PSYCH 710

1-way Within-Subjects ANOVA

Prof. Patrick Bennett

Typical Experimental Design

- treatment is a fixed factor
- subjects is a random factor
- typically, 1 observation per cell
  - not possible to measure within-cell variance
  - cannot distinguish contributions of error and subject x treatment interaction to between-cell variation
- residuals are correlated

	treatment 1	treatment 2	treatment 3
subject 1	n=1	n=1	n=1
subject 2	n=1	n=1	n=1
subject 3	n=1	n=1	n=1
subject 4	n=1	n=1	n=1

linear model

$$Y_{ij} = \mu + \alpha_j + \pi_i + (\pi\alpha)_{ij} + \epsilon_{ij}$$

- most within-S designs have 1 measure per "cell"
- model with subject x treatment interaction has too many parameters
- so, we drop interaction term
  - variation becomes part of residuals

$$Y_{ij} = \mu + \alpha_j + \pi_i + \epsilon_{ij}$$
$$Y_{ij} = \mu + \pi_i + \epsilon_{ij}$$
$$H0 \; : \; \alpha_1 = \alpha_2 = \cdots = \alpha_j = 0$$

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

linear model

$$Y_{ij} = \mu + \alpha_j + \pi_i + \epsilon_{ij}$$
$$Y_{ij} = \mu + \pi_i + \epsilon_{ij}$$
$$H0 \; : \; \alpha_1 = \alpha_2 = \cdots = \alpha_j = 0$$

- best-fitting coefficients =>
- evaluate H0 with standard F test
- df_F = (n-1)(a-1)

Least-squares coefficients:

$$\hat{\mu} = \bar{Y}_{..}$$
$$\hat{\alpha}_j = \bar{Y}_{.j} - \bar{Y}_{..}$$
$$\hat{\pi}_i = \bar{Y}_{i.} - \bar{Y}_{..}$$

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

Expected Mean Squares & F tests

- E(MS) for treatment & residuals includes subject x treatment interaction
- interaction & error are confounded in this design
- F (treatment) = MS_{treat} / MS_{resid}
  - F = MS_{treat} / $MS_{treat\ x\ subj}$
- Var Comp (subjects) = (MS_S - MS_{resid}) / a
  - assumes that interaction is zero

Effect	Type	E(Mean Square)
treatment	fixed	$\sigma_e^2 + \sigma_{\pi\alpha}^2 + n\sum_{j=1}^{a}\alpha_j^2/(a-1)$
subjects	random	$\sigma_e^2 + a\sigma_\pi^2$
residuals		$\sigma_e^2 + \sigma_{\pi\alpha}^2$

Compound Symmetry

- Each dependent variable has a variance, and each pair of dependent variables has a covariance
  - VAR(X) = E(X^2) - E(X)^2 [E == "Expectation"]
  - COV(XY) = E(XY) - E(X)E(Y)
- when all variances are equal (VAR(X) = a) and all covariances are equal (COV(XY)=b), the variance-covariance matrix has <u>**compound symmetry**</u>
  - when COV(XY) = 0, then X & Y are independent and compound symmetry holds
  - when COV(XY) = k (k≠0), then X & Y are not independent but compound symmetry holds
- the data will exhibit compound symmetry when the correlations between each dependent variable are equal and/or when within-S factor has 1 df
- when compound symmetry exists, then the F statistic for *a* treatments will follow F distribution with df=[(*a*-1), (*a*-1)(n-1)]

Sphericity

- Compound symmetry is <u>sufficient</u> condition for F statistic to follow F distribution
- but, compound symmetry is not <u>necessary</u>
- a more lenient condition is <u>sphericity</u>:
  - sphericity implies that all <u>pairwise differences</u> between dependent variables have the same variance
- when sphericity holds, the F statistic will follow the appropriate F distribution

