Random effects ANOVA

Peter Hoff
Duke STA 610

ANOVA limitations

Hierarchical normal model

Estimation and inference

Classical data analysis and estimation

The "classical" testing and estimation procedure is as follows:

If the p-value < 0.05,

- reject H_0, and conclude there are group differences,
- estimate θ_j with $\bar{y}_{\cdot j}$.

$$\hat{\theta}_j = \bar{y}_{\cdot j}$$

If the p-value > 0.05,

- accept H_0, and conclude there is no evidence of group differences,
- estimate θ_j with $\bar{y}_{\cdot\cdot}$.

$$\hat{\theta}_j = \bar{y}_{\cdot\cdot}$$

Note that the estimator of θ_j can be written as

$$\hat{\theta}_j = w\bar{y}_j + (1 - w)\bar{y}_{\cdot\cdot}$$

Classical data analysis and estimation

Advantages of classical procedure:

- controls the type I error rate of rejecting H_0;
- is easy to implement and report.

Disadvantages:

- rejecting H_0 doesn't mean no similarities across groups $\Rightarrow \bar{y}_{\cdot j}$ is an inefficient estimate of θ_j
- accepting H_0 doesn't mean no differences between groups $\Rightarrow \bar{y}_{\cdot\cdot}$ is an inaccurate estimate of θ_j.

An alternative strategy

$$\hat{\theta}_j = w\bar{y}_j + (1 - w)\bar{y}_{..}$$

Classical approach: w is the indicator of rejecting H_0.

Multilevel approach: $w = \frac{n/\hat{\sigma}^2}{n/\hat{\sigma}^2 + 1/\hat{\tau}^2}$

The multilevel approach will allow for

- sharing of information across groups,
- the amount of sharing to be estimated from the data.

Example: Test scores

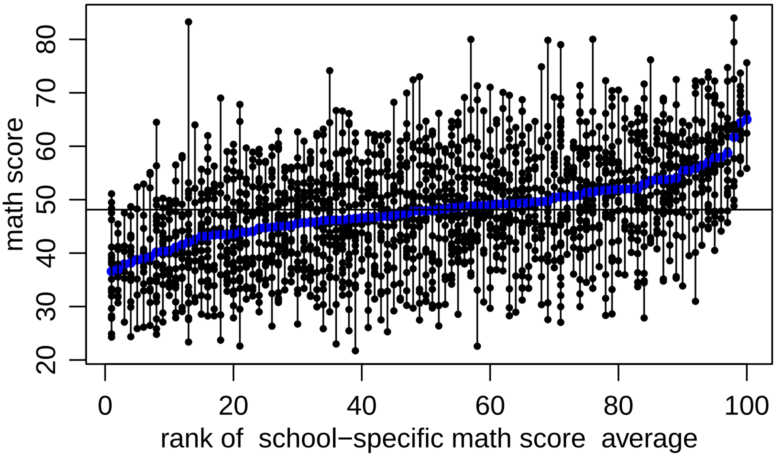

Example: Test scores

```
y.3122<-ndat$mathscore[ndat$school=="3122"]
y.2832<-ndat$mathscore[ndat$school=="2832"]

y.3122

## [1] 75.62 55.86 66.16 62.43

y.2832

##  [1] 66.26 66.12 71.22 54.90 61.98 69.42 61.22 62.99 57.99 61.33 66.85 67.87
## [13] 63.94 73.70 70.36 64.01 57.35 68.25 57.39

mean(ndat$mathscore)

## [1] 48.07446

mean(y.3122)

## [1] 65.0175

mean(y.2832)

## [1] 64.37632
```

Example: Test scores

$$
\begin{array}{ccccc}
48.0744556 & < & 64.3763158 & < & 65.0175 \\
\bar{y}_{\cdot\cdot} & < & \bar{y}_{2832} & < & \bar{y}_{3122}
\end{array}
$$

but

$$n_{3122} = 4 < 19 = n_{2832}$$

Based on the data $\{y_{i,j}\}$, how would you estimate θ_{3122} and θ_{2832}?

