Random Effects Likelihood RatioTest Examples

The result of maximum likelihood estimation is a -2 log likelihood value, which is a summary of the fit of the observed to the expected values. These values can be used for comparing different models that are nested (see the "Significance Testing in Multilevel Regression" handout). The difference in likelihood values can be evaluated against the chi-square distribution for significance—the likelihood ratio test. If the models differ only in the random effects, REML estimation is fine. If the fixed effects differ at all, then full ML should be used.

SPSS

To illustrate the likelihood ratio test approach, I use the HSB data to compare the model with SES as a level-1 predictor (uncentered) with varying slopes.[1]

Test of Slope Variance using the Wald Test

Estimates of Covariance Parameters[a]

Parameter		Estimate	Std. Error	Wald Z	Sig.	95% Confidence Interval	
						Lower Bound	Upper Bound
Residual		36.830165	.629312	58.524	.000	35.617160	38.084480
Intercept + ses [subject = schoolid]	UN (1,1)	4.828637	.672123	7.184	.000	3.675711	6.343190
	UN (2,1)	-.154275	.298837	-.516	.606	-.739984	.431434
	UN (2,2)	.412929	.235014	1.757	.079	.135340	1.259873

a. Dependent Variable: mathach.

The UN() notation refers to the rows and columns of the variance-covariance matrix. The UN(1,1,) row refers to the intercept variance, because β_{j0} is the first parameter and an element with the same row and column number refers to the variance, called τ_0^2 in the text. The UN(2,2) row refers to the second row and second column of the variance-covariance matrix, so if β_{j0} is the first row and first column and β_{j1} is the second row and second column, then UN(2,2) refers to the variance of the β_{j1} slope, known as τ_1^2. UN(2,1) then refers to the covariance between the intercept β_{j0} and the slope β_{j1}, known as τ_{01}. So, the table indicates that the intercept variance is significant, even after halving the p-value (one-tailed *p*-value = .000/2 = <.01 at least), the variance of the slope is significant with a one-tailed test (.079/2 = .0395). The *p*-value for the covariance between the intercept and slope p = .606 (two-tailed test that is not halved) is not significant.

Likelihood Ratio Test of the Just the Covariance

It may be of interest, particular with lower power, to examine the likelihood ratio test of the covariance between intercept and slope even though a Wald test is automatically provided in the "Estimates of Covariance Parameters" box when `COVTYPE(UN)` is used on the `RANDOM` subcommand. A model without the covariance is tested by specifying `COVTYPE(VC)`, for variance components, on the `RANDOM` subcommand.

Here is the original model with the random effect for slope and the intercept-slope covariance estimated because `COVTYPE(UN)` was used

```
MIXED mathach WITH ses
  /METHOD = REML
  /PRINT = SOLUTION TESTCOV
  /FIXED =  ses | SSTYPE(3)
  /RANDOM = INTERCEPT ses  | SUBJECT(schoolid) COVTYPE(UN).
```

[1] I do not generally recommend uncentered predictors in most circumstances, but SES in the HSB data set was pre-standardized so it has a mean of zero.

Information Criteria[a]

-2 Restricted Log Likelihood	46640.398
Akaike's Information Criterion (AIC)	46648.398
Hurvich and Tsai's Criterion (AICC)	46648.404
Bozdogan's Criterion (CAIC)	46679.916
Schwarz's Bayesian Criterion (BIC)	46675.916

The information criteria are displayed in smaller-is-better forms.

a. Dependent Variable: mathach.

The model is retested using a different specification on the `RANDOM` subcommand `COVTYPE(VC)` which requests a diagonal matrix assuming covariance between intercept and slope is equal to 0 (not a reasonable assumption usually).

```
MIXED mathach WITH ses
  /CRITERIA=MXITER(1000) SCORING(1)
  /METHOD = REML
  /PRINT = SOLUTION TESTCOV HISTORY
  /FIXED =  ses | SSTYPE(3)
  /RANDOM = INTERCEPT  ses | SUBJECT(schoolid) COVTYPE(VC).
```

In the results, notice that there is no covariance estimate—the UN(2,1) value is absent.

