

16.4: Assumption Checking

As with one-way ANOVA, the key assumptions of factorial ANOVA are homogeneity of variance (all groups have the same standard deviation), normality of the residuals, and independence of the observations. The first two are things we can test for. The third is something that you need to assess yourself by asking if there are any special relationships between different observations. Additionally, if you aren't using a saturated model (e.g., if you've omitted the interaction terms) then you're also assuming that the omitted terms aren't important. Of course, you can check this last one by running an ANOVA with the omitted terms included and see if they're significant, so that's pretty easy. What about homogeneity of variance and normality of the residuals? As it turns out, these are pretty easy to check: it's no different to the checks we did for a one-way ANOVA.

16.4.1 Levene test for homogeneity of variance

To test whether the groups have the same variance, we can use the Levene test. The theory behind the Levene test was discussed in Section 14.7, so I won't discuss it again. Once again, you can use the `leveneTest()` function in the `car` package to do this. This function expects that you have a saturated model (i.e., included all of the relevant terms), because the test is primarily concerned with the within-group variance, and it doesn't really make a lot of sense to calculate this any way other than with respect to the full model. So we try either of the following commands:

```
leveneTest( model.2 )
leveneTest( mood.gain ~ drug + therapy, clin.trial )
```

R will spit out the following error:

```
Error in leveneTest.formula(formula(y), data = model.frame(y), ...) :
  Model must be completely crossed formula only.
```

Instead, if you want to run the Levene test, you need to specify a saturated model. Either of the following two commands would work:²³⁷

```
library(car)
leveneTest( model.3 )
```

```
## Levene's Test for Homogeneity of Variance (center = median)
##      Df F value Pr(>F)
## group  5  0.0955 0.9912
##      12
```

```
leveneTest( mood.gain ~ drug * therapy, clin.trial )
```

```
## Levene's Test for Homogeneity of Variance (center = median)
##      Df F value Pr(>F)
## group  5  0.0955 0.9912
##      12
```


The fact that the Levene test is non-significant means that we can safely assume that the homogeneity of variance assumption is not violated.

16.4.2 Normality of residuals


As with one-way ANOVA, we can test for the normality of residuals in a straightforward fashion (see Section 14.9). First, we use the `residuals()` function to extract the residuals from the model itself, and then we can examine those residuals in a few different ways. It's generally a good idea to examine them graphically, by drawing histograms (i.e., `hist()` function) and QQ

plots (i.e., `qqnorm()` function. If you want a formal test for the normality of the residuals, then we can run the Shapiro-Wilk test (i.e., `shapiro.test()`). If we wanted to check the residuals with respect to `model.2` (i.e., the model with both main effects but no interactions) then we could do the following:

```
resid <- residuals( model.2 ) # pull the residuals
hist( resid )                # draw a histogram
```



```
qqnorm( resid )              # draw a normal QQ plot
```



```
shapiro.test( resid )        # run the Shapiro-Wilk test
```

```
##
## Shapiro-Wilk normality test
##
## data:  resid
## W = 0.95635, p-value = 0.5329
```

I haven't included the plots (you can draw them yourself if you want to see them), but you can see from the non-significance of the Shapiro-Wilk test that normality isn't violated here.

This page titled [16.4: Assumption Checking](#) is shared under a [CC BY-SA 4.0](#) license and was authored, remixed, and/or curated by [Danielle Navarro](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.