

STA 610L: MODULE 2.9

RANDOM EFFECTS ANCOVA (HOLISTIC ANALYSIS)

DR. OLANREWaju MICHAEL AKANDE

NELS DATA: TAKING A STEP BACK

Until now, we have used the NELS data to illustrate different aspects of model fitting for the multilevel model.

Now let's step back and think about model selection for the data more holistically, as if we're seeing them for the first time (for the most part).

NELS VARIABLES

Here are our variables of interest in the NELS:

- Math score (individual-level outcome)
- SES (individual-level socio-economic status)
- FLP (school level % of kids eligible for free or reduced-price lunch -- think of this as school-level SES)
 - 1: 0-5% eligible
 - 2: 5-30% eligible
 - 3: >30% eligible
- Enrollment (school level # of kids in 10th grade, rounded and measured in hundreds, so 0= \leq 100, 1=around 100, ..., 5=around 500)
- Public (school level, takes value 1 if public school and 0 if private school)
- Urbanicity (school level factor with levels rural, suburban, and urban)

MODEL SELECTION

As we think about our model selection process, we'll keep in mind a couple of methods for comparison.

- Likelihood ratio test for nested models
 - For tests involving fixed effects only, we can use a χ_d^2 for testing whether d fixed effects all equal 0 (ML, not ok for REML)
 - For tests involving random effects only, we can use a 50-50 mixture of χ_{p-1}^2 and χ_p^2 , where p is the number of random effect variances in the larger model
 - Non-nested models or testing both fixed and random effects, not so simple.

MODEL SELECTION

One option is to rely on the metrics we already often use for model comparison, like AIC, BIC, etc.





For BIC in particular, we have the following

- smaller is better
- it already adjusts for model complexity
- there is an approximation to posterior model probability
- model selection is consistent
- nesting between models is not required

NELS DATA

```
load('data/nels.Rdata')
avmscore.schools <- tapply(nels$mscore,nels$school,mean,na.rm=TRUE)
id.schools <- names(avmscore.schools)
m <- length(id.schools)
nels$sesstd <- nels$ses/sd(nels$ses)
nels$enroll <- factor(nels$enroll)
nels$flp <- factor(nels$flp)
nels$public <- factor(nels$public)
nels$urban <- factor(nels$urban)
```

DESCRIPTIVE STATISTICS



WHAT'S WRONG WITH ANOVA?

Suppose I don't really care about school effects one way or the other. Why not just use ANOVA (or other fixed effects model) here?

Under a fixed effects model,

$$\text{Cov}(y_j) = \begin{pmatrix} \sigma^2 & 0 & \dots & 0 \\ 0 & \sigma^2 & \dots & 0 \\ \vdots & & \ddots & \\ 0 & 0 & \dots & \sigma^2 \end{pmatrix}$$

WHAT'S WRONG WITH ANOVA?

Under a random intercepts model,

$$\text{Cov}(y_j) = \begin{pmatrix} \sigma^2 + \tau^2 & \tau^2 & \dots & \tau^2 \\ \tau^2 & \sigma^2 + \tau^2 & \dots & \tau^2 \\ \vdots & & & \vdots \\ \tau^2 & \tau^2 & \dots & \sigma^2 + \tau^2 \end{pmatrix},$$

and

$$\text{Corr}(y_{ij}, y_{i'j}) = \frac{\tau^2}{\tau^2 + \sigma^2}$$

We generally don't believe independence within the same school environment holds.

This type of covariance structure is often called *exchangeable* or *compound symmetric*.

