## Two way (factorial) Analysis of Variance :

Two way ANOVA: introduction.

Suppose we want to take a look at two factors at once:

What is a factor? It is similar to our "treatment". It's what we want to investigate.

Here's an example:

We investigate the "effect" of estrogen on bats. We inject a group of bats with estrogen, and we use another group of bats as controls.

We also want to know the difference in estrogen levels between male and female bats (suppose for the moment we didn't know anything about estrogen levels in males vs. females).

## Approach I:

Compare four groups and do a one way ANOVA:

male control	male estrogen	female control	female estrogen
XX	XX	XX	XX
XX	XX	XX	XX
			etc.

Yes, we'd compare the four groups, but this design doesn't really tell us what effect estrogen has and what effect sex has (the two different things would be difficult to tease apart).

## Approach II:

Do two t-tests and compare estrogen bats vs. controls, then males vs. females.

This approach might tell us if estrogen has an effect, or if sex has an effect.

What problems are there here?

The effect of estrogen might be covered up by the difference in estrogen due to sex, and we might not get any results here.

#### Approach III:

We do one ANOVA, but this time we do a two way ANOVA:

	males	females
estrogen	XX,XX,	XX,XX,
control	XX,XX,	xx,xx,etc.

Note that this looks just a bit like a contingency table (which it is NOT), so don't get confused.

This design will test three things all at once:

i) The effect of estrogen injections

ii) The effect of sex

iii) Any interactions between sex and estrogen injections

So how do we do it?

Two way ANOVA: the math.

We won't actually go into too many details here (it may look like a lot, but it's not), but here are the basics:

There are actually several ways to do this, but let's use the same method as in your text.

Step 1: calculate the between group SS. Do this assuming each combination of factors above is actually a group. So you'd have four groups above (one group for each combination of factors).

Your text calls this the "cells" SS:

$$SS_{cells} = n \sum_{i=1}^{a} \sum_{j=1}^{b} (\bar{x}_{ij} - \bar{x})^2$$

where "i" refers to group "i" of factor "a", and "j" refers to group "j" of factor "b". Note that if you want to refer to actual values, these are now indexed by "l".

(The constantly changing subscripts are a bit annoying, but we're stuck with them).

This equation is identical to what we used before for "between" groups and would give you the same result if you re-arranged the data into a one-way ANOVA with four groups.

Step 2: calculate the within group SS. This, again, will be identical to what we had before, though the equation will look a bit different:

$$SS_{within} = \sum_{i=1}^{a} \sum_{j=1}^{b} \sum_{l=1}^{n} (x_{ijl} - \bar{x}_{ij})^2$$

Step 3: now we subdivide SS<sub>cells</sub> (what we calculated in step 1):

a) We get a  $SS_{between}$  for factor a:

We pretend factor (b) didn't exist, and calculate SS<sub>between</sub> for two groups (in this example):

$$SS_{factor\,a} = bn \sum_{i=1}^{a} \left( \bar{x}_{i} - \bar{x} \right)^2$$

The "." above means we ignore factor b and calculate the average for group 1 of factor a and then for group 2 of factor a.

b) Now repeat for factor b:

$$SS_{factor b} = an \sum_{i=1}^{b} (\bar{x}_{j} - \bar{x})^2$$

Step 4: we now notice the following relationship:

$$SS_{cells} = SS_{factor a} + SS_{factor b} + ????$$

It turns out the relationship is virtually never perfect. In other words, the variation in group a + the variation in group b will not give you the total variation.

There is a third source of variation:

Step 5: Variation due to interaction:

The two factors when considered at the same time will increase the overall variation.

So we have:

 $SS_{cells} = SS_{factor a} + SS_{factor b} + SS_{interaction}$ 

Calculating  $SS_{interaction}$  directly is a pain, but fortunately we don't have to do that. We already have  $SS_{cells}$ ,  $SS_{factor a}$ , and  $SS_{factor b}$ . So we can get  $SS_{interaction}$  from the above equation:

 $SS_{interaction} = SS_{cells} - SS_{factor a} - SS_{factor b}$ 

Step 6: Degrees of freedom. For details, see your text, but here they are:

$df_{total} = N-1$	$df_{cells} = ab - 1$	$df_{\text{within}} = ab(n-1)$
$df_{factor a} = a - 1$	$df_{factor b} = b - 1$	$df_{interaction} = (a-1)(b-1)$

Step 7 (optional): You can also calculate  $SS_{total}$  (the equation looks different, but gives the same result as with a one way ANOA). If you're interested, see the text (p. 251).

