One-way Analysis of Variance

The additive model
One-way ANOVA procedure
Multiple comparison procedures (MCPs)
- Fisher's LSD
- Tukey's W
- Student-Newman-Keuls SNK
- Duncan's multiple range test
- Scheffe's method

Overview of MCPs

One-way ANOVA

So far we have discussed group comparison tests for two independent samples that came from normal populations with possibly different means but similar variances; i.e., we tested the null hypothesis:

$$H_o : \mu_1 = \mu_2$$

We now need to expand our repertoire to include the comparison of three, four, or more independent samples taken from normal populations; i.e., to test the null hypothesis:

$$H_o : \mu_1 = \mu_2 = \mu_3 \ldots \mu_i$$

The Additive Model

The underlying basis of all ANOVA procedures is the *additive model*.

By way of example, consider the spread of a pathogen through orange groves in Florida. You obtain an independent random sample of 3 oranges from 4 orange groves & score the number of fungal infections per orange. The data are as follows:

A	B	C	D
11	7	6	5
9	9	5	3
10	8	7	4

The Additive Model

As always, first plot the data:

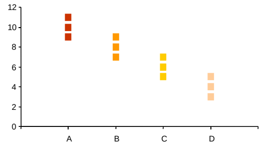

The graph does seem to indicate that there are differences among samples, but we need to apply a formal test:

$$H_o: \mu_1 = \mu_2 = \mu_3 = \mu_4$$
$$H_a: \mu_1 \neq \mu_2 \neq \mu_3 \neq \mu_4$$

The Additive Model

<u>CAUTION</u>: One might be tempted to proceed using methods we have already learned (e.g., using multiple *t*-tests to make all comparisons), but this is a gross violation of assumption and does not perform the same hypothesis test.

Repeatedly performing statistical tests on the same data set increases the chance of committing a Type-I Error. *Probability values are additive!* Thus, if we do 6 separate t-tests (as would be needed in our example), the overall P = 6 × 0.05 = 0.30...hardly a level at which inference can be made.

The Additive Model

In order to speak more precisely about individual variates we need to develop a lexicon for ANOVA using I & j subscripts to designate the j_{th} observation (row) in the i_{th} column (or vice versa):

y_{1j}	y_{2j}	y_{3j}	y_{4j}	
$y_{11} = 11$	$y_{21} = 7$	$y_{31} = 6$	$y_{41} = 5$	
$y_{12} = 9$	$y_{22} = 9$	$y_{32} = 5$	$y_{42} = 3$	
$y_{13} = 10$	$y_{23} = 8$	$Y_{33} = 7$	$y_{43} = 4$	
$\Sigma_{1j} = 30$	$\Sigma_{2j} = 24$	$\Sigma_{3j} = 18$	$\Sigma_{4j} = 12$	$\Sigma\Sigma\ y_{ij} = 84$

The Additive Model

In our orange example, the number of groups is $A = 4$ and the number of observations is $N = 3$. Based upon our knowledge of the raw data and graphs, we know the variances to be homogeneous.

Now, we know the appropriate procedure is the ANOVA, but our null hypothesis is framed in the context of means not variances. Why?

The test is possible because if the null hypothesis is true, both statistics are estimates of the variance. Let's return to our example to develop this notion...

The Additive Model

Grand mean: $\bar{y} = \sum_i \sum_j y_{ij} / an = 84/12 = 7$

Group means: $\bar{y}_1 = \sum_j y_{1j} / n = 30/3 = 10$

$\bar{y}_2 = \sum_j y_{2j} / n = 24/3 = 8$

$\bar{y}_3 = \sum_j y_{3j} / n = 18/3 = 6$

$\bar{y}_4 = \sum_j y_{4j} / n = 12/3 = 4$

The Additive Model

If we consider the population parameters related to these sample means, each observation can be thought of in terms of an additive model consisting of three terms:

$$y_{ij} = \mu + \alpha_i + \varepsilon_{ij}$$

In which:
μ = the mean of all oranges
α_i = the mean treatment effect for oranges in the i_{th} group
ε_{ij} = the "random effect" due to individual oranges

The Additive Model

We can re-write the data to demonstrate the additive components of each observation to the overall model:

$$y_{ij} = \mu + \alpha_i + \varepsilon_{ij}$$
$$y_{ij} = \mu + (y_i - \mu) + (y_{ij} - \bar{y}_i)$$

11 = 7 + (10-7) + 1
9 = 7 + (10-7) + (-1)
10 = 7 + (10-7) + 0

7 = 7 + (8-7) + (-1)
9 = 7 + (8-7) + 1
8 = 7 + (8-7) + 0

6 = 7 + (6-7) + 0
5 = 7 + (6-7) + (-1)
7 = 7 + (6-7) + 1

5 = 7 + (4-7) + 1
3 = 7 + (4-7) + (-1)
4 = 7 + (4-7) + 0

One-way ANOVA Procedure

In terms of the additive model, we can rewrite the null hypothesis in a different manner where we use α (the mean treatment effect estimated as the difference between grand mean and treatment mean):

$$H_o: \alpha_1 = \alpha_2 = \alpha_3 = \alpha_4 = 0$$
$$H_a: \alpha_i \neq 0, \text{ for some } i$$

The development of the F-test that follows, comparing the variance among groups with the variance within groups to test the above null hypothesis, assumes this additive model. The F-test also assumes that the data in each group is normally distributed and all groups have homogeneous variances.

