Random Effects ANOVA

James H. Steiger

Department of Psychology and Human Development
Vanderbilt University

Random Effects ANOVA

Introduction

So far, in our coverage of ANOVA, we have dealt only with the case of fixed-effects models.

With fixed-effects factors, the effects of the different levels of the independent variable are treated as fixed constants to be estimated.

In this module, we introduce the idea of a *random-effect* factor. The treatment effects for random effects factors are treated as random variables, or, equivalently, random samples from a population of possible treatment effects.

When models include random effects, the Expected Mean Squares will often differ from the same model with fixed effects.

In most cases, this will affect how F-tests are performed, and the distribution of the F-statistic and its degrees of freedom.

A One-Way Random Effects ANOVA

Suppose you are interested in the natural degree of variation in sodium content across brands of beer.

There are hundreds of brands of beer being sold in the U.S., and you only have time to test 8.

You select your 8 brands randomly from a list of all brands available, and you test 6 bottles of beer for each brand.

A One-Way Random Effects ANOVA

Although we can visualize an "effect" for each brand, we recognize that this effect need not be conceptualized as a fixed value — in an important sense, the effect of the first brand is a random variable, since that brand was sampled from a larger set.

If we were only interested in generalizing to the 8 brands in the study, we could choose to regard the effect of each beer as a fixed constant.

But because we randomly sampled the beer brands from a larger population, we actually can generalize back to the entire population. Let's see how.

The first thing we have to realize is that the basic ANOVA model has changed. We now have

$$\begin{aligned} Y_{ij} &= \mu_j + \epsilon_{ij} \\ &= \mu + \alpha_j + \epsilon_{ij} \quad 1 \leq i \leq n, \; 1 \leq j \leq a \end{aligned} \tag{1}$$

with

$$\epsilon_{ij} \sim \text{i.i.d. } N(0, \sigma_e^2) \tag{2}$$

$$\alpha_{ij} \sim \text{i.i.d } N(0, \sigma_A^2) \tag{3}$$

I should note immediately that many books distinguish between fixed and random effects by using Greek letters for the former and standard Arabic letters for the latter.

We are not using that convention here, as neither textbook used recently in this course employs it.

However, there are advantages to a notation that explicitly "types" its effects as fixed or random.

A One-Way Random Effects ANOVA

The Basic Model

This shift in models means that now there are two sources of random variation on the right side of the model equation, while before there was only one.

One immediate consequence of this fact is that now observations are correlated within group!

Let's use standard linear combination theory to derive the variances and covariances of scores.

Since the effects α_j and the errors ϵ_{ij} are independent, it follows that

$$\text{Var}(Y_{ij}) = \sigma_A^2 + \sigma_e^2 \tag{4}$$

A One-Way Random Effects ANOVA

The Basic Model

On the other hand, consider two observations Y_{11} and Y_{21}, both in group 1. What is their covariance?

We can solve directly for this, using results from Psychology 310. Since $Y_{11} = \mu + \alpha_1 + \epsilon_{11}$, $Y_{21} = \mu + \alpha_1 + \epsilon_{21}$, and covariances are unaffected by additive constants, we can see that the covariance between Y_{11} and Y_{21} is the covariance between $\alpha_1 + \epsilon_{11}$ and $\alpha_1 + \epsilon_{21}$.

Remember the heuristic rule? We simply multiply the two expressions and apply a conversion rule. Here is the multiplication.

$$(\alpha_1 + \epsilon_{11})(\alpha_1 + \epsilon_{21}) = \alpha_1^2 + \epsilon_{11}\epsilon_{21} + \alpha_1\epsilon_{11} + \alpha_1\epsilon_{21} \tag{5}$$

Next we apply the conversion rule

$$\begin{aligned} \mathrm{Cov}(Y_{11}, Y_{21}) &= \mathrm{Var}(\alpha_1) + \mathrm{Cov}(\epsilon_{11}, \epsilon_{21}) + \mathrm{Cov}(\alpha_1, \epsilon_{11}) + \mathrm{Cov}(\alpha_1, \epsilon_{21}) \\ &= \sigma_A^2 + 0 + 0 + 0 \\ &= \sigma_A^2 \end{aligned} \tag{6}$$

A One-Way Random Effects ANOVA

The Basic Model

So while the observations within any group are independent in the fixed-effects model, they are correlated in the random effects model. Since the correlation coefficient is the ratio of the covariance to the product of standard deviations, and each observation has a variance of $\sigma_A^2 + \sigma_e^2$, it follows that, within any group, pairs of observations observations have a correlation of

$$\rho = \frac{\sigma_A^2}{\sigma_A^2 + \sigma_e^2} \tag{7}$$

This correlation is sometimes called the *population intraclass correlation*.

