

PSYCH 710

Comparing Means in a 1-Way Design

Sept, 2023

Prof. Patrick Bennett

2

1-way ANOVA

1-way ANOVA

overview

- Comparisons of nested linear models
- Interpretations of ANOVA tables
- Mean Squares are estimates of variance
- effect size & association strength
- Assumptions of ANOVA & alternatives

nested linear models

## Linear Models

- ANOVA fits & compares several nested, linear models

- $Y_{ij} = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \dots + e_{ij}$
- $Y_{ij}$  is score "i" in group "j"
- $X_i$  are a predictor variables (e.g., representing groups/conditions)
- $\beta_i$  are coefficients adjusted to minimize  $\sum e_{ij}^2$

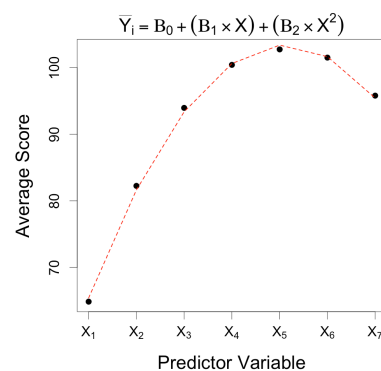

## Linear Models

- Linear models are a very broad class of models
- can be used to characterize curvilinear associations between Y and X's

- $Y_{ij} = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \dots + e_{ij}$
- $Y_{ij} = \beta_0 + \beta_1 X_1 + \beta_2 X_2^2 + \dots + e_{ij}$
- $Y_{ij} = \beta_0 + \beta_1 \exp(X_1) + \beta_2 \log(X_2) + \dots + e_{ij}$

## Linear Models

include linear regression & multiple regression



## Linear Models

Least-squares estimate of model parameters

$$Y_{ij} = \underbrace{\beta_0}_{\text{intercept}} + \beta_1 X_1 + \beta_2 X_2 + \dots + \beta_p X_p + \underbrace{e_{ij}}_{\text{residual}}$$

$\hat{Y}_i$  predicted score

$e_{ij} = Y_{ij} - \hat{Y}_{ij}$  residual = difference between observed & predicted scores

$\sum e_{ij}^2$  Least-squares measure of goodness-of-fit: sum of squared-residuals

ANOVA models are fit using the criterion of least-squares

## Comparing Nested Linear Models

- Compare nested linear models
  - models vary in complexity and how well they fit the data
  - select model that provides best fit with fewest parameters
- Define “complexity” as number of parameters (i.e., predictor variables)
- Define “goodness-of-fit” as sum of squared residuals
  - measure effect of removing free parameters on goodness-of-fit
  - if change is small, keep simpler, *reduced* model
  - if change is large, do not remove parameters & keep *full* model

## Best-fitting (least-squares) parameters

( for a 1-way design )

### Nested Linear Models

$$\text{full model: } Y_{ij} = \mu + \alpha_j + e_{ij}$$

$$\text{reduced model: } Y_{ij} = \mu + e_{ij}$$

### Group effects (alphas)

- Defined as difference between group mean & overall mean
  - with this definition, alphas MUST sum to zero
  - often called “sigma” or “sum-to-zero” definition of alphas
- There are **other definitions** of group effects & intercept that yield the same, minimum value of the sum of squared residuals.

$$\sum_j \alpha_j = 0$$

## Best-fitting (least-squares) parameters

( for a 1-way design )

### Nested Linear Models

$$\text{full model: } Y_{ij} = \mu + \alpha_j + e_{ij}$$

$$\text{reduced model: } Y_{ij} = \mu + e_{ij}$$

- Estimate parameters of models from sample

- which values of  $\mu$  and  $\alpha_j$  minimize  $\sum e_{ij}^2$ ?

