows 1, 11, 21, 31 of rust
```

The above four observations matched the corresponding data values in the external dataset in the file `CH17TA02.txt`.

Q#2A asks for summary statistics, and for them to be displayed in a chart.

mean, sd, sample size by group

Next we obtain the usual summary statistics separately by group, and then display the results in a chart. (R output is not shown.)
[R code can use the function `aggregate()`, which is illustrated in `Lec1.r` for computing the means and sds of the Kenton data. Notes: (1) You need to compute the group sample sizes also. (2) Use at most two significant decimal places]
[Show the *clean R code.*]

Rough idea about "aggregate" —

> cell.sizes <- aggregate (rust$r, $\frac{\text{designate}}{\text{group variable}}$, FUN = length)

Table 1. Summary statistics for Rust Inhibitor Brands

Brand	Mean	SD	n
A	43.14	3.00	10
B	89.44	2.22	10
C	67.95	2.17	10
D	40.47	2.44	10

3.0
2.2
2.2
2.4

43
89
68

better to use rounding to 2 signif digits

ALWAYS PLOT YOUR DATA

Dotplots of rust inhibition scores for each of the four brands:

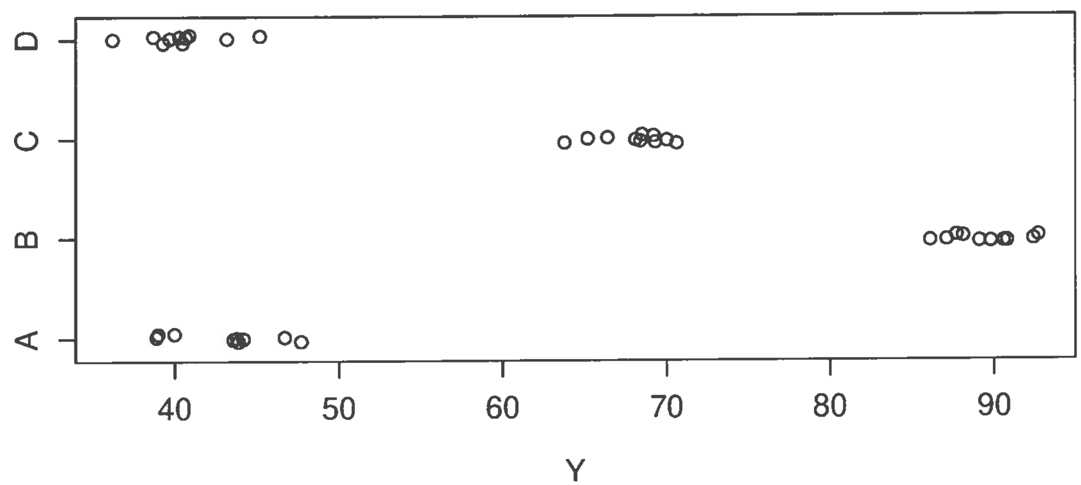

Do the means appear different?

Are the variances within groups roughly equal? (informal check of equal-variances assumption)

Ex. Rust data Here's a model of how to answer Q#s 2BC in Project 1:

Next we show R code for the plot of factor level means and the comparative boxplots; the resulting graphs are below the code.
[Use the built-in `plot.design()` and `plot()` functions in R, as shown for the Rust data in `Lec1.r`. Plot needs to be properly sized, and should have a title.]

From the summary statistics and the plot of factor level means (not shown here for Rust example; do show the plot of factor level means in Project 1), rust inhibitors A and D have similar mean effectiveness . Rust inhibitors B and C both do considerably better than A and D; the aligned dot plots show that these two distributions do not even have any overlap with the distributions of A and D scores. Rust inhibitor B appears to be the best of all with mean score of 89.

We also note that from the dotplots, the assumptions of normality and constant variance appear reasonably well-satisfied, since there are no outliers.

Here's a model of how to answer Q#s 3AC in Project 1:

Next we will carry out the F test to see if the observed differences among the means of the four groups extend to the population.

Below is the R code for fitting the one-way ANOVA model and for obtaining the ANOVA table; then the ANOVA table is shown as R output:

```
> rust.aov <- aov(Y~brand, data=rust)
> anova(rust.aov)
Analysis of Variance Table

Response: Y
          Df Sum Sq Mean Sq F value    Pr(>F)
brand      3  15954  5317.8  866.12 < 2.2e-16 ***
Residuals 36    221     6.1
```

Example, to imitate for Project 1 Q#3C

Let μ_A be the population mean rust inhibition score for the population of rusty objects (or: experimental units), if all units received an application of Rust Inhibitor A; define μ_B, μ_C, and μ_D similarly.

We are testing the null hypothesis $H_0 : \mu_A = \mu_B = \mu_C = \mu_D$, versus the alternative H_a : For at least one pair of rust inhibitors, the two population mean rust inhibition scores are not equal.

The value of the test statistic from the ANOVA table is $F_{\text{obs}} = 866.12$ with a P-value of $P = 2.2 \times 10^{-16}$. With such a huge test statistic and tiny P, we reject H_0 [maybe emphasize the strength of this conclusion a bit more!].

We conclude that for at least two of the inhibitors, their population mean rust inhibition values do differ. We next need to examine which means differ from each other, and by how much.

based on the null distribution, which is $F_{3,36}$

Let μ_A be the population mean rust inhibition score for the population of rusty objects (or: experimental units), if all units received an application of Rust Inhibitor A; define μ_B, μ_C, and μ_D similarly.

We are testing the null hypothesis $H_0 : \mu_A = \mu_B = \mu_C = \mu_D$, versus the alternative H_a : For at least one pair of means, the two means are not equal.

The value of the test statistic from the ANOVA table is $F_{\text{obs}} = 866.12$ with a P-value of $P = 2.2 \times 10^{-16}$. With such a huge test statistic and tiny P, we reject H_0 [maybe emphasize the strength of this conclusion a bit more!].

We conclude that for at least two of the inhibitors, their population mean effectiveness values do differ. We next need to examine which means differ from each other, and by how much. Below is the R code for obtaining the plot of results from Tukey's studentized range method of all pairwise comparisons, followed by the graph, which shows the six pairwise simultaneous 95% confidence intervals.