

# Lab Tutorial 8: Random Effects & Repeated-Measures ANOVA

Chris M. Fiacconi

In all designs covered so far in this course, we have examined *fixed-effect* factors. The levels of a factor are said to be *fixed* when subsequent replications of the study would include the same fixed levels. If I were conducting a drug study with three different levels corresponding to different dose amounts (e.g., 100mg, 200mg, 300mg), dose would be a fixed factor if I were only interested in how those particular dose amounts affected behaviour. So if I were to run a replication, I would use those same levels.

## Random Factors

However, some factors in an experiment are better thought of as *random-effect* factors, meaning that upon subsequent replications, different levels of the factor would be chosen. The levels of such factors can be thought of as representing a sample from a larger population of possible levels. For example, if I wanted to test whether a computer-based training program or a paper-and-pencil package of materials better prepared high school students for college, I could randomly select 4 different high schools from the region and assign students from each school to either the computer or paper-and-pencil training programs and then examine ACT scores for both programs. In this case, the Program factor would be *fixed* because I'm interested in whether there is a difference between those two particular training methods. However, we could think of the particular high schools chosen as a *random* factor because if I were to run the study again, I might select different high schools. Designs with both a fixed and a random factor are sometimes known as *mixed-designs*.

Analysis of mixed-designs is similar to that of designs consisting of only fixed factors, but there are some important differences in the way we conceptualize random factors, and in how we evaluate the associated null hypothesis, which itself is slightly different. Recall that the null hypothesis for a fixed factor is that all  $\alpha_j$ 's = 0, and that the alternative hypothesis is that at least one  $\alpha_j \neq 0$ , where  $\alpha_j$  represents the effect associated with being in level  $j$ . However, with random factors, we're not so much interested in whether there is a difference among the particular levels chosen (e.g., school 1 vs. school 2), but rather, whether in general there are differences among all possible levels. The notation for the null and alternative hypotheses is thus as follows:

$$H_0 : \sigma_\alpha^2 = 0$$
$$H_a : \sigma_\alpha^2 > 0$$

Note that for random factors, this change in notation means that our results can generalize to all possible levels of the factor, even those that we did not choose in this particular study. In contrast, for fixed factors, our results allow us to draw conclusions that are restricted to only those levels that were selected in the study. Going back to our example, if Program is treated as a fixed factor, any difference that we observe between the two programs would not allow us to draw conclusions to other training programs that were not directly investigated in this study. However, because School is treated as a random factor and the schools themselves were selected randomly, any systematic differences in ACT scores between these schools are indicative of the differences that would be seen if a different sample of schools were chosen.

## Model Comparisons for Mixed-Designs

Since designs that include only a single random factor are seldom used in the behavioural sciences (and as it turns out, are analyzed exactly the same way as a single fixed factor design), we will not cover the analysis

of such designs in this course. Mixed-designs are, however, fairly common and so I will demonstrate how to analyze such a design using the example described above (see MDK Table 10.5).

Before we begin, it is important to understand that the effect of the fixed-factor in a mixed-design is evaluated using a different error term than if the design consisted of two fixed-factors (see Lab Tutorial 6 on factorial ANOVA). In all designs we've covered so far, we've been using  $MS_W$ , as our error term for each effect in the model. However, in a mixed-design with Factor A fixed and Factor B random, the appropriate error term for evaluating the effect of the fixed-factor is **not**  $MS_W$ . Rather, the appropriate error term in this case is actually  $MS_{AB}$ , or the A  $\times$  B interaction variance. Because B is a random factor, the A  $\times$  B interaction variance can artificially inflate  $MS_A$ , making it easier to reject the null hypothesis that all  $\alpha_j$ 's = 0.

The models being compared to evaluate the effect of a fixed factor in a mixed-design look very similar to those we've seen before, with the following full model

$$Y_{ij} = \mu + \alpha_j + \beta_k + (\alpha\beta)_{jk} + \epsilon_{ij}$$

being compared with the following restricted model

$$Y_{ij} = \mu + \beta_k + (\alpha\beta)_{jk} + \epsilon_{ij}$$

You might think that the appropriate method to calculate the  $F$ -statistic evaluating the main effect of fixed-factor A would be:

$$F = \frac{(E_R - E_F)/(df_R - df_F)}{E_F/df_F} = \frac{MS_A}{MS_W}$$

but it is not. In this case,  $MS_A$  has the following expected mean-square value:

$$(E_R - E_F)/(df_R - df_F) = MS_A = \sigma_\epsilon^2 + n\sigma_{(\alpha\beta)}^2 + bn \sum_j \alpha_j^2/(a-1)$$

whereas the expected mean-square value for  $MS_W$  is:

$$E_F/df_F = MS_W = \sigma_\epsilon^2$$

.

