Spock Example

David Gerard

2018-12-07

Objectives

- Demonstrate how to interact with ANOVA in R.
- Analyze the Spock trial data in R.

Load in Data

```
library(Sleuth3)
library(ggplot2)
data("case0502")
```

Spock EDA

```
qplot(Judge, Percent, data = case0502)
```

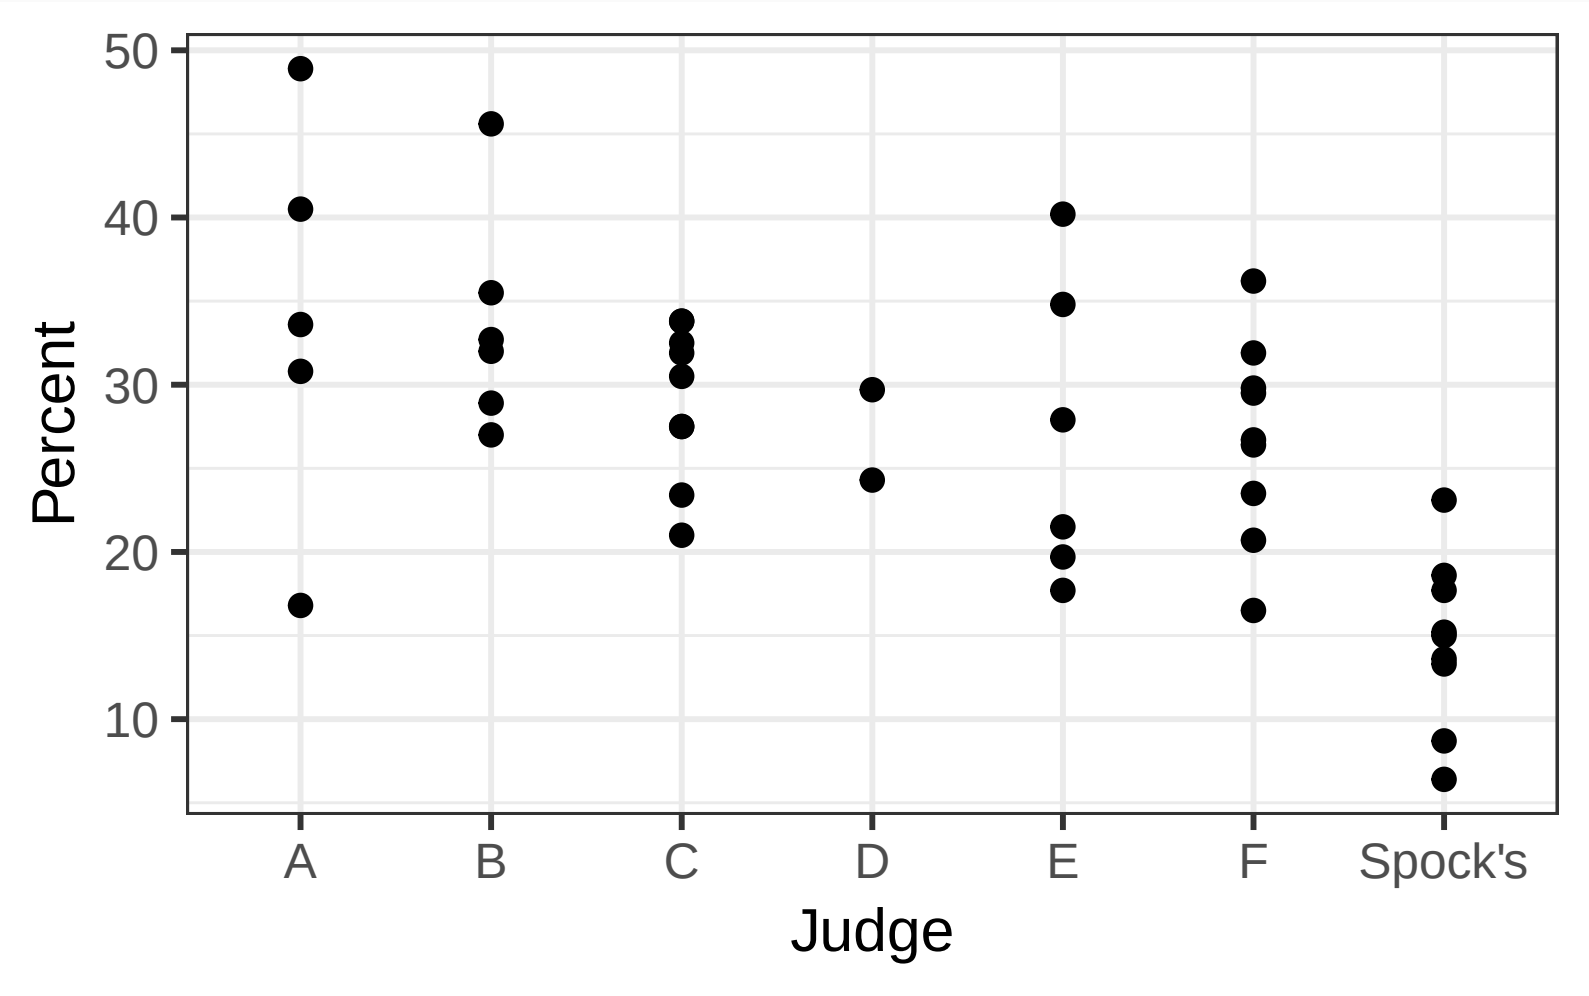

Spock EDA

```
qplot(Judge, Percent, data = case0502, geom = "boxplot")
```

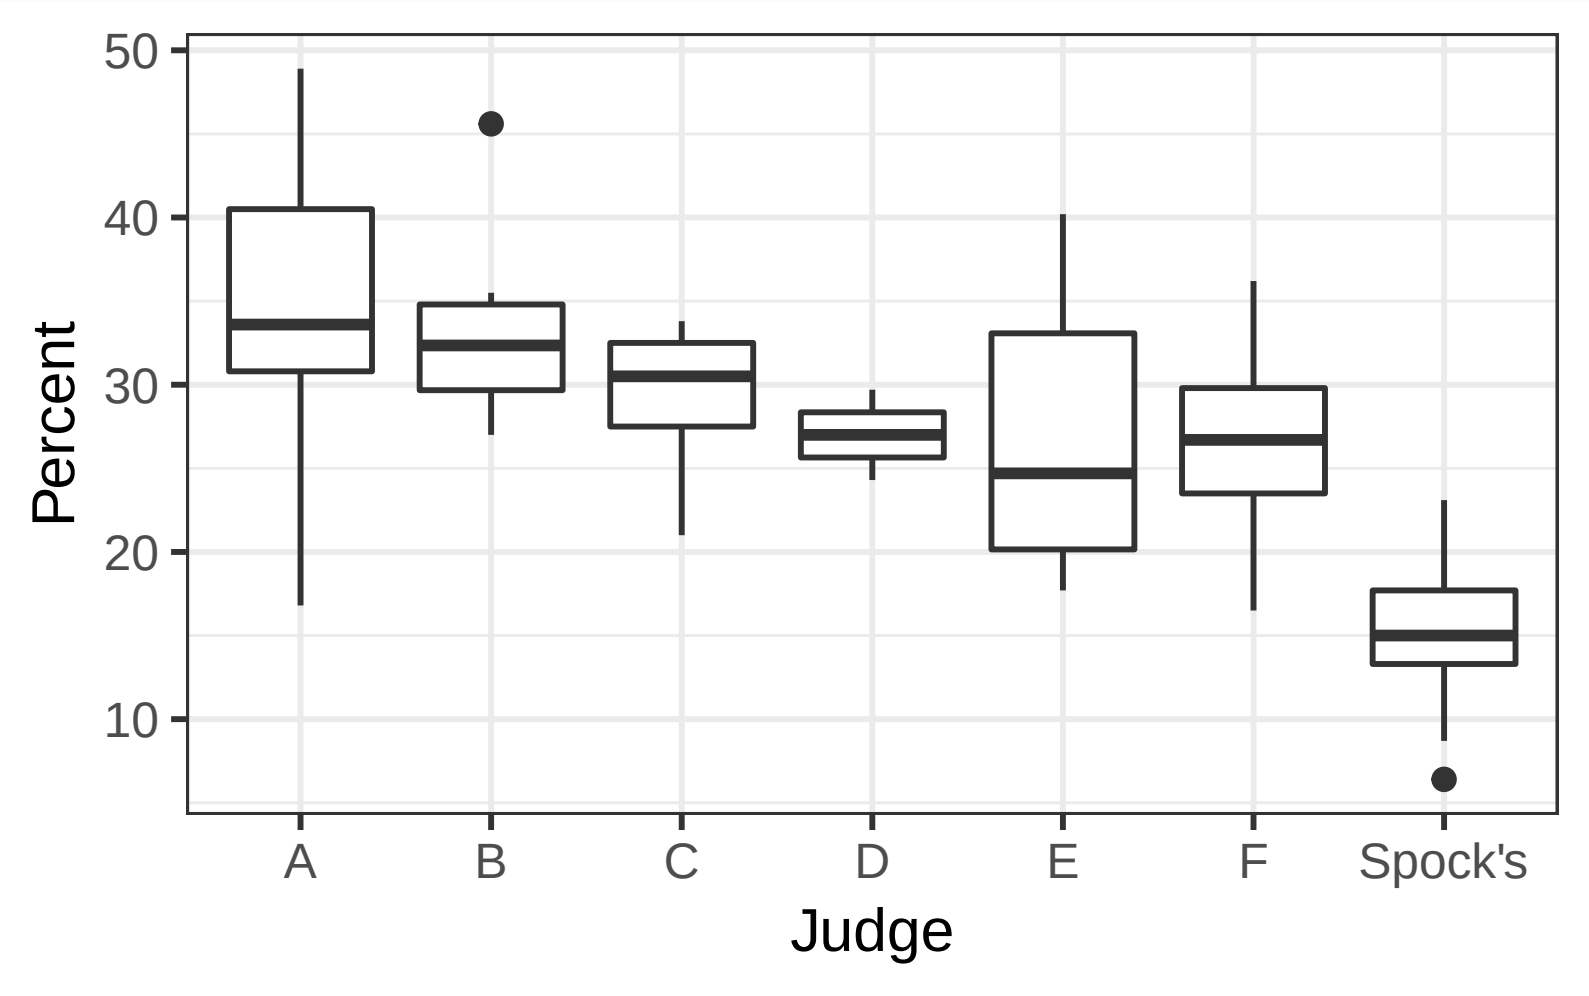

Before Fitting

- Always make sure the grouping variable (the explanatory variable) is either a "`factor`", a "`character`", or a "`logical`" with the `class()` function.

```
class(case0502$Judge)
```

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

- Things will go wrong if this is any other type.
- You can force a variable to be a `factor` with the `as.factor()` function:

```
case0502$Judge <- as.factor(case0502$Judge)
```

Fit the full model

- Use `aov()` function (for **A**nalysis **O**f **V**ariance) to fit the model that assumes μ_1, μ_2, ..., μ_I are all *different*.
- Always save this output.
- The **response** variable goes on the left of the tilde (~) and the **explanatory** variable goes to the right of the tilde.

Fit the full model

```
aout_alldiff <- aov(Percent ~ Judge, data = case0502)
aout_alldiff
```

```
## Call:
##    aov(formula = Percent ~ Judge, data = case0502)
##
## Terms:
##                    Judge Residuals
## Sum of Squares    1927.2    1863.8
## Deg. of Freedom        6        39
##
## Residual standard error: 6.914
## Estimated effects may be unbalanced
```

Get p-values

- Apply `summary()` to this output to run the omnibus F-test

```
summary(aout_alldiff)
```

```
##              Df Sum Sq Mean Sq F value  Pr(>F)
## Judge         6   1927     321    6.72 6.1e-05
## Residuals    39   1864      48
```

What is that Table?