Ignoring across-group information :

- $\hat{\theta}_{2832} = \bar{y}_{2832} = 64.3763158$
- $\hat{\theta}_{3122} = \bar{y}_{3122} = 65.0175$
- $\hat{\theta}_{2832} < \hat{\theta}_{3122}$

Considering across-group information and sample size: $\bar{y}_{\cdot\cdot} = 48.0744556$.

- $\bar{y}_{\cdot\cdot} < \hat{\theta}_{2832} < \bar{y}_{2832} = 64.3763158$
- $\bar{y}_{\cdot\cdot} < \hat{\theta}_{3122} < \bar{y}_{3122} = 65.0175$
- $\hat{\theta}_{2832} \gtrless \hat{\theta}_{3122}$?

Example: Test scores

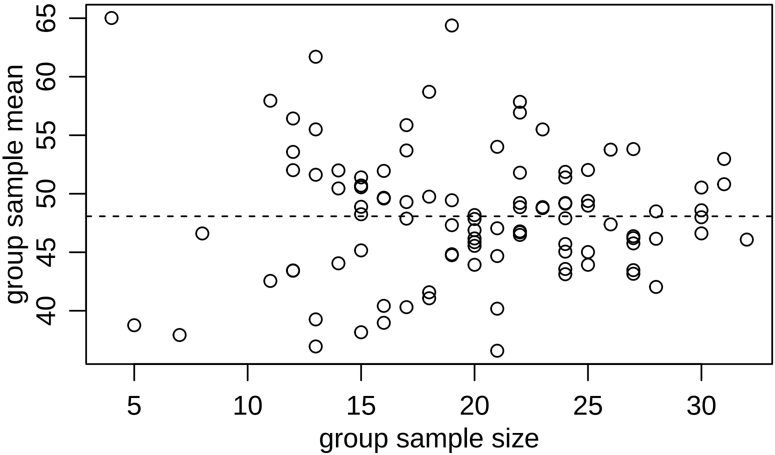

Example: Test scores

Possible explanations for $\bar{y}_{3122}$:

- $\bar{y}_{3122}$ is large because θ_{3122} is large;
- $\bar{y}_{3122}$ is large because $\text{sd}(\bar{y}_{3122})$ is large.

Possible explanations for $\bar{y}_{2832}$:

- $\bar{y}_{2832}$ is large because θ_{2832} is large;
- $\bar{y}_{2832}$ is large because $\text{sd}(\bar{y}_{2832})$ is large (but is smaller than $\text{sd}(\bar{y}_{3122})$).

The plausibility of the explanations will depend on

- the group specific sample sizes, $n_1, \ldots, n_m$;
- the observed across-group heterogeneity.

Example: Free throws

```
ftdat[1:20,]

##      PLAYER1     PLAYER2 TEAM  MIN FTM FTA   FT.
## 1        Sam    Jacobson  LAL   12   2   2 1.000
## 2      Steve      Henson  DET   25   2   2 1.000
## 3   Radoslav Nesterovic  MIN   30   2   2 1.000
## 4      Bryce        Drew  HOU  441   8   8 1.000
## 5    Charles   O'bannon  DET  165   8   8 1.000
## 6      Marty      Conlon  MIA   35   2   2 1.000
## 7      Mikki       Moore  DET    6   2   2 1.000
## 8       John      Crotty  POR   19   3   3 1.000
## 9     Gerald     Wilkins  ORL   28   2   2 1.000
## 10  Korleone       Young  DET   15   2   2 1.000
## 11     Brian       Evans  MIN  145   4   4 1.000
## 12      Pooh  Richardson  LAC  130   4   4 1.000
## 13   Michael     Hawkins  SAC  203   3   3 1.000
## 14     Randy  Livingston  PHO   22   2   2 1.000
## 15     Rusty       Larue  CHI  732  17  17 1.000
## 16      Fred     Hoiberg  IND   87   6   6 1.000
## 17      Herb    Williams  NYK   34   2   2 1.000
## 18      Ryan       Stack  CLE  199  19  20 0.950
## 19       Sam     Cassell  MIL  199  47  50 0.940
## 20    Reggie      Miller  IND 1787 226 247 0.915
```

Who does Indiana pick to shoot its technical foul free throws?

