R11 - ANOVA

HCI/PSYCH 522
Iowa State University

April 28, 2022

Outline

Mouse dataset

```
mouse <- read_csv('mouse.csv') %>%
  mutate(Mouse = factor(Mouse),
         Mouse = relevel(Mouse, ref="Dell"))
head(mouse)
```

```
## # A tibble: 6 x 2
##   Skill Mouse
##   <dbl> <fct>
## 1  35.5 Dell
## 2  35.4 Dell
## 3  34.9 Dell
## 4  34.8 Dell
## 5  33.8 Dell
## 6  33.5 Dell
```

```
summary(mouse)
```

```
##      Skill                        Mouse
##  Min.   : 6.4   Dell                :49
##  1st Qu.:31.8   Basilisk (Wired)    :57
##  Median :39.5   Mamba (Wired)       :71
##  Mean   :38.8   Mamba (Wireless)    :56
##  3rd Qu.:46.9   Viper (Wired, light):60
##  Max.   :54.6   Viper (Wired)       :56
```

Mouse graphically

```
ggplot(mouse, aes(x = Mouse, y = Skill)) + geom_jitter(width=0.1)
```

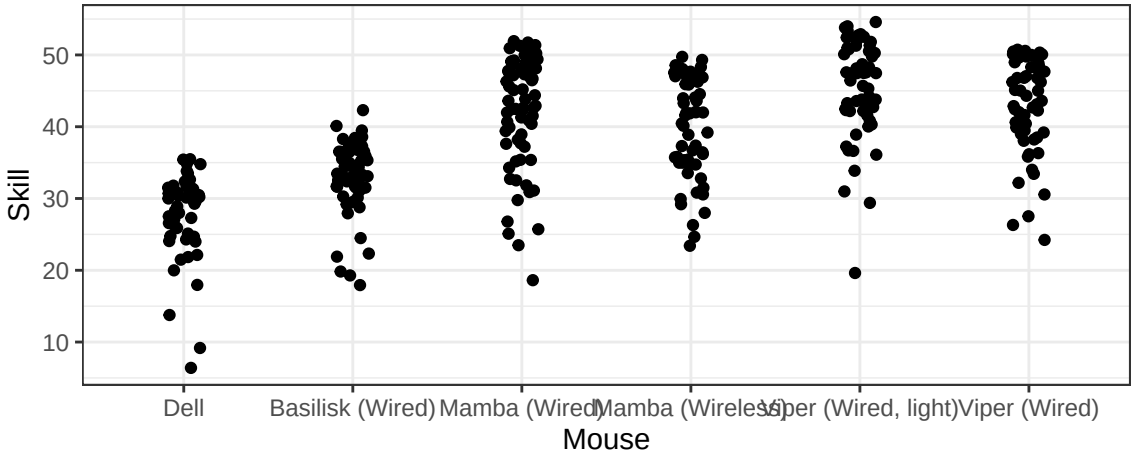

Regression model

```
m <- lm(Skill ~ Mouse, data = mouse)
summary(m)
```

```
##
## Call:
## lm(formula = Skill ~ Mouse, data = mouse)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -25.5167  -3.3857   0.8143   5.1833  10.0143
##
## Coefficients:
##                            Estimate Std. Error t value Pr(>|t|)
## (Intercept)                  27.402      0.954  28.722  < 2e-16 ***
## MouseBasilisk (Wired)         5.289      1.301   4.065 5.95e-05 ***
## MouseMamba (Wired)           14.895      1.240  12.009  < 2e-16 ***
## MouseMamba (Wireless)        12.284      1.306   9.403  < 2e-16 ***
## MouseViper (Wired, light)    17.715      1.286  13.776  < 2e-16 ***
## MouseViper (Wired)           15.484      1.306  11.852  < 2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 6.678 on 343 degrees of freedom
## Multiple R-squared:  0.4543,Adjusted R-squared:  0.4463
## F-statistic:  57.1 on 5 and 343 DF,  p-value: < 2.2e-16
```