Therefore, the resulting  $F$ -statistic for evaluating the main effect of fixed-factor A will be inflated because the numerator incorporates variability not only from  $bn \sum_j \alpha_j^2/(a-1)$  and  $\sigma_\epsilon^2$ , but also from  $n\sigma_{(\alpha\beta)}^2$ . This is because when a random factor consists of only a sample of the population of levels, the  $(\alpha\beta)_{jk}$  term will sum-to-zero *only across the levels of the fixed factor*. This means that  $MS_A$  will be influenced by interaction variance as  $MS_A$  and  $MS_{AB}$  are no longer orthogonal.

In other words, this  $F$ -test is biased and will tend to inflate the Type I Error rate. To properly evaluate the main effect of fixed-factor A, we need to choose a denominator that differs in only one variance component. The correct choice is  $MS_{A \times B}$ , which has the following expected mean-square value:

$$MS_{A \times B} = \sigma_\epsilon^2 + n\sigma_{(\alpha\beta)}^2$$

.

Therefore, the appropriate  $F$ -statistic is now:

$$F = \frac{(E_R - E_F)/(df_R - df_F)}{MS_{AB}} = \frac{MS_A}{MS_{A \times B}}$$

Despite this change in the composition of the  $F$ -statistic for the main effect of fixed-factor A, the main effect of random-factor B and the A  $\times$  B interaction are assessed using the familiar  $MS_W$  error term.

For the main effect of random-factor B, we would compare the following full and restricted models:

$$Y_{ij} = \mu + \alpha_j + \beta_k + (\alpha\beta)_{jk} + \epsilon_{ij}$$

$$Y_{ij} = \mu + \alpha_j + (\alpha\beta)_{jk} + \epsilon_{ij}$$

The appropriate  $F$ -statistic to evaluate the main effect of random-factor B is then:

$$F = \frac{(E_R - E_F)/(df_R - df_F)}{E_F/df_F} = \frac{MS_B}{MS_W}$$

Finally, for the the  $A \times B$  interaction, we would compare the following full and restricted models:

$$Y_{ij} = \mu + \alpha_j + \beta_k + (\alpha\beta)_{jk} + \epsilon_{ij}$$

$$Y_{ij} = \mu + \alpha_j + \beta_k + \epsilon_{ij}$$

The appropriate  $F$ -statistic to evaluate the  $A \times B$  interaction is then:

$$F = \frac{(E_R - E_F)/(df_R - df_F)}{E_F/df_F} = \frac{MS_{A \times B}}{MS_W}$$

## Analyzing Mixed-Designs in R

### Example #1

Let's now conduct the appropriate analyses for the example I described above using *R*. This example is taken from your textbook (MDK Table 10.5).

```
library(tidyverse)
act.data<-read.table(file="CH10T5.txt",header=T) # Load in data file
str(act.data)

## 'data.frame':  40 obs. of  3 variables:
## $ Program: Factor w/ 2 levels "Computer","Pencil": 1 1 1 1 1 1 1 1 1 1 ...
## $ School : int  1 1 1 1 1 2 2 2 2 2 ...
## $ ACT    : int  28 26 26 23 17 26 26 25 21 17 ...

act.data$School<-as.factor(act.data$School) # Make sure School is treated as a factor

# Get means for each condition
cell.means<-act.data %>% group_by(Program,School) %>% summarize(Mean=mean(ACT))

# Transform cell.means dataframe into matrix for plotting
cell.means<-xtabs(Mean~Program+School,data=cell.means)
print(cell.means)

##           School
## Program    1  2  3  4
## Computer  24 23 19 18
## Pencil    26 29 25 28

# Make a bar graph
means.bp<-barplot(cell.means,beside=TRUE,las=1,ylab="ACT Scores",xlab="High School",
  ylim=c(0,40),legend.text=c("Computer","Pencil"))

print(means.bp) # Print bar graph and x-coordinates of bars
```