$$\sum_j \alpha_j = 0$$

## Best-fitting (least-squares) parameters

( for a 1-way design )

$$\text{Full Model: } Y_{ij} = \mu + \alpha_j + e_{ij}$$

$$\mu = \bar{Y}_u = \sum_{j=1}^a \bar{Y}_j / a \quad \text{unweighted mean of group means}$$

$$\hat{\alpha}_j = \bar{Y}_j - \bar{Y}_u \quad \text{difference between group \& unweighted means} \quad \sum_j \alpha_j = 0$$

$$\text{Reduced Model: } Y_{ij} = \mu + e_{ij}$$

$$\mu = \bar{Y} \quad \text{intercept equals the grand mean}$$

## Comparing Full & Reduced Models

- We expect full model to fit the data better
- Why?
  - because it has more free parameters (is more complex)
  - adjusting values of alpha is expected to reduce  $\sum e_{ij}^2$
  - even when true values of alphas in population are zero
- $H_0: \alpha_j = 0$  (for all j's)      $H_1: \alpha_j \neq 0$  (for at least 1 group, j)
  - when  $H_0$  is true, all groups are selected from populations with same mean  $\mu$
  - variation among  $\bar{Y}_j$  is due to sampling variation
  - reduced model should provide a reasonably good fit to data
  - difference between  $\sum e_{ij}^2$  for full & reduced model should not be unusually large
- Question: is change in  $\sum e_{ij}^2$  unusually large assuming  $H_0$  is true (and true alphas are 0)?

## Measure of relative goodness-of-fit

$$F = \frac{(E_R - E_F) / (df_R - df_F)}{E_F / df_F}$$

$E_R$ :  $\sum e_{ij}^2$  for reduced model

$E_F$ :  $\sum e_{ij}^2$  for full model

$df_R$ : degrees-of-freedom reduced model

$df_F$ : degrees-of-freedom full model

- $df = N - 1$  (the number of estimated parameters)
- $df_R = N - 1$
- $df_F = N - a$  [ $N$  = total number scores;  $a$  = number of groups]
- $(df_R - df_F) = (a - 1)$  = change in number of free parameters in full & reduced models

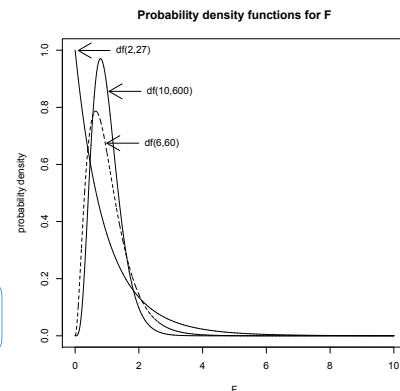
## F distribution when $H_0$ is true

$$F = \frac{(E_R - E_F) / (df_R - df_F)}{E_F / df_F}$$

$$H_0 : \alpha_1 = \alpha_2 = \dots = \alpha_a = 0$$

$$H_1 : \alpha_j \neq 0$$


When  $H_0$  is true, F statistic follows F distribution with  $(df_R - df_F)$  &  $(df_F)$  degrees of freedom



## R example (mood induction experiment)

## mood induction experiment

```
> library(sciplot)
> par(mfrow=c(1,1),cex=1.5)
> with(data=mood.data,
      bargraph.CI(x.factor=group,
                 response=mood,
                 main="Mood Data") )
```



```
> summary(mood.data)
      group      mood
neutral  :10   Min.   :1.00
pleasant :10   1st Qu.:3.25
unpleasant:10 Median :4.00
              Mean  :4.33
              3rd Qu.:5.00
              Max.  :7.00

> # grand mean
> mean(mood.data$mood)
[1] 4.3333
> # group means
> with(mood.data, tapply(mood, group, mean))
      neutral pleasant unpleasant
           4         6           3

> # group SD
> with(mood.data, tapply(mood, group, sd))
      neutral pleasant unpleasant
      0.66667  1.15470  1.05409
```

## Comparison of full & restricted models

```
> # load data:
> load(file=url("http://pnb.mcmaster.ca/bennett/psy710/datasets/mood_data.rda"))
> # fit models:
> mood.full <- lm(mood~1+group, data=mood.data)
> mood.restricted <- lm(mood~1, data=mood.data)
```