Estimating Compound Symmetry & Sphericity

- Compound symmetry & sphericity are rarely met perfectly
- when neither holds, the F statistic will follow (approximately) an F distribution with adjusted degrees of freedom
  - df adjustment depends on degree of sphericity $(\hat{\epsilon}, \tilde{\epsilon})$
  - two common measures of sphericity: epsilon-hat & epsilon-tilde
  - both are derived from data and vary from 1 (perfect sphericity) to 1/(a-1) (no sphericity)
  - adjusted df = epsilon x (a-1), epsilon x (a-1)(n-1)

df adjustment [Conservative F test]

$$\epsilon = \frac{1}{(a-1)} \text{ (minimum value)}$$

$$df = (\epsilon \times (a-1), \epsilon \times (a-1)(n-1))$$

$$df = (1, (n-1))$$

Geisser/Greenhouse & Huynh-Feldt df adjustment

Geisser Greenhouse : $\hat{\epsilon}$

Huynh Feldt : $\tilde{\epsilon}$

- GG & HF estimates of ϵ are derived from variance-covariance matrix
- $df_{adj} = \epsilon \times (a-1), \epsilon \times (a-1)(n-1)$
- GG slightly more conservative than HF

R Example

aov, aov_car, aov_ez, aov_5, lmer, & lme

aov & aov_car

R Example

aov command (assumes sphericity)

```
> load(url("http://pnb.mcmaster.ca/bennett/psy710/datasets/mw11_5.rda") )
> summary(mw115L)
      subj       age          score      
 s1     : 4   a30:12   Min.   : 84.0  
 s2     : 4   a36:12   1st Qu.: 98.2  
 s3     : 4   a42:12   Median :107.0  
 s4     : 4   a48:12   Mean   :108.0  
 s5     : 4            3rd Qu.:117.2  
 s6     : 4            Max.   :133.0  
 (Other):24                           
```

R Example

aov command (assumes sphericity)

```
> # aov command:
> # following anova table ASSUMES SPHERICITY:
> mw115.aov.01b <- aov(Y~age+Error(subj),data=mw115.long)
> summary(mw115.aov.01b)

Error: subj
          Df Sum Sq Mean Sq F value Pr(>F)
Residuals 11   6624   602.2               

Error: Within
          Df Sum Sq Mean Sq F value Pr(>F)  
age        3    552  184.00   3.027 0.0432 *
Residuals 33   2006   60.79                 
```

R Example

aov command (assumes sphericity)

```
> # aov command:
> # following anova table ASSUMES SPHERICITY:
> # better because within subj x age made explicit:
> mw115.aov.01c <- aov(Y~age+Error(subj/age),data=mw115.long)
> summary(mw115.aov.01c)

Error: subj
          Df Sum Sq Mean Sq F value Pr(>F)
Residuals 11   6624   602.2               

Error: subj:age
          Df Sum Sq Mean Sq F value Pr(>F)  
age        3    552  184.00   3.027 0.0432 *
Residuals 33   2006   60.79                 
```

R Example

aov_car command in afex package

can affect measure of association strength

```
> library(afex)
> mw115.aov.02a <- aov_car(Y~age+Error(subj/age),
+                          data=mw115.long, observed="age")
> summary(mw115.aov.02a)

Univariate Type III Repeated-Measures ANOVA Assuming Sphericity

            Sum Sq num Df Error SS den Df  F value    Pr(>F)    
(Intercept) 559872      1     6624     11 929.7391 5.586e-12 ***
age            552      3     2006     33   3.0269   0.04322 *  
```

R Example

aov_car command in afex package

```
Univariate Type III Repeated-Measures ANOVA Assuming Sphericity

            Sum Sq num Df Error SS den Df  F value    Pr(>F)
(Intercept) 559872      1     6624     11 929.7391 5.586e-12 ***
age            552      3     2006     33   3.0269   0.04322 *

Mauchly Tests for Sphericity
    Test statistic  p-value
age        0.24265 0.017718

Greenhouse-Geisser & Huynh-Feldt Corrections
     GG eps Pr(>F[GG])
age 0.60954    0.074

       HF eps Pr(>F[HF])
age 0.7248502 0.0635
```

