

Notes on Maxwell & Delaney

PSYCH 710

6 Chapter 6 - Trend Analysis

Previously, we discussed how to use linear contrasts, or comparisons, to test specific hypotheses about differences among group means. Those discussions implicitly assumed that the grouping variable, or factor, was *qualitative*, in the sense that the levels of the factor were not ordered. In such cases, the levels of the factor are said to form a **nominal scale**. In some situations, however, the levels of the factor clearly *are* ordered: the ages of children in a developmental study, for example, or the amount of time given to study lists of words in a memory experiment. It makes sense to say that one level is greater than another, and therefore these factors are *quantitative* rather than qualitative. In such cases, we often are interested in particular kinds of comparisons across groups. For example, we may want to know if performance in a memory test is a *linear* or *nonlinear* function of study time. As we shall see, we can use linear comparisons to answer such questions. Using linear comparisons to examine the effects of a quantitative factor is often referred to as **trend analysis**, or the **method of orthogonal polynomials**. The mathematical methods for computing these contrasts are exactly the same as those used to compare the levels of qualitative factor.

6.1 rose data

6.1.1 omnibus test

We will use the data stored in the data frame, `rose` to illustrate trend analysis. The following commands set up R to use sum-to-zero coding for the effects (i.e., the α 's) associated with different levels of factors, polynomial contrasts for ordered factors (more on this in a moment), and loads the data:

```
> options(contrasts=c("contr.sum","contr.poly") )
> options(digits=4) # formats the output
> load(url("http://psycserv.mcmaster.ca/bennett/psy710/datasets/rose.Rdata") )
```

`rose` contains the integer variables `age` and `score` from a hypothetical score that recorded some measure of cognitive performance in children who were 8, 9, 10, 11, and 12 years of age. I am going to create a grouping variable (i.e., a factor) from `age`, and then plot the mean score as a function of age group in Figure 1.

```
> rose$ageGroup <- factor(rose$age, labels = "g", ordered = FALSE)
> y.means <- tapply(rose$score, rose$ageGroup, mean)
> x.group <- sort(unique(rose$age))$
```

Just for fun, let's compute the correlation between age and the mean:

```
> cor(x.group, y.means)
```

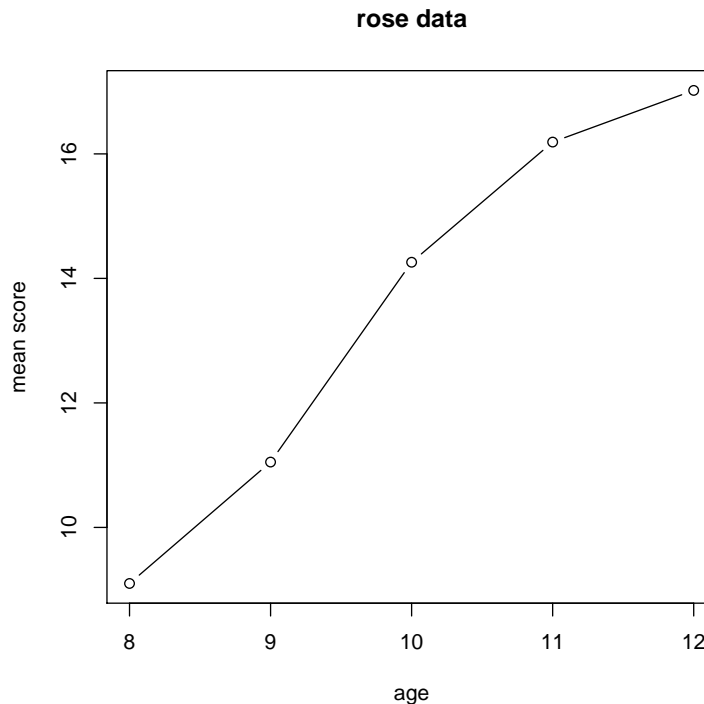


Figure 1: Rose data.

```
[1] 0.98287
```

The correlation is quite high and, together with the graph, suggests that there is an association between `score` and `age`. Now we do an ANOVA using the `aov()` and `summary()` commands:

```
> rose.aov.01 <- aov(score ~ ageGroup, data = rose)
> summary(rose.aov.01)
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
ageGroup	4	455.6	113.9	1.672	0.173
Residuals	45	3064.9	68.1		

We get the same results if we use `lm()` and print the summary table with `anova()`:

```
> rose.lm.01 <- lm(score ~ ageGroup, data = rose)
> anova(rose.lm.01)
```

Response: score

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
ageGroup	4	456	114	1.67	0.17
Residuals	45	3065	68		

Surprisingly, the effect of `ageGroup` is not statistically significant. What is going on here?

The problem is that the omnibus F test evaluates a very general question, namely whether there are *any* differences across groups. However, the very high correlation, as well as the data in Figure 1, suggest that there is a *linear* relation between `score` and `ageGroup`. What we need is a method for testing for the presence of specific trends that may exist between `score` and `ageGroup`. This method is called, not surprisingly, trend analysis.

6.1.2 ordered factors

Trend analysis involves nothing more than the calculation of linear contrasts, or comparisons, among levels of a quantitative factor. In R, a quantitative factor is called an **ordered factor**: it is a grouping variable that preserves the order among levels. Although linear contrasts can be used to test for arbitrary differences across levels of an ordered factor, in trend analysis we are interested in only a few possible patterns of differences, or trends, across the levels of an ordered factor. Several trends are shown in Figure 2. Each plot in Figure 2 show the weights for a single trend “contrast”. The weights are shown in the columns of the following matrix:

```
> contr.poly(n = 5,scores=c(8,9,10,11,12))
```

	.L	.Q	.C	\sim^4
[1,]	-0.6325	0.5345	-3.162e-01	0.1195
[2,]	-0.3162	-0.2673	6.325e-01	-0.4781
[3,]	0.0000	-0.5345	-4.096e-16	0.7171
[4,]	0.3162	-0.2673	-6.325e-01	-0.4781
[5,]	0.6325	0.5345	3.162e-01	0.1195

These contrast weights shown in Figure 2 with the assumption that there are an **equal number of subjects in each group**. A careful inspection of this matrix, and the plots in Figure 2, will show that the trend contrasts form an orthogonal set. (Note that these values are not the same as those listed in your textbook. For example, the book uses the following weights to test for a linear trend across 5 groups: [-1.5, -0.5, 0.5, 1.5]. The difference in weights will not alter the results of F tests, but it will alter the value of ψ for each comparison.)

To perform a trend analysis, we first need to create an ordered factor for our `rose` data set:

```
> rose$orderedAge <- factor(rose$age, labels = "a", ordered = TRUE)
> class(rose$orderedAge)
```

```
[1] "ordered" "factor"
```

```
> class(rose$ageGroup)
```

```
[1] "factor"
```

```
> unique(rose$orderedAge)
```

```
[1] a1 a2 a3 a4 a5
```

```
Levels: a1 < a2 < a3 < a4 < a5
```

```
> unique(rose$ageGroup)
```

```
[1] g1 g2 g3 g4 g5
```

```
Levels: g1 g2 g3 g4 g5
```

Note the difference in the levels for `orderedAge` and `ageGroup`.

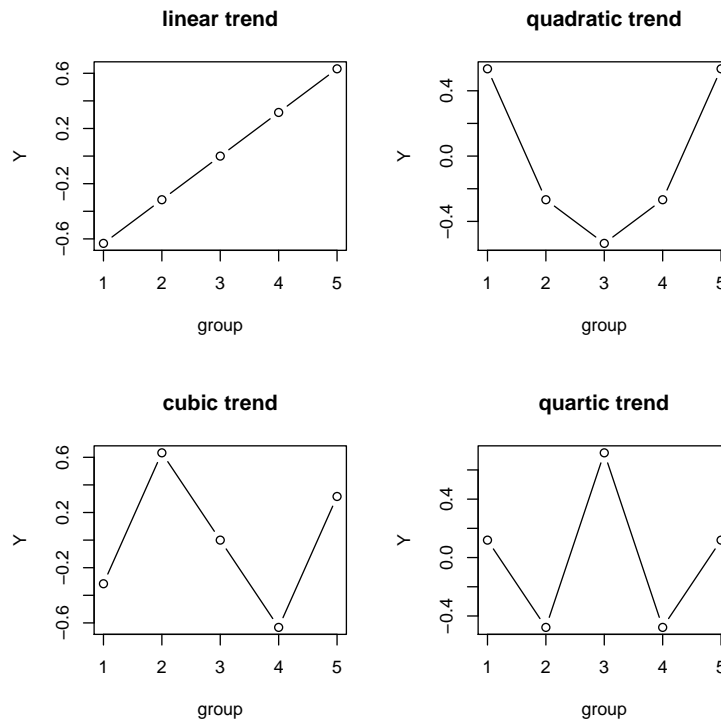


Figure 2: Linear, quadratic, cubic, and quartic trends.

6.1.3 trend analysis using aov

Now we simply do an ANOVA with our new ordered variable using `aov`:

```
> rose.aov.02 <- aov(score ~ orderedAge, data = rose)
> summary(rose.aov.02)
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
orderedAge	4	455.6	113.9	1.672	0.173
Residuals	45	3064.9	68.1		

The results are exactly the same as the previous one. Why? Because the ANOVA table shows the overall effect of group, which does not depend on the code used to represent the group variable. In other words, the overall, omnibus F is exactly the same regardless of whether group is represented as a factor or an ordered factor. The effect of using an ordered factor becomes apparent when we partition $SS_{orderedAge}$. Notice how the command `split` is used in the next command to break $SS_{orderedAge}$ into four components. Why four components? Because there are five age groups, so there are $5 - 1 = 4$ orthogonal trends.

```
> summary(rose.aov.02, split = list(orderedAge = list(1, 2, 3, 4)))
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
orderedAge	4	455.6	113.9	1.672	0.1730
orderedAge: C1	1	440.2	440.2	6.463	0.0145 *
orderedAge: C2	1	8.9	8.9	0.130	0.7202

```

orderedAge: C3  1    5.6    5.6  0.082 0.7762
orderedAge: C4  1    1.1    1.1  0.016 0.9014
Residuals      45 3064.9    68.1

```

```
---
```

```
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

This table is not very informative, so I am going to give names to each component.

```
> summary(rose.aov.02, split = list(orderedAge = list(linear = 1, quadratic = 2,
+   cubic = 3, quartic = 4)))
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
orderedAge	4	456	114	1.67	0.173
orderedAge: linear	1	440	440	6.46	0.015 *
orderedAge: quadratic	1	9	9	0.13	0.720
orderedAge: cubic	1	6	6	0.08	0.776
orderedAge: quartic	1	1	1	0.02	0.901
Residuals	45	3065	68		

```
---
```

```
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

The table shows SS_{contrast} , F , and p for each linear comparison. The first contrast is the linear trend, the second is the quadratic trend, the third is the cubic trend, and the fourth contrast is the quartic trend. We can see that the test for linear trend is significant, which means that we reject the null hypothesis

$$-0.6325\mu_1 - 0.3162\mu_2 + 0\mu_3 + 0.3162\mu_4 + 0.6325\mu_5 = 0 \quad (1)$$

Compare $SS_{\text{orderedAge}}$ to the sums of squares for each comparison. Because the trends are derived from orthogonal contrasts, $SS_{\text{orderedAge}} = SS_{\text{linear}} + SS_{\text{quadratic}} + SS_{\text{cubic}} + SS_{\text{quartic}}$. Also, because each trend is a linear contrast, there is 1 df for each trend. Finally, because the trends are orthogonal, $F_{\text{omnibus}} = (1/4)(F_{\text{linear}} + F_{\text{quadratic}} + F_{\text{cubic}} + F_{\text{quartic}})$.

6.1.4 testing for significance of nonlinear trends

Sometimes we want to know if a the relationship between our independent and dependent variables is or is not nonlinear, regardless of the form of the nonlinearity. This type of question can be addressed by doing a single test on the grouped nonlinear trends. The following commands perform such an analysis in R. Notice the 2:4 syntax used to specify that contrasts 2, 3, and 4 are combined.

```
> summary(rose.aov.02, split = list(orderedAge = list(linear = 1, nonlinear = 2:4)))
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
orderedAge	4	456	114	1.67	0.173
orderedAge: linear	1	440	440	6.46	0.015 *
orderedAge: nonlinear	3	15	5	0.08	0.973
Residuals	45	3065	68		

```
---
```

```
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

The results again indicate that the linear trend is significant. The test for nonlinearity is not significant. This test for nonlinearity can be thought of as an omnibus F test on the nonlinear trends.

Finally, in some situations we may be interested in evaluating the linear and quadratic trends and ignoring any remaining nonlinear trends. This analysis can be done by listing the linear and quadratic trends in the split command and grouping the remaining nonlinear trends:

```
> summary(rose.aov.02, split = list(orderedAge = list(lin = 1, quad=2, other = 3:4)))
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
orderedAge	4	456	114	1.67	0.173
orderedAge: lin	1	440	440	6.46	0.015 *
orderedAge: quad	1	9	9	0.13	0.720
orderedAge: other	2	7	3	0.05	0.953
Residuals	45	3065	68		

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

6.1.5 calculating ψ for each trend

Finally, it is a simple matter to list the value of ψ for each comparison:

```
> coef(rose.aov.02)
```

(Intercept)	orderedAge.L	orderedAge.Q	orderedAge.C	orderedAge^4
13.5240	6.6345	-0.9408	-0.7463	0.3251

```
> confint(rose.aov.02) # confidence intervals
```

	2.5 %	97.5 %
(Intercept)	11.173	15.875
orderedAge.L	1.378	11.891
orderedAge.Q	-6.197	4.316
orderedAge.C	-6.003	4.510
orderedAge^4	-4.931	5.581

6.1.6 interpretation of ψ values

What do these values of ψ mean? Each ψ can be turned into a regression weight, b :

$$b = \frac{\hat{\psi}}{\sum_{j=1}^a c_j^2} \quad (2)$$

where c_j is the j -th weight in the linear contrast. These regression weights can be used to construct a regression equation that expresses the observed scores in terms of a polynomial contrast weights

$$Y_{ij} = \beta_0 + b_L c_{Lj} + b_Q c_{Qj} + b_C c_{Cj} + \cdots + e_{ij} \quad (3)$$

where B_0 is the intercept, the b 's are regression weights, and the c 's are the contrast weights, and j represents the level of the ordered factor. For example, c_{L1} , c_{Q1} , and c_{C1} represent the first weights of the linear, quadratic, and cubic contrasts; c_{L2} , c_{Q2} , and c_{C2} represent the second weights of the

linear, quadratic, and cubic contrasts; and so on. This interpretation is valid *provided* the ordered levels are equally spaced *and* there are equal numbers of subjects in each group.

In the present example, the values of the intercept and four ψ 's are 13.524, 6.6345, -0.9409, -0.7463, and 0.3251. For the polynomial contrasts used by R, the sum of the squared weights is 1. Therefore, the regression coefficients equal the values of the comparisons ($b_j = \psi_j = [6.6345, -0.9409, -0.7463, 0.3251]$). Combining the b 's with the values of the contrast weights (shown above) allows us to use Equation 3 to calculate the mean scores for the four groups:

$$\begin{aligned}\bar{Y}_1 &= 13.524 + 6.6345(-0.6325) - 0.9408(0.5345) - 0.7463(-0.3162) + 0.3251(0.1195) = 9.10 \\ \bar{Y}_2 &= 13.524 + 6.6345(-0.3162) - 0.9408(-0.2673) - 0.7463(0.6325) + 0.3251(-0.4781) = 11.05 \\ \bar{Y}_3 &= 13.524 + 6.6345(0) - 0.9408(-0.5345) - 0.7463(0) + 0.3251(0.7171) = 14.26 \\ \bar{Y}_4 &= 13.524 + 6.6345(0.3162) - 0.9408(-0.2673) - 0.7463(-0.6325) + 0.3251(-0.4781) = 16.19 \\ \bar{Y}_5 &= 13.524 + 6.6345(0.6325) - 0.9408(0.5345) - 0.7463(0.3162) + 0.3251(0.1195) = 17.02\end{aligned}$$

These values correspond exactly to the group means:

```
> with(rose, tapply(score, orderedAge, mean) ) # calc group means
      a1      a2      a3      a4      a5
9.10 11.05 14.26 16.19 17.02
```

6.2 planned vs. post-hoc comparisons

I have been acting as though the previous analyses were planned comparisons. However, it certainly is permissible to do post-hoc trend analyses. In such cases, you should use Scheffé's method to control familywise Type I errors. However, I think that we often do ourselves a disservice by not doing planned trend analyses on ordered factors. Why? Because psychological theories rarely are developed to the point where they predict a quantitative relation between independent and dependent variables that includes anything higher than a quadratic component. In other words, it is unlikely that a theory would have the form shown in Equation 3. It is more likely to find a theory that looks like this

$$Y_{ij} = \beta_0 + \alpha_L X_j + e_{ij} \quad (4)$$

or this

$$Y_{ij} = \beta_0 + \alpha_L X_j + \alpha_Q X_j^2 + e_{ij} \quad (5)$$

Sometimes the "theory" is implicit. For instance, if we are measuring some ability in children in grades 1 through 5, we may strongly suspect that the ability increases from grades 1 through 4 and levels off. This expectation is equivalent to saying that we expect to find significant linear and quadratic trends. In this case, it makes sense to test for linear and quadratic trends and ignore the higher-order nonlinear trends. Such comparisons would be more powerful than an omnibus F test and provide more direct tests of our hypothesis. Of course, I am not saying that you should *always* do trend analyses on ordered factors. Instead, I am suggesting that you think about how your dependent variable may be related to the levels of your grouping variable and then use an appropriate test to evaluate *your* hypothesis.