Further limitations of ANOVA

In the wheat yield example we might be interested in

(1) what the yield might be in other plots of land in these 10 regions, or

(2) what the yield might be in other regions.

For general hierarchical data, these questions translate into

(1) making inference about units within groups in our study;

(2) making inference about groups that weren't in our study.

Inference for (1) can be obtained with ANOVA.

Inference for (2) requires

- treating the m groups as a sample from a larger population;
- a statistical model for this larger population.

The hierarchical normal model

$$
\begin{aligned}
y_{i,j} &= \mu + a_j + \epsilon_{i,j} && (1)\\
\{\epsilon_{1,1},\ldots,\epsilon_{n_1,1}\},\ldots,\{\epsilon_{1,m},\ldots,\epsilon_{n_m,m}\} &\sim \text{i.i.d. normal}(0,\sigma^2) && (2)\\
a_1,\ldots,a_m &\sim \text{i.i.d. normal}(0,\tau^2) && (3)
\end{aligned}
$$

The classical ANOVA model consists of (1) and (2).

The HNM assumes the sampling model (3) for the groups.

- $\{a_1,\ldots,a_m\}$ represent differences across groups
- $\{\epsilon_{i,j}\}$ represent differences within groups

The HNM represents this heterogeneity in terms of population variances:

$$
\begin{aligned}
\text{Var}[a] &= \tau^2 = \text{across-group variance}\\
\text{Var}[\epsilon] &= \sigma^2 = \text{within-group variance}
\end{aligned}
$$

Marginal and conditional variation

Two levels of heterogeneity require two versions of variance and covariance:

Within-group variance:

- The variance of $y_{i,j}$ around θ_j;
- Describes heterogeneity/variance within a particular group;
- Mathematically, is calculated *conditionally* on group-level parameters.

Population-level variance:

- Variance of $y_{i,j}$ around μ;
- Describes heterogeneity/variance across the population;
- Mathematically, is calculated *marginally* over group-level parameters.

Conditional variance and covariance

For a fixed group j,

$$\begin{aligned} \{y_{1,j},\ldots,y_{n_j,j}\}|\mu,a_j,\sigma^2 &\sim \text{i.i.d. normal}(\mu+a_j,\sigma^2) \\ \{y_{1,j},\ldots,y_{n_j,j}\}|\theta_j,\sigma^2 &\sim \text{i.i.d. normal}(\theta_j,\sigma^2) \end{aligned}$$

Variation *around the group mean θ_j* is as follows

$$\begin{aligned} \text{E}[y_{i,j}|\mu,a_j] &= \mu+a_j=\theta_j \\ \text{Var}[y_{i,j}|\mu,a_j] &= \sigma^2, \\ \text{Cov}[y_{i_1,j},y_{i_2,j}|\mu,a_j] &= 0. \end{aligned}$$

In words,

- sample observations *from the group* are centered around θ_j;
- the variation of the sample *around θ_j* is σ^2;
- the observations within a group are uncorrelated *around θ_j*.

Regarding correlation: Knowing how far $y_{1,j}$ is from θ_j doesn't inform you about about how far $y_{2,j}$ is from θ_j.

Within-group variance and covariance

$$
\begin{aligned}
y_{i,j} &= \mu + a_j + \epsilon_{i,j} \\
y_{i,j} &= \theta_j + \epsilon_{i,j}
\end{aligned}
$$

$$
\begin{aligned}
\text{Var}[y_{i,j}|\theta_j] &\equiv \text{E}[(y_{i,j} - \text{E}[y_{i,j}|\theta_j])^2|\theta_j] \\
&= \text{E}[(y_{i,j} - \theta_j)^2|\theta_j] \\
&= \text{E}[(\theta_j + \epsilon_{i,j} - \theta_j)^2|\theta_j] \\
&= \text{E}[\epsilon_{i,j}^2|\theta_j] = \sigma^2
\end{aligned}
$$

$$
\begin{aligned}
\text{Cov}[y_{i_1,j}, y_{i_2,j}|\theta_j] &\equiv \text{E}[(y_{i_1,j} - \text{E}[y_{i_1,j}|\theta_j]) \times (y_{i_2,j} - \text{E}[y_{i_2,j}|\theta_j])|\theta_j] \\
&= \text{E}[(y_{i_1,j} - \theta_j) \times (y_{i_2,j} - \theta_j)|\theta_j] \\
&= \text{E}[\epsilon_{i_1,j}\epsilon_{i_2,j}|\theta_j] = 0
\end{aligned}
$$

