ANCOVA

ANCOVA (Analysis of covariance) is a linear model that allows you to compare two (or more) groups while adjusting for one or more quantitative covariates.
Similarly, you might be interested in comparing the relationship between the quantitative variables while accounting for group differences.

ANCOVA

As a made up example, suppose you have a sample of books of different lengths, some of which are hardcover, some softcover. Maybe these are books you buy at the beginning of a semester. You want to model price as a function of book length and book cover (hard cover or soft cover).

- You might think that longer books tend to be more expensive. This might be true within categories (longer novels might be more expensive than shorter novels, longer hardback textbooks tend to be more expensive than shorter hardback textbooks).
- But if you ignore the category, it might be hard to compare.

title	category	length	price
Applied Linear Statistical Models	soft	1396	54.64
It	soft	1000	12.95
War & Peace	soft	1300	14.95
Wizard of Oz (5 volumes)	hard	1200	7.98

ANCOVA

As another example, suppose you recruit 20 patients for a study

- wish to determine whether treatment A or treatment B is more effective.
- patients are assigned randomly to treatments, however, we notice that due to the small sample size, patients receiving treatment A tend to be younger than those receiving treatment B.
- In this case, although randomization should have helped, you might wish to additionally take into account age information of patients in determining the effects of the two treatments.

Example: Suppose that you are interested in comparing the typical lifetime (hours) of two tool types (A and B).

- A simple analysis is to make side-by-side boxplots followed by a two-sample test of equal means (or medians).
- The standard two-sample test using the pooled variance estimator is a special case of the one-way ANOVA with two groups.
- The summaries suggest that the distribution of lifetimes for the tool types are different.
  —— In the output below, μ_i is population mean lifetime for tool type $i (i = A, B)$.

```
tools <- read.table("http://statacumen.com/teach/ADA2/ADA2_not
                  , header = TRUE)
tools
  lifetime  rpm type
1     18.73  610    A
2     14.52  950    A
3     17.43  720    A
4     14.54  840    A
5     13.44  980    A
6     24.39  530    A
7     13.34  680    A
8     22.71  540    A
9     12.68  890    A
10    19.32  730    A
11    30.16  670    B
12    27.09  770    B
13    25.40  880    B
14    26.05 1000    B
```

Figure: Boxplot of the two tool types

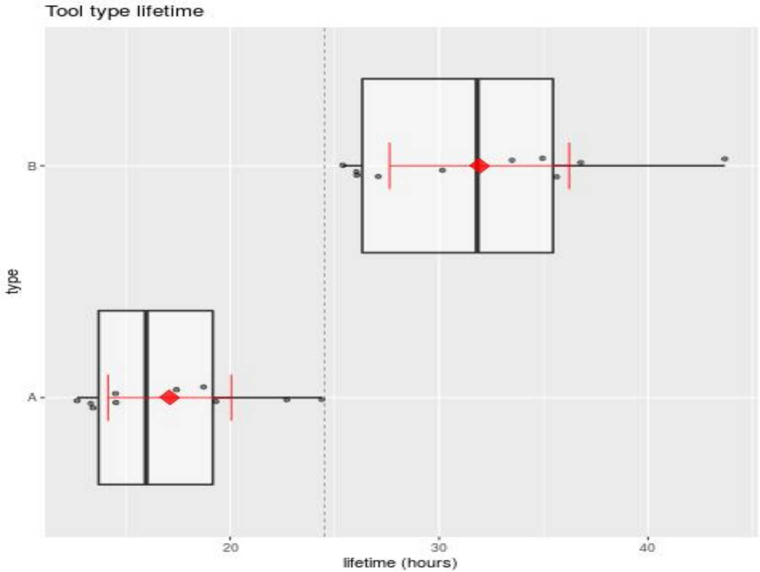

```
> t.summary <- t.test(lifetime ~ type, data = tools)
> t.summary

Welch Two Sample t-test

data:  lifetime by type
t = -6.435, df = 15.93, p-value = 8.422e-06
alternative hypothesis: true difference in means is not equal
to 0
95 percent confidence interval:
 -19.70128  -9.93472
sample estimates:
mean in group A mean in group B
         17.110          31.928
```

Comments:

- Side by side box plot and two sample t-test comparing mean lifetimes of tool types indicates a difference between means.
- These comparisons are potentially misleading because the samples may not be com- parable.
  —— A one-way ANOVA is most appropriate for designed experiments where **all the factors influencing the response, other than the treatment (tool type), are controlled by the experimenter**.
- The tools were operated at different speeds.
  —— If speed influences lifetime, then the observed differences in lifetimes could be due to differences in speeds at which the two tool types were operated.
- Need to compare groups (tools) after ad- justing the lifetimes to account for the influence of a measurement variable, speed.
  ——-The appropriate statistical technique for handling this problem is called analysis of covariance (ANCOVA).

