

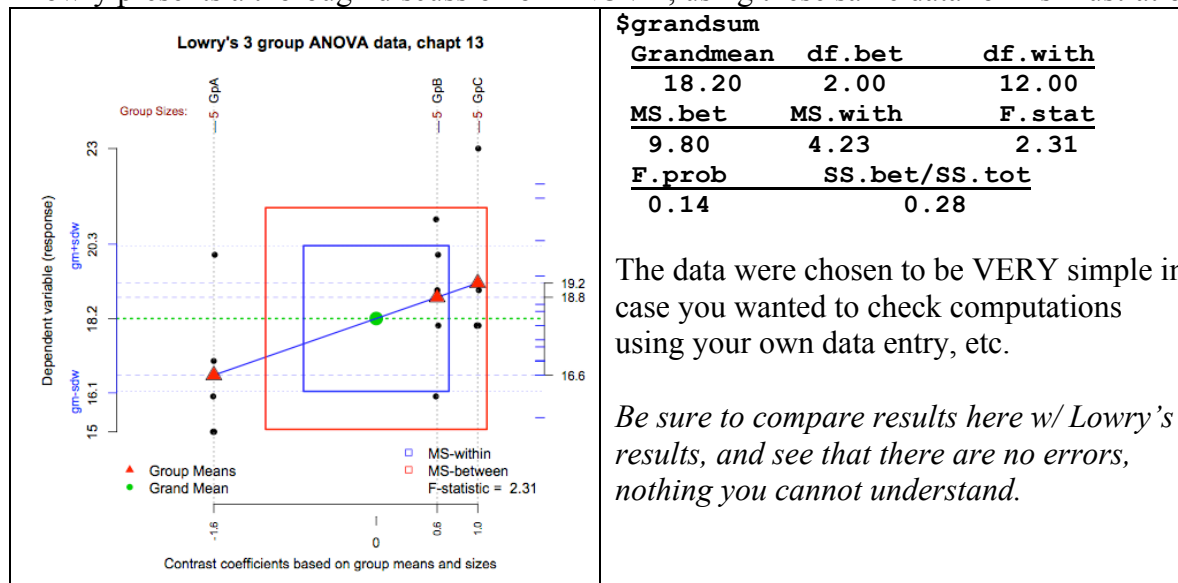
## ONE WAY ANOVA, and some notes re: contrasts:

One way *anova* can be seen as a straight-forward generalization of the two independent sample comparison to accommodate a comparison of several groups. This is often taught as a hypothesis testing method. For J groups, with a single continuous response variate  $y$  for the groups, the *null hypothesis* is that means of the J populations are equal:

$H_0: \mu_j = \mu_{j^*}$  for all  $j \neq j^*$ , versus the alternative,

$H_1: \mu_j \neq \mu_{j^*}$  for at least one  $j, j^*$  pair.

First, a quick example, I used data from <http://faculty.vassar.edu/lowry/ch13pt1.html> because Lowry presents a thorough discussion of ANOVA, using these same data for his illustrations:



The data were chosen to be VERY simple in case you wanted to check computations using your own data entry, etc.

*Be sure to compare results here w/ Lowry's results, and see that there are no errors, nothing you cannot understand.*

The null hypothesis in the one-way set up is an *omnibus* hypothesis. In other words, one does not prespecify anything as to what groups are to be compared, or what differences among (between) groups are to be tested; we only ask if there is evidence in sample data to reject the hypothesis that all means of the set of J populations are equal. In effect, this means that the analyst asks whether the estimate of the variance (sigma squared =  $\sigma^2$ , which is assumed to be the same in all J populations), derived from the variation among the J sample means, is 'notably' larger than the estimate of the same variance derived from variation within groups. The numerator and denominator of the statistic  $F = MS(\text{between})/MS(\text{within})$  that each ratio is of the form  $SS/df$ . **Some theoretical considerations:** Note that the variance of J sample means ( $\text{var}(\text{means})$ ) can be seen as estimating  $\sigma^2/n$  when the population means are the same and n has same value in each group,  $n_j = n$  for all j. This follows from basic facts – & your knowledge? – about the sampling distribution of the mean. Thus, the **mean square between** groups is of the form  $n \cdot \text{var}(\text{means})$ . The denominator of F on the other hand is computed as the simple average of the variances **within** all J groups (because all n's are equal). So, *under a true NULL hypothesis*, the numerator and denominator of the F statistic estimate the same quantity. Note that the variance of J groups means would show ONLY simple random sampling variation IF populations really had the same means (*i.e.*, null  $H_0$  true). But when *all population means are NOT the same* (*i.e.*,  $H_1$  is true) then the variance of the sample means estimates a quantity LARGER than  $\sigma^2/n$ . Hence one should **expect** a systematically larger variance estimate between groups (*i.e.*, between means of the J groups) for the numerator of F when the ALTERNATIVE hypothesis is true (when at least one population mean differs from the J-1 others).

