RANDOM EFFECTS MODELS (Chapter 17)

So far we have studied experiments and models with only *fixed effect* factors: factors whose levels have been specifically fixed by the experimenter, and where the interest is in comparing the response for just these fixed levels.
A *random effect* factor is one that has many possible levels, and where the interest is in the variability of the response over the entire population of levels, but we only include a random sample of levels in the experiment.

Examples: Classify as fixed or random effect.

1. The purpose of the experiment is to compare the effects of three specific dosages of a drug on response.

2. A textile mill has a large number of looms. Each loom is supposed to provide the same output of cloth per minute. To check whether this is the case, five looms are chosen at random and their output is noted at different times.

3. A manufacturer suspects that the batches of raw material furnished by his supplier differ significantly in zinc content. Five batches are randomly selected from the warehouse and the zinc content of each is measured.

4. Four different methods for mixing Portland cement are economical for a company to use. The company wishes to determine if there are any differences in tensile strength of the cement produced by the different mixing methods.

Note: The theory behind the techniques we discuss assumes that the population of levels of the random effect factor is infinite. However, the techniques fit well as long as the population is at least 100 times the size of the sample being observed. Situations where the population/sample size ratio is smaller than 100 require "finite population" methods which we will not cover in this class.

The Random-Effects One-Way Model

For a completely randomized design, with v randomly selected levels of a single treatment factor T, and r_i observations for level i of T, we can use the model

$$Y_{it} = \mu + T_i + \varepsilon_{it},$$

where:

Each $\varepsilon_{it} \sim N(0, \sigma^2)$
The ε_{it}'s are independent random variables
The T_i's are independent random variables with distribution $N(0, \sigma_T^2)$
The T_i's and ε_{it}'s are independent of each other.

Note: If all r_i have the same value r, then we have a *balanced* design.

Caution: The terminology can be confusing. Here is how to think of the model:

1. Each possible level of the factor T might have a different effect. "Effect of level i" is thus a random variable, hence has a certain distribution – the "distribution of effects". One of our model assumptions is that this distribution is normal. By adjusting μ if necessary, we may assume this distribution of effects has mean 0. We call the variance of this distribution σ_T^2.

2. The effect of level i of T is called (confusingly) T_i.

3. Since the levels i are randomly chosen (that is, we have a simple random sample of levels), we can say that the T_i's are independent random variables with the same distribution $N(0, \sigma_T^2)$.

4. The conditions on the ε_{it}'s are the same as in previous models. They say that the observations are independently taken within each level and between levels, and all have the same normal distribution.

5. However, we also assume that the observations are made independently of the choice of levels, hence the last condition.

(Think about what these mean in, e.g., Example 2 or Example 3 above.)

Consequences:

$E[Y_{it}] =$

$Var(Y_{it}) =$

$Y_{it} \sim$

$Cov(Y_{it}, Y_{is}) =$

$\rho(Y_{it}, Y_{is}) =$

Thus: Observations within the same treatment level are correlated. Does this make sense?

Terminology: σ_T^2 and σ^2 are called *variance components*. (Why?)

Hypothesis test: If we wish to test whether or not the level of the factor T makes a difference in the response, what should the null and alternate hypotheses be?

H_0^T:

H_a^T:

Least squares estimates: Given data y_{it}, i = 1, 2, … , v, t = 1, 2, … , r_i, we can still use the method of least squares to obtain "fitted values" $\hat{y}_{it} = \bar{y}_{i\bullet}$. However, our interest here will not be the fits, but the sums of squares obtained from them. In particular, we can still form

$$\text{ssE} = \sum_{i=1}^{v}\sum_{t=1}^{r_i}\left(y_{it} - \bar{y}_{i\bullet}\right)^2$$

and the corresponding random variable

$$\text{SSE} = \sum_{i=1}^{v}\sum_{t=1}^{r_i}\left(Y_{it} - \bar{Y}_{i\bullet}\right)^2$$

Also as with the fixed effects model, we obtain the least squares estimate (or "fit") $\bar{y}_{\bullet\bullet}$ for the submodel (assuming H_0 is true)

$$Y_{it} = \mu + \varepsilon_{it},$$

and can form its error sum of squares $\text{ssE}_0 = \sum_{i=1}^{v}\sum_{t=1}^{r_i}\left(y_{it} - \bar{y}_{\bullet\bullet}\right)^2$ and the sum of squares for treatment $\text{ssT} = \text{ssE}_0 - \text{ssE} = \sum_{i=1}^{v}\sum_{t=1}^{r_i}\left(y_{it} - \bar{y}_{\bullet\bullet}\right)^2 - \sum_{i=1}^{v}\sum_{t=1}^{r_i}\left(y_{it} - \bar{y}_{i\bullet}\right)^2$, and the corresponding random variable

$$\text{SST} = \sum_{i=1}^{v}\sum_{t=1}^{r_i}\left(Y_{it} - \bar{Y}_{\bullet\bullet}\right)^2 - \sum_{i=1}^{v}\sum_{t=1}^{r_i}\left(Y_{it} - \bar{Y}_{i\bullet}\right)^2$$