Information Criteria[a]

-2 Restricted Log Likelihood	46640.398
Akaike's Information Criterion (AIC)	46648.398
Hurvich and Tsai's Criterion (AICC)	46648.404
Bozdogan's Criterion (CAIC)	46679.916
Schwarz's Bayesian Criterion (BIC)	46675.916

The information criteria are displayed in smaller-is-better form.

a. Dependent Variable: mathach.

Estimates of Covariance Parameters[a]

Parameter		Estimate	Std. Error	Wald Z	Sig.	95% Confidence Interval Lower Bound	95% Confidence Interval Upper Bound
Residual		36.822286	.628871	58.553	.000	35.610122	38.075712
Intercept [subject = schoolid]	Variance	4.852841	.673828	7.202	.000	3.696621	6.370702
ses [subject = schoolid]	Variance	.424034	.235016	1.804	.071	.143097	1.256521

a. Dependent Variable: mathach.

The likelihood ratio test subtracts the -2 log likelihood value for the previous model with the covariance estimated (same as D1 below), from this more restricted model 46640.398 with the covariance not estimated (set to 0), 46640.663. The resulting chi-square test can be compared to a standard chi-square table. The difference in this case is not significant, $\chi^2(1) = .265$, ns.

Likelihood Ratio Test of the Slope Variance and Intercept-Slope Covariance Together

D_0 model in which slope is non-varying

```
MIXED mathach WITH ses
  /METHOD = REML
  /PRINT = SOLUTION TESTCOV
  /FIXED =  ses | SSTYPE(3)
  /RANDOM = INTERCEPT   | SUBJECT(schoolid)
COVTYPE(UN).
```

Information Criteria[a]

-2 Restricted Log Likelihood	46645.169
Akaike's Information Criterion (AIC)	46649.169
Hurvich and Tsai's Criterion (AICC)	46649.171
Bozdogan's Criterion (CAIC)	46664.928
Schwarz's Bayesian Criterion (BIC)	46662.928

The information criteria are displayed in smaller-is-better forms.

a. Dependent Variable: mathach.

D_1 model in which slope is allowed to vary

```
MIXED mathach WITH ses
  /METHOD = REML
  /PRINT = SOLUTION TESTCOV
  /FIXED =  ses | SSTYPE(3)
  /RANDOM = INTERCEPT ses  | SUBJECT(schoolid)
COVTYPE(UN).
```

Information Criteria[a]

-2 Restricted Log Likelihood	46640.398
Akaike's Information Criterion (AIC)	46648.398
Hurvich and Tsai's Criterion (AICC)	46648.404
Bozdogan's Criterion (CAIC)	46679.916
Schwarz's Bayesian Criterion (BIC)	46675.916

The information criteria are displayed in smaller-is-better forms.

a. Dependent Variable: mathach.

Because *p*-values are halved for variances but not for covariances a mixture of the two p-values is needed. I did the adjustments using a spreadsheet (available on the class website).
$D_0 - D_1 = 46645.163 - 46640.398 = 4.765$ This value does not exceed the 5.14 chi-square cutoff value from the mixture ("chi-bar") chi-square critical value in Snijders & Bosker (2012, p. 99), so the test of the slope and covariance between slope and intercept, tested together, is not significant.

R

```
> library(lme4)
> #install.packages("lmerTest")
> library(lmerTest)  #lmerTest generates Satterthwaitte df with summary function
> model1 <- lmer(mathach ~ ses + (ses|schoolid), data = mydata, REML = TRUE)
> summary(model1)
Linear mixed model fit by REML
t-tests use  Satterthwaite approximations to degrees of freedom ['lmerMod']
Formula: mathach ~ ses + (ses | schoolid)
   Data: mydata

REML criterion at convergence: 46640.4

Scaled residuals:
     Min       1Q   Median       3Q      Max
-3.12272 -0.73046  0.02144  0.75610  2.94356

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 schoolid (Intercept)  4.8286  2.1974
          ses          0.4129  0.6426   -0.11
 Residual             36.8302  6.0688
Number of obs: 7185, groups:  schoolid, 160

Fixed effects:
            Estimate Std. Error       df t value            Pr(>|t|)
(Intercept)  12.6650     0.1898 145.5500   66.71 <0.0000000000000002
ses           2.3938     0.1181 157.5300   20.27 <0.0000000000000002

Correlation of Fixed Effects:
    (Intr)
ses -0.045
```

