STAT 516 hw 4

Solutions

Chp 6 Ex 1

The completed table is:

Source	df	SS	MS	F value
Between	4	25	6.25	1.41
Within	35	155	4.43	
Total	39	180		

To test H_0: $\mu_1 = \cdots = \mu_5$, we compare the F value to the critical value $F_{4,35,0.05} = 2.6414652$. Since the F value is smaller than the critical value, we fail to reject the null hypothesis of no treatment effect, i.e. equal means.

Chp 6 Ex 7

The data can be read in with this code:

```
stay <- c(rep(3,6),rep(4,14),
          rep(4,18),rep(5,2),
          rep(4,10),rep(5,9),rep(6,1),
          rep(4,8),rep(5,12))
grp <- as.factor(c(rep("A",6+14),rep("B",18+2),rep("C",10+9+1),rep("D",8+12)))
```

(a)

This code produces the ANOVA table:

```r
lm_out <- lm(stay ~ grp)
anova(lm_out)
```

```
Analysis of Variance Table

Response: stay
          Df Sum Sq Mean Sq F value    Pr(>F)    
grp        3 10.738  3.5792  15.325 6.86e-08 ***
Residuals 76 17.750  0.2336                      
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Provided the assumptions are satisfied, we would reject the null hypothesis of no effect due to the treatment; we conclude that the type of anesthesia has an effect on the average length of hospital stay.

(b), (c), (d)

We check the residuals versus fitted values plot as well as the Normal quantile-quantile plot of the residuals:

```r
plot(lm_out,which = 1)
```

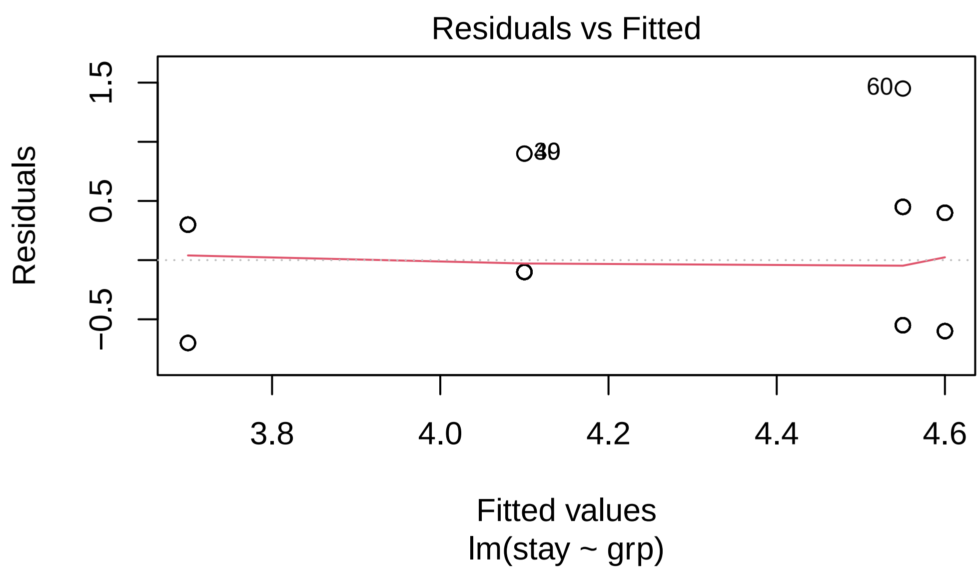

```
plot(lm_out,which = 2)
```

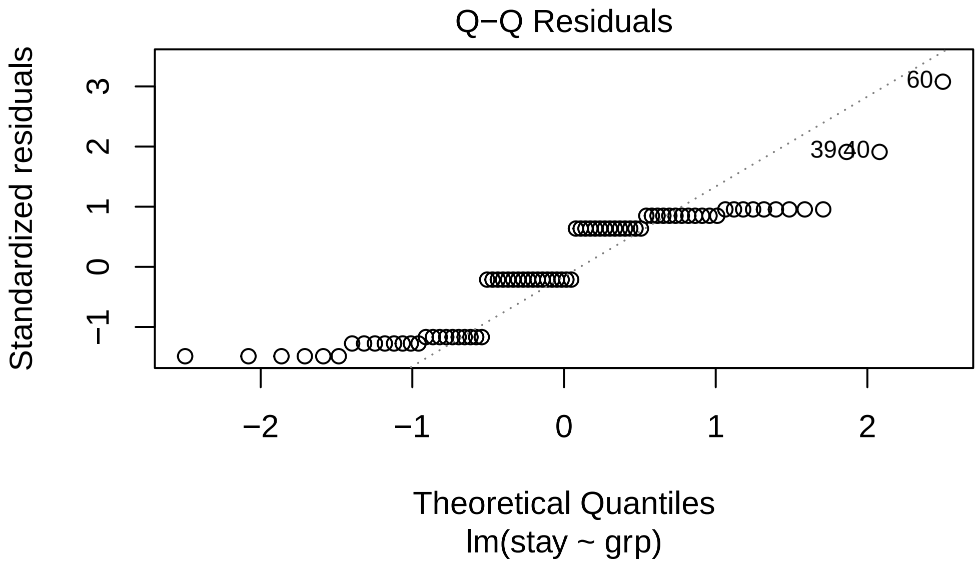

The Normal quantile-quantile plot shows that the responses do not follow Normal distributions around the treatment means. This is no surprise, since the number of days a patients stays in the hospital following a procedure is a discrete, rather than a continuous random variable. Moreover, the response only takes two values in three of the treatment groups and only three values in group C. Because of this we should therefore interpret our results with caution. What may make the results nevertheless trustworthy is that the sample sizes in each group are somewhat large: We have $n = 20$ in each group. The distributions of the group means might be regarded as approximately Normal.

Chp 6 Ex 8

Read in the data:

```
time <- c(9,11,10,9,15,20,21,23,17,30,6,5,8,14,7)
color <- as.factor(c(rep("Red",5),rep("Green",5),rep("Black",5)))
```

(a)

Print the ANOVA table:

```r
lm_out <- lm(time ~ color)
anova(lm_out)
```

```
Analysis of Variance Table

Response: time
          Df Sum Sq Mean Sq F value    Pr(>F)    
color      2 565.73 282.867  20.014 0.0001505 ***
Residuals 12 169.60  14.133                      
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Check the diagnostic plots:

```r
plot(lm_out,which = 1)
```