aov_ez & aov_4

R Example

aov_ez command in afex package

```
> library(afex)
> mw115.aov.02b <- aov_ez(id=“subj",dv="Y",data=mw115.long,
+                         between=NULL, within=“age", observed="age")
> summary(mw115.aov.02b)
Univariate Type III Repeated-Measures ANOVA Assuming Sphericity

            Sum Sq num Df Error SS den Df  F value    Pr(>F)
(Intercept) 559872      1     6624     11 929.7391 5.586e-12 ***
age            552      3     2006     33   3.0269   0.04322 *

Mauchly Tests for Sphericity
    Test statistic  p-value
age        0.24265 0.017718

Greenhouse-Geisser and Huynh-Feldt Corrections
     GG eps Pr(>F[GG])
age 0.60954    0.07479 .

      HF eps Pr(>F[HF])
age 0.7248502 0.06353773
```

R Example

aov_4 command in afex package

```
> # aov_4:
> mw115.aov.02c <- aov_4(Y~age+(1+age|subj),mw115.long)
> summary(mw115.aov.02c)
Univariate Type III Repeated-Measures ANOVA Assuming Sphericity
            Sum Sq num Df Error SS den Df  F value    Pr(>F)
(Intercept) 559872      1     6624     11 929.7391 5.586e-12 ***
age            552      3     2006     33   3.0269   0.04322 *

Mauchly Tests for Sphericity
   Test statistic  p-value
age       0.24265 0.017718

Greenhouse-Geisser and Huynh-Feldt Corrections
     GG eps Pr(>F[GG])
age 0.60954    0.07479 .

      HF eps Pr(>F[HF])
age 0.7248502 0.06353773
```

lmer & lme

lmer()

```
> library(lmerTest)
> cog.lmer.01 <- lmer(score~age+(1|subj),data=mw115L)

> anova(cog.lmer.01) # assumes sphericity
Type III Analysis of Variance Table with Satterthwaite's method
    Sum Sq Mean Sq NumDF DenDF F value  Pr(>F)
age    552     184     3    33  3.0269 0.04322
```

assumes independence/sphericity

```
> ranova(cog.lmer.01)
ANOVA-like table for random-effects: Single term deletions

Model:
score ~ age + (1 | subj)
           npar  logLik    AIC    LRT Df Pr(>Chisq)
<none>        6 -171.76 355.53
(1 | subj)    5 -184.92 379.85 26.318  1  2.896e-07
```

chi square tests are approximate/conservative

lmer()

evaluating fixed effect with chi square test [does not assume sphericity]

```
> anova(cog.lmer.01) # assumes sphericity
Type III Analysis of Variance Table with Satterthwaite's method
    Sum Sq Mean Sq NumDF DenDF F value  Pr(>F)
age    552     184     3    33  3.0269 0.04322

> cog.lmer.02 <- lmer(score~1+(1|subj),data=mw115L) # remove age

> anova(cog.lmer.02,cog.lmer.01) # evaluate change in deviance
refitting model(s) with ML (instead of REML)
Data: mw115L
Models:
cog.lmer.02: score ~ 1 + (1 | subj)
cog.lmer.01: score ~ age + (1 | subj)
            npar    AIC    BIC  logLik deviance Chisq Df Pr(>Chisq)
cog.lmer.02    3 371.47 377.08 -182.73   365.47
cog.lmer.01    6 368.71 379.94 -178.36   356.71 8.751  3    0.03279
```

chi square tests are approximate/conservative

lme() in nlme package

defining variance-covariance matrix

```
> library(nlme)

> # assume independence:
> cog.nlme.00 <- lme(score~age,data=mw115L,
+                random=~1|subj)

> # assume compound symmetry:
> cog.nlme.01 <- lme(score~age,data=mw115L,
+                random=~1|subj,
+                correlation=corCompSymm(value=0.3,form=~1|subj))

> # no constraints on between-level correlations:
> cog.nlme.02 <- lme(score~age,data=mw115L,
+                random=~1|subj,
+                correlation=corSymm(value=c(.3,.3,.3,.3,.3,.3),form=~1|subj))
```