A One-Way Random Effects ANOVA

Calculations

We are very fortunate, in that the sums of squares for the random effects model are calculated in exactly the same way they are calculated for the fixed effects model.

A One-Way Random Effects ANOVA

Expected Mean Squares and F Test

Remember that earlier in the course, I mentioned that ultimately Expected Mean Squares would play an important role in deciding how to perform an F test on a particular model. We have almost reached that point. Notice in the table below that the expected mean squares for the random effects model are almost identical to those for the fixed effects model.

A One-Way Random Effects ANOVA

Expected Mean Squares and F Test

The F test construction principle says that, to construct a test for an effect of interest:

1. Take the $E(MS)$ for the effect. Examine which component(s) of the $E(MS)$ involve the effect of interest. Under the null hypothesis, these will be zero.
2. Imagine that the null hypothesis is true. This will cause the terms identified in the previous step to drop out. This revised formula is your "Null Effect Formula." It will be the numerator of your F statistic for this effect.
3. Scan the list of $E(MS)$ formulas, and find a formula that is identical to the "Null Effect Formula." This will be the denominator (error) term for your F test.

A One-Way Random Effects ANOVA

Expected Mean Squares and F Test

Below is a table with Expected Mean Squares for 1-Way ANOVA with fixed effects, and with random effects.

In this case, how do we compute the F statistic for the A effect? Let's do it step by step for the random effects model:

	Expected MS	
Factor	Fixed Effects Model	Random Effects Model
A	$\sigma_e^2 + n\theta_A^2$	$\sigma_e^2 + n\sigma_A^2$
S/A	σ_e^2	σ_e^2

1. The MS for your effect (MS_A) will be the numerator of the F statistic. Take the $E(MS)$ for the effect. ($E(MS_A) = \sigma_e^2 + n\sigma_A^2$) Examine which component(s) of the $E(MS)$ involve the effect of interest. ($n\sigma_A^2$) Under the null hypothesis, these will be zero.
2. Imagine that the null hypothesis ($H_0: \ \sigma_A^2 = 0$), is true. This will cause the term identified in the previous step to drop out. This revised formula is your "Null Effect Formula." ($E(MS_A) = \sigma_e^2 + n\sigma_A^2 = \sigma_e^2$)
3. Scan the list of $E(MS)$ formulas, and find a formula that is identical to the "Null Effect Formula." This will be the denominator (error) term for your F test. Since $E(MS_{S/A}) = \sigma_e^2$, $MS_{S/A}$ will be our denominator (error) term.

A One-Way Random Effects ANOVA

Expected Mean Squares and F Test

You can quickly see that the F test for the A effect in the fixed effects model also uses $MS_{S/A}$ as the error term.

Thus, by our F-test construction principle, the F statistic is computed as $MS_A/MS_{S/A}$ in both models.

With the F test computed the same way, one might be tempted to think that it has the same distribution. It does when the null hypothesis is true, but it does not when the null hypothesis is false.

A One-Way Random Effects ANOVA

Expected Mean Squares and F Test

The general distribution of the F-statistic in the one-way random effects model is

$$(1 + \frac{n\sigma_A^2}{\sigma_e^2})F_{a-1,a(n-1)}$$

Note that when H_0 is true, the distribution has a central F distribution identical to the fixed effects model.

However, when H_0 is false, the distribution is not a noncentral F, but rather a constant multiplied by a central F.