ANCOVA

The simplest approach to ANCOVA is to adjust the intercept for each group. The model is

$$\text{response} = \text{Grand mean} + \text{group} + \text{covariate}$$

or

$$y_{ij} = \mu + \alpha_i + \beta x_{ij} + \varepsilon_{ij}$$

where $\mu_i = \mu + \alpha_i$ is the intercept for group i.

- The effect of the model is that there is a separate regression line for each group, but the regression lines are assumed to be parallel.
- The effect of belonging to a particular group is to shift the regression line up or down.
  ——Typically, the main interest is in testing whether $\alpha_i = 0$ for each i, meaning that there is no difference in regression line among the groups.

Example continued

A natural way to account for the effect of speed is through a multiple regression model with lifetime as the response and two predictors,

——speed and tool type.

—— A binary categorical variable, here tool type, is included in the model as a dummy variable or indicator variable

$$X_i = \begin{cases} 1 & \text{if tool i is from type B} \\ 0 & \text{if tool i is from type A} \end{cases}$$

Consider the model

$$\text{Tool lifetime} = \beta_0 + \beta_1 X_i + \beta_2 rpm + \epsilon_i,$$

where

$$X_i = \begin{cases} 1 & \text{if tool i is from type B} \\ 0 & \text{if tool i is from type A} \end{cases}$$

For type A tools, the model simplifies to:

$$\text{Tool lifetime} = \beta_0 + \beta_2 rpm + \epsilon_i,$$

For type B tools, the model simplifies to:

$$\text{Tool lifetime} = (\beta_0 + \beta_1) + \beta_2 rpm + \epsilon_i.$$

This ANCOVA model fits two regression lines, one for each tool type, but restricts the slopes of the regression lines to be identical.

$$\text{Tool lifetime} = \beta_0 + \beta_1 X_i + \beta_2 rpm + \epsilon_i,$$

- $\beta_2 =$ slope of population regression lines for tool types A and B.
- β_0=intercept of population regression line for tool A (called the reference group).
- $\beta_0 + \beta_1$ is the intercept of population regression line for tool B.
  ——-β_1 is the difference between tool B and tool A intercepts.
  ——A test of $H_0 : \beta_1 = 0$ is the primary interest, and is interpreted as a comparison of the tool types, after adjusting or allowing for the speeds at which the tools were operated.