As with the fixed effects model, we can get an alternate expression for SSE (cf Equation 3.4.5):

$$\text{SSE} = \sum_{i=1}^{v}\sum_{t=1}^{r_i}Y_{it}^2 - \sum_{i=1}^{v}r_i\bar{Y}_{i\bullet}^2 \qquad (*)$$

To find an estimator of σ^2, we first need to find E(SSE). We can use (*) if we find $E(Y_{it}^2)$ and $E(\bar{Y}_{i\bullet}^2)$. To do this:

Recall: $\text{Var}(Y) = E(Y^2) - [E(Y)]^2$, so
$E(Y^2) = \text{Var}(Y) + [E(Y)]^2$

Applying this to Y_{it}^2:
$E(Y_{it}^2) = \text{Var}(Y_{it}) + [E(Y_{it})]^2 =$ ______________________

To find $E(\bar{Y}_{i\bullet}^2)$, first note that

$$\bar{Y}_{i\bullet} = \frac{1}{r_i}\left[\sum_{t=1}^{r_i}\left(\mu + T_i + \varepsilon_{it}\right)\right]$$

$$= \mu + T_i + \frac{1}{r_i}\left[\sum_{t=1}^{r_i} \varepsilon_{it}\right]$$

Thus: $E(\overline{Y}_{i\bullet}) =$

$\mathrm{Var}(\overline{Y}_{i\bullet}) =$

$E(\overline{Y}_{i\bullet}{}^2) =$

Finally, we get

$$E(SSE) = \sum_{i=1}^{v}\sum_{t=1}^{r_i} E(Y_{it}{}^2) - \sum_{i=1}^{v} r_i E(\overline{Y}_{i\bullet}^2)$$

$=$

We define MSE = SSE/(n-v) as before, and so E(MSE) = __________

Thus:

To do inference, we also need an unbiased estimator of $\sigma_T{}^2$. To this end, look at E(SST). As with the fixed effects model,

$$SST = \sum_i r_i \overline{Y}_{i\bullet}{}^2 - n\overline{Y}_{\bullet\bullet}^2$$

Now $\overline{Y}_{\bullet\bullet} =$

$=$

So $E(\overline{Y}_{\bullet\bullet}) =$

and $\mathrm{Var}(\overline{Y}_{\bullet\bullet}) =$

Thus E($\overline{Y}_{\bullet\bullet}{}^2$) = Var($\overline{Y}_{\bullet\bullet}$) + E($\overline{Y}_{\bullet\bullet}$)^2

=

Then

E(SST) = $\sum_i r_i E(\overline{Y}_{i\bullet}{}^2) - nE(\overline{Y}_{\bullet\bullet}^2)$ =

So (defining MST = SST/(v-1) as usual)

E(MST) = =

where c =

Recalling that E(MSE) = σ^2, we can calculate

E([MST – MSE]/c) =

i.e., [MST – MSE]/c is an unbiased estimator of _________

Note: 1. If we have a balanced design (all $r_i = r$), then n = vr, and

c =

2. In general, since the r_i's are positive and sum to n,

Testing Equality of Treatment Effects:

Recall: H_0^T: $\sigma_T^2 = 0$ (i.e., T ≡ 0)
H_a^T: $\sigma_T^2 > 0$

If H_0^T is true, then E(MST) = $c\sigma_T^2 + \sigma^2 = 0 + \sigma^2$ = E(MSE), so we expect MST/MSE ≈ 1. If H_a^T is true, then E(MST) > E(MSE), so if σ_T^2 is large enough, we expect MST/MSE >1. Thus MST/MSE will be a reasonable test statistic, if it has a known distribution.

Under the model assumptions, the following can be proved:

i. SST/($c\sigma_T^2 + \sigma^2$) ~ χ^2(v-1)

ii. SSE/ $\sigma^2 \sim \chi^2$(n-v)

iii. SST and SSE are independent random variables

Thus

$$\frac{SST \big/ (c\sigma_T^2 + \sigma^2)(v-1)}{SSE \big/ \sigma^2 (n-v)} \sim \text{F(v-1, n-v)}.$$

This fraction can be re-expressed as

$$\frac{MST}{MSE}\frac{\sigma^2}{c\sigma_T^2 + \sigma^2}.$$

Thus if H_0^T is true, MST/MSE ~ F(v-1, n-v). Thus MST/MSE is indeed a suitable test statistic for H_0^T.