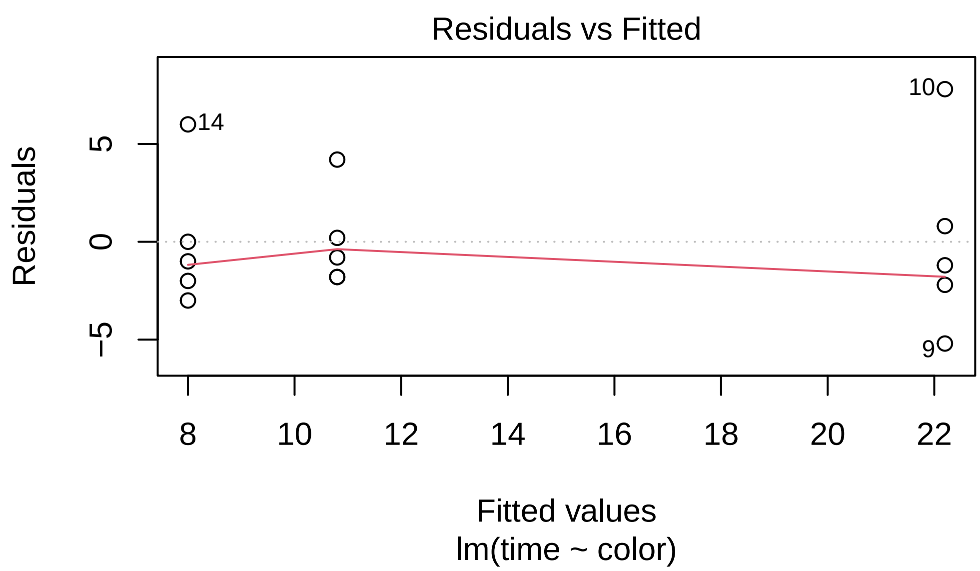


```r
plot(lm_out,which = 2)
```

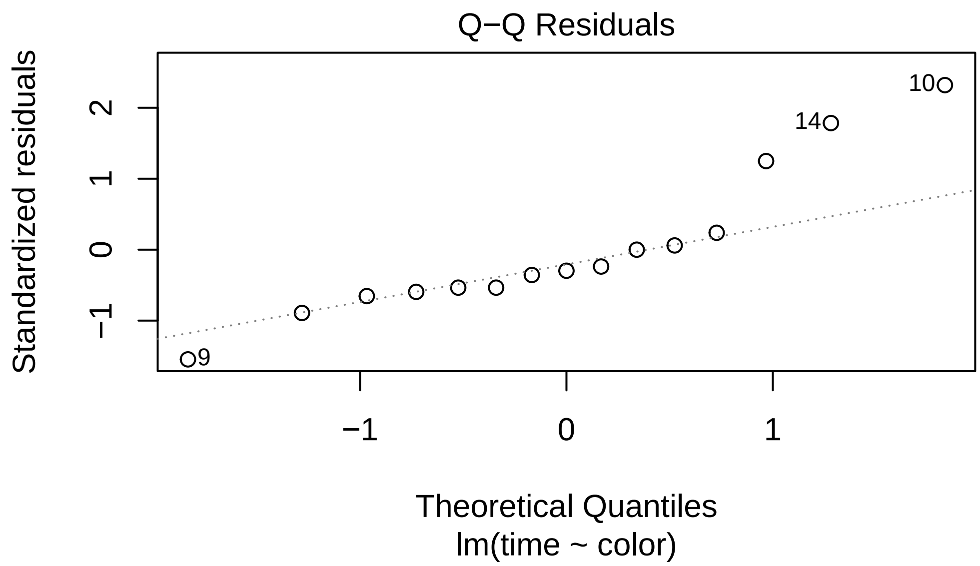

While the Normal quantile-quantile plot exhibits some departure from Normality in the residuals, the residuals versus fitted values plot shows an approximately equal spread of the residuals in each group. It is probably okay to trust our inferences.

The ANOVA table shows a p value for the effect of color of 0.0001505. Since this p value is quite small, there seems to be sufficient evidence in the data to claim that the color of the doors in the maze makes a difference in the average time taken by mice to complete it.

(b)

Tukey's method for comparing all pairs of means is the appropriate method.

```
Tukey_out <- TukeyHSD(aov(time ~ color))
Tukey_out
```

```
  Tukey multiple comparisons of means
    95% family-wise confidence level

Fit: aov(formula = time ~ color)

$color
             diff       lwr      upr     p adj
Green-Black  14.2   7.85669 20.54331 0.0001774
Red-Black     2.8  -3.54331  9.14331 0.4879085
Red-Green   -11.4 -17.74331 -5.05669 0.0011754
```

```r
plot(Tukey_out)
```