Figure: Tool data, lifetime by rmp with categorical types

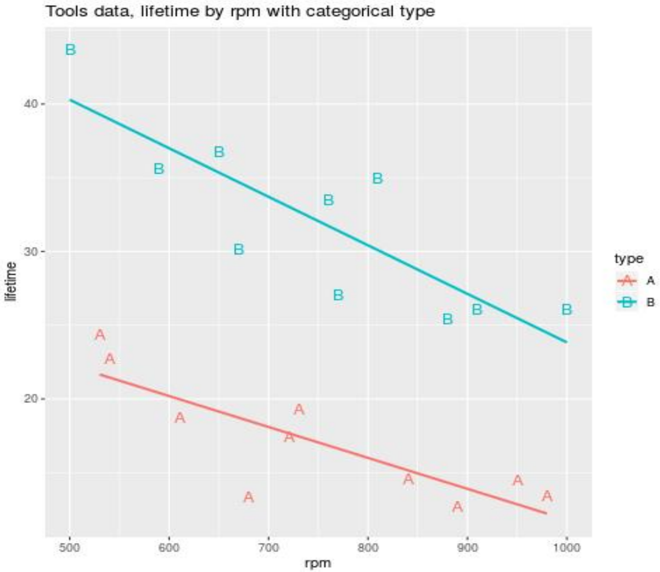

The ANCOVA model is plausible. The relationship between lifetime and speed is roughly linear within tool types, with similar slopes but unequal intercepts across groups.

```
 > lm.l.r.t <- lm(lifetime ~ rpm + type, data = tools)
> summary(lm.l.r.t)

Call:
lm(formula = lifetime ~ rpm + type, data = tools)
Residuals:
    Min      1Q  Median      3Q     Max
-5.5527 -1.7868 -0.0016  1.8395  4.9838

Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept) 36.98560    3.51038  10.536 7.16e-09 ***
rpm         -0.02661    0.00452  -5.887 1.79e-05 ***
typeB       15.00425    1.35967  11.035 3.59e-09 ***

Residual standard error: 3.039 on 17 degrees of freedom
Multiple R-squared:  0.9003,Adjusted R-squared:  0.8886
F-statistic: 76.75 on 2 and 17 DF,  p-value: 3.086e-09
```

- The t-test of
$$H_0 : \beta_1 = 0$$
checks whether the intercepts for the population regression lines for the two tool types are equal, assuming equal slopes.
——-The t-test $p - value < 0.0001$ suggests that there is a significant difference between the mean lifetimes of the two tool types, after adjusting for the effect of the speeds at which the tools were operated.
——The estimated difference in average lifetime between the two tool types is 15 hours, regardless of the speed.
- The LS lines indicate that the average lifetime of either type tool decreases by 0.0266 hours for each increase in 1 RPM. Regardless of the the type.

```
 Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept) 36.98560    3.51038  10.536 7.16e-09 ***
rpm         -0.02661    0.00452  -5.887 1.79e-05 ***
typeB       15.00425    1.35967  11.035 3.59e-09 ***
```

For tool type A, the fitted relationship is

$$\text{Predicted Lifetime} = 36.99 - 0.0266 rpm.$$

For tool type B, the fitted relationship is

$$\text{Predicted Lifetime} = (36.99 + 15.00) - 0.0266 rpm = 51.99 - 0.0266 rpm$$

——-the model predicts that type B tools will last 15 hours longer than type A tools.

Figure: Tool data, diagnostic plot

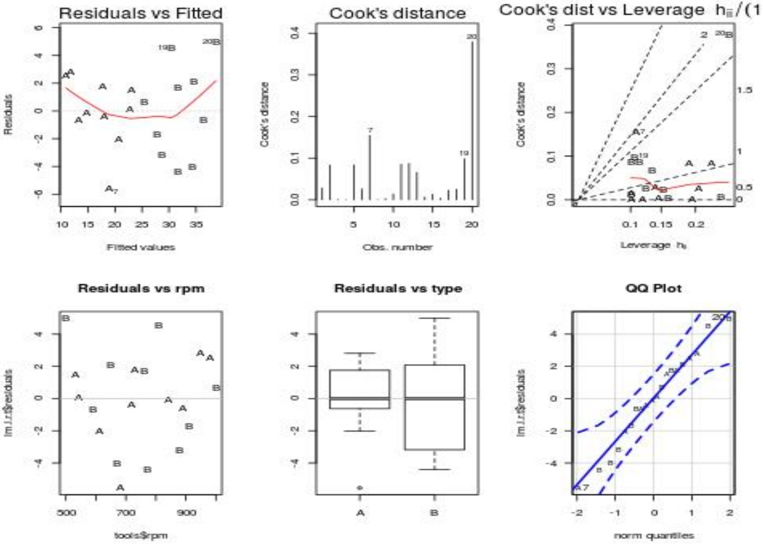

Comments:

- The plot of residuals against the fitted values shows no gross abnormalities, but suggests that the variability about the regression line for tool type A is somewhat smaller than the variability for tool type B.
- Cook's distance are all less than 1, no influential point observed.
- The QQ-plot does not show any gross deviations from a straight line.

Generalizing the ANCOVA Model to Allow Unequal Slopes

- Introduce a flexible approach for checking equal slopes and equal intercepts in ANCOVA-type models.
- The algorithm also provides a way to build regression models in studies where the primary interest is comparing the regression lines across groups rather than comparing groups after adjusting for a regression effect.
- The approach can be applied to an arbitrary number of groups and predictors.