lme() in nlme package

defining variance-covariance matrix

```
> summary(cog.nlme.00) # independent
Random effects:
 Formula: ~1 | subj
        (Intercept) Residual
StdDev:    11.63394 7.796658

Correlation Structure: Independence
 Formula: ~1 | subj
 Parameter estimate(s):
 Rho
 0
```

```
> anova(cog.nlme.00) # independent
            numDF denDF  F-value p-value
(Intercept)     1    33 929.7391  <.0001
age             3    33   3.0269  0.0432
```

lme() in nlme package

defining variance-covariance matrix

```
> summary(cog.nlme.01) # compound symmetry
Random effects:
 Formula: ~1 | subj
        (Intercept) Residual
StdDev:    10.59335  9.16064

Correlation Structure: Compound symmetry
 Formula: ~1 | subj
 Parameter estimate(s):
 Rho
 0.2756218
```

```
> anova(cog.nlme.01) # compound symmetry
            numDF denDF  F-value p-value
(Intercept)     1    33 929.7391  <.0001
age             3    33   3.0269  0.0432
```

lme() in nlme package

defining variance-covariance matrix

```
> summary(cog.nlme.02) # no constraints
Random effects:
 Formula: ~1 | subj
        (Intercept) Residual
StdDev:    11.46896 8.859535
Correlation Structure: General
 Formula: ~1 | subj
 Parameter estimate(s):
 Correlation:
  1      2      3
2  0.496
3  0.231  0.469
4  0.023 -0.200  0.644
```

```
> anova(cog.nlme.02) # no constraints
            numDF denDF  F-value p-value
(Intercept)     1    33 916.3946  <.0001
age             3    33   2.6743  0.0633
```

Note that this p value is similar to corrected p values obtained with anova.

```
> anova(cog.nlme.00,cog.nlme.01,cog.nlme.02) # 3rd model fits better
            Model df      AIC      BIC    logLik   Test  L.Ratio p-value
cog.nlme.00     1  6 355.5288 366.2340 -171.7644
cog.nlme.01     2  7 357.5288 370.0182 -171.7644 1 vs 2  0.00000  1.0000
cog.nlme.02     3 12 352.6203 374.0306 -164.3102 2 vs 3 14.90853  0.0108
```

variance components

Variance Components

calculating from ANOVA table

```
> summary(mw115.aov.01c)

Error: subj
          Df Sum Sq Mean Sq F value Pr(>F)
Residuals 11   6624   602.2

Error: subj:age
          Df Sum Sq Mean Sq F value Pr(>F)
age        3    552  184.00   3.027 0.0432 *
Residuals 33   2006   60.79

> # variance component for subj:
> (602.2 - 60.79) / 4 # divide by number of levels of within-S factor
[1] 135.4
```

Variance Components

```
> # anova variance components
> library(VCA)
> cog.aov.vca <- anovaMM(score~age+(subj),Data=mw115L)
> print(cog.aov.vca,digits=3)
ANOVA-Type Estimation of Mixed Model:
--------------------------------------
[Variance Components]
 Name  DF     SS   MS     VC     %Total SD     CV[%]
1 total 18.117             196.13  100   14.005 12.96
2 subj  11     6624 602.18  135.34  69.02 11.634 10.77
3 error 33     2006  60.79   60.79  30.99  7.797  7.22
Mean: 108 (N = 48)
Experimental Design: balanced  |  Method: ANOVA
```

```
> # lmer
> cog.lmer.vca <- VarCorr(cog.lmer.01) # independence
> print(cog.lmer.vca,comp=c("Variance","Std.Dev."))
 Groups   Name        Variance Std.Dev.
 subj     (Intercept) 135.348  11.6339
 Residual              60.788   7.7967
```

