

8.3: Calculating the RM ANOVA

Now that you are familiar with the concept of an ANOVA table (remember the table from last chapter where we reported all of the parts to calculate the F -value?), we can take a look at the things we need to find out to make the ANOVA table. The figure below presents an abstract for the repeated-measures ANOVA table. It shows us all the thing we need to calculate to get the F -value for our data.

	df	SS	MSE	F	P
EFFECT	df1 effect = conditions-1	SS Effect= SS Total - SS Error (within-conditions)	MSE Effect = $\frac{SS \text{ Effect}}{df1 \text{ Effect}}$	F = $\frac{MSE \text{ Effect}}{MSE \text{ Error}}$	p = From Sampling distribution of F(df1,df2)
ERROR	df2 Effect = (n-1) x (conditions-1) n=# of subjects	SS Error (left-over)= SS Error (within-conditions) - SS Subjects	MSE Error (left-over) = $\frac{SS \text{ Error}}{df2 \text{ Error}}$		

Figure 8.3.1: Equations for computing the ANOVA table for a repeated measures design.

So, what we need to do is calculate all the SS s that we did before for the between-subjects ANOVA. That means the next three steps are identical to the ones you did before. In fact, I will just basically copy the next three steps to find SS_{TOTAL} , SS_{Effect} , and $SS_{Error (within-conditions)}$. After that we will talk about splitting up $SS_{Error (within-conditions)}$ into two parts, this is the new thing for this chapter. Here we go!

SS Total

The total sums of squares, or SS_{Total} measures the total variation in a set of data. All we do is find the difference between each score and the grand mean, then we square the differences and add them all up.

```
suppressPackageStartupMessages(library(dplyr))
scores <- c(20,11,2,6,2,7,2,11,2)
conditions <- as.character(rep(c("A","B","C"), each=3))
subjects <- rep(1:3,3)
diff <- scores - mean(scores)
diff_squared <- diff^2
df <- data.frame(subjects, conditions, scores, diff, diff_squared)
df$conditions <- as.character(df$conditions)
df$subjects <- as.character(df$subjects)
```

```
df <- df %>%
  rbind(c("Sums", "", colSums(df[1:9,3:5]))) %>%
  rbind(c("Means", "", colMeans(df[1:9,3:5])))
knitr::kable(df)
```

run restart restart & run all

subjects	conditions	scores	diff	diff_squared
1	A	20	13	169
2	A	11	4	16
3	A	2	-5	25
1	B	6	-1	1
2	B	2	-5	25
3	B	7	0	0
1	C	2	-5	25
2	C	11	4	16
3	C	2	-5	25
Sums		63	0	302
Means		7	0	33.5555555555556

The mean of all of the scores is called the Grand Mean. It's calculated in the table, the Grand Mean = 7.

We also calculated all of the difference scores from the Grand Mean. The difference scores are in the column titled `diff`. Next, we squared the difference scores, and those are in the next column called `diff_squared`.

When you add up all of the individual squared deviations (difference scores) you get the sums of squares. That's why it's called the sums of squares (SS).

Now, we have the first part of our answer:

$$SS_{\text{total}} = SS_{\text{Effect}} + SS_{\text{Error}}$$

$$SS_{\text{total}} = 302$$

and

$$302 = SS_{\text{Effect}} + SS_{\text{Error}}$$

SS Effect

SS_{Total} gave us a number representing all of the change in our data, how they all are different from the grand mean.

What we want to do next is estimate how much of the total change in the data might be due to the experimental manipulation. For example, if we ran an experiment that causes change in the measurement, then the means for each group will be different from other, and the scores in each group will be different from each. As a result, the manipulation forces change onto the numbers, and this will naturally mean that some part of the total variation in the numbers is caused by the manipulation.

The way to isolate the variation due to the manipulation (also called effect) is to look at the means in each group, and the calculate the difference scores between each group mean and the grand mean, and then the squared deviations to find the sum for SS_{Effect} .