Although the profile likelihood confidence intervals can be obtained with the `confint()` function (demonstrated previously) to test the covariance for significance, the `anova` function from the `lmerTest` package can be used to perform a likelihood ratio test to compare two nested models. Note that it retests the models using full maximum likelihood, which is not necessary when just random effects differ but should be ok with larger sample sizes.

```
> library(lmerTest)
> #sets the covariance equal 0 (diagonal matrix of random effects)
> model2 <- lmer (mathach ~ ses + (1|schoolid) + (0+ses|schoolid), data=mydata)
> summary(model2)

Random effects:
 Groups     Name        Variance Std.Dev.
 schoolid   (Intercept)  4.853   2.2029
 schoolid.1 ses          0.424   0.6511
 Residual               36.822   6.0681
Number of obs: 7185, groups:  schoolid, 160

> #conducts a LR comparison of the unconstrained and constrained models
> anova(model,model2)
refitting model(s) with ML (instead of REML)
Data: mydata
Models:
model2: mathach ~ ses + (1 | schoolid) + (0 + ses | schoolid)
model: mathach ~ ses + (ses | schoolid)
       Df   AIC   BIC logLik deviance  Chisq Chi Df Pr(>Chisq)
model2  5 46647 46681 -23318    46637
model   6 46648 46690 -23318    46636 0.2762      1     0.5992
```

Likelihood Ratio Test of the Slope Variance and Intercept-Slope Covariance Together

```
> library(lmerTest)
> #rand function reports LR comparison to intercept only model using mixture chi-square
> rand(model1)
Analysis of Random effects Table:
            Chi.sq Chi.DF p.value
ses:schoolid   4.77      2    0.09
```

HLM

The HLM package provides a method of likelihood ratio tests by requesting a comparison of a current model with a prior model. The first step is to test the full model, which is not shown but was exactly as the model demonstrated earlier with random slopes for SES specified and the covariance between intercept and slope estimated by default.

Likelihood Ratio Test of the Intercept-Slope Covariance

Under the Other Settings menu and then Estimation Settings one can check the box Diagonalize Tau, which sets all covariances equal to zero. When that model is estimated, the covariance between the intercept and slope will be set to 0. HLM automatically generates the likelihood ratio test when this box is checked.

τ

INTRCPT1,β_0	4.85319	0.00000
SES,β_1	0.00000	0.42574

Variance-Covariance components test

χ^2 statistic = 0.26428
Degrees of freedom = 3
p-value = >.500

Likelihood Ratio Test of the Slope Variance and Intercept-Slope Covariance Together

To obtain the likelihood ratio test below, I first tested the model with non-varying slopes. The deviance and number of parameters from this model (D_0) was:

Deviance = 46643.331427
Number of estimated parameters = 2

I then enter these values into the **Hypothesis Testing** dialog box under the **Other Settings** menu when I test a new model (D_1) that allows the SES slopes to vary across schools.

Test of Slope Variance Using the HLM Chi-square Parameter Test

Here are the results:
σ^2 = 36.82835
τ

INTRCPT1,β_0	4.82978	-0.15399
SES,β_1	-0.15399	0.41828

Final estimation of variance components

Random Effect	Standard Deviation	Variance Component	*d.f.*	χ^2	*p*-value
INTRCPT1, u_0	2.19768	4.82978	159	905.26472	<0.001
SES slope, u_1	0.64675	0.41828	159	216.21178	0.002
level-1, *r*	6.06864	36.82835			

The result of the model with varying slopes had a significant intercept variance, χ^2 = 905.26472, $p <$ 0.001 and slope variance χ^2 = 216.21178, p = 0.002. *p*-values are not halved in HLM. No test of the covariance between the intercept and the slope is printed.

Statistics for current covariance components model
Deviance = 46638.560929
Number of estimated parameters = 4

Variance-Covariance components test
χ^2 statistic = 4.77050
Degrees of freedom = 2
p-value = 0.090

The chi-square represents the $D_0 – D_1$ difference, which this test indicates is non-significant. No mixture adjustment is needed as the HLM tests already makes a comparable adjustment.