Population level variance and covariance

Across all groups,

$$\begin{aligned} a_1, \ldots, a_m &\sim \text{i.i.d. normal}(0, \tau^2) \\ \{y_{1,j}, \ldots, y_{n_j,j}\} &\sim \text{i.i.d. normal}(\mu + a_j, \sigma^2) \end{aligned}$$

For a randomly sampled observation i from a randomly sampled group j,

$$\begin{aligned} \text{E}[y_{i,j}|\mu] &= \text{E}[\mu + a_j + \epsilon_{i,j}|\mu] \\ &= \text{E}[\mu|\mu] + \text{E}[a_j|\mu] + \text{E}[\epsilon_{i,j}|\mu] \\ &= \mu + 0 + 0 = \mu \end{aligned}$$

This is the *population mean*.

Population level variance and covariance

Variation *around the population mean* μ is as follows:

$$\begin{aligned} \mathrm{E}[y_{i,j}|\mu] &= \mathrm{E}[\mu + a_j|\mu] = \mu + 0 = \mu, \\ \mathrm{Var}[y_{i,j}|\mu] &= \sigma^2 + \tau^2, \\ \mathrm{Cov}[y_{i_1,j}, y_{i_2,j}|\mu] &= \tau^2. \end{aligned}$$

In words,

- sampled observations *across groups* are centered around μ;
- the variation of the sample *around* μ is $\sigma^2 + \tau^2$;
- the observations within a group are correlated *around* μ.

Regarding correlation: Knowing how far $y_{1,j}$ is from μ *does* inform you about how far $y_{2,j}$ is from μ.

Population level variance

$$
\begin{aligned}
\text{Var}[y_{i,j}|\mu] &\equiv \text{E}[(y_{i,j} - \text{E}[y_{i,j}|\mu])^2|\mu] \\
&= \text{E}[(y_{i,j} - \mu)^2|\mu] \\
&= \text{E}[(\mu + a_j + \epsilon_{i,j} - \mu)^2|\mu] \\
&= \text{E}[(a_j + \epsilon_{i,j})^2|\mu] \\
&= \text{E}[a_j^2 + 2a_j\epsilon_{i,j} + \epsilon_{i,j}^2|\mu] \\
&= \tau^2 + 0 + \sigma^2 = \sigma^2 + \tau^2
\end{aligned}
$$

Exercise: Draw a picture of within and across group sampling.

Population level covariance and correlation

$$
\begin{aligned}
\text{Cov}[y_{i_1,j}, y_{i_2,j}|\mu] &\equiv \text{E}[(y_{i_1,j} - \text{E}[y_{i_1,j}|\mu]) \times (y_{i_2,j} - \text{E}[y_{i_2,j}])|\mu] \\
&= \text{E}[(y_{i_1,j} - \mu) \times (y_{i_2,j} - \mu)|\mu] \\
&= \tau^2
\end{aligned}
$$

$$
\begin{aligned}
\text{Cor}[y_{i_1,j}, y_{i_2,j}|\mu] &\equiv \frac{\text{Cov}[y_{i_1,j}, y_{i_2,j}|\mu]}{\sqrt{\text{Var}[y_{i_1,j}|\mu]\text{Var}[y_{i_2,j}|\mu]}} \\
&= \frac{\tau^2}{\tau^2 + \sigma^2} \equiv \rho
\end{aligned}
$$

The correlation ρ is the *intraclass correlation coefficient*.