Confidence/credible intervals

```
coef(m)
```

```
##               (Intercept)     MouseBasilisk (Wired)        MouseMamba (Wired)     MouseMamba (Wireless)
##                 27.402041                  5.289187                 14.895142                 12.283673
## MouseViper (Wired, light)        MouseViper (Wired)
##                 17.714626                 15.483673
```

```
confint(m)
```

```
##                                2.5 %    97.5 %
## (Intercept)                25.525547 29.278535
## MouseBasilisk (Wired)       2.730232  7.848142
## MouseMamba (Wired)         12.455599 17.334686
## MouseMamba (Wireless)       9.714178 14.853169
## MouseViper (Wired, light) 15.185417 20.243835
## MouseViper (Wired)         12.914178 18.053169
```

Regression model

```
em <- emmeans(m, pairwise ~ Mouse, adjust = "none")
confint(em)
```

```
## $emmeans
##  Mouse                emmean    SE  df lower.CL upper.CL
##  Dell                   27.4 0.954 343     25.5     29.3
##  Basilisk (Wired)       32.7 0.885 343     31.0     34.4
##  Mamba (Wired)          42.3 0.793 343     40.7     43.9
##  Mamba (Wireless)       39.7 0.892 343     37.9     41.4
##  Viper (Wired, light)   45.1 0.862 343     43.4     46.8
##  Viper (Wired)          42.9 0.892 343     41.1     44.6
##
## Confidence level used: 0.95
##
## $contrasts
##  contrast                                estimate   SE  df lower.CL upper.CL
##  Dell - Basilisk (Wired)                   -5.289 1.30 343   -7.848   -2.730
##  Dell - Mamba (Wired)                     -14.895 1.24 343  -17.335  -12.456
##  Dell - Mamba (Wireless)                  -12.284 1.31 343  -14.853   -9.714
##  Dell - Viper (Wired, light)              -17.715 1.29 343  -20.244  -15.185
##  Dell - Viper (Wired)                     -15.484 1.31 343  -18.053  -12.914
##  Basilisk (Wired) - Mamba (Wired)          -9.606 1.19 343  -11.942   -7.270
##  Basilisk (Wired) - Mamba (Wireless)       -6.994 1.26 343   -9.466   -4.523
##  Basilisk (Wired) - Viper (Wired, light)  -12.425 1.24 343  -14.855   -9.996
##  Basilisk (Wired) - Viper (Wired)         -10.194 1.26 343  -12.666   -7.723
##  Mamba (Wired) - Mamba (Wireless)           2.611 1.19 343    0.264    4.959
##  Mamba (Wired) - Viper (Wired, light)      -2.819 1.17 343   -5.123   -0.516
##  Mamba (Wired) - Viper (Wired)             -0.589 1.19 343   -2.936    1.759
##  Mamba (Wireless) - Viper (Wired, light)   -5.431 1.24 343   -7.872   -2.990
```

Regression model

```
summary(m)

##
## Call:
## lm(formula = Skill ~ Mouse, data = mouse)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -25.5167  -3.3857   0.8143   5.1833  10.0143
##
## Coefficients:
##                           Estimate Std. Error t value Pr(>|t|)
## (Intercept)                 27.402      0.954  28.722  < 2e-16 ***
## MouseBasilisk (Wired)        5.289      1.301   4.065 5.95e-05 ***
## MouseMamba (Wired)          14.895      1.240  12.009  < 2e-16 ***
## MouseMamba (Wireless)       12.284      1.306   9.403  < 2e-16 ***
## MouseViper (Wired, light)   17.715      1.286  13.776  < 2e-16 ***
## MouseViper (Wired)          15.484      1.306  11.852  < 2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 6.678 on 343 degrees of freedom
## Multiple R-squared:  0.4543,Adjusted R-squared:  0.4463
## F-statistic:  57.1 on 5 and 343 DF,  p-value: < 2.2e-16
```

Analysis of variance (ANOVA)

```
anova(m)

## Analysis of Variance Table
##
## Response: Skill
##            Df Sum Sq Mean Sq F value    Pr(>F)
## Mouse       5  12734  2546.8  57.104 < 2.2e-16 ***
## Residuals 343  15297    44.6
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Analysis of variance (ANOVA)