One-way ANOVA Procedure

Thus, we can say that the random effects due to individual oranges are independently normally distributed with a mean of zero and a common variance. Formally...

$$y_{ij} = \mu + \alpha_i + \varepsilon_{ij}$$
$$\sum_i \alpha_i = 0$$
$$\varepsilon_{ij} \text{ IND}(0, \sigma^2)$$

One-way ANOVA Procedure

Total Variance	$\frac{\sum_i \sum_j (y_{ij} - \bar{y})^2}{na - 1}$	The mean squared deviation of the observations from the grand mean
Within-group Variance	$\frac{\sum_i \sum_j (y_{ij} - \bar{y}_i)^2}{a(n-1)}$	The mean squared deviation of the observations from their respective group mean (the pooled variance)
Among-group Variance	$n\left[\frac{\sum_i (y_i - \bar{y})^2}{a-1}\right]$	The mean squared deviation of the group mean from the grand mean multiplied by the number of observations in each group

One-way ANOVA Procedure

Graphically, we can look at the data as:

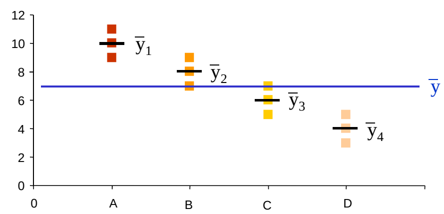

One-way ANOVA Procedure

If the null hypothesis is true, the group means will not differ from the overall mean and the within-group variance will be approximately the same as the among-group variance.

However, if the null hypothesis is false, then the among-group variance will be larger because of the significant deviations of the group means from the grand mean.

Now let's compute the variances (sums-of-squares) for each group...

One-way ANOVA Procedure

Total SS $\sum_i \sum_j (y_{ij} - \bar{y})^2 = (11-7)^2 + (9-7)^2 + (10-7)^2 + ... + (4-7)^2 = 68$

Within SS $\sum_i \sum_j (y_{ij} - \bar{y}_i)^2 = [1^2 + (-1)^2 + 0^2] + ... [1^2 + (-1)^2 + 0^2] = 8$

Among SS $n\sum_j (\bar{y}_i - \bar{y})^2 = 3[(10-7)^2 + (8-7)^2 + (6-7)^2 + (4-7)^2] = 60$

This illustrates how the total sums-of-squares can be partitioned into two parts: among-group and within-group; i.e.,

Total SS = Among SS + Within SS or 68 = 8 + 60

One-way ANOVA Procedure

Now, in order to turn our SS in to actual variances (mean squares = MS) we must divide each by their appropriate df.

The degrees of freedom (df) are partitioned much like SS:

Total df	=	Among df	+	Within df
na-1	=	a-1	+	a(n-1)
11	=	3	+	8

One-way ANOVA Procedure

We can now start to arrange all of the data in to a conventional ANOVA table:

Source	df	SS	MS
Among groups	3	60	20
Within groups	8	8	1
Total	11	68	

One-way ANOVA Procedure

The F-test is then performed as follows:

$$F = \frac{\text{Among MS}}{\text{Within MS}} = \frac{20}{1} = 20$$

$F_{table} = 4.066$ at alpha = 0.05 and df = 3,8

$F_{calc} > F_{table}$ therefore reject H_o & conclude there is at least one inequality among the four samples of oranges (i.e., one or more samples is significantly different from the other samples).

One-way ANOVA Procedure

The complete and final ANOVA table can then be summarized:

Source	df	SS	MS	F	P
Among groups	3	60	20	20	<0.05
Within groups	8	8	1		
Total	11	68			

Multiple Comparison Procedures

The null hypothesis we just rejected permitted us to conclude that *there is at least one inequality* amongst the four samples of oranges.

However, we still do not know *which* samples are unequal with respect to the others.

Typically, the investigator wishes to make further decisions, particularly if this is a carefully designed experiment with a control and one or more treatments.

We will now examine the various methods (multiple comparison procedures) by which this is accomplished.

Fisher's Least Significant Difference

The LSD test utilizes the *t*-statistic to perform pairwise comparisons via a form of confidence interval.

$$\text{LSD}_{ij}(\text{F}) = t(\nu, \alpha/2)\sqrt{\left(\frac{1}{n_i}+\frac{1}{n_j}\right)\text{MS[E]}}$$

where $\nu = a(n-1)$

ν can also be obtained from the ANOVA table (MS[E]).