## Comparison of full & restricted models

```
> # extract residuals:
> (E.full<-sum(residuals(mood.full)^2))
[1] 26
> (E.restricted<-sum(residuals(mood.restricted)^2))
[1] 72.667
> (E.restricted-E.full)
[1] 46.667
> (df.full<-mood.full$df.residual)
[1] 27
> (df.restricted<-mood.restricted$df.residual)
[1] 29
```

$$F = \frac{(E_R - E_F) / (df_R - df_F)}{E_F / df_F}$$

```
> # compare SSresids with F test:
> (F <- ( (E.restricted-E.full)/(df.restricted-df.full)/(E.full/df.full) ) )
[1] 24.231
> (p.value <- 1-pf(F, df1=(df.restricted-df.full), df2=df.full) )
[1] 9.4214e-07
```

## ANOVA tables

## ANOVA tables

Source	df	Sum Sq	Mean Sq	F value	Pr(>F)
Group	$a - 1$	$SS_B$	$MS_B$	$MS_B/MS_W$	$p$
Residuals	$a(n - 1)$	$SS_W$	$MS_W$		

Table 1: A standard ANOVA table for a one-way design.

Source	df	Sum Sq	Mean Sq	F value	Pr(>F)
Between-Group	$a - 1$	$SS_B$	$MS_B$	$MS_B/MS_W$	$p$
Within-Residuals	$a(n - 1)$	$SS_W$	$MS_W$		

Table 1: A standard ANOVA table for a one-way design.

## ANOVA tables

lm()

```
options(contrasts=c("contr.sum", "contr.poly")) # IMPORTANT!
> # use lm()
> mood.lm <- lm(mood~1+group, data=mood.data)
```

```
> class(mood.lm)
[1] "lm"
```

```
> anova(mood.lm)
Analysis of Variance Table
Response: mood
      Df Sum Sq Mean Sq F value Pr(>F)
group   2  46.7   23.33   24.2 9.4e-07 ***
Residuals 27  26.0    0.96
      E.full
```

## ANOVA tables

lm()

```
options(contrasts=c("contr.sum", "contr.poly")) # IMPORTANT!
> # use lm()
> mood.lm <- lm(mood~1+group, data=mood.data)
```

```
> class(mood.lm)
[1] "lm"
```

```
> anova(mood.lm)
Analysis of Variance Table
Response: mood
      Df Sum Sq Mean Sq F value Pr(>F)
group   2  46.7   23.33   24.2 9.4e-07 ***
Residuals 27  26.0    0.96
      E.restricted - E.full
```

## ANOVA tables

interpret as changes in goodness of fit in nested models

```
> anova(mood.lm)
Analysis of Variance Table
Response: mood
      Df Sum Sq Mean Sq F value Pr(>F)
group   2  46.7   23.33   24.2 9.4e-07 ***
Residuals 27  26.0    0.96
```

$SS_{\text{group}}$  is change in goodness-of-fit when alphas are set to zero

Is this change in goodness-of-fit unusual when  $H_0$  is true?

Calculate F statistic...

## ANOVA tables

$$F = \frac{(E_R - E_F)/(df_R - df_F)}{E_F/df_F}$$

lm()

```
options(contrasts=c("contr.sum", "contr.poly")) # IMPORTANT!
> # use lm()
> mood.lm <- lm(mood~1+group, data=mood.data)
```

```
> class(mood.lm)
[1] "lm"
```

```
df_restricted = df_total = N - 1 = 29
df_group = Δdf = (a-1) = 2
df_full = df_resid = N - 1 - df_group = 27
```

```
> anova(mood.lm)
```

Analysis of Variance Table  
Response: mood

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
group	2	46.7	23.33	24.2	9.4e-07 ***
Residuals	27	26.0	0.96		

## ANOVA tables

$$F = \frac{(E_R - E_F)/(df_R - df_F)}{E_F/df_F}$$

lm()

```
options(contrasts=c("contr.sum", "contr.poly")) # IMPORTANT!
> # use lm()
> mood.lm <- lm(mood~1+group, data=mood.data)
```

```
> class(mood.lm)
[1] "lm"
```

```
> anova(mood.lm)
```

Analysis of Variance Table  
Response: mood

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
group	2	46.7	23.33	24.2	9.4e-07 ***
Residuals	27	26.0	0.96		

$$MS = \frac{SS}{df}$$

## ANOVA tables

$$F = \frac{(E_R - E_F)/(df_R - df_F)}{E_F/df_F}$$

lm()

```
options(contrasts=c("contr.sum", "contr.poly")) # IMPORTANT!
> # use lm()
> mood.lm <- lm(mood~1+group, data=mood.data)
```

```
> class(mood.lm)
[1] "lm"
```

```
> anova(mood.lm)
```