```
anova(m)

## Analysis of Variance Table
##
## Response: Skill
##            Df Sum Sq Mean Sq F value    Pr(>F)
## Mouse       5  12734  2546.8  57.104 < 2.2e-16 ***
## Residuals 343  15297    44.6
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

ANOVA F-test: Comparison of models

Regression model with categorical variable:

$$Y_i \overset{ind}{\sim} N(\beta_0 + \beta_1 X_{i,1} + \cdots + \beta_p X_{i,p}, \sigma^2)$$

where $\beta_p, p > 0$ is the difference between mean response in the reference level compared to the level associated with the pth level

F-test:

- Reduced model: no categorical variable $\beta_1 = \cdots = \beta_p = 0$
- Full model: with categorical variable (see above)

ANOVA F-test: Summary

```
anova(m)

## Analysis of Variance Table
##
## Response: Skill
##            Df Sum Sq Mean Sq F value    Pr(>F)
## Mouse       5  12734  2546.8  57.104 < 2.2e-16 ***
## Residuals 343  15297    44.6
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

There is evidence of a difference in mean player skill using different mice ($F_{5,343} = 57, p \approx 0$).

YouTube videos

(hyperlinks)

Playlists:

- Probability
- Inference
- Regression
  - One-way ANOVA
  - F-tests

Bias in jury selection

```
case0502 <- Sleuth3::case0502 %>% mutate(Judge = relevel(Judge, ref="Spock's"))
head(case0502)
```

```
##   Percent   Judge
## 1     6.4 Spock's
## 2     8.7 Spock's
## 3    13.3 Spock's
## 4    13.6 Spock's
## 5    15.0 Spock's
## 6    15.2 Spock's
```

```
summary(case0502)
```

```
##     Percent          Judge
##  Min.   : 6.40   Spock's:9
##  1st Qu.:19.95   A      :5
##  Median :27.50   B      :6
##  Mean   :26.58   C      :9
##  3rd Qu.:32.38   D      :2
##  Max.   :48.90   E      :6
##                  F      :9
```

Bias in jury selection - Plot

```
ggplot(case0502, aes(x = Judge, y = Percent)) + geom_jitter(width=0.1)
```

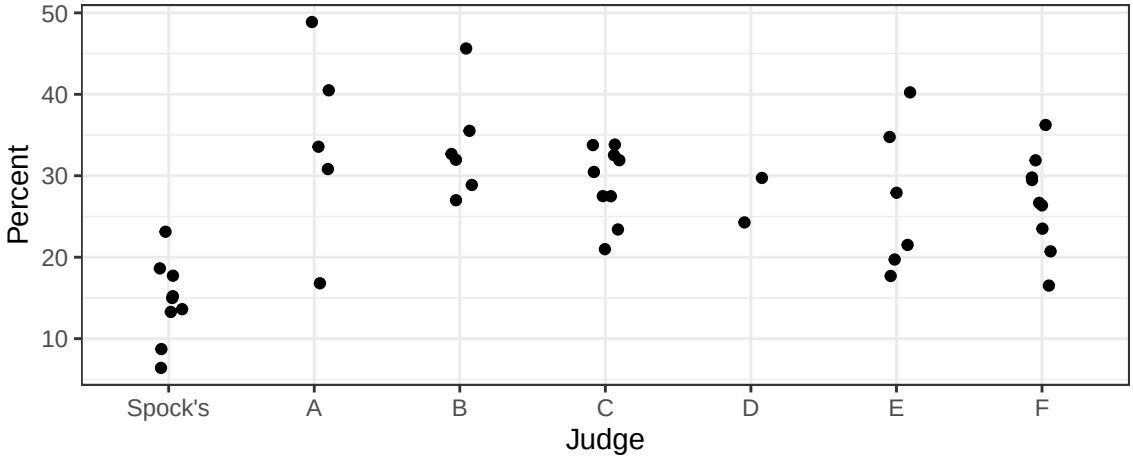

Bias in jury selection - One-way ANOVA

```
m <- lm(Percent ~ Judge, data = case0502)
anova(m)

## Analysis of Variance Table
##
## Response: Percent
##           Df Sum Sq Mean Sq F value    Pr(>F)
## Judge      6 1927.1  321.18  6.7184 6.096e-05 ***
## Residuals 39 1864.5   47.81
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Manuscript statement: There is evidence of a difference in mean percent women on juries amongst the judges ($F_{6,39} = 7, p \approx 0$).