Now that we have all the mathematical details taken care of, we arrange things into our ANOVA table as follows

#### Two way ANOVA: the table:

Source	SS	df	MS	F
cells factor a factor b interaction within	$\begin{array}{l} SS_{cells} \\ SS_{factor \ a} \\ SS_{factor \ b} \\ SS_{interaction} \\ SS_{within} \end{array}$	ab-1 a-1 b-1 (a-1)(b-1) ab(n-1)	$\frac{SS_{cells}/(ab - 1)}{SS_{factor a}/(a - 1)}$ $\frac{SS_{factor b}}{(b - 1)}$ $\frac{SS_{interaction}}{(a-1)(b-1)}$	N/A MS <sub>factor a</sub> /MS <sub>within</sub> MS <sub>factor b</sub> /MS <sub>within</sub> MS <sub>interaction</sub> /MS <sub>within</sub>
TOTAL	SS <sub>total</sub>	N-1		

Let's look at this a bit more closely.

What we get is three "F's" in the last column. Each F represents a hypothesis test, but because of the way ANOVA works, we don't need to worry about  $\alpha$  changing.

The first F represents a test for difference between groups for factor a

The second F represents a test for difference between groups for factor b

The third F represents a test for interactions (more on this below)

In R (or another software package) the probabilities would be printed in a last column at the end.

We are actually testing three hypotheses at the same time:

H<sub>0</sub>: there is no effect of factor a

H<sub>0</sub>: there is no effect of factor b

H<sub>0</sub>: there is no interaction between factors a and b

Note that if the interaction is significant, this causes all kinds of problems (again, more below).

Two way ANOVA: example:

Let's do an example. We'll go through 12.1 (p. 251) from the text. This example is similar to the example mentioned above.

We'll skip the math details (see 12.1a on p. 253 if you're interested) and look at the data (p. 251) and the ANOVA table (p. 255).

(Please *ignore* the "machine formulas" when you read through this section).

End result:

Hormone does affect plasma calcium levels

Sex does not affect plasma calcium levels

The two factors do not interact

Two way ANOVA: comments:

As presented here, the ANOVA assumes equal replication. That is, each combination of factors has the same number of measurements in it.

In the example above, each combination of sex and hormone has 5 replicates (i.e., 5 birds were measured at each combination of factors.

If the replications are not the same, this creates problems for the calculations as well as for the theory.

See section 12.2 (p. 265) for more on this. We'll skip this for now.

We can graph the results. To do this, we plot all of our means on the same graph:

The y axis represents the actual value (measurement)

The x axis represents one of the factors (e.g., male, female)

Then we connect the means that are the same for the other factor.

For our bird example we would get the graph on p. 259.

This is also called an "interaction plot" and can be easily generated in R (the R code for the bird example is below).

What about this interaction?

Interaction - an interaction is when two things come together and give you an effect that's greater (or less) than either effect alone would predict.

A classic example is in medicine. Suppose some medication for blood pressure "interacted" with a medication for heart rate in such a way as to cause a dangerous loss of blood pressure. That would be an interaction.

Or perhaps, the opposite happens - quite unexpectedly - when both medications are given together, heart rate skyrockets to dangerous levels (say 160 - 180 beats/minute, while just sitting still). Again, this is an interaction.

Interactions are difficult to deal with:

If an interaction is significant, that means that the effect of either one of the factors may or may not be significant regardless of what the results are:

The change in factor A might have been caused by an interaction with factor B; by itself, factor A might not do anything!

If interactions are present, you need to do further analyses to determine the effect(s) of your factors.

Model I vs. Model II vs. Model III ANOVA

Everything we've done so far has assumed that the factors we're interested in are "fixed". That is, we have control over them.

We select "sex", we select the amount of "hormone" in the example above.

The effects are "fixed" by us in the sense that we set these factors.

This is a Model I or (better) fixed effects model.

What happens if we don't have control over one (or both) of the factors?

The good news is that most of the math stays the same.

The bad news is that the ANOVA table does change.

Let's examine a "mixed" effects ANOVA (or Model III ANOVA):

In this case, one of the effects is fixed, the other is random (you don't pick the "levels" deliberately):

You want to find out if the effect of your hormone is the same in all birds (let's ignore sex since it wasn't significant anyway).

You are NOT interested in robins vs. blue jays (for example). Any birds will do.

You just happen to have access to robins and blue jays, however, so you want to find out if there is any variation between these birds.

Does species make a difference? Note that you didn't "pick" robins or blue jays. Any birds would do, but you just happen to have access to these.

What happens to the analysis?

Everything stays the same. Factor a is hormone, factor b is species.

Calculate everything and set up your ANOVA table as always.

However, your F's change.

Notice that hormone is still "fixed", but species is now a "random" effect.

Source	SS	df	MS	F
cells	$SS_{cells}$	ab-1	$SS_{cells}/(ab - 1)$	N/A
factor a	SS <sub>factor a</sub>	a-1	$SS_{factor a}/(a - 1)$	MS <sub>factor a</sub> /MS <sub>interactioin</sub>
factor b	$SS_{factor b}$	b-1	$SS_{factor b}/(b - 1)$	MS <sub>factor b</sub> /MS <sub>within</sub>
interaction	$SS_{interaction}$	(a-1)(b-1)	$SS_{interaction}/(a-1)(b-1)$	$MS_{interaction}/MS_{within}$
within	SSwithin	ab(n-1)		
TOTAL	$SS_{total}$	N-1		

So what's the difference?