OTHER CONSIDERATIONS

Why not treat school as a fixed effect? That should handle the school heterogeneity.

```
m10 <- lm(mscore~school+enroll+flp+public+
          urbanicity, data=nels)
#summary(m10)
coef(m10)[(length(coef(m10))-30):length(coef(m10))]
```

OTHER CONSIDERATIONS

##	school4513	school4521	school4522	school4531
##	1.771737	3.085000	4.330590	-5.556333
##	school4532	school4541	school4542	school4551
##	1.619069	1.912625	4.158000	1.240200
##	school4552	school4553	school4561	school4562
##	2.027769	7.574857	8.552385	1.357000
##	school4571	school4572	school4582	school4591
##	-3.348000	4.821000	9.443250	6.169727
##	school4592	school4601	school4602	school4611
##	12.405182	-13.559667	3.622333	5.820846
##	school4612	enroll1	enroll2	enroll3
##	-7.980692	NA	NA	NA
##	enroll4	enroll5	flp2	flp3
##	NA	NA	NA	NA
##	public1	urbanicitysuburban	urbanicityurban	
##	NA	NA	NA	

What happened to the estimates for enrollment, eligibility for free lunch, public/private status, and urbanicity?

OTHER CONSIDERATIONS

The school-specific fixed effects explain approximately *all* heterogeneity in means across schools, leaving basically no room for the other factors (which we care more about in terms of learning about patterns in the data) to explain any heterogeneity.

So this approach does not allow us to evaluate school-level predictors, and it is also very expensive in terms of spending degrees of freedom (estimating a lot of parameters).

This is a relatively common phenomenon when dealing with categorical group-level predictors.

HETEROGENEITY ACROSS SCHOOLS

Let's take a more detailed look at the heterogeneity across schools and how much of that can be explained by measured school-level factors including urbanicity, public/private status, free lunch percentage, and school size.

In a model with only a random intercept, let's calculate the intraclass correlation -- the correlation between two kids in the same school.

$$y_{ij} = \beta_{0j} + \varepsilon_{ij}, \quad \beta_{0j} \stackrel{iid}{\sim} N(\beta_0, \tau^2) \perp \varepsilon_{ij} \stackrel{iid}{\sim} N(0, \sigma^2)$$

```
fit0 <- lmer(mscore~(1|school),data=nels, REML=FALSE)
sigma2hat <- sigma(fit0)*sigma(fit0) #pick off estimate of sigma2
tau2hat <- as.numeric(VarCorr(fit0)$school) #pick off est of tau2
c(sigma2hat,tau2hat,tau2hat/(tau2hat+sigma2hat)) #show vars and correlation
```

```
## [1] 73.7084447 23.6341046 0.2427932
```

HETEROGENEITY ACROSS SCHOOLS

How much of the heterogeneity across schools is explained by enrollment?

```
fit1 <- lmer(mscore~enroll+(1|school),data=nels, REML=FALSE)
sigma2hat <- sigma(fit1)*sigma(fit1) #pick off estimate of sigma2
tau2hat <- as.numeric(VarCorr(fit1)$school) #pick off est of tau2
c(sigma2hat,tau2hat,tau2hat/(tau2hat+sigma2hat)) #show vars and correlation
```

```
## [1] 73.7202995 23.0948621 0.2385459
```

Not much!

HETEROGENEITY ACROSS SCHOOLS

How much of the remaining heterogeneity across schools is explained by the percentage of kids eligible for free or reduced price lunch?

```
fit2=lmer(mscore~enroll+flp+(1|school),  
          data=nels, REML=FALSE)  
sigma2hat <- sigma(fit2)*sigma(fit2) #pick off estimate of sigma2  
tau2hat <- as.numeric(VarCorr(fit2)$school) #pick off est of tau2  
c(sigma2hat,tau2hat,tau2hat/(tau2hat+sigma2hat)) #show vars and correlation
```

```
## [1] 73.7655328 13.5097521 0.1547947
```

Wow, "school-level SES" explained a lot of that heterogeneity.