Analysis of Variance Table  
Response: mood

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
group	2	46.7	23.33	24.2	9.4e-07 ***
Residuals	27	26.0	0.96		

$$F = \frac{MS_{\text{group}}}{MS_{\text{Residuals}}}$$

$$p = P(F \geq F_{\text{observed}} | H_0)$$

## ANOVA tables

$$F = \frac{(E_R - E_F)/(df_R - df_F)}{E_F/df_F}$$

lm()

```
options(contrasts=c("contr.sum", "contr.poly")) # IMPORTANT!
> # use lm()
> mood.lm <- lm(mood~1+group, data=mood.data)
```

```
> anova(mood.lm)
```

Analysis of Variance Table  
Response: mood

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
group	2	46.7	23.33	24.2	9.4e-07 ***
Residuals	27	26.0	0.96		

- Our observed F is unusual when H<sub>0</sub> is true.
- Reject H<sub>0</sub> in favor of H<sub>1</sub> α<sub>j</sub> ≠ 0 for at least one group, j

## ANOVA tables

aov()

$$F = \frac{(E_R - E_F)/(df_R - df_F)}{E_F/df_F}$$

```
options(contrasts=c("contr.sum", "contr.poly")) # IMPORTANT!
```

```
> # use aov()
```

```
> mood.aov <- aov(mood~1+group, data=mood.data)
```

```
> class(mood.aov)
```

```
[1] "aov" "lm"
```

```
> anova(mood.aov)
```

Analysis of Variance Table

Response: mood

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
group	2	46.7	23.33	24.2	9.4e-07 ***
Residuals	27	26.0	0.96		

## ANOVA tables

Source	df	Sum Sq	Mean Sq	F value	Pr(>F)
Group	$a - 1$	$SS_B$	$MS_B$	$MS_B/MS_W$	$p$
Residuals	$a(n - 1)$	$SS_W$	$MS_W$		

Table 1: A standard ANOVA table for a one-way design.

## ANOVA tables

Source	df	Sum Sq	Mean Sq	F value	Pr(>F)
Group	$a - 1$	$SS_B$	$MS_B$	$MS_B/MS_W$	$p$
Residuals	$a(n - 1)$	$SS_W$	$MS_W$		

Table 1: A standard ANOVA table for a one-way design.

$$E_F = SS_W$$

$E_F$  and  $E_R$  are  $\sum e_{ij}^2$  for full & reduced models

## ANOVA tables

Source	df	Sum Sq	Mean Sq	F value	Pr(>F)
Group	$a - 1$	$SS_B$	$MS_B$	$MS_B/MS_W$	$p$
Residuals	$a(n - 1)$	$SS_W$	$MS_W$		

Table 1: A standard ANOVA table for a one-way design.

$$E_R = SS_B + SS_W = SS_{Total} \quad E_F = SS_W$$

$E_F$  and  $E_R$  are  $\sum e_{ij}^2$  for full & reduced models



## ANOVA tables

$$E_R - E_F = SS_{Total} - SS_W = SS_B$$

Source	df	Sum Sq	Mean Sq	F value	Pr(>F)
Group	$a - 1$	$SS_B$	$MS_B$	$MS_B/MS_W$	$p$
Residuals	$a(n - 1)$	$SS_W$	$MS_W$		

Table 1: A standard ANOVA table for a one-way design.

$$E_R = SS_B + SS_W = SS_{Total} \quad E_F = SS_W$$

$E_F$  and  $E_R$  are  $\sum e_{ij}^2$  for full & reduced models

## ANOVA tables

$$E_R - E_F = SS_{Total} - SS_W = SS_B$$

Source	df	Sum Sq	Mean Sq	F value	Pr(>F)
Group	$a - 1$	$SS_B$	$MS_B$	$MS_B/MS_W$	$p$
Residuals	$a(n - 1)$	$SS_W$	$MS_W$		