Estimation of τ^2 and ρ

The easiest way to estimate τ^2 is using the method-of-moments. Recall,

$$
\begin{aligned}
MSA &= \frac{1}{m-1}\sum_j\sum_i(\bar{y}_j - \bar{y}_{\cdot\cdot})^2 \\
&= \frac{n}{m-1}\sum(\bar{y}_j - \bar{y}_{\cdot\cdot})^2 \\
\mathrm{E}[MSA|a_1,\ldots,a_m] &= \frac{n}{m-1}\left(\frac{m-1}{n}\sigma^2 + \sum a_j^2\right) \\
&= \sigma^2 + n \times \frac{1}{m-1}\sum a_j^2.
\end{aligned}
$$

The expectation of *MSA* over samples *and groups* is given by

$$
\begin{aligned}
\mathrm{E}[\mathrm{E}[MSA|a_1,\ldots,a_m]] &= \mathrm{E}[\sigma^2 + n \times \frac{1}{m-1}\sum a_j^2] \\
&= \sigma^2 + n \times \mathrm{E}[\frac{1}{m-1}\sum a_j^2] \\
&= \sigma^2 + n\tau^2.
\end{aligned}
$$

(In the ANOVA parameterization, $\sum a_j^2 = \sum(a_j - \bar{a})^2$ becuase $\bar{a} = 0$)

Estimation of τ^2 and ρ

The result suggests

$$\widehat{\sigma^2 + n\tau^2} = MSA.$$

How to estimate τ^2? Recall $\mathrm{E}[MSW] = \sigma^2$, so we can use

$$\hat{\sigma^2} = MSW.$$

This suggests

$$\widehat{n\tau^2} = MSA - MSW$$
$$\hat{\tau}^2 = (MSA - MSW)/n.$$

Comments:

- $MSA - MSW$ *could* be negative. If so, it is standard to set $\hat{\tau}^2 = 0$.
- If sample sizes are unequal, the formula must be modified slightly:

$$\hat{\tau}^2 = (MSA - MSW)/\tilde{n}$$

  where there is a horrible formula for $\tilde{n}$.

Unequal sample sizes

$$\hat{\tau}^2 = (MSA - MSW)/\tilde{n}$$

$$\tilde{n} = \bar{n} - \frac{\text{sample variance}(n_1, \ldots, n_m)}{m\bar{n}}$$

where $\bar{n} = \sum_j n_j/m = \text{sample mean}(n_1, \ldots, m_m)$.

Estimation of τ^2 and ρ

It is common to use a "plug-in" estimate of ρ:

$$\hat{\rho} = \widehat{\frac{\tau^2}{\tau^2 + \sigma^2}} = \frac{\hat{\tau}^2}{\hat{\tau}^2 + \hat{\sigma}^2}.$$

A standard error for ρ (with which we can get a CI) is

$$\text{se}(\hat{\rho}) = (1 - \hat{\rho}) \times (1 + (n-1)\hat{\rho})\sqrt{\frac{2}{n(n-1)(m-1)}}.$$

Example: Wheat

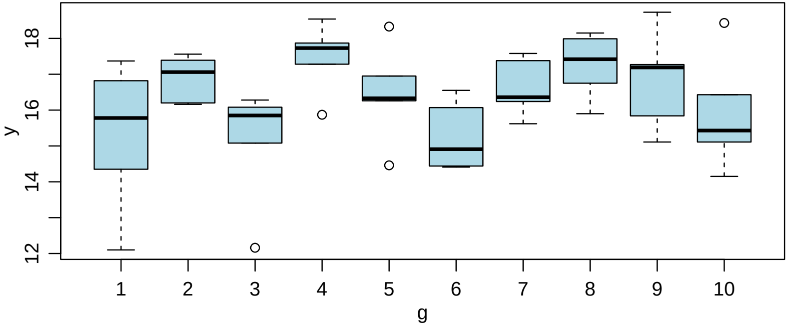


```
anova(lm(y~as.factor(g)))

## Analysis of Variance Table
##
## Response: y
##              Df Sum Sq Mean Sq F value Pr(>F)
## as.factor(g)  9 33.368  3.7076  2.0745 0.0555 .
## Residuals    40 71.488  1.7872
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Example: Wheat

```
fit<-anova( lm(y~as.factor(g)) )

MSA<-fit[1,3]
MSW<-fit[2,3]