Example: IQ scores of identical twins

- One raised in a foster home (IQ score is denoted as IQF) and the other raised by natural parents (IQ score is denoted as IQN).
- The 27 pairs are divided into three groups by social status of the natural parents (H=high, M=medium, L=low).
- Want to examine the regression of IQF on IQN for each of the three social classes.

```
> twins
   IQF IQN status
1   82  82      H
2   80  90      H
3   88  91      H
4  108 115      H
5  116 115      H
6  117 129      H
7  132 131      H
8   71  78      M
9   75  79      M
10  93  82      M
11  95  97      M
12  88 100      M
13 111 107      M
14  63  68      L
15  77  73      L
16  86  81      L
```

Indicator variable and coding

- in order that such a qualitative variable can be used in a regression model, indicator variables that take on the values 0 and 1 for the classes of the qualitative variable must be employed
- if there are c categories, need to use $c-1$ indicator variables to distinguish them

Let $I_1 = \begin{cases} 1 & \text{if H status families} \\ 0 & \text{otherwise} \end{cases}$, $I_2 = \begin{cases} 1 & \text{if M status families} \\ 0 & \text{otherwise} \end{cases}$

The indicators I_1 and I_2 jointly assume 3 values:

Status	I_1	I_2
L	0	0
M	0	1
H	1	0

The most general model allows separate slopes and intercepts for each group:

$$IQF = \beta_0 + \beta_1 I_1 + \beta_2 I_2 + \beta_3 IQN + \beta_4 I_1 IQN + \beta_5 I_2 IQN + \epsilon.$$

If status = L, then $I_1 = I_2 = 0$. For these families

$$IQF = \beta_0 + \beta_3 IQN + \epsilon.$$

If status = M, then $I_1 = 0$ and $I_2 = 1$. For these families

$$\begin{aligned} IQF &= \beta_0 + \beta_2(1) + \beta_3 IQN + \beta_5 IQN + \epsilon \\ &= (\beta_0 + \beta_2) + (\beta_3 + \beta_5) IQN + \epsilon \end{aligned}$$

If status = H, then $I_1 = 1$ and $I_2 = 0$. For these families

$$\begin{aligned} IQF &= \beta_0 + \beta_1(1) + \beta_3 IQN + \beta_4 IQN + \epsilon \\ &= (\beta_0 + \beta_1) + (\beta_3 + \beta_4) IQN + \epsilon \end{aligned}$$

- The regression coefficients β_0 and β_3 are the intercept and slope for the L status population regression line. L status families is treated as a baseline or reference group.
- The other parameters measure differences in intercepts and slopes across the three groups, using L status families as a baseline or reference group.
- $\beta_1 =$ difference between the intercepts of the H and L population regression lines.

  $\beta_2 =$ difference between the intercepts of the M and L population regression lines.

  $\beta_4 =$ difference between the slopes of the H and L population regression lines.

  $\beta_5 =$ difference between the slopes of the M and L population regression lines.

Figure: Twin data

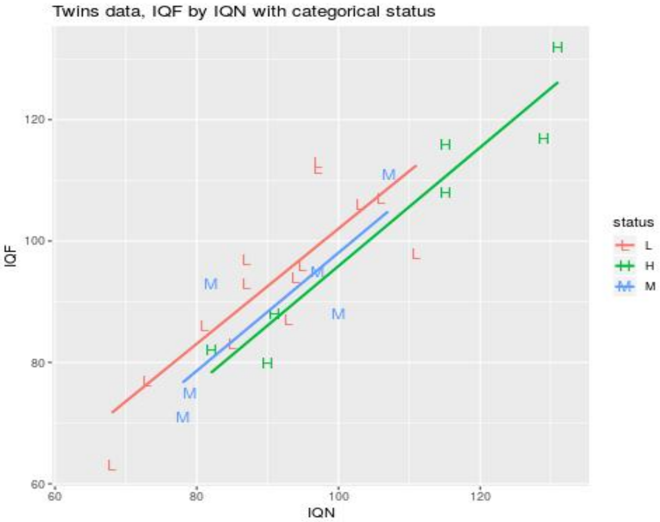

```
 twins <-read.table("http://statacumen.com/teach/ADA2/ADA2_not
 # set "L" as baseline level
 twins$status <-relevel(twins$status, "L’’)
 lm.f.n.s.ns <- lm(IQF ~ IQN*status, data = twins)
summary(lm.f.n.s.ns)
> summary(lm.f.n.s.ns)
Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept)  7.20461   16.75126   0.430    0.672
IQN          0.94842    0.18218   5.206 3.69e-05 ***
statusH     -9.07665   24.44870  -0.371    0.714
statusM     -6.38859   31.02087  -0.206    0.839
IQN:statusH  0.02914    0.24458   0.119    0.906
IQN:statusM  0.02414    0.33933   0.071    0.944
---
Residual standard error: 7.921 on 21 degrees of freedom
Multiple R-squared:  0.8041,Adjusted R-squared:  0.7574
F-statistic: 17.24 on 5 and 21 DF,  p-value: 8.31e-07
```