Figure 1: Plot of cell means from data taken from MDK Chapter 10 Table 5

```
##      [,1] [,2] [,3] [,4]
## [1,]  1.5  4.5  7.5 10.5
## [2,]  2.5  5.5  8.5 11.5
```

```
# Compute +/- SEM Error Bars
all.SEM<-act.data %>% group_by(Program,School) %>%
  summarize(SEM=sd(ACT)/sqrt(length(ACT))) # Calculate SEM for each condition

SEM.matrix<-xtabs(SEM~Program+School,data=all.SEM) # Transform SEM data.frame into matrix for plotting

# Draw Error Bars
arrows(x0=means.bp,y0=cell.means,x1=means.bp,y1=cell.means+SEM.matrix,
       angle=90,length=.05,lwd=2) # Draw upper half

arrows(x0=means.bp,y0=cell.means,x1=means.bp,y1=cell.means-SEM.matrix,
       angle=90,length=.05,lwd=2) # Draw lower half
```

If we were to (incorrectly I might add) treat School as a fixed-factor, then we could simply go ahead and conduct a 2 X 4 factorial ANOVA as we did in Lab Tutorial 6:

```
options(contrasts=c("contr.sum","contr.poly")) # Set sum-to-zero contrasts

school.fixed.aov<-aov(ACT~Program*School,data=act.data)
print(summary(school.fixed.aov))
```

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

```
## Program          1    360    360.0  19.931  9.35e-05 ***
## School           3    100     33.3   1.845   0.159
## Program:School   3     80     26.7   1.476   0.240
## Residuals       32    578     18.1
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

From these results it appears as though there is a highly significant effect of Program, with the pencil-and-paper students outperforming the computer-based training students.

Note however, that this analysis uses  $MS_W$  to evaluate the main effect of Program, which is inappropriate because School is a random-factor, not a fixed-factor. Instead, we should be using  $MS_{AB}$  in the denominator of our  $F$ -statistic when evaluating the main effect of Program:

```
MS.program<-360
MS.interaction<-26.7

program.F.correct<-MS.program/MS.interaction
print(program.F.correct)
```

```
## [1] 13.48315
```

```
p.value<-1-pf(program.F.correct,df1=1,df2=3)
print(p.value)
```

```
## [1] 0.03495279
```

So after treating School as a random effect and re-calculating  $F$ , we see that there is still a significant effect of Program,  $F(1, 3) = 13.48, p = .035$ .

From comparing the first analysis with the second, it is clear that by using  $MS_W$  as our error term when evaluating the effect of Program, there was an obvious inflation of our  $F$ -statistic, yielding a  $p$ -value that was far smaller than it should have been.

Another point to note here is that, because there were only 4 schools involved in this study, there were only 3 degrees of freedom in the denominator. In general, so few degrees of freedom will reduce our power to detect a true effect. So, in this case, we could substantially increase our power by selecting a larger number of schools to participate in this study.

## One-Way Repeated-Measures ANOVA

Our preceding discussion regarding random factors is important when we consider within-subject or *repeated-measures* designs in which each subject in our study receives all levels of the experimental manipulation. Indeed, a standard one-way repeated-measures design can be thought of as a particular kind of mixed-design, with the primary factor of interest being fixed, and subjects as a random factor. Importantly, testing each subject in all conditions means that the Subject factor is *crossed* with the primary manipulation of interest, allowing us to measure the variation in our dependent measure that can be attributed to our fixed-factor, and the random-factor Subject. Note that it is impossible to disentangle these different sources of variability in a one-way fixed-factor between subjects design because each subject is *nested* within a particular condition. Therefore, we have no way of isolating variability that is due to individual subject-level differences vs. random error variability.

Let's formalize this idea by constructing the models that are compared in a one-way repeated-measures design. You can conceive of the full model as:

$$Y_{ij} = \mu + \alpha_j + \pi_i + (\alpha\pi)_{ij} + \epsilon_{ij}$$

where  $\mu$  is the grand mean of all scores on the DV,  $\alpha_j$  is the effect of being in condition  $j$ ,  $\pi_i$  is the effect associated with the  $i$ th subject (or participant),  $(\alpha\pi)_{ij}$  is the interaction between the  $i$ th subject in condition  $j$ , and  $\epsilon_{ij}$  is the prediction error associated with subject  $i$  in condition  $j$ .

However, before we proceed it is necessary to recognize that because each subject contributes only one mean to each condition in a repeated-measures design, it is not possible disentangle the variability in our dependent measure that is associated with the interaction between the fixed-factor and Subjects from random error variability. In other words, the interaction variance is the error variance in such a design, and as a consequence, we can remove the  $(\alpha\pi)_{ij}$  parameters from our model as they get incorporated into the error term (i.e., residuals). Therefore, our new full model becomes:

$$Y_{ij} = \mu + \alpha_j + \pi_i + \epsilon_{ij}$$

where each term is calculated as follows:

$$\begin{aligned}\mu &= \bar{Y}_{..} \\ \alpha_j &= \bar{Y}_{.j} - \bar{Y}_{..} \\ \pi_i &= \bar{Y}_{i.} - \bar{Y}_{..} \\ \epsilon_{ij} &= Y_{ij} - (\bar{Y}_{.j} + \bar{Y}_{i.} - \bar{Y}_{..})\end{aligned}$$

The null hypothesis tested here is the same as we've seen before for a fixed-factor A:

$$H_0 : \alpha_1 = \alpha_2 = \dots = \alpha_j = 0$$

Similarly, the alternative hypothesis is that at least one  $\alpha_j \neq 0$ .

To test this null hypothesis, the full model would be compared to the following restricted model:

$$Y_{ij} = \mu + \pi_i + \epsilon_{ij}$$

using the following  $F$ -statistic:

$$F = \frac{(E_R - E_F)/(df_R - df_F)}{E_F/df_F}$$

where

$$\begin{aligned}(df_R - df_F) &= a - 1 \\ df_F &= (n - 1)(a - 1)\end{aligned}$$

Note that according to our previous discussion on evaluating the effect of a fixed-factor when a random-factor is included in the same design, we could also express the  $F$ -statistic for the effect of the fixed-factor A as:

$$F = \frac{MS_A}{MS_{A \times S}}$$

Remember that the Subject factor is treated as a random-factor here because we randomly selected only a subset of the full population to participate in the study. If we were to replicate this study, we would not select the same subjects.

You might now be wondering how we test for a significant effect of Subject. It turns out that because each subject contributes only one mean to each condition, such a test is not possible because there is no appropriate error term to test for this effect. This isn't really to much of a problem because we're not usually interested in knowing that there are in fact overall differences between individuals.




Figure 2: Boxplot of data take from MDK Chapter 11 Table 5

## Example #2

Let's illustrate how to conduct a one-way repeated-measures ANOVA in *R* using the data from MDK Table 11.5. In this example, 12 children are measured at 4 different ages (30-, 36-, 42-, 48-months) on their cognitive abilities. The researcher wants to examine whether their cognitive performance changes with age.

```
age.data<-read.table(file="Ch11T5.txt",header=T)
str(age.data)
```

```
## 'data.frame': 48 obs. of 3 variables:
## $ Age : Factor w/ 4 levels "Months30","Months36",...: 1 1 1 1 1 1 1 1 1 1 ...
## $ Score : int 108 103 96 84 118 110 129 90 84 96 ...
## $ Subject: int 1 2 3 4 5 6 7 8 9 10 ...
```

```
# Boxplot of age.data
boxplot(Score~Age,data=age.data,las=1,ylab="Cognitive Ability",xlab="Age (months)")
```

```
# Get descriptives for each Age group
age.descriptives<-age.data %>% group_by(Age) %>% summarize(Mean=mean(Score),S.Dev=sd(Score))
print(age.descriptives)
```

```
## # A tibble: 4 x 3
## Age Mean S.Dev
## <fct> <dbl> <dbl>
```

```
## 1 Months30 103 13.7
## 2 Months36 107 14.2
## 3 Months42 110 13.3
## 4 Months48 112 14.8
```

Let's now perform a repeated-measures ANOVA in *R* using the basic `aov` command:

```
age.data$Subject<-as.factor(age.data$Subject) # Ensure Subject is treated as a factor!
age.aov<-aov(Score~Age+Error(Subject/Age),data=age.data) # Perform RM-ANOVA
print(summary(age.aov))
```

```
##
## Error: Subject
##           Df Sum Sq Mean Sq F value Pr(>F)
## Residuals 11  6624   602.2
##
## Error: Subject:Age
##           Df Sum Sq Mean Sq F value Pr(>F)
## Age         3    552   184.00  3.027 0.0432 *
## Residuals 33   2006    60.79
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

There appears to be a significant effect of Age,  $F(3, 33) = 3.03, p = .043$ . Notice that within the `aov` command I specified the appropriate error term,  $MS_{A \times S}$ , using the *Error* argument. This tells *R* that we're performing a repeated-measures ANOVA, and that Age is manipulated within-subjects.