Variance Components

depend on within-S variance-covariance matrix

```
> # lme
> VarCorr(cog.nlme.00) # independence
subj = pdLogChol(1)
            Variance  StdDev
(Intercept) 135.348   11.633
Residual     60.787    7.796
> VarCorr(cog.nlme.01) # compound symmetry
subj = pdLogChol(1)
            Variance  StdDev
(Intercept) 112.219   10.593
Residual     83.917    9.160
> VarCorr(cog.nlme.02) # no constraints
subj = pdLogChol(1)
            Variance  StdDev
(Intercept) 131.536   11.468
Residual     78.491    8.859
```

checking residuals for normality

Checking residuals

residuals() does not work with aov() objects

```
> shapiro.test(residuals(cog.aov.02)) # aov_car in afex
Data was changed during ANOVA calculation. Thus, residuals cannot be
added to original data. residuals(..., append = TRUE) will return data and residuals.

Shapiro-Wilk normality test
data:  residuals(cog.aov.02)
W = 0.96965, p-value = 0.2455

> shapiro.test(residuals(cog.lmer.01)) # lmer in lme4

Shapiro-Wilk normality test
data:  residuals(cog.lmer.01)
W = 0.98053, p-value = 0.6008

> shapiro.test(residuals(cog.nlme.02)) # lme in nlme

Shapiro-Wilk normality test
data:  residuals(cog.nlme.02)
W = 0.97417, p-value = 0.3648
```

Checking residuals

residuals() does not work with aov() objects

```
> par(mfrow=c(1,3))
> qqnorm(residuals(cog.aov.02),main="aov");qqline(residuals(cog.aov.02))
> qqnorm(residuals(cog.lmer.01),main="lmer/lme4");qqline(residuals(cog.lmer.01))
> qqnorm(residuals(cog.nlme.02),main="lme/nlme");qqline(residuals(cog.nlme.02))
```

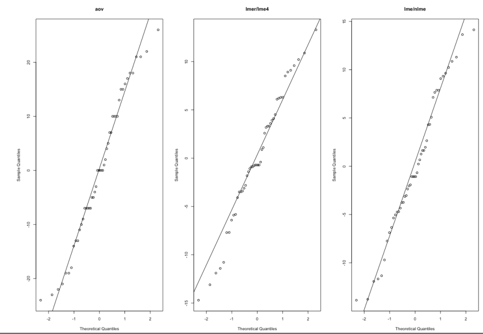

linear contrasts

Linear Contrasts

- create contrast weights
- contrast transforms multiple within-S measures to single composite score for each subject
  - use %*% operator to create composite scores for each subject
- perform t test or F test on composite scores

Linear Contrasts

```
> dat.mat <- with(mw115,cbind(age.30,age.36,age.42,age.48) )
> dat.mat

      age.30 age.36 age.42 age.48
 [1,]    108     96    110    122
 [2,]    103    117    127    133
 [3,]     96    107    106    107
 [4,]     84     85     92     99
 [5,]    118    125    125    116
 [6,]    110    107     96     91
 [7,]    129    128    123    128
 [8,]     90     84    101    113
 [9,]     84    104    100     88
[10,]     96    100    103    105
[11,]    105    114    105    112
[12,]    113    117    132    130
```

- use %*% operator to create composite scores for each subject
- perform t test on composite scores

Linear Contrasts

```
> lin.trend<-c(-1.5,-0.5,0.5,1.5);
> lin.scores<-dat.mat %*% lin.trend;
> lin.scores
      [,1]
 [1,]   28
 [2,]   50
 [3,]   16
 [4,]   26
 [5,]   -3
 [6,]  -34
 [7,]   -4
 [8,]   43
 [9,]    4
[10,]   15
[11,]    6
[12,]   33
```

$\sum w_i d_i$

- composite score is weighted sum of data points for each subject
- positive values means <u>increasing</u> linear trend

- use %*% operator to create composite scores for each subject
- perform t test on composite scores

Linear Contrasts

```
> t.test(lin.scores)

        One Sample t-test

data:  lin.scores
t = 2.2414, df = 11, p-value = 0.04659
alternative hypothesis: true mean is not equal to 0
95 percent confidence interval:
  0.2701827 29.7298173
sample estimates:
mean of x
       15
```

linear trend is significant

with our weights, positive value means that the linear trend is increasing

- use %*% operator to create composite scores for each subject
- perform t test on composite scores

