

**Stat 414 – Day 5**  
**Random Effects ANOVA**

**Last Time:**

- Likelihood Ratio Tests
  - Compare the likelihood values of two nested models
  - Reference distribution:  $\chi^2$  with  $df =$  difference in number of parameters
- Restricted Maximum Likelihood Estimation
  - Restricts likelihood space based on estimated regression coefficients
  - Adjusts for  $\hat{\sigma}^2$  for degrees of freedom of model
  - Make sure making consistent comparisons
    - One clue not nested is whether have same  $df$
    - Can still compare likelihoods “informally”

**Example:** Recall the chip melting study we carried out in class. Each student measured the melting time of semi-sweet chocolate, peanut butter, and butterscotch chips under identical conditions. Of interest is whether there is a difference in the melting times for these three types of chips. (ChipData.txt)

(a) Explore the data: What do you learn from the boxplots? Do you think these differences will be statistically significant?

(b) State null and alternative hypotheses. What statistical method could you use to analyze these data?

(c) Produce a regression model for testing the hypotheses in (b). Report the test statistic, degrees of freedom, and p-value for your (one) test. What do you conclude?

(d) Also report the  $R^2$  and  $s$  (residual standard error) values. Interpret these values in context. Are they consistent with your p-value? Also report “MSTotal” =  $Var(time)$ .

(e) Examine the residual plots, do you consider the model assumptions met?

A huge problem with this analysis is we have not accounted for the repeated measures on each individual. Our study was considered a *randomized block design* and we need to adjust for the blocking variable in the analysis. Similar to running a matched pairs *t*-test rather than an independent samples *t*-test if there had only been two groups – we want to focus on the differences within each block. This can make a difference if the blocking variable is significantly related to the response variable.

(f) Does chip melting time appear to vary across the individuals in the study?

(g) Look at the data: If we were to “adjust” for the individual doing the melting, is there evidence of a chip effect?

```
with(chipdata, coplot(time~chip | person))
```

(h) Fit a regression model including chip type and person. What do you learn from each of the “effects” tests? What can you tell me about the  $R^2$  for each variable? What is “MSTotal”?

(i) Interpret the coefficient of “chipchoc” in context.

(j) Interpret the coefficient of “personCC” in context.

(k) Include the interaction between chip type and person in the model. What do you learn?

There are a few issues with the above, mostly the “boundary conditions” involved with estimating the likelihood for variance parameters that must be positive. A quick modification is to cut the p-value in half.

Some downsides to including the person variable in the model

- We aren't really interested in the person-to-person comparisons
- With lots of people, we use a lot of degrees of freedom in our model
- We can only generalize our conclusions to the 9 people in the study

But we really need to account for (and even measure) the amount of person to person variation in order to better assess our treatment effects of interest.

In a situation like this, one option is to treat *person* as a *random effect* rather than a *fixed effect*. This means we are going to treat these 9 people not as 9 levels of a factor (if I did the study again, I would get 9 different people) but as a random sample from a population. We are going to assume that the "person effects" follow a normal distribution with mean zero and a variance, call this variance  $\tau^2$ ,  $u_j \sim N(0, \tau^2)$ . So our regression model becomes

$$y_{ij} = \beta_0 + \beta_{choc} + \beta_{pb} + u_j + \epsilon_{ij} \text{ w/ } u_j \sim iid N(0, \tau^2), \epsilon_{ij} \sim iid N(0, \sigma^2), cov(u_i, \epsilon_{ij}) = 0$$

(l) What do  $i$  and  $j$  refer to? What are  $n_i$  and  $n_j$ ? What does "iid" refer to? What is  $E(Y_{ij} | person_j)$ ? What is  $E(Y_{ij})$ ? What is  $Var(Y_{ij})$ ? How many parameters are there to estimate in this model?

To fit this model in R:

```
library(nlme)
model4 = lme(fixed = time ~ chip, random = ~1 | person, data = chipdata)
summary(model4)
```

(m) Did this procedure use ML or REML estimation?

(n) Report the AIC and logLik values.

(o) *Fixed effects*: How do the intercept, and coefficients for chip type compare to (h)? Interpret  $\hat{\beta}_0$  in context.

*Random effects*: Notice you no longer see estimates for each individual. Instead you should see

```
Random effects:
Formula: ~1 | person
      (Intercept) Residual
StdDev:   18.56445  7.358543
```

This gives  $\hat{\sigma}_\epsilon$  and  $\hat{\tau}$ . Which is which?

Notice that no p-values are given for these variance components. How do we decide whether the person-to-person variation is meaningful? Statistically significant?

(p) Sum the two *variances* together and compare this to MSTotal.

(q) How much (what percentage) of this total variation is due to the person-to-person variability?

**Definition:** This ratio, called the *intraclass correlation coefficient*, is the proportion of the total variance that is explained by the between-group variability. It can also be interpreted as the correlation between the responses within the same group.

For statistical significance, we want to test  $H_0: \tau^2 = 0$  vs.  $H_a: \tau^2 > 0$ .

(r) Why are we using a one-sided test here?

One answer to this question is we already did that in (f) (with an ANOVA on just the person effects). This is treating the variable as fixed. Alternatively, we can compare the model with and without these random effects using REML.

(s) Fit the model without random effects using REML and find the log likelihood value.

```
gls(time ~ chip, data= chipdata)
```

(t) What do you find for the likelihood ratio test statistic, df, and p-value?