Note that although we can't test whether there is a significant effect of Subject, the sums-of-squares for Subject ( $SS_{Subject} = 6624$ ) can still be calculated and are provided in the top portion of the output from the `aov` command.

## Assumption of Sphericity

In order for the previous analysis to be correct, we have to assume that our data conform to the sphericity assumption. I won't go into too much detail here, but I want to explain the concept of sphericity, how it arises in repeated-measures designs with three or more levels, and what to do if the assumption of sphericity is violated.

Because we are measuring the *same* subject in multiple conditions, the scores are often correlated across conditions. In order for RM-ANOVA to produce valid results, the variance of the difference scores for all pairs of conditions must be equal. This assumption, known as the *sphericity assumption*, will necessarily be true when the variances across conditions are equal, and the covariances between all pairs of conditions are equal. When the data conform to these criteria, the sphericity assumption is upheld. **IMPORTANT: When there are only two conditions, we don't have to worry, as only one set of difference scores can be computed and therefore sphericity cannot be violated.** When there are three or more conditions, however, it is possible that the variances of the pairwise difference scores between conditions are not equal and violations of the sphericity assumption can occur.

Fortunately, we can explicitly test whether the sphericity assumption holds using *Mauchley's Test for Sphericity*. This test is automatically included in the output from the `ezANOVA` function, which can be found in the package *ez*. If this test is significant, then the sphericity assumption has been violated, and we need to introduce corrections to account for this violation (some statisticians argue you should always perform the corrections regardless of whether Mauchley's test is significant).



There are two different correction procedures, with one (Greenhouse-Geisser, or  $\hat{\epsilon}$ ) being slightly more conservative than the other (Huynh-Feldt, or  $\tilde{\epsilon}$ ). Luckily, each correction factor along with its associated  $p$ -adjustment is also included in the output of the **ezANOVA** command.

Let's redo the last analysis using the **ezANOVA** command and set the *detailed* argument to TRUE so that we can evaluate whether the sphericity assumption is violated and obtain all  $SS$  values, the appropriate correction factors ( $\hat{\epsilon}$  and  $\tilde{\epsilon}$ ) and associated  $p$ -values.

```
library(ez) # Ensure 'ez' is installed first - install.packages("ez") (Make sure you're online)

age.ezAOV<-ezANOVA(data=age.data,dv=Score,wid=Subject,within=Age,between=NULL,detailed=TRUE)
print(age.ezAOV)

## $ANOVA
##      Effect DFn DFd      SSn  SSd          F          p p<.05      ges
## 1 (Intercept)   1  11 559872 6624 929.739130 5.586316e-12 * 0.98481975
## 2      Age      3  33   552 2006   3.026919 4.321863e-02 * 0.06011762
##
## $`Mauchly's Test for Sphericity`
##      Effect      W          p p<.05
## 2      Age 0.2426472 0.01771762 *
##
## $`Sphericity Corrections`
##      Effect      GGe      p[GG] p[GG]<.05      HFe      p[HF] p[HF]<.05
## 2      Age 0.6095445 0.0747874          0.7248502 0.06353773
```

In this case Mauchley's test is significant ( $p = .018$ ). As a consequence, we need use either the  $\hat{\epsilon}$  or  $\tilde{\epsilon}$  correction factors to adjust the  $p$ -value associated with the effect of Age. My own preference is to use the Huynh-Feldt (HF;  $\tilde{\epsilon}$ ) correction when this assumption is violated. Notably, the adjusted  $p$ -value associated with the  $\tilde{\epsilon}$  correction is .064, and the effect of Age is no longer significant after adjusting for the apparent violation of sphericity. This example illustrates that violations of sphericity left uncorrected can yield biased results that increase the probability of a Type I Error.