MSA

## [1] 3.70759

MSW

## [1] 1.787206

t2<-(MSA-MSW)/n

t2

##         1
## 0.3840768
```

Example: Wheat

```
rho<-t2/(t2+MSW)

rho

##         1
## 0.1768894

se.rho<- (1-rho)*(1+(n-1)*rho)*sqrt( 2/(  n*(n-1)*(m-1)))

rho + c(-2,2)*se.rho

## [1] -0.1194179  0.4731966
```

Two-stage sampling

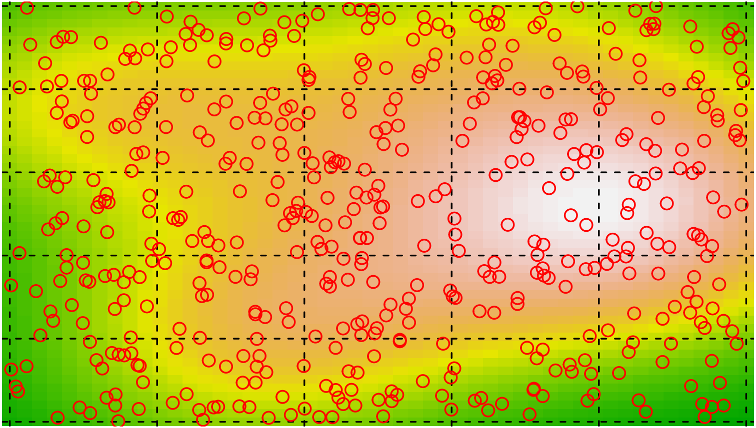

μ=2.1124814

Ignoring across-group heterogeneity

Task: Construct a 95% CI for the population mean.

***t*-interval for SRS:**
If $y_1, \ldots, y_n$ is an iid sample with $\mathrm{E}[y_i] = \mu$ and $\mathrm{Var}[y_i] = \sigma^2$,

$$\mathrm{E}[\bar{y}] = \mu \ , \ \mathrm{Var}[\bar{y}] = \sigma^2/n.$$

By the central limit theorem,

$$\bar{y} \dot{\sim} N(\mu, \sigma^2/n) \ , \ \frac{\bar{y} - \mu}{\sigma/\sqrt{n}} \dot{\sim} N(0,1).$$

As σ^2 is generally unknown, we use

$$\frac{\bar{y} - \mu}{s/\sqrt{n}} \dot{\sim} t_{n-1}, \ \text{, where } s^2 = \frac{1}{n-1} \sum (y_i - \bar{y})^2.$$

From this, we have

$$\bar{y} \pm t_{n-1,.975} \times s/\sqrt{n} \ \text{ is a 95\% CI for } \mu.$$

Ignoring across-group heterogeneity

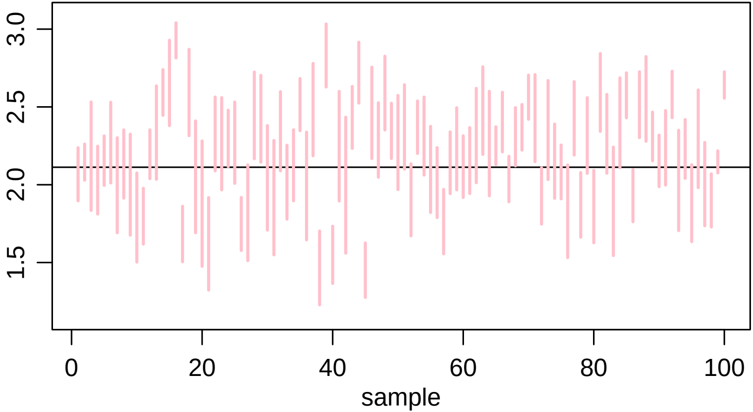

Building an accurate t-interval

Recall that an approximate 95% CI for μ is given by

$$\bar{y} \pm 2 \times \text{se}(\bar{y}),$$

where $\text{se}(\bar{y})$ is an approximation to the standard deviation of $\bar{y}$.

How to find $\text{se}(\bar{y})$:

1. compute the variance v of $\bar{y}$ based on the model;
2. find an estimate $\hat{v}$ of v;
3. let $\text{se}(\bar{y}) = \sqrt{v}$.

So the first step is to find $\text{Var}[\bar{y}]$:

Variance of the grand mean around population mean

$$
\begin{aligned}
\text{Var}[\bar{y}] &= \text{Var}[\frac{1}{mn}\sum_j \sum_i y_{i,j}] \\
&= \text{Var}[\frac{1}{m}\sum_j \frac{1}{n}\sum_i y_{i,j}] \\
&= \text{Var}[\frac{1}{m}\sum_j \bar{y}_j] \\
&= \frac{1}{m^2}\text{Var}[\sum_j \bar{y}_j] \\
&= \frac{1}{m^2}\sum_j \text{Var}[\bar{y}_j] \\
&= \frac{1}{m^2} m \text{Var}[\bar{y}_1] \\
&= \frac{1}{m}\text{Var}[\bar{y}_1]
\end{aligned}
$$

Variance of a group mean around population mean

What is $\text{Var}[\bar{y}_1]$? We've shown

$$\text{Var}[y_{i,1}] = \sigma^2 + \tau^2,$$

but generally,

$$\text{Var}[\bar{y}_1] \neq [\sigma^2 + \tau^2]/n.$$

Quiz: What is the smallest that $\text{Var}[\bar{y}_1]$ could be for fixed σ^2 and n? Recall