- H_0: $\mu_1 = \mu_2 = \mu_3 = \mu_4 = \mu_5 = \mu_6 = \mu_7$
- H_A: At least two means are different.

```
##              Df Sum Sq Mean Sq F value  Pr(>F)
## Judge         6   1927     321    6.72 6.1e-05
## Residuals    39   1864      48
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
Judge	df_{extra}	ESS	ESS/df_{extra}	F-statistic	p-value
Residuals	df_{full}	RSS_{full}	$RSS_{full}/df_{full} = s_p^2$		

Pairwise Comparisons

- Use `pairwise.t.test()` to perform all pairwise comparisons, and then just extract the ones that you want.

Pairwise Comparisons

```
ptout <- pairwise.t.test(x = case0502$Percent,
                         g = case0502$Judge,
                         p.adjust.method = "none")
ptout
```

```
##
##  Pairwise comparisons using t tests with pooled SD
##
## data:  case0502$Percent and case0502$Judge
##
##          A     B     C     D     E     F
## B        0.905 -     -     -     -     -
## C        0.201 0.223 -     -     -     -
## D        0.226 0.248 0.700 -     -     -
## E        0.095 0.104 0.562 0.995 -     -
## F        0.065 0.069 0.485 0.971 0.964 -
## Spock's 1e-05 6e-06 7e-05 0.028 0.002 6e-04
##
## P value adjustment method: none
```

General Nested Comparisons

- To run tests that do not include the "all equal model", you must first *fit* both models using `aov()`, then run the `anova()` command.
- To compare the full model where all means are different to the reduced model where the non-Spock-judges have the same mean, fit a new model of Spock's judge vs the rest of the judges.

```
case0502$isSpock <- case0502$Judge == "Spock's"
case0502$isSpock
```

```
##  [1]  TRUE  TRUE  TRUE  TRUE  TRUE  TRUE  TRUE  TRUE  TR
## [12] FALSE FALSE FALSE FALSE FALSE FALSE FALSE FALSE FAL
## [23] FALSE FALSE FALSE FALSE FALSE FALSE FALSE FALSE FAL
## [34] FALSE FALSE FALSE FALSE FALSE FALSE FALSE FALSE FAL
## [45] FALSE FALSE
```

General Nested Comparisons

- Now fit the reduced model where $\mu_2 = \mu_3 = \cdots = \mu_7$

```
aout_otherssame <- aov(Percent ~ isSpock, data = case0502)
```

General Nested Comparisons

- Use `anova()` with both the full and reduced models to get the appropriate ANOVA table.

```
anova(aout_otherssame, aout_alldiff)
```

```
## Analysis of Variance Table
##
## Model 1: Percent ~ isSpock
## Model 2: Percent ~ Judge
##   Res.Df  RSS Df Sum of Sq    F Pr(>F)
## 1     44 2191
## 2     39 1864  5       326 1.37   0.26
```

What is that Table?

- H_0: $\mu_1 \neq \mu_2 = \mu_3 = \mu_4 = \mu_5 = \mu_6 = \mu_7$
- H_A: At least two means are different among judges 2 through 7.

```
## Analysis of Variance Table
##
## Model 1: Percent ~ isSpock
## Model 2: Percent ~ Judge
##   Res.Df  RSS Df Sum of Sq    F Pr(>F)
## 1     44 2191
## 2     39 1864  5       326 1.37   0.26
```

	Res.Df	RSS	Df	Sum of Sq	F	Pr(>F)
1	$df_{reduced}$	$RSS_{reduced}$				
2	dr_{full}	RSS_{full}	df_{extra}	ESS	F-statistic	p-value

More nested comparisons

- use `anova()` with more than 2 models to get the nested sum of squares.

```
aout_allequal <- aov(Percent ~ 1, data = case0502)
anova(aout_allequal, aout_otherssame, aout_alldiff)
```

```
## Analysis of Variance Table
##
## Model 1: Percent ~ 1
## Model 2: Percent ~ isSpock
## Model 3: Percent ~ Judge
##   Res.Df  RSS Df Sum of Sq     F Pr(>F)
## 1     45 3792
## 2     44 2191  1      1601 33.48  1e-06
## 3     39 1864  5       326  1.37   0.26
```