HETEROGENEITY ACROSS SCHOOLS

What if we add public/private status?

```
fit3=lmer(mscore~enroll+flp+public+
          (1|school),data=nels, REML=FALSE)
sigma2hat=sigma(fit3)*sigma(fit3) #pick off estimate of sigma2
tau2hat=as.numeric(VarCorr(fit3)$school) #pick off est of tau2
c(sigma2hat,tau2hat,tau2hat/(tau2hat+sigma2hat)) #show vars and correlation
```

```
## [1] 73.7749179 13.2366759 0.1521254
```


HETEROGENEITY ACROSS SCHOOLS

Now we add urbanicity.

```
fit4=lmer(mscore~enroll+flp+public+
          urbanicity+(1|school),data=nels, REML=FALSE)
sigma2hat=sigma(fit4)*sigma(fit4) #pick off estimate of sigma2
tau2hat=as.numeric(VarCorr(fit4)$school) #pick off est of tau2
c(sigma2hat,tau2hat,tau2hat/(tau2hat+sigma2hat)) #show vars and correlation
```

```
## [1] 73.779034 12.908770 0.148911
```

SUMMARY

As we add more group-level predictors,

- $\hat{\tau}^2$ decreases
- $\hat{\sigma}^2$ stays about the same
- the within-group correlation is nonincreasing (and with the addition of some variables decreases substantially)

NELS DATA

Let's return to our data from a data analysis perspective (rather than just illustrating aspects of the multi-level model), considering the hypotheses regarding the role of school-specific and individual-specific factors in math test scores.

We'll start with a simple model and build from there, using the BIC as our primary selection criterion.

$$y_{ij} = \beta_{0j} + \beta_{1j}\text{ses}_{ij} + \varepsilon_{ij}, \quad \beta_{0j} = \beta_0 + b_{0j} \quad \beta_{1j} = \beta_1 + b_{1j}$$

$$\begin{pmatrix} b_{0j} \\ b_{1j} \end{pmatrix} \sim N \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \tau_{11} & \tau_{12} \\ \tau_{12} & \tau_{22} \end{pmatrix} \right) \quad \varepsilon_{ij} \sim N(0, \sigma^2)$$

This model allows random intercepts and slopes across schools.

NELS DATA

We saw previously that the random slope did explain additional heterogeneity in a model without school-level predictors.

We'll come back to that question again once we add a few school level predictors to the model.

Let's first compare our starting model to models that add enrollment to the mix, so that

$$\beta_{0j} = \beta_0 + \alpha_{0k}I(\text{enroll}_j = k) + b_{0j}$$

$$\beta_{1j} = \beta_1 + \alpha_{1k}I(\text{enroll}_j = k) + b_{1j}$$

$$k = 1, \dots, 5$$

We'll use ML estimation because we may wish to consider likelihood ratio tests of the mean parameters.

NELS DATA

First, check out the base model.

```
mod1=lmer(mscore~sesstd+(sesstd|school),data=nels, REML=FALSE)
summary(mod1)
```

```
## Linear mixed model fit by maximum likelihood ['lmerMod']
## Formula: mscore ~ sesstd + (sesstd | school)
## Data: nels
##
##      AIC      BIC   logLik deviance df.resid
## 92553.1 92597.9 -46270.5 92541.1    12968
##
## Scaled residuals:
##      Min       1Q   Median       3Q      Max
## -3.8910 -0.6382  0.0179  0.6669  4.4613
##
## Random effects:
##  Groups   Name                Variance Std.Dev. Corr
##  school   (Intercept) 12.2231   3.4961
##           sesstd      0.8562   0.9253   0.11
## Residual                67.3451   8.2064
## Number of obs: 12974, groups: school, 684
##
## Fixed effects:
##              Estimate Std. Error t value
## (Intercept) 50.67670    0.15511   326.7
## sesstd      3.27708    0.09256    35.4
##
## Correlation of Fixed Effects:
##          (Intr)
## sesstd 0.007
```