Consider this table, showing the calculations for SS_{Effect} .

```
suppressPackageStartupMessages(library(dplyr))
scores <- c(20,11,2,6,2,7,2,11,2)
conditions <- as.character(rep(c("A","B","C"), each=3))
subjects <- rep(1:3,3)
```

```
means <-c(11,11,11,5,5,5,5,5,5)
diff <-means-mean(scores)
diff_squared <-diff^2
df<-data.frame(subjects,conditions,scores,means,diff, diff_squared)
df$conditions<-as.character(df$conditions)
df$subjects<-as.character(df$subjects)
df <- df %>%
  rbind(c("Sums", "", colSums(df[1:9,3:6]))) %>%
  rbind(c("Means", "", colMeans(df[1:9,3:6])))
knitr::kable(df)
```

run

restart

restart & run all

subjects	conditions	scores	means	diff	diff_squared
1	A	20	11	4	16
2	A	11	11	4	16
3	A	2	11	4	16
1	B	6	5	-2	4
2	B	2	5	-2	4
3	B	7	5	-2	4
1	C	2	5	-2	4
2	C	11	5	-2	4
3	C	2	5	-2	4
Sums		63	63	0	72
Means		7	7	0	8

Notice we created a new column called `means`, these are the means for each condition, A, B, and C.

SS_{Effect} represents the amount of variation that is caused by differences between the means. The `diff` column is the difference between each condition mean and the grand mean, so for the first row, we have $11 - 7 = 4$, and so on.

We found that $SS_{\text{Effect}} = 72$, this is the same as the ANOVA from the previous chapter

SS Error (within-conditions)

Great, we made it to SS Error. We already found SS Total, and SS Effect, so now we can solve for SS Error just like this:

$$SS_{\text{total}} = SS_{\text{Effect}} + SS_{\text{Error (within-conditions)}}$$

switching around:

$$SS_{\text{Error}} = SS_{\text{total}} - SS_{\text{Effect}}$$

$$SS_{\text{Error (within conditions)}} = 302 - 72 = 230$$

Or, we could compute $SS_{\text{Error (within conditions)}}$ directly from the data as we did last time:

```
suppressPackageStartupMessages(library(dplyr))
scores <- c(20,11,2,6,2,7,2,11,2)
conditions <- as.character(rep(c("A","B","C"), each=3))
subjects <- rep(1:3,3)
means <-c(11,11,11,5,5,5,5,5,5)
diff <-means-scores
```

```
diff_squared <- diff^2
df<-data.frame(subjects,conditions,scores,means,diff, diff_squared)
df$conditions<-as.character(df$conditions)
df$subjects<-as.character(df$subjects)
df <- df %>%
  rbind(c("Sums", "", colSums(df[1:9,3:6]))) %>%
  rbind(c("Means", "", colMeans(df[1:9,3:6])))
knitr::kable(df)
```

run restart restart & run all

subjects	conditions	scores	means	diff	diff_squared
1	A	20	11	-9	81
2	A	11	11	0	0
3	A	2	11	9	81
1	B	6	5	-1	1
2	B	2	5	3	9
3	B	7	5	-2	4
1	C	2	5	3	9
2	C	11	5	-6	36
3	C	2	5	3	9
Sums		63	63	0	230
Means		7	7	0	25.5555555555556

When we compute $SS_{\text{Error (within conditions)}}$ directly, we find the difference between each score and the condition mean for that score. This gives us the remaining error variation around the condition mean, that the condition mean does not explain.

SS Subjects

Now we are ready to calculate new partition, called SS_{Subjects} . We first find the means for each subject. For subject 1, this is the mean of their scores across Conditions A, B, and C. The mean for subject 1 is 9.33 (repeating). Notice there is going to be some rounding error here, that's OK for now.

