R Code: F tests for ANOVA, More Nested Models

(Sleuth3 Sections 5.3 and 5.4)

Spock Trials Example

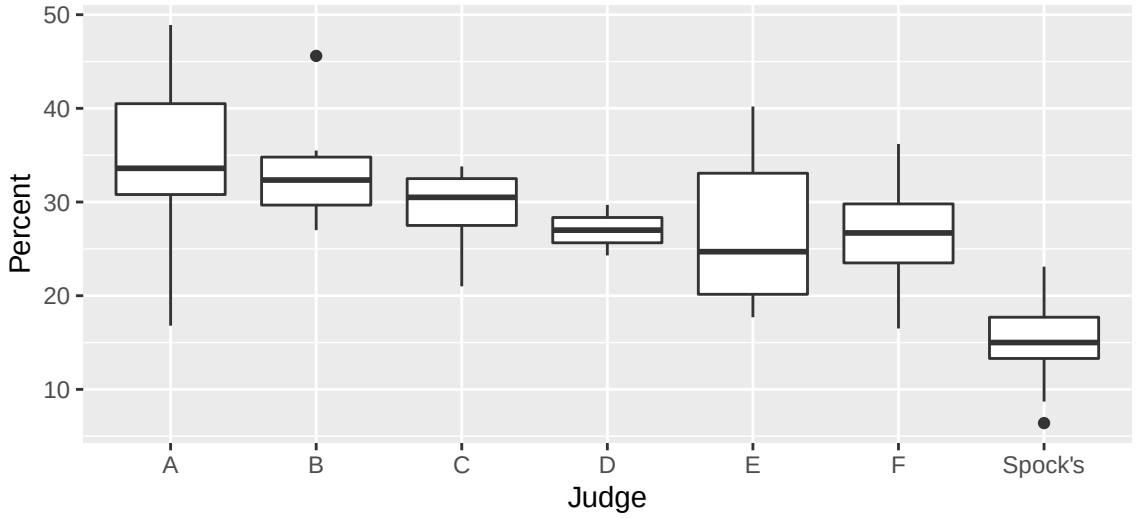

Question: are the means for judges A through F equal?

$H_0 : \mu_1 = \mu_2 = \mu_3 = \mu_4 = \mu_5 = \mu_6$

H_A : At least one mean for a judge other than Spock is not equal to the others.

Note: if the null hypothesis is correct, there are effectively two groups:

1. All judges other than Spock's (they all have the same mean!)
2. Spock's judge

Tasks we need to do:

1. Fit the full model
2. Create a new variable in the data set representing the groups in the reduced model
3. Fit the reduced model
4. Call `anova` to compare the reduced model to the full model

Step 1: Fit the full model

```
fit_full <- lm(Percent ~ Judge, data = juries)
summary(fit_full)
```

```
##
## Call:
## lm(formula = Percent ~ Judge, data = juries)
##
## 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)   34.1200     3.0921  11.034 1.47e-13 ***
## JudgeB        -0.5033     4.1868  -0.120   0.9049
## JudgeC        -5.0200     3.8566  -1.302   0.2007
## JudgeD        -7.1200     5.7848  -1.231   0.2258
## JudgeE        -7.1533     4.1868  -1.709   0.0955 .
## JudgeF        -7.3200     3.8566  -1.898   0.0651 .
## JudgeSpock's -19.4978     3.8566  -5.056 1.05e-05 ***
## ---
## 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
```

Step 2: Create a new variable in the data set representing the groups in the reduced model

The mutate statement below creates a new variable in the `juries` data frame called `judge_reduced`. For each observation in the data set, `judge_reduced` is `"Other"` if the Judge was in `"A"`, `"B"`, `"C"`, `"D"`, `"E"`, or `"F"`. Otherwise, `judge_reduced` is `"Spock's"`.

```
juries <- juries %>%
  mutate(
    judge_reduced = ifelse(Judge %in% c("A", "B", "C", "D", "E", "F"), "Other", "Spock's")
  )
head(juries, 15) # just to check and make sure our new variable was created correctly
```

```
## # A tibble: 15 x 3
##    Percent Judge   judge_reduced
##      <dbl> <chr>   <chr>
##  1     6.4 Spock's Spock's
##  2     8.7 Spock's Spock's
##  3    13.3 Spock's Spock's
##  4    13.6 Spock's Spock's
##  5    15   Spock's Spock's
##  6    15.2 Spock's Spock's
##  7    17.7 Spock's Spock's
##  8    18.6 Spock's Spock's
##  9    23.1 Spock's Spock's
## 10    16.8 A       Other
## 11    30.8 A       Other
## 12    33.6 A       Other
## 13    40.5 A       Other
## 14    48.9 A       Other
## 15    27   B       Other
```

Step 3: Fit the reduced model

```
fit_reduced <- lm(Percent ~ judge_reduced, data = juries)
summary(fit_reduced)
```

```
##
## Call:
## lm(formula = Percent ~ judge_reduced, data = juries)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -12.9919  -4.6669   0.2581   3.7854  19.4081
##
## Coefficients:
##                       Estimate Std. Error t value Pr(>|t|)
## (Intercept)             29.492      1.160   25.42  < 2e-16 ***
## judge_reducedSpock's  -14.870      2.623   -5.67 1.03e-06 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 7.056 on 44 degrees of freedom
## Multiple R-squared:  0.4222, Adjusted R-squared:  0.409
## F-statistic: 32.15 on 1 and 44 DF,  p-value: 1.03e-06
```