NELS DATA

```
mod2a=lmer(mscore~enroll+sesstd+(sesstd|school),
            data=nels, REML=FALSE)
mod2b=lmer(mscore~enroll+sesstd+enroll:sesstd+
            (sesstd|school),data=nels, REML=FALSE)
anova(mod2b,mod2a)
```

```
## Data: nels
## Models:
## mod2a: mscore ~ enroll + sesstd + (sesstd | school)
## mod2b: mscore ~ enroll + sesstd + enroll:sesstd + (sesstd | school)
##      npar   AIC   BIC logLik deviance Chisq Df Pr(>Chisq)
## mod2a   11 92557 92639 -46267    92535
## mod2b   16 92559 92678 -46263    92527 7.9798  5      0.1574
```

```
anova(mod2a,mod1)
```

```
## Data: nels
## Models:
## mod1: mscore ~ sesstd + (sesstd | school)
## mod2a: mscore ~ enroll + sesstd + (sesstd | school)
##      npar   AIC   BIC logLik deviance Chisq Df Pr(>Chisq)
## mod1     6 92553 92598 -46271    92541
## mod2a   11 92557 92639 -46267    92535 6.1315  5      0.2936
```

Here we don't see much evidence that enrollment is useful, so we don't need to use it.

NELS DATA

Next we consider eligibility for free and reduced lunch.

Here we'll explore a variety of models, including the one above, a model without the interaction with flp, a model that has the flp main effect but drops the SES random effect, and a model that drops all the school random effects given that flp is in the model ($\tau = 0$).

```
mod3a=lmer(mscore~flp+sesstd+(sesstd|school),
           data=nels, REML=FALSE)
mod3b=lmer(mscore~flp+sesstd+flp*sesstd+
           (sesstd|school),data=nels, REML=FALSE)
mod3c=lmer(mscore~flp+sesstd+(1|school),
           data=nels, REML=FALSE)
mod3d=lm(mscore~flp+sesstd,data=nels)
anova(mod3b,mod3a)
anova(mod3a,mod1)
anova(mod3c,mod3a) #just look at BIC here
BIC(mod3d) #check if random intercept needed by comparing to BIC from 3c
```

MODEL SELECTION

```
mod3a=lmer(mscore~flp+sesstd+(sesstd|school),  
            data=nels, REML=FALSE)  
mod3b=lmer(mscore~flp+sesstd+flp*sesstd+(sesstd|school),  
            data=nels, REML=FALSE)  
anova(mod3b,mod3a)
```

```
## Data: nels  
## Models:  
## mod3a: mscore ~ flp + sesstd + (sesstd | school)  
## mod3b: mscore ~ flp + sesstd + flp * sesstd + (sesstd | school)  
##      npar   AIC    BIC logLik deviance Chisq Df Pr(>Chisq)  
## mod3a     8 92395 92454 -46189    92379  
## mod3b    10 92384 92459 -46182    92364 14.384  2  0.0007525
```


MODEL SELECTION

```
anova(mod3a,mod1)
```

```
## Data: nels
## Models:
## mod1: mscore ~ sesstd + (sesstd | school)
## mod3a: mscore ~ flp + sesstd + (sesstd | school)
##      npar   AIC   BIC logLik deviance  Chisq Df Pr(>Chisq)
## mod1      6 92553 92598 -46271    92541
## mod3a      8 92395 92454 -46189    92379 162.51  2 < 2.2e-16
```

MODEL SELECTION

```
mod3c=lmer(mscore~flp+sesstd+(1|school),  
           data=nels, REML=FALSE)  
anova(mod3c,mod3a) #just look at BIC here
```

```
## Data: nels  
## Models:  
## mod3c: mscore ~ flp + sesstd + (1 | school)  
## mod3a: mscore ~ flp + sesstd + (sesstd | school)  
##      npar   AIC    BIC logLik deviance Chisq Df Pr(>Chisq)  
## mod3c     6 92404 92449 -46196    92392  
## mod3a     8 92395 92454 -46189    92379 13.371  2    0.001249
```

```
mod3d=lm(mscore~flp+sesstd,data=nels)  
anova(mod3c,mod3d) #check if random intercept needed by comparing to BIC from 3c
```

```
## Data: nels  
## Models:  
## mod3d: mscore ~ flp + sesstd  
## mod3c: mscore ~ flp + sesstd + (1 | school)  
##      npar   AIC    BIC logLik deviance Chisq Df Pr(>Chisq)  
## mod3d     5 93115 93152 -46552    93105  
## mod3c     6 92404 92449 -46196    92392 712.63  1 < 2.2e-16
```

```
BIC(mod3d) #also use BIC
```

```
## [1] 93151.9
```