Two-Way Model with Both Effects Random

The model for the CRF-*pq* model with both effects random is

$$
\begin{aligned}
Y_{ijk} &= \mu_{jk} + \epsilon_{ijk} && (8)\\
&= \mu + \alpha_j + \beta_k + (\alpha\beta)_{jk} + \epsilon_{ijk} && (9)\\
& && (10)
\end{aligned}
$$

The assumptions are that

$$
\begin{aligned}
\epsilon_{ij} &\sim \text{ i.i.d } N(0, \sigma_e^2) && (11)\\
\alpha_j &\sim \text{ i.i.d } N(0, \sigma_A^2) && (12)\\
\beta_k &\sim \text{ i.i.d } N(0, \sigma_B^2) && (13)\\
(\alpha\beta)_{jk} &\sim \text{ i.i.d } N(0, \sigma_{AB}^2). && (14)
\end{aligned}
$$

Moreover, there is the assumption that all of the above terms are independent of each other.

A Typical Two-Way Model with One Random Effect

Suppose we were interested in the effect of a type of Study Program on student learning, but we were also interested in the effect of the school environment.

We go to a large local school district and select 4 schools at random from a list of potential participating schools.

Next, we sample 10 student volunteers from each school, and randomly assign them to two training methods, "Computer" and "Standard." In this design, Study Program is Factor A, and School is Factor B.

A Typical Two-Way Model with One Random Effect

Factor B

Consider Factor B (School), and how we might model it. Actually, we have a choice of models that we might apply to this situation:

1. One model views the 4 selected schools as the only schools of interest. Since these are the only schools of interest, the effects of school 1, for example, on learning can be viewed as a fixed quantity to be estimated. In this model, the school factor, factor B, is a *fixed effect*.
2. The other model, probably more in line with our substantive goal, views the schools as simply a sample from a larger population of interest. In this model, we can make inferences about the entire population of schools, and school is a *random effect*.

A Typical Two-Way Model with One Random Effect
Factor A

Factor A, in this study, is Study Program.

There are two study programs, and we are interested in comparing these two specific programs.

These programs have not been sampled from some population of interest. Consequently, we treat Factor A as a fixed-effects factor.

A Typical Two-Way Model with One Random Effect

The Basic Model

The model, with Factor A fixed and B random, is, for $1 \leq i \leq n$, $1 \leq j \leq a$, and $1 \leq k \leq b$,

$$\begin{aligned} Y_{ijk} &= \mu_{jk} + \epsilon_{ijk} && (15) \\ &= \mu + \alpha_j + \beta_k + (\alpha\beta)_{jk} + \epsilon_{ijk} && (16) \end{aligned}$$

A Typical Two-Way Model with One Random Effect

The Basic Model

The assumptions are that

$$\epsilon_{ij} \sim \text{i.i.d } N(0, \sigma_e^2) \tag{17}$$

$$\beta_k \sim \text{i.i.d } N(0, \sigma_B^2) \tag{18}$$

$$(\alpha\beta)_{jk} \sim \text{i.i.d } N(0, \frac{a-1}{a}\sigma_{AB}^2). \tag{19}$$

Moreover, there are independence assumptions:

1. $(\alpha\beta)_{jk}$ are independent of the β_k.
2. Different $(\alpha\beta)_{jk}$ in different columns are independent, but will be dependent within column.
3. ϵ_{ijk} are independent of β_k and $(\alpha\beta)_{jk}$

There are also identification restrictions:

1. $\sum_j \alpha_j = 0$
2. $\sum_j (\alpha\beta)_{jk} = 0 \quad \forall k$

When testing the fixed main effect, an additional assumption is that of *sphericity* across the levels of the fixed factor A. The assumption requires all pairwise difference between levels to have equal variance.

A Typical Two-Way Model with One Random Effect

Computations and Expected Mean Squares

Again, we *compute* the sums of squares and mean squares exactly as in the fixed effects case.

However, the expected mean squares and F tests are now not all the same.

Below is a compact table from the textbook by Maxwell and Delaney, showing the Expected Mean Squares for two way factorial models with both effects fixed, both random, and one fixed and one random.

Recall that, although it is not stated in the table, in all models, $E(MS_{S/AB}) = \sigma_e^2$.

Examine the table and see if you can determine how to test the null hypotheses for A, B, and the AB interaction.