Local vs Global Estimates of Error

```
> wLinear <- contr.poly(n=4)[,1]
> y.mat <- data.matrix(mw115[,1:4])
> linTrends <-  y.mat %*% wLinear
> # uses local error estimate
> t.test(linTrends)

        One Sample t-test

data:  linTrends
t = 2.2414, df = 11, p-value = 0.04659
H1: true mean is not equal to 0
95 percent confidence interval:
  0.1208294 13.2955785
sample estimates:
mean of x
 6.708204
```

```
> # uses local error estimate
> summary(aov(linTrends~1),intercept=T)
            Df Sum Sq Mean Sq F value Pr(>F)
(Intercept)  1    540   540.0   5.024 0.0466
Residuals   11   1182   107.5
```

Local vs Global Estimates of Error

- Significance tests for contrasts may use local or global error term
  - global error term comes from original ANOVA
  - local error term comes from t test
- Global error term has more degrees of freedom
  - may provide more powerful test of null hypothesis

Local vs Global Estimates of Error

```
> wLinear <- contr.poly(n=4)[,1]
> y.mat <- data.matrix(mw115[,1:4])
> linTrends <-  y.mat %*% wLinear

> # uses local error estimate
> summary(aov(linTrends~1),intercept=T)
            Df Sum Sq Mean Sq F value Pr(>F)
(Intercept)  1    540   540.0   5.024 0.0466
Residuals   11   1182   107.5
```

MS-contrast / sum(wLinear^2)
MS-contrast / 1 = 540

```
> # uses global error estimate
> summary(cog.aov.01)

Error: subj
          Df Sum Sq Mean Sq F value Pr(>F)
Residuals 11   6624   602.2

Error: subj:age
          Df Sum Sq Mean Sq F value Pr(>F)
age        3    552  184.00   3.027 0.0432
Residuals 33   2006   60.79

> MS.err <- 60.79
> df.err <- 33
> ( F.global <- 540/MS.err )
[1] 8.883
> 1-pf(F.global,1,df.err)
[1] 0.00537
```

contrasts with emmeans

uses local error estimate with aov_car() objects

```
> wLinear <- c(-1.5,-0.5,0.5,1.5)
> # aov_car object:
> # mw115L.aov.car.01 <- aov_car(score~1+age+Error(subj/age),data=mw115L)
> mw115L.aov.car.emm <- emmeans(mw115L.aov.car.01,specs="age")
> # uses local estimate of error:
> contrast(mw115L.aov.car.emm,method=list(linear=wLinear))
 contrast estimate   SE df t.ratio p.value
 linear         15 6.69 11   2.241  0.0466

> 2.241^2 # square t to get value of F statistic
[1] 5.022
```

contrasts with emmeans

uses global error estimate with aov() objects

```
> wLinear <- c(-1.5,-0.5,0.5,1.5)
> # aov object:
> # mw115L.aov.01 <- aov(score~1+age+Error(subj/age),data=mw115L)
> mw115L.aov.emm <- emmeans(mw115L.aov.01,specs="age")
> # contrast uses the uses global estimate of error:
> contrast(mw115L.aov.emm,method=list(linear=wLinear))
 contrast estimate   SE df t.ratio p.value
 linear         15 5.03 33   2.981  0.0054

> 2.981^2 # square t to get value of F statistic
[1] 8.886
```

contrasts with emmeans

uses global error estimate with lmer() objects

```
> wLinear <- c(-1.5,-0.5,0.5,1.5)
> # lmer object:
> # mw115L.lmer.01 <- lmer(score~age+(1|subj),data=mw115L)
> mw115L.lmer.emm <- emmeans(mw115L.lmer.01,specs="age")
> contrast(mw115L.lmer.emm,method=list(linear=wLinear))
 contrast estimate   SE df t.ratio p.value
 linear         15 5.03 33   2.981  0.0054

Degrees-of-freedom method: kenward-roger
> 2.981^2 # square t to get value of F statistic
[1] 8.886
```