Bias in jury selection - Model summary

```
summary(m)
```

```
##
## Call:
## lm(formula = Percent ~ Judge, data = case0502)
##
## Residuals:
##     Min      1Q  Median      3Q     Max
## -17.320  -4.367  -0.250   3.319  14.780
##
## Coefficients:
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept)   14.622      2.305   6.344 1.72e-07 ***
## JudgeA        19.498      3.857   5.056 1.05e-05 ***
## JudgeB        18.994      3.644   5.212 6.39e-06 ***
## JudgeC        14.478      3.259   4.442 7.15e-05 ***
## JudgeD        12.378      5.405   2.290 0.027513 *
## JudgeE        12.344      3.644   3.388 0.001623 **
## JudgeF        12.178      3.259   3.736 0.000597 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 6.914 on 39 degrees of freedom
## Multiple R-squared:  0.5083,Adjusted R-squared:  0.4326
## F-statistic: 6.718 on 6 and 39 DF,  p-value: 6.096e-05
```

Bias in jury selection - Treatment vs Control

```
em <- emmeans(m, trt.vs.ctrl ~ Judge, adjust = "none")
confint(em)

## $emmeans
##  Judge   emmean   SE df lower.CL upper.CL
##  Spock's   14.6 2.30 39     9.96     19.3
##  A         34.1 3.09 39    27.87     40.4
##  B         33.6 2.82 39    27.91     39.3
##  C         29.1 2.30 39    24.44     33.8
##  D         27.0 4.89 39    17.11     36.9
##  E         27.0 2.82 39    21.26     32.7
##  F         26.8 2.30 39    22.14     31.5
##
## Confidence level used: 0.95
##
## $contrasts
##  contrast    estimate   SE df lower.CL upper.CL
##  A - Spock's     19.5 3.86 39    11.70     27.3
##  B - Spock's     19.0 3.64 39    11.62     26.4
##  C - Spock's     14.5 3.26 39     7.89     21.1
##  D - Spock's     12.4 5.41 39     1.44     23.3
##  E - Spock's     12.3 3.64 39     4.97     19.7
##  F - Spock's     12.2 3.26 39     5.59     18.8
##
## Confidence level used: 0.95
```

Bias in jury selection - Custom contrast

Average of all other judge's percent women minus Spock's.

```
em <- emmeans(m, ~ Judge)
co <- contrast(em, list(`Mean(others) - Spock` = c(-7, rep(1,6))/7))
confint(co)

##  contrast             estimate   SE df lower.CL upper.CL
##  Mean(others) - Spock     10.7 2.56 39     5.58     15.9
##
## Confidence level used: 0.95
```

Seaweed grazers

```
head(case1301)
```

```
##   Cover Block Treat
## 1    14    B1     C
## 2    23    B1     C
## 3    22    B2     C
## 4    35    B2     C
## 5    67    B3     C
## 6    82    B3     C
```

```
summary(case1301)
```

```
##      Cover           Block    Treat
##  Min.   : 1.00   B1     :12   C  :16
##  1st Qu.: 9.00   B2     :12   L  :16
##  Median :22.50   B3     :12   Lf :16
##  Mean   :28.62   B4     :12   LfF:16
##  3rd Qu.:42.25   B5     :12   f  :16
##  Max.   :95.00   B6     :12   fF :16
##                  (Other):24
```

Seaweed grazers

```
ggplot(case1301, aes(x = Block, y = Cover, shape=Treat, color=Treat)) + geom_jitter(width=0.1, height=0)
```