The numerators will always be the MS for the effect of interest. The denominator ("error") term for a test can change, depending on the model.

A Typical Two-Way Model with One Random Effect

Computations and Expected Mean Squares

Expected Mean Squares
CRF-ab Design

Effect	*A fixed* *B fixed*	*A fixed* *B random*	*A random* *B fixed*	*A random* *B random*
A	$\sigma_e^2 + bn\theta_A^2$	$\sigma_e^2 + bn\theta_A^2 + n\sigma_{AB}^2$	$\sigma_e^2 + bn\sigma_A^2$	$\sigma_e^2 + bn\sigma_A^2 + n\sigma_{AB}^2$
B	$\sigma_e^2 + an\theta_B^2$	$\sigma_e^2 + an\sigma_B^2$	$\sigma_e^2 + an\theta_B^2 + n\sigma_{AB}^2$	$\sigma_e^2 + an\sigma_B^2 + n\sigma_{AB}^2$
AB	$\sigma_e^2 + n\theta_{AB}^2$	$\sigma_e^2 + n\sigma_{AB}^2$	$\sigma_e^2 + n\sigma_{AB}^2$	$\sigma_e^2 + n\sigma_{AB}^2$
S/AB	σ_e^2	σ_e^2	σ_e^2	σ_e^2

A Typical Two-Way Model with One Random Effect

Computations and Expected Mean Squares

Expected Mean Squares
CRF-ab Design

Effect	*A fixed B fixed*	*A fixed B random*	*A random B fixed*	*A random B random*
A	$\sigma_e^2 + bn\theta_A^2$	$\sigma_e^2 + bn\theta_A^2 + n\sigma_{AB}^2$	$\sigma_e^2 + bn\sigma_A^2$	$\sigma_e^2 + bn\sigma_A^2 + n\sigma_{AB}^2$
B	$\sigma_e^2 + an\theta_B^2$	$\sigma_e^2 + an\sigma_B^2$	$\sigma_e^2 + an\theta_B^2 + n\sigma_{AB}^2$	$\sigma_e^2 + an\sigma_B^2 + n\sigma_{AB}^2$
AB	$\sigma_e^2 + n\theta_{AB}^2$	$\sigma_e^2 + n\sigma_{AB}^2$	$\sigma_e^2 + n\sigma_{AB}^2$	$\sigma_e^2 + n\sigma_{AB}^2$
S/AB	σ_e^2	σ_e^2	σ_e^2	σ_e^2

From the above, we see that in the design with both factors fixed, our test construction principle would lead us to test all effects with $MS_{S/AB}$ as the error term.

or the mixed model with A fixed and B random, the B and AB effects would be tested with $MS_{S/AB}$ as the error term, but the A effect would be tested with MS_{AB} as the error term. We'll work through that in the next slide, in case this is still giving you difficulties. With A random and B fixed, the labels are simply reversed.

or the model with both A and B random, both A and B main effects are tested with MS_{AB} in the denominator error term, while the AB interaction is tested with $MS_{S/AB}$ as the error term.

A Typical Two-Way Model with One Random Effect

Computations and Expected Mean Squares

Expected Mean Squares
CRF-ab Design

Effect	*A fixed B fixed*	*A fixed B random*	*A random B fixed*	*A random B random*
A	$\sigma_e^2 + bn\theta_A^2$	$\sigma_e^2 + bn\theta_A^2 + n\sigma_{AB}^2$	$\sigma_e^2 + bn\sigma_A^2$	$\sigma_e^2 + bn\sigma_A^2 + n\sigma_{AB}^2$
B	$\sigma_e^2 + an\theta_B^2$	$\sigma_e^2 + an\sigma_B^2$	$\sigma_e^2 + an\theta_B^2 + n\sigma_{AB}^2$	$\sigma_e^2 + an\sigma_B^2 + n\sigma_{AB}^2$
AB	$\sigma_e^2 + n\theta_{AB}^2$	$\sigma_e^2 + n\sigma_{AB}^2$	$\sigma_e^2 + n\sigma_{AB}^2$	$\sigma_e^2 + n\sigma_{AB}^2$
S/AB	σ_e^2	σ_e^2	σ_e^2	σ_e^2

Suppose we wish to calculate how to do the F test in the mixed model with factor A fixed and B random. Which mean squares do we use to perform F_A, the F test for the A main effect?