$$\text{Cor}[y_{i,1}, y_{i,2}] = \frac{\tau^2}{\tau^2 + \sigma^2}$$

Answer: When τ^2 is zero the within group samples are independent and so

$$\text{Var}[\bar{y}_1] \geq \sigma^2/n$$

Variance of a group mean around population mean

Quiz: what is the smallest that $\text{Var}[\bar{y}_1]$ could be for fixed σ^2 and τ^2?

Answer: Increasing *n* reduces variation of $\bar{y}_1$ around θ_1, but across group heterogeneity remains:

$$\begin{aligned} \text{for large } n, \bar{y}_1 &\approx \theta_1 \\ \text{Var}[\theta_1] &= \tau^2 \\ \text{Var}[\bar{y}_1] &\geq \tau^2 \end{aligned}$$

Variance of a group mean around population mean

Let's compute $\text{Var}[\bar{y}_1]$. For notational convenience, we'll drop the group index, and assume $\mu = 0$, so

$$\text{E}[y_i] = 0 \ , \ \text{E}[y_i^2] = \sigma^2 + \tau^2 \ , \ \text{E}[y_i y_k] = \tau^2$$

In this case,

$$\begin{aligned}
\text{Var}[\bar{y}] &= \text{E}[\bar{y}^2] \\
&= \text{E}[\frac{1}{n^2}(\sum y_i)^2] \\
&= \frac{1}{n^2}\text{E}[\sum y_i^2 + \sum_{i \neq k} y_i y_j] \\
&= \frac{1}{n^2}(n[\sigma^2 + \tau^2] + n(n-1)\tau^2) \\
&= \frac{\sigma^2}{n} + \frac{1}{n}\tau^2 + \frac{n-1}{n}\tau^2 \\
&= \frac{\sigma^2}{n} + \tau^2
\end{aligned}$$

Exercise: Make sure the answer makes sense to you intuitively.

Variance of the sample grand mean

$$\text{Var}[\bar{y}_{\cdot\cdot}] = \frac{1}{m}\text{Var}[\bar{y}_j]$$

$$\text{Var}[\bar{y}_j] = \frac{1}{n}\sigma^2 + \tau^2$$

$$\text{Var}[\bar{y}_{\cdot\cdot}] = \frac{1}{nm}\sigma^2 + \frac{1}{m}\tau^2$$

What happens as

- $n \to \infty$ and m stays fixed?
- $m \to \infty$ and n stays fixed?

In this sense, m is the "sample size" for the population-level parameter μ.

Standard error and CI

$$\widehat{\mathrm{Var}}[\bar{y}_{\cdot\cdot}] = \frac{1}{nm}\hat{\sigma}^2 + \frac{1}{m}\tau^2$$

- $\hat{\sigma}^2 = MSW$
- $\hat{\tau}^2 = (MSA - MSW)/n$

$$\widehat{\mathrm{Var}}[\bar{y}_{\cdot\cdot}] = \frac{1}{mn} MSA$$

This should make sense, because previously we claimed

$$\mathrm{E}[MSA] = \sigma^2 + n \times \tau^2,$$

so

$$\mathrm{E}[\frac{1}{mn} MSA] = \frac{1}{mn}\sigma^2 + \frac{1}{m}\tau^2 = \mathrm{Var}[\bar{y}_{\cdot\cdot}]$$

Confidence interval

$$\bar{y}_{\cdot\cdot} \pm 2 \times \sqrt{MSA/mn}$$