In general, Expected value of  $MS(\text{betw}) = \sigma^2 + n \cdot \sigma^2(\text{pop. means})$ ; but for the denominator, Expected value of  $MS(\text{within}) = \sigma^2$ , regardless of whether the null or alternative hypothesis is true. It follows that when the sample F exceeds some *theoretical or tabular*  $F(df_{\text{num}}/df_{\text{den}})$ , we reject

$H_0$  in favor of the alternative  $H_1$ . All this is useful to understand, but in practice, it is often better NOT do ONE WAY ANOVA, for several reasons that we should discuss.

*One of the problems w/ one-way ANOVA when used for inference is that the omnibus hypothesis merely speaks to generic differences in population means; if one rejects the  $H_0$ , then post-hoc (multiple) comparisons are often advised to learn which means may be statistically different from one another. Planned comparisons are far better in principle, since they ask that the investigator think carefully in advance about which comparisons would be most interesting or informative. For example, suppose we had set Group C as a control group, where A and B were novel treatments to be compared first with one another, and then, as a combined group, w/ C.*

*We can make this pair of planned comparisons using contrasts, using the following matrix: `contr.helmert(3)` [ see the help, i.e. documentation]*

A -1 -1 ← Note that column 1 compares A and B; column 2 compares (A,B) with C.

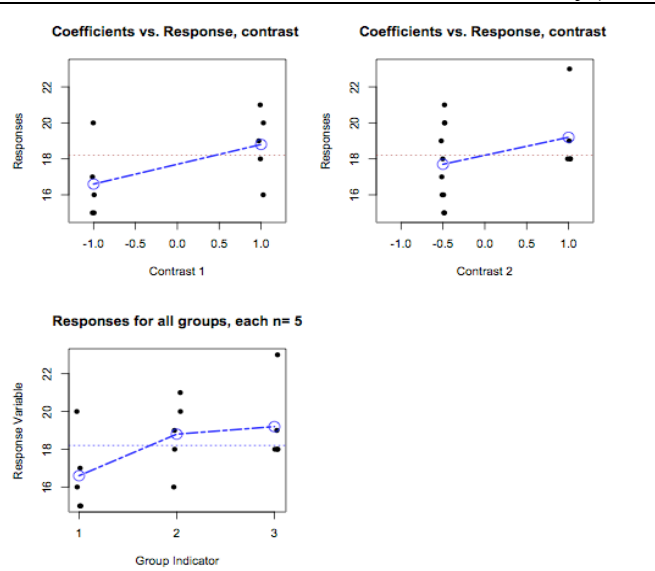
B 1 -1 This general idea of *superimposing contrasts on group comparisons* works

C 0 2 best, and most simply (especially for interpretations) when group sizes are equal.

I shall use the function `granova.contr` to carry out the analysis; I also use the `stack` function again since it facilitates putting all three columns of the Lowry3gp matrix in a column vector. The same ideas work quite generally. (I needed to set `Lowry3gp=data.frame(Lowry3gp)` initially.)

```
>granova.contr(stack(Lowry3gp)[,1],
contr.helmert(3))
$summary.lm
Call: lm(formula = resp ~ contrast)
Coefficients:
      Estimate Std. Error t value Pr(>|t|)
(Intercept) 18.200    0.531   34.26  ~0.0 ***
contrast1    1.100    0.651    1.69   0.12
contrast2    1.000    0.751    1.33   0.21
-----
Residual standard error: 2.06 on 12
degrees of freedom.
Multiple R-squared: 0.278,
Adjusted R-squared: 0.158
F-statistic: 2.31 on 2 and 12 DF,
p-value: 0.141
$means.pos.neg.coeff (Study this!)
      neg pos diff stESftSze
[1,] 16.6 18.8 2.2 1.07 Rows index contrasts
[2,] 17.7 19.2 1.5 0.73
$contrasts (so-called standard contrasts-
will discuss)
      [,1] [,2]
1     -1  -0.5
2      1  -0.5
3      0   1.0
$group.means.sds
      [,1] [,2] [,3]
Means  16.60 18.80 19.20
S.D.s   2.07  1.92  2.17

$data
...snip...
```



*Note that the first row panels above show comparisons for contrasts of A vs. B, then for (A,B) vs. C, where the dashed blue line connects the means above, be sure that you understand (and that there are no errors). The final panel might be seen as a crude version of the `granova.lw` plot (here that is especially clear since the group means are ordered just as are A, B and C). **You may want to TRY TO REPRODUCE ALL THESE results on your own machine.***

**A final note:** As the `$means.pos.neg.coeff` array above suggests, it is easy to define standardized effect sizes for each contrast by dividing mean differences by the pooled s.d. for each contrast. Each test statistic (actually a  $t$  [1.69, 1.33 above]) is just a constant times such a stES; further the *average of these squared  $t$ 's* is itself equal to the *omnibus F statistic* (try it to see).