To get the numerator MS, we scan the list of Expected Mean Squares and find that only MS_A has a term involving θ_A^2. So MS_A will be the numerator of our F statistic.

Next, we ask what will happen to $E(MS_A)$ if the null hypothesis is true. We see that, under $H_0 : \theta_A^2 = 0$, we have

$$E(MS_A) = \sigma_e^2 + bn\theta_A^2 + n\sigma_{AB}^2 = \sigma_e^2 + n\sigma_{AB}^2 \qquad (20)$$

Scanning the table of Expected Mean Squares, we see that MS_{AB} has the same $E(MS)$, and so it will serve as the error term for the F test.

Two-Way Mixed Model ANOVA: An Example

Maxwell and Delaney (p.481, Table 10.5) present data for a 2-way ANOVA exploring the effects of *Program* and *School* on ACT *Score*. The data are in a file called *MD10_05.txt*.

Reading in the file, and examining the interaction plot, we see definite signs of a main effect for *Program*.

Two-Way Mixed Model ANOVA: An Example

```
> set.seed(12345)
> two.way <- read.csv("MD10_05.csv")
> attach(two.way)
> interaction.plot(School, Program, Score, col = c("red", "blue"))
```

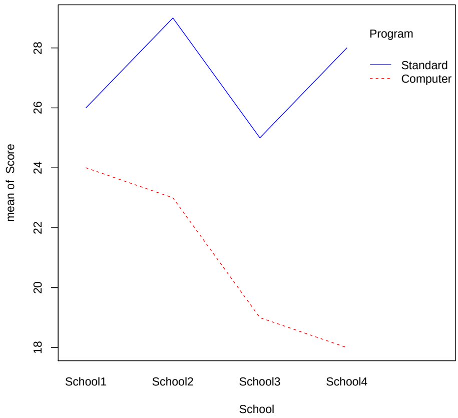

Two-Way Mixed Model ANOVA: An Example

Unfortunately, unlike some commercial software, R does not currently possess a facility whereby one can indicate whether an effect is fixed or random, and have the correct F statistic generated automatically.

Some user intervention is required, although one may easily write a function to automate the process for simple designs such as a two-way.

Since the computations for the sums of squares and mean squares are identical for fixed, random, and mixed effects models, we need only perform the computations the standard way, and then change the F tests only for those effects with a different error term than the standard fixed-effects model.

What that boils down to, for a two-way completely randomized factorial design, is:

1. If the model is mixed, the fixed effect F test is performed using the interaction mean square as the (denominator) error term, and
2. If the model has random effects for *both* main effects, then both main effect F tests are performed using the interaction mean square as the error term.

We demonstrate the calculations for the current example on the next slide. Notice how the code performs the standard fixed-effects ANOVA, then replaces the test for the *Program* factor with the correct F statistic and reconstitutes the table.

Two-Way Mixed Model ANOVA: An Example

```
> fit <- lm(Score ~ Program * School)
> results <- anova(fit)
> Df <- results$Df
> SumSq <- results$"Sum Sq"
> MeanSq <- results$"Mean Sq"
> Fvalue <- results$"F value"
> Pvalue <- results$"Pr(>F)"
> Error.Term <- MeanSq[3]
> df.error <- Df[3]
> Fvalue[1] <- MeanSq[1]/Error.Term
> Pvalue[1] <- 1 - pf(Fvalue[1], Df[1], df.error)
> Ftable <- cbind(Df, SumSq, MeanSq, Fvalue, Pvalue)
> rownames(Ftable) <- c("Program", "School", "Program:School", "Residuals")
```

Two-Way Mixed Model ANOVA: An Example

```
> Ftable

               Df SumSq     MeanSq    Fvalue     Pvalue
Program         1   360 360.00000 13.500000 0.03489698
School          3   100  33.33333  1.845444 0.15880542
Program:School  3    80  26.66667  1.476355 0.23951513
Residuals      32   578  18.06250        NA         NA
```