```
round(y,2)

##  [1] 0.55 0.56 0.48 0.85 0.81 2.76 2.71 2.47 2.43 2.43 2.68 2.52 2.97 2.92 2.60
## [16] 2.42 1.90 1.99 2.37 1.87

g

##  [1] 1 1 1 1 1 2 2 2 2 2 3 3 3 3 3 4 4 4 4 4

anova(lm(y~as.factor(g)))

## Analysis of Variance Table
##
## Response: y
##              Df  Sum Sq Mean Sq F value    Pr(>F)
## as.factor(g)  3 13.4751  4.4917   110.5 6.603e-11 ***
## Residuals    16  0.6504  0.0406
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

MSA<-anova(lm(y~as.factor(g)))[1,3]

mean(y) + c(-2,2)*sqrt( MSA/(m*n)  )

## [1] 1.066935 2.962551

mean(y) + c(-2,2)*sqrt( var(y)/(m*n)  )

## [1] 1.629141 2.400345
```

Accounting for across-group heterogeneity

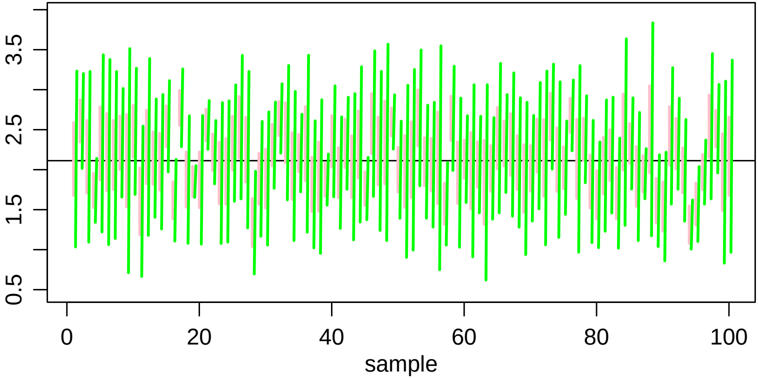


```
mean(  CI.tss0[,1] < mu  & mu < CI.tss0[,2] )

## [1] 0.729

mean(  CI.tss1[,1] < mu  & mu < CI.tss1[,2] )

## [1] 0.933
```

Summary

$$y_{i,j} = \mu + a_j + \epsilon_{i,j}$$
$$\text{Var}[\epsilon_{i,j}] = \sigma^2$$
$$\text{Var}[a_j] = \tau^2$$

Variation around the group mean: $\theta_j = \mu + a_j$

- $\text{Var}[y_{i,j}|\theta_j] = \sigma^2$
- $\text{Cov}[y_{i_1,j}, y_{i_2,j}|\theta_j] = 0$
- $Exp\bar{y}_j|\theta_j = \theta_j$, $\text{Var}[\bar{y}_j|\theta_j] = \sigma^2/n$

Variation around the grand mean:

- $\text{Var}[y_{i,j}|\mu] = \sigma^2 + \tau^2$
- $\text{Cov}[y_{i_1,j}, y_{i_2,j}|\mu] = \tau^2$
- $\text{E}[\bar{y}_j|\mu] = \mu$, $\text{Var}[\bar{y}_j|\mu] = \sigma^2/n + \tau^2$
- $\text{E}[\bar{y}_{\cdot\cdot}|\mu] = \mu$, $\text{Var}[\bar{y}_{\cdot\cdot}|\mu] = \sigma^2/(mn) + \tau^2/m$