Table 1: A standard ANOVA table for a one-way design.

$$E_R = SS_B + SS_W = SS_{Total} \quad E_F = SS_W$$

$$MS = \frac{SS}{df} \quad \text{between-group \& within-group variance}$$

$E_F$  and  $E_R$  are  $\sum e_{ij}^2$  for full & reduced models

## ANOVA tables

$$E_R - E_F = SS_{Total} - SS_W = SS_B$$

$$F = \frac{(E_R - E_F)/(df_R - df_F)}{E_F/df_F} = \frac{MS_B}{MS_W}$$

Source	df	Sum Sq	Mean Sq	F value	Pr(>F)
Group	$a - 1$	$SS_B$	$MS_B$	$MS_B/MS_W$	$p$
Residuals	$a(n - 1)$	$SS_W$	$MS_W$		

Table 1: A standard ANOVA table for a one-way design.

$$E_R = SS_B + SS_W = SS_{Total} \quad E_F = SS_W$$

$$MS = \frac{SS}{df}$$

$E_F$  and  $E_R$  are  $\sum e_{ij}^2$  for full & reduced models

MS values are estimates of variance

Models

full:  $Y_{ij} = \mu + \alpha_j + e_{ij}$

reduced:  $Y_{ij} = \mu + e_{ij}$

$\alpha_j = \mu_j - \mu_G$

- within-group variation due only to population error variance
- between-group variation due to group EFFECTS and sampling variation
  - we expect to find some Between-Group variation even when effects are zero
  - amount of variation depends on error variance & sample size

### Estimates of Population Error Variance

$$MS_W = \frac{E_F}{df_F} = \hat{\sigma}_e^2$$

$MS_W$  is weighted average of within-group variances

$$\sqrt{\frac{E_F}{df_F}} = \hat{\sigma}_e$$

Residual Standard Error

### ANOVA tables

$MS_{resid} = \text{AVG}(\text{within-group variance})$

```
> anova(mood.lm)
Analysis of Variance Table
Response: mood
      Df Sum Sq Mean Sq F value Pr(>F)
group  2  46.7   23.33   24.2  <.001
Residuals 27  26.0    0.96
```

```
> # mean of group variances:
> mean(with(mood.data,
+ tapply(mood,group,var)) )
[1] 0.96296
```

### ANOVA tables

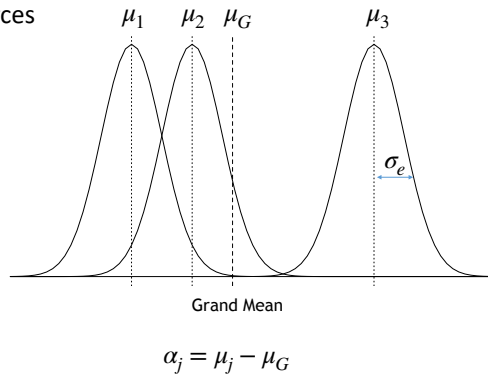
$MS_{group} = n \times (\text{between-group variance})$

```
> anova(mood.lm)
Analysis of Variance Table
Response: mood
      Df Sum Sq Mean Sq F value Pr(>F)
group  2  46.7   23.33   24.2  <.001
Residuals 27  26.0    0.96
```

```
> n <- 10
> ( groupMeans <- with(mood.data,tapply(mood,group,mean)) )
  neutral pleasant unpleasant
      4         6             3
> n * var(groupMeans)
[1] 23.33333
```

## MS<sub>Group</sub> is the sum of 2 components

- between-group variance has 2 sources
  - real differences among  $\mu$  values
  - sampling variation
- if H0 is true:  $\alpha_j = 0$
- between-group variance is caused only by sampling variation



## Estimates of Population Error Variance

when H0 is TRUE variation among group means is pure SAMPLING VARIANCE

- recall that  $MS_B = n \times \hat{\sigma}_Y^2$
- if H0 is true:  $\alpha_j = 0$
- between-group variance is caused only by sampling variation
  - $\hat{\sigma}_Y^2 = \frac{\hat{\sigma}_e^2}{n}$ , so  $\hat{\sigma}_e^2 = n \times \hat{\sigma}_Y^2$ ,
- when H0 is true
  - $MS_B = \hat{\sigma}_e^2$  and  $MS_W$  &  $MS_B$  are independent estimates of  $\sigma_e^2$
  - in the long run, averages of  $MS_W$  and  $MS_B$  equal  $\sigma_e^2$  when H0 is true, and
  - $F = \frac{MS_B}{MS_W}$  is a random variable that follows the so-called F distribution