Notice how the F's are now calculated.

The interaction and random effects are calculated just as before.

The fixed effect is done by comparing  $MS_{factor a}/MS_{interaction}$  (by comparing the MS for the fixed factor with the MS for the interaction)

This might be a bit counterintuitive.

Why the difference?

Because now we're interested in how variable birds are NOT in if there's a difference between robins and blue jays.

That's not a very good explanation. But the details require some math and an introduction to what are called "expected values".

For now, think of it this way:

If our fixed effect is significant, the variation in our fixed effect (factor a) should be higher than the variation our interaction.

Again, the reason would require another hour or two worth of lecture.

A model II (or random effects) ANOVA:

In this case, both effects are random. This is rather unusual and isn't used that much.

Just as above, the table is set up the same way, but now your calculations for F become:

factor a	MS <sub>factor a</sub> /MS <sub>interaction</sub>
factor b	$MS_{factor b}/MS_{interaction}$
interaction	$MS_{interaction}/MS_{within}$

For details see your text on p. 262

We won't discuss this any further here.

Let's do the two way ANOVA example using R:

Organize your file like this:

16.3	female	e no	You want three columns here. The first has the plamsa
20.4	female	e no	concentration. The second identifies the sex, and the
			third identifies whether or not the bird got the hormone
15.3	male	no	treatment.
 38.1	 female	 e yes	
 34.0	 male	 yes	
•••			

etc.

Either enter your data directly into R or import it.

Once your data is in R, you can type the following:

sex <- factor(sex)
hormone <- factor(hormone)</pre>

to make sure R recognizes these as factors.

then do:

birds2way <- aov(conc ~ sex \* hormone)</pre>

note the "\*" (don't do "+")

then look at the results:

```
summary(birds2way)
```

and we get:

```
Df Sum Sq Mean Sq F value Pr(>F)

sex 1 70.3 70.3 3.733 0.0713 .

hormone 1 1386.1 1386.1 73.585 2.22e-07 ***

sex:hormone 1 4.9 4.9 0.260 0.6170

Residuals 16 301.4 18.8

---

Signif. codes: 0 `***' 0.001 `**' 0.01 `*' 0.05 `.' 0.1 ` ' 1
```

Which matches what the book does on page 255.

You can also do the following if you want:

TukeyHSD(birds2way)

This will give you confidence intervals for sex, hormone, and then all possible pairs:

```
Tukey multiple comparisons of means
    95% family-wise confidence level
Fit: aov(formula = conc ~ sex * hormone)
$sex
     diff
                lwr
                          upr
                                  p adj
m-f -3.75 -7.864692 0.3646918 0.0712638
$hormone
        diff
                  lwr
                           upr p adj
yes-no 16.65 12.53531 20.76469 2e-07
$`sex:hormone`
             diff
                         lwr
                                           p adj
                                   upr
            -2.76 -10.613382
m:no-f:no
                              5.093382 0.7485244
            17.64
                    9.786618 25.493382 0.0000452
f:yes-f:no
                    5.046618 20.753382 0.0012416
m:yes-f:no
            12.90
f:yes-m:no
            20.40 12.546618 28.253382 0.0000078
m:yes-m:no
            15.66 7.806618 23.513382 0.0001729
m:yes-f:yes -4.74 -12.593382 3.113382 0.3428025
```

This is more interesting if one of your factors has more than two levels (here it doesn't tell you much - the pairwise comparisons aren't very useful).

Two more things to do with R:

1) You can easily get an interaction plot (put your two factors first, then your measurement variable):

interaction.plot(sex,hormone,conc)



2) You can check the normal distribution assumption with one (yes one!) qq plot by doing:

```
qqnorm(birds2way$residuals)
qqline(birds2way$residuals)
```

How does this work?

It would take a while to explain. Basically it combined all the data into a single sample and figures out how far away each point is from what is expected in each group for each factor (or combination) of factors.

All these "distances" are then plotted on the same graph. A more thorough explanation will have to wait until regression.

Note that up until now you would need to do one plot for each group within a factor or combination of factors (i.e, you'd have to do four plots).

(This also works with a regular one way ANOVA).

Final comment for R:

Doing a mixed model or random effects model is a bit more complicated. The easiest way would be to just use the ANOVA table (for example the one above) and re-calculate the F's that you need by hand. It shouldn't take more than a minute or two. Then you can get the p-values by doing:

pf(F,df1,df2,lower.tail = FALSE)

where "F" is your F, df1 is numerator degrees of freedom, and df2 is denominator degrees of freedom.

If you're really want to do this directly in R, you can (you can do almost anything in R), but you'll need to install the "nlme" package and check up how to use it on the web.

# Finally:

Three (or more) way ANOVA's are handled in a similar way to two way ANOVA's. The basic idea should be pretty obvious. See chapter 14 in your text if you're interested.

But there's a lot more to ANOVA than just one, two or three way designs.....