- For the baseline group with status = L,

  $$\text{Predicted IQF} = 7.20 + 0 + (0.948 + 0)IQN = 7.20 + 0.948IQN.$$

- For the M status group with indicator I_2 and product effect I_2IQN:

  $$\text{Predicted IQF} = 7.20 - 6.39 + (0.948 + 0.024)IQN = 0.81 + 0.972IQN.$$

- For the H status group with indicator I_1 and product effect I_1IQN:

  $$\text{Predicted IQF} = 7.20 - 9.08 + (0.948 + 0.029)IQN = -1.88 + 0.977IQN.$$

- The LS lines are identical to separately fitting simple linear regressions to the three groups.

Figure: Twin data, diagnostic plot

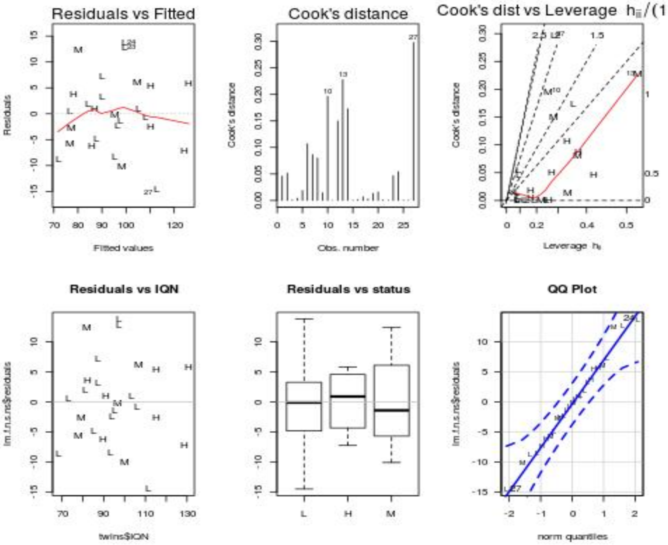

General Model

$$IQF = \beta_0 + \beta_1 I_1 + \beta_2 I_2 + \beta_3 IQN + \beta_4 I_1 IQN + \beta_5 I_2 IQN + \epsilon.$$

IQF = Grand Mean + Status Effect + IQN effect +StatusIQN interaction + Residual.

- β_0: grand mean
- $\beta_1 I_1 + \beta_2 I_2$: status effect (i.e., the two indicators I_1 and I_2 allow you to differentiate among social classes)
- $\beta_3 IQN$: the IQN effect
- $\beta_4 I_1 IQN + \beta_5 I_2 IQN$: status by IQN interaction

The general model is a special case of the two-factor interaction ANOVA model because it restricts the means to change linearly with IQN.

ANCOVA model

The ANCOVA model has main effects for status and IQN but no interaction:
IQF = Grand Mean + Status Effect + IQN effect + Residual.

- The ANCOVA model is a special case of the additive two-factor ANOVA model because the plot of the population means has parallel profiles, but is not equivalent to the additive two-factor ANOVA model.

Equal slopes and intercepts model

The model with equal slopes and intercepts has no main effect for status nor an interaction between status and IQN:
IQF = Grand Mean + IQN effect + Residual.

One-way ANOVA model

The one-way ANOVA model has no main effect for IQN nor an interaction between status and IQN:
IQF = Grand Mean + Status Effect + Residual.