## Estimates of Population Error Variance

when H0 is FALSE variation among groups caused by SAMPLING VARIANCE plus group effects

- if H0 is false:  $\alpha_j \neq 0$  (for at least group)
- between-group variance is caused by sampling variation & alpha's

$$\hat{\sigma}_Y^2 = \frac{\hat{\sigma}_e^2}{n} + \frac{\sum \alpha_j^2}{a-1}$$

$$E(MS_B) = n \times \sigma_\mu^2 = \sigma_e^2 + \frac{n \sum \alpha_j^2}{a-1}$$


- and we expect  $MS_B > MS_W$  as  $\sum \alpha_j^2$  increases, and  $F \gg 1$
- evaluate with 1-tailed tests... look for unusually large values of F

## Estimates of Population Error Variance

value of  $MS_B$  reflects error variance and group effects

$$E(MS_W) = E\left(\frac{E_F}{df_F}\right) = \sigma_e^2$$

$$E(MS_B) = \sigma_e^2 + \frac{\sum_j n_j \alpha_j^2}{a-1}$$



When H0 is true &  $\alpha_j = 0$

$$E(MS_B) = E(MS_W)$$

$$F \approx 1$$

When H0 is false &  $\alpha_j \neq 0$

$$E(MS_B) > E(MS_W)$$

$$F > 1$$

## ANOVA tables

MS values are estimates of variance

```
> anova(mood.lm)
```

Analysis of Variance Table

Response: mood

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
group	2	46.7	23.33	24.2	9.4e-07 ***
Residuals	27	26.0	0.96		

### When H0 is true:

$MS_{resid}$  &  $MS_{group}$  are independent estimates of  $VAR(error)$

$F = MS_{group}/MS_{resid} \approx 1$  (distributed as  $F(2,27)$ )

### When H1 is true:

$MS_{resid}$  = estimate of  $VAR(error)$

$MS_{group}$  = estimate of  $VAR(error) + [\text{positive number that depends on } \alpha\text{'s}]$

$F = MS_{group}/MS_{resid} \gg 1$

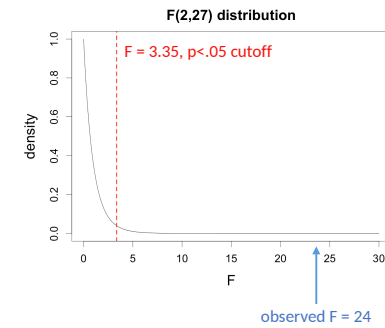
## ANOVA tables

Using F to test null hypothesis

### When H0 is true:

$MS_{resid}$  &  $MS_{group}$  are independent estimates of  $VAR(error)$

$F = MS_{group}/MS_{resid} \approx 1$  (distributed as  $F(2,27)$ )



## ANOVA tables

interpretations of H0 & H1

```
> anova(mood.lm)
```

Analysis of Variance Table

Response: mood

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
group	2	46.7	23.33	24.2	9.4e-07 ***
Residuals	27	26.0	0.96		

H0: all group effects (alpha's) are zero  
H1: not all group effects (alpha's) are zero

H0: all populations have the same mean  
H1: not all populations have the same mean


Assuming all of the effects are zero, is the change in goodness of fit (SS error) unusually large?