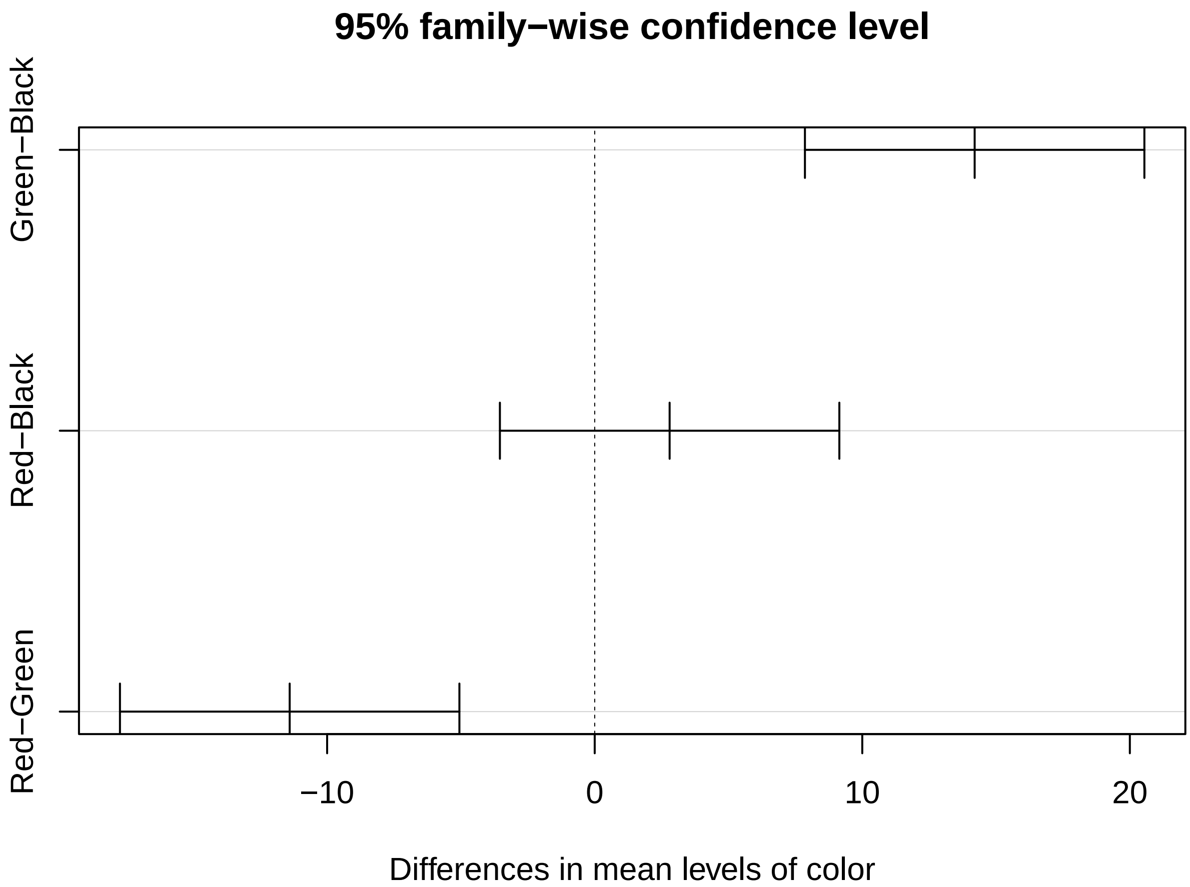

According to the output, there is no significant difference in the completion times of the maze with red versus black doors. However, green doors resulted in a longer average completion time than black doors, and red doors resulted in a shorter average completion time than green doors.

(c)

One should build Dunnett's CIs for comparing all the means to the mean completion time of the maze with green doors.

```r
library(DescTools)
Dunnett_out <- DunnettTest(time ~ color,control = "Green")
Dunnett_out
```

```
  Dunnett's test for comparing several treatments with a control :
    95% family-wise confidence level

$Green
             diff    lwr.ci    upr.ci    pval
Black-Green -14.2 -20.15148 -8.248518 0.00012 ***
Red-Green   -11.4 -17.35148 -5.448518 0.00083 ***

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

```r
plot(Dunnett_out)
```

Based on the output, we conclude that both door colors—black and red—result in faster average maze completion times than green.

Chp 6 Ex 14

We need to make all pairwise comparisons of means to know which paint is the best. Since we have $n = 6$ in each of the $a = 4$ groups, we compute $\hat{\sigma}^2$ as $\hat{\sigma}^2 = \text{MS}_{\text{Error}} = \text{SS}_{\text{Error}}/20$. We obtain this with the code

```
MSE <- 5*(82.7 + 77.9 + 91.0 + 105.2)/(6*4 - 4)
```

We now build confidence intervals of the form

$$\bar{Y}_{1.} - \bar{Y}_{j.} \pm q_{4,20,0.05}\hat{\sigma}/\sqrt{n},$$

where $q_{4,20,0.05}$ is found in Table A.6 of the book to be $q_{4,20,0.05} = 3.96$. Or we can use `qtukey(.95,4,20)` which gives $q_{4,20,0.05} = 3.9582935$.

So the intervals are

```
me <- 3.96 * sqrt(MSE) / sqrt(6)
y1bar <- 48.6
y2bar <- 51.2
y3bar <- 60.1
y4bar <- 55.2
tab <- rbind(c(y1bar - y2bar - me,y1bar - y2bar + me),
             c(y1bar - y3bar - me,y1bar - y3bar + me),
             c(y1bar - y4bar - me,y1bar - y4bar + me),
             c(y2bar - y3bar - me,y2bar - y3bar + me),
             c(y2bar - y4bar - me,y2bar - y4bar + me),
             c(y3bar - y4bar - me,y3bar - y4bar + me))
colnames(tab) <- c("lower","upper")
rownames(tab) <- c("1-2","1-3","1-4","2-3","2-4","3-4")
tab
```

```
        lower     upper
1-2 -17.8687 12.668697
1-3 -26.7687  3.768697
1-4 -21.8687  8.668697
2-3 -24.1687  6.368697
2-4 -19.2687 11.268697
3-4 -10.3687 20.168697
```

We do not find any significant differences, so we cannot say which paint is the best.

Another way to think about this question is to just test H_0: $\mu_1 = \cdots = \mu_6$, where these are the means of the paints.

We can compute the p value of the F test as

```
y.bar <- mean(c(y1bar,y2bar,y3bar,y4bar))
SSA <- 6*((y1bar - y.bar)^2+(y2bar - y.bar)^2+(y3bar - y.bar)^2+(y4bar - y.bar)^2)
```

```r
MSA <- SSA / (4 - 1)
SSE <- MSE * (6*4 - 4)
SST <- SSA + SSE
FA <- MSA / MSE
pA <- 1 - pf(FA,4-1,6*4 - 4)
pA
```

```
[1] 0.2008665
```

Since the p value is quite large, we fail to reject the null hypothesis that all the means are the same. This means we cannot say that any paint is better than any other.