The `means` column now shows all of the subject means. We then find the difference between each subject mean and the grand mean. These deviations are shown in the `diff` column. Then we square the deviations, and sum them up.

```
suppressPackageStartupMessages(library(dplyr))
scores <- c(20,11,2,6,2,7,2,11,2)
conditions <- as.character(rep(c("A","B","C"), each=3))
subjects <- rep(1:3,3)
means <-c(9.33,8,3.66,9.33,8,3.66,9.33,8,3.66)
diff <-means-mean(scores)
diff_squared <-diff^2
df<-data.frame(subjects,conditions,scores,means,diff, diff_squared)
df$conditions<-as.character(df$conditions)
df$subjects<-as.character(df$subjects)
df <- df %>%
  rbind(c("Sums", "", colSums(df[1:9,3:6]))) %>%
  rbind(c("Means", "", colMeans(df[1:9,3:6])))
```

```
knitr::kable(df)
```

[run](#)
[restart](#)
[restart & run all](#)

subjects	conditions	scores	means	diff	diff_squared
1	A	20	9.33	2.33	5.4289
2	A	11	8	1	1
3	A	2	3.66	-3.34	11.1556
1	B	6	9.33	2.33	5.4289
2	B	2	8	1	1
3	B	7	3.66	-3.34	11.1556
1	C	2	9.33	2.33	5.4289
2	C	11	8	1	1
3	C	2	3.66	-3.34	11.1556
Sums		63	62.97	-0.029999999999999994	52.7535
Means		7	6.996666666666667	-0.0033333333333333326	5.8615

We found that the sum of the squared deviations $SS_{\text{Subjects}} = 52.75$. Note again, this has some small rounding error because some of the subject means had repeating decimal places, and did not divide evenly.

We can see the effect of the rounding error if we look at the sum and mean in the `diff` column. We know these should be both zero, because the Grand mean is the balancing point in the data. The sum and mean are both very close to zero, but they are not zero because of rounding error.

SS Error (left-over)

Now we can do the last thing. Remember we wanted to split up the $SS_{\text{Error (within conditions)}}$ into two parts, SS_{Subjects} and $SS_{\text{Error (left-over)}}$. Because we have already calculate $SS_{\text{Error (within conditions)}}$ and SS_{Subjects} , we can solve for $SS_{\text{Error (left-over)}}$:

$$SS_{\text{Error (left-over)}} = SS_{\text{Error (within conditions)}} - SS_{\text{Subjects}}$$

$$SS_{\text{Error (left-over)}} = SS_{\text{Error (within conditions)}} - SS_{\text{Subjects}} = 230 - 52.75 = 177.25$$

Check our work

Before we continue to compute the MSEs and F-value for our data, let's quickly check our work. For example, we could have R compute the repeated measures ANOVA for us, and then we could look at the ANOVA table and see if we are on the right track so far.

```
suppressPackageStartupMessages(library(dplyr))
library(xtable)
scores <- c(20,11,2,6,2,7,2,11,2)
conditions <- as.character(rep(c("A","B","C"), each=3))
subjects <- rep(1:3,3)
means <- c(9.33,8,3.66,9.33,8,3.66,9.33,8,3.66)
diff <- means - mean(scores)
diff_squared <- diff^2
df <- data.frame(subjects, conditions, scores, means, diff, diff_squared)
df$conditions <- as.character(df$conditions)
```

```
df$subjects<-as.character(df$subjects)
```

```
summary_out <- summary(aov(scores~conditions + Error(subjects/conditions),df[1:9,]))
knitr::kable(xtable(summary_out))
```

[run](#)
[restart](#)
[restart & run all](#)

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
Residuals	2	52.66667	26.33333	NA	F)" style="vertical-align:middle;" class="lt-stats-7935">NA
conditions	2	72.00000	36.00000	0.8120301	F)" style="vertical-align:middle;" class="lt-stats-7935">0.505848
Residuals	4	177.33333	44.33333	NA	F)" style="vertical-align:middle;" class="lt-stats-7935">NA

OK, looks good. We found the SS_{Effect} to be 72, and the SS for the conditions (same thing) in the table is also 72. We found the SS_{Subjects} to be 52.75, and the SS for the first residual (same thing) in the table is also 53.66 repeating. That's close, and our number is off because of rounding error. Finally, we found the $SS_{\text{Error (left-over)}}$ to be 177.25, and the SS for the bottom residuals in the table (same thing) in the table is 177.33 repeating, again close but slightly off due to rounding error.