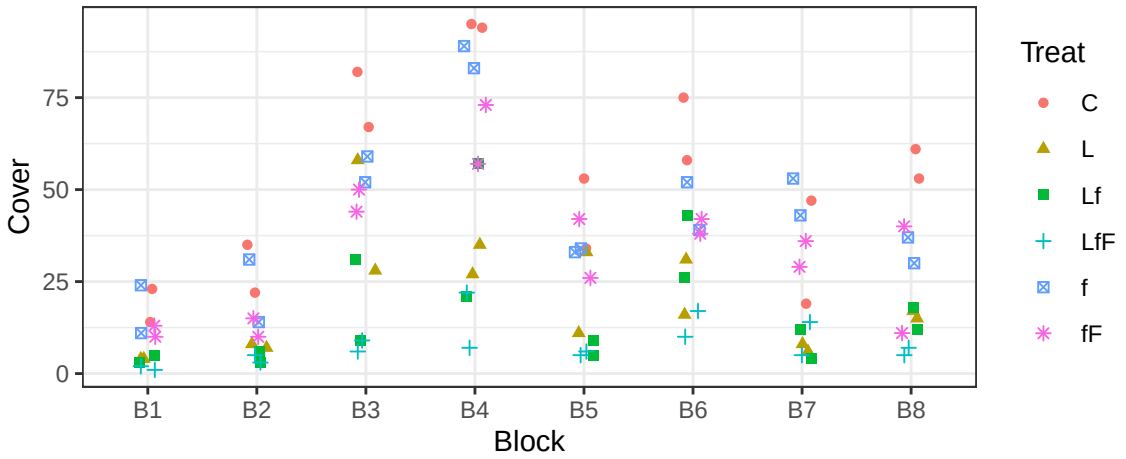

Seaweed grazers

```
m <- lm(Cover ~ Block + Treat, data = case1301)
summary(m)

##
## Call:
## lm(formula = Cover ~ Block + Treat, data = case1301)
##
## Residuals:
##     Min      1Q  Median      3Q     Max
## -27.375  -5.812   0.625   5.438  26.125
##
## Coefficients:
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept)   32.875      4.217   7.795 1.66e-11 ***
## BlockB2        3.750      4.679   0.801 0.425132
## BlockB3       31.750      4.679   6.786 1.59e-09 ***
## BlockB4       45.500      4.679   9.725 2.32e-15 ***
## BlockB5       14.750      4.679   3.153 0.002253 **
## BlockB6       27.750      4.679   5.931 6.67e-08 ***
## BlockB7       13.500      4.679   2.885 0.004980 **
## BlockB8       16.000      4.679   3.420 0.000974 ***
## TreatL       -32.750      4.052  -8.083 4.45e-12 ***
## TreatLf      -35.500      4.052  -8.761 1.96e-13 ***
## TreatLfF     -44.250      4.052 -10.921  < 2e-16 ***
## Treatf        -9.250      4.052  -2.283 0.024995 *
## TreatfF      -18.500      4.052  -4.566 1.71e-05 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
```

Seaweed grazers

```
anova(m)

## Analysis of Variance Table
##
## Response: Cover
##           Df Sum Sq Mean Sq F value    Pr(>F)
## Block      7  19106  2729.4  20.780 6.977e-16 ***
## Treat      5  23046  4609.1  35.092 < 2.2e-16 ***
## Residuals 83  10902   131.3
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Two-way ANOVA

Regression model with two categorical variable:

$$Y_i \overset{ind}{\sim} N(\beta_0 + \underbrace{\beta_1 X_{i,1} + \cdots + \beta_p X_{i,p}}_{\text{variable 1}} + \underbrace{\beta_{p+1} X_{i,p+1} + \cdots + \beta_{p+q} X_{i,p+q}}_{\text{variable 2}}, \sigma^2)$$

where

- $\beta_r, 1 \leq r \leq p$ is the difference between mean response in the reference level compared to the level associated with the rth level
- $\beta_r, p+1 \leq r \leq p+q$ is the difference between mean response in the reference level compared to the level associated with the rth level

F-tests:

1. Variable 1
   - Reduced model: no categorical variables $\beta_1 = \cdots = \beta_{p+q} = 0$
   - Full model: with first variable $\beta_{p+1} = \cdots = \beta_{p+q} = 0$
2. Variable 2
   - Reduced model: with first variable $\beta_{p+1} = \cdots = \beta_{p+q} = 0$
   - Full model: with both variables (see model above)

Seaweed grazers