Assuming all of the effects are zero, is the variation among group means unusually large?

effect size & association strength

## effect size (Cohen's f)

$$d_i = \frac{(\mu_1 - \mu_G)}{\sigma_e} \quad f = \sqrt{\frac{\sum_i d_i^2}{3}}$$



```
> library(effectsize)
> cohens_f(mood.lm)
Parameter | Cohen's f | 90% CI
-----|-----|-----
group    | 1.34 | [0.88, 1.75]

> cohens_f_squared(mood.lm)
Parameter | Cohen's f2 | 90% CI
-----|-----|-----
group    | 1.79 | [0.77, 3.08]
```

## association strength

proportion of variance accounted for by group

adjusted R-squared

omega-squared

$$\tilde{R}^2 \quad \omega^2 = \frac{\sigma_\alpha^2}{\sigma_\alpha^2 + \sigma_e^2}$$

$\omega^2 = 0.01$  is a small association  
 $\omega^2 = 0.06$  is a medium association  
 $\omega^2 = 0.14$  or larger is a large association

## association strength

proportion of variation that is accounted for by group

```
> summary(mood.lm)
Coefficients:
              Estimate Std. Error t value Pr(>|t|)
(Intercept)  4.333      0.179    24.19 < .0001
group1       -0.333      0.253    -1.32  0.2
group2        1.667      0.253     6.58 < .0001
Residual standard error: 0.981 on 27 df
Multiple R^2: 0.642, Adjusted R^2: 0.616
F: 24.2 on 2 and 27 DF, p-value: 9.42e-07
```

```
> eta_squared(mood.lm)
Parameter | Eta2 | 90% CI
-----|-----|-----
group    | 0.64 | [0.43, 0.75]
```

$$\eta^2 = R^2$$

```
> omega_squared(mood.lm)
Parameter | Omega2 | 90% CI
-----|-----|-----
group    | 0.61 | [0.39, 0.73]
```

$$\omega^2 \approx \tilde{R}^2$$

- VAR(residuals) = 0.896
- VAR(original scores) = 2.506
- so we accounted 64%, of VAR(original scores)

## assumptions

## Assumptions of ANOVA

- The scores must be statistically independent of each other
- Population of scores, Y, distributed normally within each group
  - equivalent to assuming that error follows normal distribution
- Error variance is constant across groups
- If assumptions are met, F statistic follows F distribution
  - if they are not met, F statistic does not follow F distribution

## Non-normality & non-constant variance

- ANOVA reasonably robust to deviations from normality
  - if deviations are similar in all groups
  - less robust to deviations that differ across groups
    - e.g., positive skew in 1 group and negative skew in others
  - robustness also declines if n is not equal across groups
- ANOVA is reasonably robust to 3-4 fold differences in variances
  - if Y's normally distributed and equal n per group

## Tests for non-normality

- Kolmogorov-Smirnov test: `ks.test()`
- Shapiro-Wilk's test: `shapiro.test()`
- Both tests have low power, though `shapiro.test` is better
  - to compensate, could use tests with  $\alpha = 0.10$

```
shapiro.test(residuals(mood.full) )  
  
##  
## Shapiro-Wilk normality test  
##  
## data: residuals(mood.full)  
## W = 0.85, p-value = 5e-04
```

H0: residuals are distributed normally

## Bartlett.test

test for homogeneity of variance

```
bartlett.test(mood~group, data=mood.data)  
bartlett.test(mood.data$mood, mood.data$group)  
##  
## Bartlett test of homogeneity of variances  
##  
## data: mood.data$mood and mood.data$group  
## Bartlett's K-squared = 2.6, df = 2, p-value  
## = 0.3
```

H0: variance is constant across groups  
( also see `leveneTest` in `car` package )

## Alternative Analyses

- perform ANOVA on transformed data
  - square-root, log, & inverse-sine transformations common
  - conclusions apply to transformed data
- Welch correction for non-constant variance
  - `oneway.test()` [N.B. Assumes normality]

```
oneway.test(mood~group,data=mood.data)

##
## One-way analysis of means (not assuming
## equal variances)
##
## data: mood and group
## F = 18, num df = 2, denom df = 17, p-value =
## 6e-05
```

## Alternative Analyses

- Kruskal-Wallis test for group differences [`kruskal.test()`]
  - does not assume normality or constant variance
  - HO: the means of ranked data are the same in each group
  - if distributions for each group have same shape (not necessarily normal), then KW test evaluates null hypothesis that group MEDIANS are equal

```
kruskal.test(mood~group,data=mood.data)

##
## Kruskal-Wallis rank sum test
##
## data: mood by group
## Kruskal-Wallis chi-squared = 19, df = 2,
## p-value = 7e-05
```