We have finished our job of computing the sums of squares that we need in order to do the next steps, which include computing the MSEs for the effect and the error term. Once we do that, we can find the F-value, which is the ratio of the two MSEs.

Before we do that, you may have noticed that we solved for $SS_{\text{Error (left-over)}}$, rather than directly computing it from the data. In this chapter we are not going to show you the steps for doing this. We are not trying to hide anything from, instead it turns out these steps are related to another important idea in ANOVA. We discuss this idea, which is called an interaction in the next chapter, when we discuss factorial designs (designs with more than one independent variable).

Compute the MSEs

Calculating the MSEs (mean squared error) that we need for the F -value involves the same general steps as last time. We divide each SS by the degrees of freedom for the SS.

The degrees of freedom for SS_{Effect} are the same as before, the number of conditions - 1. We have three conditions, so the df is 2. Now we can compute the MSE_{Effect} .

$$MSE_{\text{Effect}} = \frac{SS_{\text{Effect}}}{df} = \frac{72}{2} = 36$$

The degrees of freedom for $SS_{\text{Error (left-over)}}$ are different than before, they are the (number of subjects - 1) multiplied by the (number of conditions -1). We have 3 subjects and three conditions, so $(3 - 1) * (3 - 1) = 2 * 2 = 4$. You might be wondering why we are multiplying these numbers. Hold that thought for now and wait until the next chapter. Regardless, now we can compute the $MSE_{\text{Error (left-over)}}$.

$$MSE_{\text{Error (left-over)}} = \frac{SS_{\text{Error (left-over)}}}{df} = \frac{177.33}{4} = 44.33$$

Compute F

We just found the two MSEs that we need to compute F . We went through all of this to compute F for our data, so let's do it:

$$F = \frac{MSE_{\text{Effect}}}{MSE_{\text{Error (left-over)}}} = \frac{36}{44.33} = 0.812$$

And, there we have it!

p-value

We already conducted the repeated-measures ANOVA using R and reported the ANOVA. Here it is again. The table shows the p -value associated with our F -value.

```
suppressPackageStartupMessages(library(dplyr))
library(xtable)
scores <- c(20,11,2,6,2,7,2,11,2)
conditions <- as.character(rep(c("A","B","C"), each=3))
subjects <- rep(1:3,3)
means <-c(9.33,8,3.66,9.33,8,3.66,9.33,8,3.66)
diff <-means-mean(scores)
diff_squared <-diff^2
df<-data.frame(subjects,conditions,scores,means,diff, diff_squared)
df$conditions<-as.character(df$conditions)
df$subjects<-as.character(df$subjects)

summary_out <- summary(aov(scores~conditions + Error(subjects/conditions),df[1:9,]))
knitr::kable(xtable(summary_out))
```

[run](#)
[restart](#)
[restart & run all](#)

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
Residuals	2	52.66667	26.33333	NA	F)" style="vertical-align:middle;" class="lt-stats-7935">NA
conditions	2	72.00000	36.00000	0.8120301	F)" style="vertical-align:middle;" class="lt-stats-7935">0.505848
Residuals	4	177.33333	44.33333	NA	F)" style="vertical-align:middle;" class="lt-stats-7935">NA

We might write up the results of our experiment and say that the main effect condition was not significant, $F(2,4) = 0.812$, $MSE = 44.33$, $p = 0.505$.

What does this statement mean? Remember, that the p -value represents the probability of getting the F value we observed or larger under the null (assuming that the samples come from the same distribution, the assumption of no differences). So, we know that an F -value of 0.812 or larger happens fairly often by chance (when there are no real differences), in fact it happens 50.5% of the time. As a result, we do not reject the idea that any differences in the means we have observed could have been produced by chance.

This page titled 8.3: Calculating the RM ANOVA is shared under a CC BY-SA 4.0 license and was authored, remixed, and/or curated by Matthew J. C. Crump via source content that was edited to the style and standards of the LibreTexts platform.