Fisher's Least Significant Difference

Let's apply the LSD procedure to the data from our orange example. Because the sample sizes are all identical, we need only calculate LSD *once* (if sample sizes were different, we would need to determine LSD separately for each pairwise *ij* comparison):

$$\text{LSD}_{ij}(\text{F}) = t(\nu, \alpha/2)\sqrt{\left(\frac{1}{n_i}+\frac{1}{n_j}\right)\text{MS[E]}}$$

$$\text{LSD} = t(0.025, 8)\sqrt{\left(\frac{1}{3}+\frac{1}{3}\right)1} = 2.306\sqrt{0.666} = 1.88$$

Fisher's Least Significant Difference

LSD = 1.88 is used as a bracket interval around each group mean. Intervals that overlap are NOT significantly different from each other. Thus, we need to examine the means ± 1.88/2 = 0.94:

Group 1: 9.06-10.0-10.94
Group 2: 7.06-8-8.94
Group 3: 5.06-6-6.94
Group 4: 3.06-4-4.94

Fisher's Least Significant Difference

Group 1: 9.06-10.0-10.94
Group 2: 7.06-8-8.94
Group 3: 5.06-6-6.94
Group 4: 3.06-4-4.94

No group overlaps, therefore, they are all significantly different from each other. The means are usually summarized in a tabular format with superscripts used to depict the results of the MCP:

Group 1	Group 2	Group 3	Group 4
10^a	8^b	6^c	4^d

Fisher's Least Significant Difference

Fisher's LSD test suffers one major draw back and that is it only controls for the comparisonwise error rate (in this case $\alpha = 0.05$).

LSD produces a much higher experimentwise error rate which is between α and $k\alpha$ (where $k = a(a-1)/2$).

For example, with $a = 4$ treatments (as with the oranges), $k = 6$ comparisons. Thus, the actual experimentwise error rate is $6 \times 0.05 = 0.30$ error rate. A far cry from 0.05.

Scheffe's Multiple Comparison

The LSD values for Scheffe's pairwise comparison procedure work similarly to Fisher's, but control for the experimentwise error rate as well (will never exceed α):

$$\mathrm{LSD}_{ij}(S) = \sqrt{(t-1)F(t-1,\nu,\alpha)\left(\frac{1}{n_i}+\frac{1}{n_j}\right)\mathrm{MS[E]}}$$

where $\nu = a(n-1)$

This test is about the most conservative of the LSD tests.

Bonferroni Multiple Comparison

The Bonferroni modification of the LSD method is also a very good way to control for the experimentwise error rate:

$$\text{LSD}_{ij}(B) = t(\nu, \alpha, k)\sqrt{\left(\frac{1}{n_i} + \frac{1}{n_j}\right)\text{MS[E]}}$$

where $\nu = a(n-1)$ and $t(\nu, \alpha, k)$ comes from Table C.10

Tukey's Multiple Comparison

If the sample sizes are equal among all groups, a procedure that is less conservative than Scheffe's MCP is Tukey's MCP:

$$\text{LSD}_{ij}(T) = q(a, \nu, \alpha)\sqrt{\frac{\text{MS[E]}}{\text{n}}}$$

where $q(a, \nu, \alpha)$ comes from Table C.11

One-way ANOVA Using R

Example: Oranges Revisited

```
> oranges<-read.csv("C:/TEMPR/Oranges.csv")
> oranges
   Infection Grove
1         11     a
2         10     a
3          9     a
4          9     b
5          8     b
6          7     b
7          5     c
8          7     c
9          6     c
10         5     d
11         3     d
12         4     d
> attach(oranges)
```

One-way ANOVA Using R
Example: Oranges Revisited

```
> summary(oranges)

   Infection  Grove
 Min.   : 3   a:3
 1st Qu.: 5   b:3
 Median : 7   c:3
 Mean   : 7   d:3
 3rd Qu.: 9
 Max.   :11

> boxplot (Infection
 ~ Grove)
```

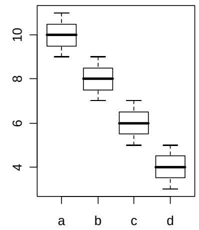

One-way ANOVA Using R
Example: Oranges Revisited

```
> anova(lm(Infection~Grove))

Analysis of Variance Table

Response: Infection
          Df Sum Sq Mean Sq F value    Pr(>F)
Grove      3     60      20      20 0.0004487 ***
Residuals  8      8       1
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05
‘.’ 0.1 ‘ ’ 1
```

One-way ANOVA Using R
Example: Oranges Revisited

```
> pairwise.t.test(Infection, Grove,
+     p.adj="bonferroni")

        Pairwise comparisons using t tests with
pooled SD

data:  Infection and Grove

  a       b       c
b 0.23981 -       -
c 0.00717 0.23981 -
d 0.00048 0.00717 0.23981

P value adjustment method: bonferroni
```

The End...

Reminder: Read Chapters 10 & 11 in your textbook.