Moreover, MST and MSE are calculated the same way as in the ANOVA table for fixed effects, so we can use the same software routine.

Model checking;

We should check model assumptions as best we can before deciding to proceed to inference. Since the least square fits are the same as for fixed effects, we can form standardized residuals and use them for some checks:

a. $\varepsilon_{it} \sim N(0, \sigma^2)$ and are mutually independent – same checks as for fixed effects model.

b. Independence of the ε_{it}'s from the T_i's – This is not easy to check, so care is needed in design and implementation of the experiment. Sometimes unequal variances of the ε_{it}'s can be a clue to a problem with independence.

c. Independence of the T_i's – Also not checkable by residuals, so care is needed in the design and implementation of the experiment.

d. $T_i \sim N(0, \sigma_T^2)$. Recall that $\text{Var}(\overline{Y}_{i\bullet}) = \sigma_T^2 + (1/r_i)\,\sigma^2$. So in the case of equal sample sizes (balanced design), the $\overline{Y}_{i\bullet}$'s should all be $\approx N(\mu, \sigma_T^2 + \sigma^2/r)$. Thus a normal plot of the $\overline{y}_{i\bullet}$'s should be approximately a straight line; however, if v is small, the normal plot may not be informative.

Note: This is an important check, since the procedure is *not* robust to departures from normality of random effects.

ONE RANDOM EFFECT EXAMPLE

A textile mill has a large number of looms. Each loom is supposed to provide the same output per minute. To test this assumption, five looms are chosen at random. Each loom's output is recorded at five different times.

Check the model: Run on Minitab to get residuals and fits. Make model-checking plots:

Standardized residuals vs factor levels and fits:

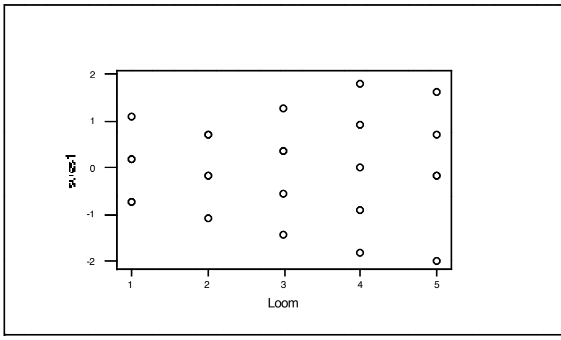

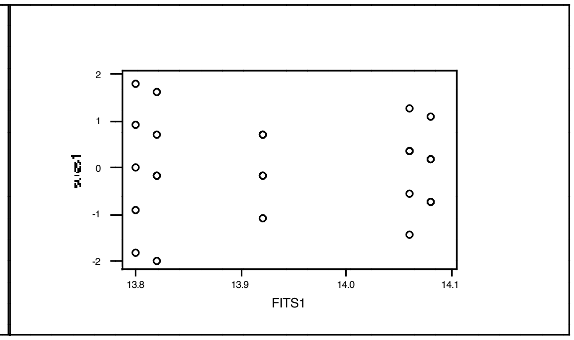

Max and min sample standard deviations by level: 0.753 and 1.424

Normal probability plots of standardized residuals and level means:

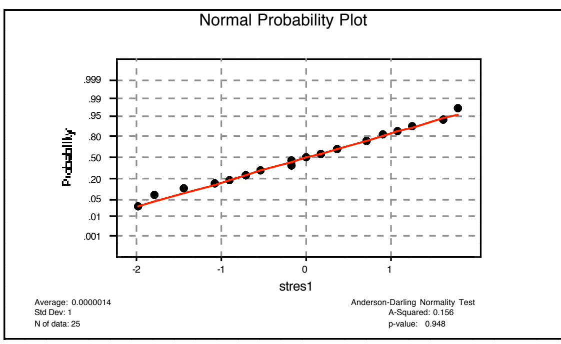

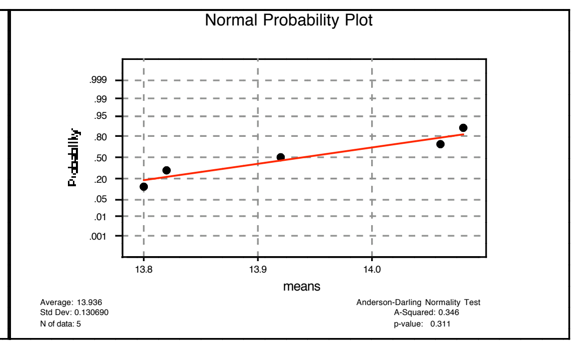

Analysis of Variance table:

Source	DF	SS	MS	F	P
Loom	4	0.34160	0.08540	5.77	0.003
Error	20	0.29600	0.01480		
Total	24	0.63760			