Step 4: Call anova to compare the reduced model to the full model

```
anova(fit_reduced, fit_full)
```

```
## Analysis of Variance Table
##
## Model 1: Percent ~ judge_reduced
## Model 2: Percent ~ Judge
##   Res.Df    RSS Df Sum of Sq      F Pr(>F)
## 1     44 2190.9
## 2     39 1864.5  5    326.46 1.3658 0.2582
```

Because interpretation is important: what is the conclusion? Does this prove that the other judges had the same mean percent of women in their jury pools?

More F Test Examples (part 1)

In all cases the full model has a separate mean for all 7 judges: μ_1 for judge A, μ_2 for judge B, $\cdots$ μ_6 for judge F, and μ_7 for Spock's judge.

We estimate this model with: `fit_full <- lm(Percent ~ Judge, data = juries)`

The sample size is n = 46, so the degrees of freedom for the full model is: 46 - 7 = 39

Null Hypothesis	Reduced Model Groups	Reduced df, Extra df	R Code and Output
$\mu_1 = \mu_2 = \mu_3 = \mu_4 = \mu_5 = \mu_6 = \mu_7$	1 group: • all judges	Reduced: 46-1=45 Extra: 45 - 39 = 6 7 - 1 = 6	**anova(fit_full)** 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
$\mu_1 = \mu_2 = \mu_3 = \mu_4 = \mu_5 = \mu_6 = \mu_7$	1 group: • all judges	Reduced: 46-1=45 Extra: 45 - 39 = 6 7 - 1 = 6	**fit_reduced <- lm(Percent ~ 1, data = juries)** **anova(fit_reduced, fit_full)** Analysis of Variance Table Model 1: Percent ~ 1 Model 2: Percent ~ Judge Res.Df RSS Df Sum of Sq F Pr(>F) 1 45 3791.5 2 39 1864.4 6 1927.1 6.7184 6.096e-05 ***
$\mu_1 = \mu_2 = \mu_3 = \mu_4 = \mu_5 = \mu_6$	2 groups: • A,B,C,D,E,F • Spock's Judge	Reduced: 46-2=44 Extra: 44 - 39 = 5 7 - 2 = 5	**juries <- juries %>%** **mutate(** **judges_grouped = ifelse(** **Judge %in% c("A", "B", "C", "D", "E", "F"),** **"grouped", Judge)** **)** **fit_reduced <- lm(Percent ~ judges_grouped, data = juries)** **anova(fit_reduced, fit_full)** Analysis of Variance Table Model 1: Percent ~ judges_grouped Model 2: Percent ~ Judge Res.Df RSS Df Sum of Sq F Pr(>F) 1 44 2190.9 2 39 1864.5 5 326.46 1.3658 0.2582

More F Test Examples (part 2)

In all cases the full model has a separate mean for all 7 judges: μ_1 for judge A, μ_2 for judge B, $\cdots$ μ_6 for judge F, and μ_7 for Spock's judge.

We estimate this model with: `fit_full <- lm(Percent ~ Judge, data = juries)`

The sample size is n = 46, so the degrees of freedom for the full model is: 46 - 7 = 39

Null Hypothesis	Reduced Model Groups	Reduced df, Extra df	R Code and Output
$\mu_1 = \mu_2 = \mu_3 = \mu_4 = \mu_5$	3 groups: • A,B,C,D,E • F • Spock's Judge	Reduced: 46-3=43 Extra: 43 - 39 = 4 7 - 3 = 4	**juries <- juries %>%** ** mutate(** ** judges_grouped = ifelse(** ** Judge %in% c("A", "B", "C", "D", "E"),** ** "grouped", Judge)** **)** **fit_reduced <- lm(Percent ~ judges_grouped, data = juries)** **anova(fit_reduced, fit_full)** Analysis of Variance Table Model 1: Percent ~ judges_grouped Model 2: Percent ~ Judge Res.Df RSS Df Sum of Sq F Pr(>F) 1 43 2104.7 2 39 1864.5 4 240.28 1.2565 0.3035
$\mu_1 = \mu_2 = \mu_3 = \mu_4$	4 groups: • A,B,C,D • E • F • Spock's Judge	Reduced: 46-4=42 Extra: 42 - 39 = 3 7 - 4 = 3	**juries <- juries %>%** ** mutate(** ** judges_grouped = ifelse(** ** Judge %in% c("A", "B", "C", "D"),** ** "grouped", Judge)** **)** **fit_reduced <- lm(Percent ~ judges_grouped, data = juries)** **anova(fit_reduced, fit_full)** Analysis of Variance Table Model 1: Percent ~ judges_grouped Model 2: Percent ~ Judge Res.Df RSS Df Sum of Sq F Pr(>F) 1 42 2016.9 2 39 1864.5 3 152.5 1.0633 0.3758