## Effect Size

There are still a couple of things we need to figure out. Although our obtained adjusted  $p$ -value does not allow us to reject the null hypothesis that the condition means for Age are equal, we can nonetheless still calculate the effect size for Age. In other words, we would like to know "how much of an effect did Age have?" For a one-way RM-ANOVA, we can again calculate eta-squared ( $\eta^2$ ) with formula:

$$\eta^2 = \frac{SS_A}{SS_A + SS_{Subject} + SS_{AS}}$$

```
SS.age<-age.ezAOV$ANOVA$SSn[2] # Get SS Age from ezANOVA
SS.subject<-6624 # Get SS Subject from original ANOVA (aov)
SS.interaction<-age.ezAOV$ANOVA$SSd[2] # Get SS A x S from ezANOVA

# Calculate eta-squared
eta.sq<-SS.age/(SS.age+SS.subject+SS.interaction)
print(round(eta.sq,digits=2))
```

```
## [1] 0.06
```

The obtained  $\eta^2$  value associated with Age is 0.06, which means that Age accounts for  $\sim 6\%$  of the total variability among exam scores. This would be considered a relatively small effect. Also note that the denominator of  $\eta^2$  contains variability due to overall subject-level differences ( $SS_{Subject}$ ) even though this variability was *not* included in the  $F$ -statistic used to evaluate the effect of Age. When we talk about effect size in a one-way repeated measures context, however, it is usually most informative to express the amount of variability associated with our manipulation (i.e., Age) as a proportion of the total variability in our dependent measure. Therefore,  $SS_{Subject}$  is included in the denominator when calculating  $\eta^2$ .

As in the case of a between-subjects ANOVA, we can use this  $\eta^2$  value to calculate Cohen's  $f$ , which will be useful for power and sample size calculations:

$$f = \sqrt{\frac{\eta^2}{1 - \eta^2}}$$

```
cohens.f<-sqrt(eta.sq/(1-eta.sq))
print(round(cohens.f,digits=2))
```

```
## [1] 0.25
```

## Comparisons Among Individual Means

As we've discussed throughout the course, a significant omnibus one-way ANOVA indicates that there likely exists a difference among the condition means included in your study, but does not indicate specifically which such means are different from each other. Therefore, it is necessary to perform follow-up tests to directly compare the individual means.

As we've done previously, we can use linear contrasts here to test hypotheses about patterns of condition means, either with planned contrasts or post-hoc contrasts. You could also perform pairwise comparisons here, but it is important that you use a separate error term for each comparison rather than the pooled  $MS_{A \times S}$ . Such comparisons could be performed using the **pairwise.t.test** function.

In the present example, given that cognitive ability could reasonably be expected to increase from 30-48 months of age, one might plan in advance to test a contrast that examines whether there exists a linear trend among condition means using the contrast weights (-1.5, -0.5, 0.5, 1.5). The procedure for performing a contrast for a repeated-measures factor is slightly different from that used for between-subject factors.

1) We need to transform our data from long to wide format using the **spread** command from the *tidyverse* package.

```
age.data.wide<-spread(age.data,key=Age,value=Score) # Convert dataframe from long to wide format
age.data.wide<-age.data.wide[,-1] # Remove Subject column
age.data.wide<-as.matrix(age.data.wide) # Convert from dataframe to matrix
```

2) We need to compute a composite score for each subject that reflects the weighted sum of the condition means for individual  $i$ :

$$\psi_i = \sum_{j=1}^a c_j \bar{Y}_{ij}$$

This computation is done using matrix multiplication with the `%*%` operator:

```
c.weights<-c(-1.5,-0.5,0.5,1.5)
comp.scores<-age.data.wide%*%c.weights # Note that the data matrix goes first here!!
print(comp.scores)
```

```
##      [,1]
## [1,]  28
## [2,]  50
## [3,]  16
## [4,]  26
## [5,]  -3
## [6,] -34
## [7,]  -4
## [8,]  43
## [9,]   4
## [10,] 15
## [11,]  6
## [12,] 33
```

3) We can test whether these composite scores differ from 0 using a one-sample *t*-test:

```
t.test(comp.scores)
```

```
##
## One Sample t-test
##
## data:  comp.scores
## t = 2.2414, df = 11, p-value = 0.04659
## alternative hypothesis: true mean is not equal to 0
## 95 percent confidence interval:
##  0.2701827 29.7298173
## sample estimates:
## mean of x
##      15
```

The results suggest that there is a significant linear trend,  $t(11) = 2.24, p = .047$ , with cognitive ability increasing with age from 30-48 months.

Since this was a planned contrast, no adjustments for Type I Error inflation are necessary here.