MODEL SELECTION

Note that BIC now likes the model without a random slope -- we evaluated that because we thought that after introducing a school-level SES variable to the model, the importance of the individual-level SES variable may change.

It also prefers a model without an interaction between individual-level SES and school-level SES (measured by flp).

MODEL SELECTION

Now our model for the coefficients is

```
summary(mod3c)

## Linear mixed model fit by maximum likelihood ['lmerMod']
## Formula: mscore ~ flp + sesstd + (1 | school)
## Data: nels
##
##      AIC      BIC   logLik deviance df.resid
## 92403.9 92448.7 -46196.0 92391.9    12968
##
## Scaled residuals:
##      Min       1Q   Median       3Q      Max
## -3.9560 -0.6434  0.0178  0.6710  4.4906
##
## Random effects:
## Groups   Name      Variance Std.Dev.
## school  (Intercept) 9.004    3.001
## Residual                67.959    8.244
## Number of obs: 12974, groups: school, 684
##
## Fixed effects:
##              Estimate Std. Error t value
## (Intercept) 52.84307    0.24462 216.020
## flp2        -1.87992    0.33042  -5.689
## flp3        -4.79607    0.36150 -13.267
## sesstd       3.10819    0.08578  36.233
##
## Correlation of Fixed Effects:
##      (Intr) flp2  flp3
## flp2   -0.739
## flp3   -0.697  0.514
## sesstd -0.202  0.143  0.237
```

MODEL SELECTION

The more students we have eligible for the free and reduced price lunch program, the lower the math scores.

In addition, the coefficient on individual-level SES did not change much in magnitude -- so SES operates both on the school level and the individual level.

Let's now add the public school indicator.

```
mod4a=lmer(mscore~flp+public+sesstd+
            (1|school),data=nels, REML=FALSE)
mod4b=lmer(mscore~flp+public +
            sesstd+public*sesstd+(1|school),
            data=nels, REML=FALSE)
anova(mod4b,mod4a)
anova(mod4b,mod3c)
#summary(mod4b)
```

MODEL SELECTION

```
mod4a=lmer(mscore~flp+public+sesstd+(1|school),  
            data=nels, REML=FALSE)  
mod4b=lmer(mscore~flp+public + sesstd+  
            public*sesstd+(1|school),data=nels, REML=FALSE)  
anova(mod4b,mod4a)
```

```
## Data: nels  
## Models:  
## mod4a: mscore ~ flp + public + sesstd + (1 | school)  
## mod4b: mscore ~ flp + public + sesstd + public * sesstd + (1 | school)  
##      npar   AIC    BIC logLik deviance  Chisq Df Pr(>Chisq)  
## mod4a     7 92406 92458 -46196    92392  
## mod4b     8 92396 92456 -46190    92380 11.948  1  0.000547
```

MODEL SELECTION

```
anova(mod4b,mod3c)
```

```
## Data: nels
## Models:
## mod3c: mscore ~ flp + sesstd + (1 | school)
## mod4b: mscore ~ flp + public + sesstd + public * sesstd + (1 | school)
##      npar   AIC    BIC logLik deviance Chisq Df Pr(>Chisq)
## mod3c     6 92404 92449 -46196    92392
## mod4b     8 92396 92456 -46190    92380 12.081  2   0.002381
```

The BIC suggests leaving public/private out of the model.