Choosing among models

Full general model:

$$IQF = \beta_0 + \beta_1 I_1 + \beta_2 I_2 + \beta_3 IQN + \beta_4 I_1 IQN + \beta_5 I_2 IQN + \epsilon.$$

Test the hypothesis of equal slopes

$$H_0 : \beta_4 = \beta_5 = 0$$

- Note that t-tests are used to test either $\beta_4 = 0$ or $\beta_5 = 0$
- Fit a reduced model

$$IQF = \beta_0 + \beta_1 I_1 + \beta_2 I_2 + \beta_3 IQN$$

- Reject $H_0 : \beta_4 = \beta_5 = 0$ if the increase in the Residual SS obtained by deleting $I_1 IQN$ and $I_2 IQN$ from the full model is significant.

Formal test

$$H_0 : \beta_4 = \beta_5 = 0$$
$$MSTest = \frac{SSE(R) - SSE(F)}{dfE(R) - dfE(F)}$$

and

$$F = \frac{\{SSE(R) - SSE(F)\}/(dfE(R) - dfE(F))}{SSE(F)/dfE(F)} = \frac{MSTest}{MSE(F)}$$

Reject H_0, if $F > F(1-\alpha; dfE(R) - dfE(F), dfE(F))$
The observed F statistic is

$$\begin{aligned} F &= \frac{\{SSE(R) - SSE(F)\}/(dfE(R) - dfE(F))}{MSE(F)} \\ &= \frac{(1318.4 - 1317.5)/(23 - 21)}{62.7} \\ &= 0.0072 \end{aligned}$$

which is much smaller than $F(2, 21, 0.95) = 3.4668$. The test suggests do not reject H_0. Therefore, equal slope model is appropriate.

```
 lm.f.n.s.ns <- lm(IQF ~ IQN*status, data = twins)
 library(car)
> Anova(aov(lm.f.n.s.ns), type=3)
Response: IQF
             Sum Sq Df F value   Pr(>F)
(Intercept)   11.61  1  0.1850   0.6715
IQN         1700.39  1 27.1035 3.69e-05 ***
status         8.99  2  0.0716   0.9311
IQN:status     0.93  2  0.0074   0.9926
Residuals   1317.47 21
 lm.f.n.s <- lm(IQF ~ IQN + status, data = twins)
> Anova(aov(lm.f.n.s), type=3)
Response: IQF
            Sum Sq Df F value    Pr(>F)
(Intercept)   18.2  1  0.3181    0.5782
IQN         4674.7  1 81.5521 5.047e-09 ***
status       175.1  2  1.5276    0.2383
Residuals   1318.4 23
```

Or you can use the following to compare the models

```
> anova(lm.f.n.s.ns,lm.f.n.s)
Analysis of Variance Table

Model 1: IQF ~ IQN * status
Model 2: IQF ~ IQN + status
  Res.Df    RSS Df Sum of Sq      F Pr(>F)
1     21 1317.5
2     23 1318.4 -2  -0.93181 0.0074 0.9926
```

When comparing the full general model with the reduced equal slope model, F-value is 0.0074, and p-value is 0.9926. We do not reject the null hypothesis, and conclude that the equal slope model is appropriate.

Test equal intercepts

Diagnostics plots of the equal slope model looks good, now we adopt the equal slope model (full model)

$$IQF = \beta_0 + \beta_1 I_1 + \beta_2 I_2 + \beta_3 IQN + \epsilon.$$

want to further test if equal intercepts model is appropriate or not

$$H_0 : \beta_1 = \beta_2 = 0$$

Reduced model: $IQF = \beta_0 + \beta_3 IQN$

$$\begin{aligned} F &= \frac{\{SSE(R) - SSE(F)\}/(dfE(R) - dfE(F))}{MSE(F)} \\ &= \frac{(1493.5 - 1318.4)/(25 - 23)}{1318.4/23} \\ &= 1.53 \end{aligned}$$

which is much smaller than $F(2, 23, 0.95) = 3.42$. The test suggests not rejecting H_0. Therefore, equal intercept model is appropriate.

```
 lm.f.n.s <- lm(IQF ~ IQN + status, data = twins)
library(car)
Anova(aov(lm.f.n.s), type=3)
Response: IQF
            Sum Sq Df F value    Pr(>F)
(Intercept)   18.2  1  0.3181    0.5782
IQN         4674.7  1 81.5521 5.047e-09 ***
status       175.1  2  1.5276    0.2383
Residuals   1318.4 23

lm.f.n <- lm(IQF ~ IQN, data = twins)
Anova(aov(lm.f.n), type=3)
Anova Table (Type III tests)
Response: IQF
            Sum Sq Df F value    Pr(>F)
(Intercept)   58.6  1  0.9802    0.3316
IQN         5231.1  1 87.5630 1.204e-09 ***
Residuals   1493.5 25
```