```
anova(m)

## Analysis of Variance Table
##
## Response: Cover
##           Df Sum Sq Mean Sq F value    Pr(>F)
## Block      7  19106  2729.4  20.780 6.977e-16 ***
## Treat      5  23046  4609.1  35.092 < 2.2e-16 ***
## Residuals 83  10902   131.3
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Manuscript statements:

- There is evidence of a difference in mean cover amongst the blocks ($F_{7,83} = 21, p \approx 0$).
- There is evidence of a difference in mean cover amongst the treatments after controlling for blocks ($F_{5,83} = 35, p \approx 0$).

Pygmalion effect

```
head(case1302)
```

```
##   Company     Treat Score
## 1      C1 Pygmalion  80.0
## 2      C1   Control  63.2
## 3      C1   Control  69.2
## 4      C2 Pygmalion  83.9
## 5      C2   Control  63.1
## 6      C2   Control  81.5
```

```
summary(case1302)
```

```
##     Company         Treat        Score
##  C1     : 3   Control  :19   Min.   :59.50
##  C10    : 3   Pygmalion:10   1st Qu.:69.20
##  C2     : 3                  Median :73.90
##  C4     : 3                  Mean   :74.07
##  C5     : 3                  3rd Qu.:78.90
##  C6     : 3                  Max.   :89.80
##  (Other):11
```

Pygmalion effect

```
ggplot(case1302, aes(x = Company, y = Score, shape=Treat, color=Treat)) + geom_jitter(width=0.1, height=0)
```