MODEL SELECTION

Now let's consider urban/suburban/rural status.

```
mod5a=lmer(mscore~flp+urbanicity+sesstd+(1|school),  
            data=nels, REML=FALSE)  
mod5b=lmer(mscore~flp+urbanicity+sesstd+  
            urbanicity*sesstd+(1|school),  
            data=nels, REML=FALSE)  
anova(mod5b,mod5a)  
anova(mod5a,mod3c)  
summary(mod5b)
```


MODEL SELECTION

```
mod5a=lmer(mscore~flp+urbanicity+sesstd+(1|school),  
            data=nels, REML=FALSE)  
mod5b=lmer(mscore~flp+urbanicity+sesstd+  
            urbanicity*sesstd+(1|school),  
            data=nels, REML=FALSE)  
anova(mod5b,mod5a)
```

```
## Data: nels  
## Models:  
## mod5a: mscore ~ flp + urbanicity + sesstd + (1 | school)  
## mod5b: mscore ~ flp + urbanicity + sesstd + urbanicity * sesstd + (1 | school)  
##      npar   AIC    BIC logLik deviance Chisq Df Pr(>Chisq)  
## mod5a     8 92400 92460 -46192    92384  
## mod5b    10 92390 92465 -46185    92370 13.373  2    0.001248
```

MODEL SELECTION

```
anova(mod5a,mod3c)
```

```
## Data: nels
## Models:
## mod3c: mscore ~ flp + sesstd + (1 | school)
## mod5a: mscore ~ flp + urbanicity + sesstd + (1 | school)
##      npar   AIC   BIC logLik deviance Chisq Df Pr(>Chisq)
## mod3c     6 92404 92449 -46196    92392
## mod5a     8 92400 92460 -46192    92384 8.0578  2    0.01779
```

BIC suggests leaving urbanicity out of the model.

SUMMARY OF SELECTION USING BIC

- Enrollment, urbanicity, and public/private status did not add much to our model using the BIC as our selection criterion
- The lower the SES status of the whole school (measured by percent eligible for free and reduced-price lunch), the lower the math scores on average
- Having higher individual-level SES was associated with higher math scores regardless of the school environment
- A random intercept for school explained significant variability across schools and controlled for lack of independence within schools

$$y_{ij} = \beta_{0j} + \beta_1 \text{ses}_{ij} + \varepsilon_{ij}$$

$$\beta_{0j} = \beta_0 + \psi_{0l} I(\text{flp}_j = l) + b_{0j}$$

$$b_{0j} \sim N(0, \tau^2) \quad \varepsilon_{ij} \sim N(0, \sigma^2)$$

FINAL MODEL AGAIN

```
summary(mod3c)
```

```
## Linear mixed model fit by maximum likelihood ['lmerMod']
## Formula: mscore ~ flp + sesstd + (1 | school)
## Data: nels
##
##      AIC      BIC   logLik deviance df.resid
## 92403.9 92448.7 -46196.0 92391.9   12968
##
## Scaled residuals:
##      Min       1Q   Median       3Q      Max
## -3.9560 -0.6434  0.0178  0.6710  4.4906
##
## Random effects:
## Groups Name Variance Std.Dev.
## school (Intercept) 9.004 3.001
## Residual 67.959 8.244
## Number of obs: 12974, groups: school, 684
##
## Fixed effects:
## Estimate Std. Error t value
## (Intercept) 52.84307 0.24462 216.020
## flp2 -1.87992 0.33042 -5.689
## flp3 -4.79607 0.36150 -13.267
## sesstd 3.10819 0.08578 36.233
##
## Correlation of Fixed Effects:
## (Intr) flp2 flp3
## flp2 -0.739
## flp3 -0.697 0.514
## sesstd -0.202 0.143 0.237
```

WHAT'S NEXT?

MOVE ON TO THE READINGS FOR THE NEXT MODULE!