Figure: Twin data, diagnostic plot for reduced model: $IQF = \beta_0 + \beta_3 IQN$

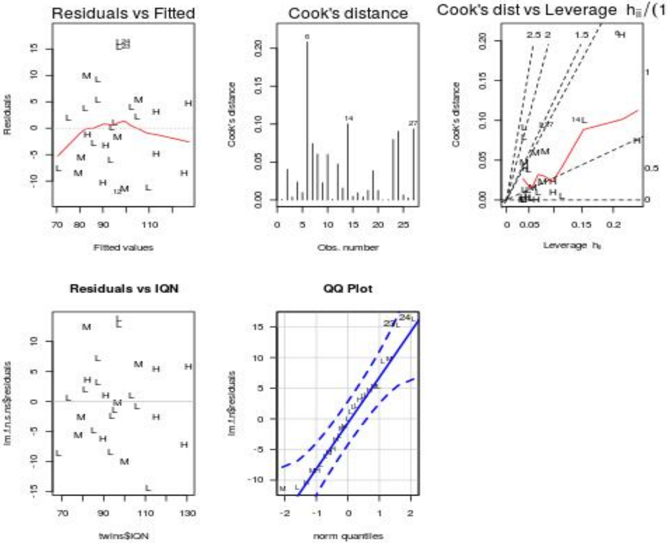

```
> shapiro.test(lm.f.n$residuals)

Shapiro-Wilk normality test
data:  lm.f.n$residuals
W = 0.96329, p-value = 0.4377

> library(lmtest)
> bptest(IQF ~ IQN, data = twins,studentize=FALSE)

Breusch-Pagan test
data:  IQF ~ IQN
BP = 0.038304, df = 1, p-value = 0.8448

#compare cook's distance to F(0.5, p, n-p)
> highcook <- which((cooks.distance(lm.f.n)) > qf(0.5,2,25))
> cooks.distance(lm.f.n)[highcook]
named numeric(0)
```

Comments:

- The plot of residuals against the fitted values shows no gross abnormalities.
- The Breusch-Pagan test show $BP = 0.038304$ with $p - value = 0.8448$. Constant variance assumption is not rejected.
- Cook's distance are all less than 1, when compared with $F(0.5, 2, 25)$, there is no influential point observed.
- The QQ-plot does not show any gross deviations from a straight line. Shapiro-Wilk normality test has $W = 0.96329$ with $p - value = 0.4377$, normality assumption is not rejected.

Suggested final model to use

```
 > summary(lm.f.n)

Call:
lm(formula = IQF ~ IQN, data = twins)

Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept)  9.20760    9.29990   0.990    0.332
IQN          0.90144    0.09633   9.358  1.2e-09 ***

Residual standard error: 7.729 on 25 degrees of freedom
Multiple R-squared:  0.7779,Adjusted R-squared:  0.769
F-statistic: 87.56 on 1 and 25 DF,  p-value: 1.204e-09
```

Suggested final model to use

The estimated regression line, regardless of social class, is:

$$\text{Predicted IQF} = 9.21 + 0.901 * IQN.$$

- There are no serious inadequacies with this model, based on a diagnostic analysis.
- The natural parents social class has no impact on the relationship between the IQ scores of identical twins raised apart.

Simultaneous testing of regression parameters

$$IQF = \beta_0 + \beta_1 I_1 + \beta_2 I_2 + \beta_3 IQN + \beta_4 I_1 IQN + \beta_5 I_2 IQN + \epsilon.$$

If status = L, $IQF = \beta_0 + \beta_3 IQN + \epsilon.$
If status = M, $IQF = (\beta_0 + \beta_2) + (\beta_3 + \beta_5) IQN + \epsilon$
If status = H, $IQF = (\beta_0 + \beta_1) + (\beta_3 + \beta_4) IQN + \epsilon$
Consider these two specific hypotheses:

1. H_0 : equal regression lines for status M and L, i.e. $\beta_2 = \beta_5 = 0$
2. H_0 : equal regression lines for status M and H, i.e. $\beta_1 = \beta_2$ and $\beta_4 = \beta_5$

```
 > lm.f.n.s.ns <- lm(IQF ~ IQN*status, data = twins)
> coef(lm.f.n.s.ns)
(Intercept)          IQN     statusH     statusM
 7.20460986  0.94842244 -9.07665352 -6.38858548
 IQN:statusH IQN:statusM
 0.02913971  0.02414450
```

-
$$IQF = \beta_0 + \beta_1 I_1 + \beta_2 I_2 + \beta_3 IQN + \beta_4 I_1 IQN + \beta_5 I_2 IQN + \epsilon.$$

- Using R default, it is modeled in ordered of $\beta_0, \beta_3, \beta_1, \beta_2, \beta_4, \beta_5$

Wald test

H_0 : equal regression lines for status M and L, i.e. $\beta_2 = \beta_5 = 0$

- Define matrix

$$mR = \left[\begin{array}{cccccc} 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{array} \right], \beta = \left[\begin{array}{c} \beta_0 \\ \beta_3 \\ \beta_1 \\ \beta_2 \\ \beta_4 \\ \beta_5 \end{array} \right]$$

so that $mR * \beta = \left[\begin{array}{c} \beta_2 \\ \beta_5 \end{array} \right]$

- We want to test $mR * \beta = \left[\begin{array}{c} 0 \\ 0 \end{array} \right]$

```
 library(aod)
 mR <- as.matrix(rbind(c(0, 0, 0, 1, 0, 0),
  c(0, 0, 0, 0, 0, 1)))
> mR
     [,1] [,2] [,3] [,4] [,5] [,6]
[1,]    0    0    0    1    0    0
[2,]    0    0    0    0    0    1
> vR <- c(0, 0)
[1] 0 0
wald.test(b = coef(lm.f.n.s.ns) , Sigma = vcov(lm.f.n.s.ns)
        , L = mR, H0 = vR)
        Wald test:
Chi-squared test:
X2 = 1.2, df = 2, P(> X2) = 0.55
```

```
wald.test(b = coef(lm.f.n.s.ns) , Sigma = vcov(lm.f.n.s.ns)
        , L = mR, H0 = vR)
        Wald test:
Chi-squared test:
X2 = 1.2, df = 2, P(> X2) = 0.55
```

Conclusion: $\chi^2 = 1.2, p - value = 0.55$, we do not reject $H_0 : \beta_2 = \beta_5 = 0$, conclude that the regression lines for status M and L are not significantly different. We suggests that M and L can be described by the same regression line.

H_0 : equal regression lines for status M and H, i.e. $\beta_1 = \beta_2$ and $\beta_4 = \beta_5$ or

$$\beta_1 - \beta_2 = 0, \text{ and } \beta_4 - \beta_5 = 0$$

- Define matrix

$$mR = \begin{bmatrix} 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & -1 \end{bmatrix}, \beta = \begin{bmatrix} \beta_0 \\ \beta_3 \\ \beta_1 \\ \beta_2 \\ \beta_4 \\ \beta_5 \end{bmatrix}$$

so that $mR * \beta = \begin{bmatrix} \beta_1 - \beta_2 \\ \beta_4 - \beta_5 \end{bmatrix}$

- We want to test $mR * \beta = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$

```
 > mR <- as.matrix(rbind(c(0, 0, 1, -1, 0, 0),
  c(0, 0, 0, 0, 1, -1)))
> mR
     [,1] [,2] [,3] [,4] [,5] [,6]
[1,]    0    0    1   -1    0    0
[2,]    0    0    0    0    1   -1
> vR <- c(0, 0)
> vR
[1] 0 0
> wald.test(b = coef(lm.f.n.s.ns)
+         , Sigma = vcov(lm.f.n.s.ns)
+         , L = mR, H0 = vR)
Wald test:
Chi-squared test:
X2 = 0.19, df = 2, P(> X2) = 0.91
```

Comments:

- The large $p - value = 0.91$ suggests that M and H can be described by the same regression line with same slope and intercept.
- The results of these tests are not surprising, given our previous analysis where we found that the status effect is not significant for all three groups.
- Any simultaneous linear combination of parameters can be tested in this way.