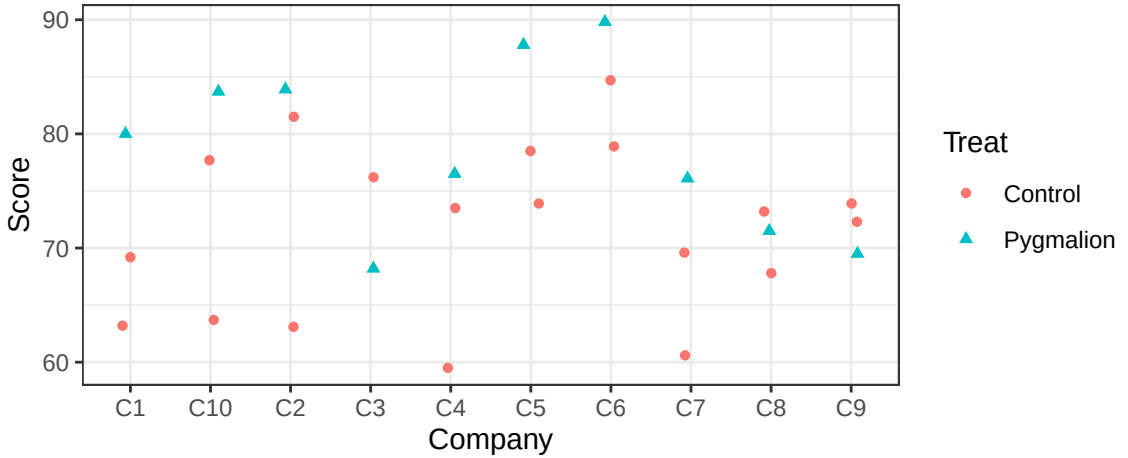

Pygmalion effect

```
m <- lm(Score ~ Company + Treat, data = case1302)
summary(m)
```

```
##
## Call:
## lm(formula = Score ~ Company + Treat, data = case1302)
##
## Residuals:
##     Min      1Q  Median      3Q     Max
## -10.660  -4.147   1.853   3.853   7.740
##
## Coefficients:
##                Estimate Std. Error t value Pr(>|t|)
## (Intercept)    68.39316    3.89308  17.568 8.92e-13 ***
## CompanyC10      4.23333    5.36968   0.788   0.4407
## CompanyC2       5.36667    5.36968   0.999   0.3308
## CompanyC3       0.19658    6.01886   0.033   0.9743
## CompanyC4      -0.96667    5.36968  -0.180   0.8591
## CompanyC5       9.26667    5.36968   1.726   0.1015
## CompanyC6      13.66667    5.36968   2.545   0.0203 *
## CompanyC7      -2.03333    5.36968  -0.379   0.7094
## CompanyC8       0.03333    5.36968   0.006   0.9951
## CompanyC9       1.10000    5.36968   0.205   0.8400
## TreatPygmalion  7.22051    2.57951   2.799   0.0119 *
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 6.576 on 18 degrees of freedom
## Multiple R-squared:  0.5647,Adjusted R-squared:  0.3228
```

Pygmalion effect

```
anova(m)

## Analysis of Variance Table
##
## Response: Score
##           Df Sum Sq Mean Sq F value  Pr(>F)
## Company    9 670.98   74.55  1.7238 0.15556
## Treat      1 338.88  338.88  7.8354 0.01186 *
## Residuals 18 778.50   43.25
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Pygmalion Effect

```
anova(m)

## Analysis of Variance Table
##
## Response: Score
##           Df Sum Sq Mean Sq F value  Pr(>F)
## Company    9 670.98   74.55  1.7238 0.15556
## Treat      1 338.88  338.88  7.8354 0.01186 *
## Residuals 18 778.50   43.25
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Manuscript statements:

- There is no evidence of a difference in mean score amongst the companies ($F_{9,18} = 2, p = 0.16$).
- There is evidence of a difference in mean score amongst the treatments after controlling for company ($F_{1,18} = 8, p = 0.01$).

Pygmalion Effect

```
drop1(m, test="F")

## Single term deletions
##
## Model:
## Score ~ Company + Treat
##         Df Sum of Sq    RSS    AIC F value  Pr(>F)
## <none>                778.5 117.41
## Company  9    682.52 1461.0 117.67  1.7534 0.14844
## Treat    1    338.88 1117.4 125.89  7.8354 0.01186 *
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Manuscript statements:

- There is no evidence of a difference in mean score amongst the companies after controlling for treatment ($F_{9,18} = 2, p = 0.15$).
- There is evidence of a difference in mean score amongst the treatments after controlling for company ($F_{1,18} = 8, p = 0.01$).

ANOVA Tables

- Sequential comparisons
  - Adds new variable to model that already includes variables above it
  - Use `anova()` in R
  - SAS Type I sums of squares
- Partial comparisons
  - Removes variable from model that includes all other variables
  - Use `drop1()` in R
  - SAS Type III sums of squares
  - SPSS default
- Suggestions
  - Results are the same for complete, balanced experiments
  - Always include variables that were part of the experimental design
  - Generally prefer `drop1()`

Three-way ANOVA

```r
m <- lm(log(Forrest) ~ Stress + factor(SO2) + factor(O3), data = case1402)
drop1(m, test="F")

## Single term deletions
##
## Model:
## log(Forrest) ~ Stress + factor(SO2) + factor(O3)
##             Df Sum of Sq     RSS      AIC F value    Pr(>F)
## <none>                   0.51159 -106.143
## Stress       1   0.00804 0.51963 -107.675  0.3456 0.5625887
## factor(SO2)  2   0.06346 0.57505 -106.635  1.3646 0.2762792
## factor(O3)   4   0.75818 1.26977  -86.871  8.1510 0.0003437 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Interactions

```
m <- lm(log(Forrest) ~ .^2, data = case1402 %>% select(-William))
drop1(m, test="F")

## Single term deletions
##
## Model:
## log(Forrest) ~ (Stress + SO2 + O3)^2
##            Df Sum of Sq     RSS     AIC F value Pr(>F)
## <none>                  0.50479 -108.54
## Stress:SO2  1  0.011613 0.51640 -109.86  0.5291 0.4743
## Stress:O3   1  0.012599 0.51739 -109.81  0.5740 0.4564
## SO2:O3      1  0.014743 0.51953 -109.68  0.6717 0.4209
```