Random Effects Likelihood RatioTest Examples

The result of maximum likelihood estimation is a -2 log likelihood value (Deviance), 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. Each likelihood ratio test is a test of whether one or more parameters (whichever parameters differ between the two models) are significantly different from zero.

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	95% Confidence Interval 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

Just as one illustration I first conduct a likelihood ratio test of the covariance between intercept and slope. It may be of interest, particularly with low power due to a small number of groups, to examine the likelihood ratio test instead of the Wald test, which is automatically provided in the "Estimates of Covariance Parameters" box when `COVTYPE(UN)` is used on the `RANDOM` subcommand.

To conduct the likelihood ratio test of the covariance, I first tested 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).
```

I omit most of the results, which were reported in the “Random Slopes Example: SPSS, R, and HLM.” We just need the -2 log likelihood value (Deviance) from that output.

[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 at or near 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 then retested using a different specification on the `RANDOM` subcommand `COVTYPE(VC)`, which stands for variance components and 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	46640.663
Akaike's Information Criterion (AIC)	46646.663
Hurvich and Tsai's Criterion (AICC)	46646.666
Bozdogan's Criterion (CAIC)	46670.301
Schwarz's Bayesian Criterion (BIC)	46667.301

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 (traditional two-tailed) can be compared to a standard chi-square table. The difference in this case is not significant, $\chi^2(1) = .265$, ns, and is the same conclusion the Wald test gave.

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

To illustrate the mixture distribution ("chi-bar") test when the test involves a variance and a covariance, which should be a mixture of a one-tailed and a two-tailed test respectively, I compare a model with and without the random effect requested (SES is not included on the random line). This restricts the variance for the slopes to 0 but also the covariance, so two parameters are different.

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).
```

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	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.

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.

Likelihood Ratio Test of Just the Variance

We could further do a comparison of the model with the covariance restricted (as in the first model above) to a model without a random slope. This test is testing the same hypothesis as the covariance test for UN(2,2) printed in our original random slope model.

I just need to test a model with no random slope (same as in the handout "ANCOVA Example (One Level-1 Predictor Assuming Homogeneous Slopes): SPSS, R, and HLM") and using REML estimation (because the only difference is a random effect). The model we will compare against has already been conducted above where I restricted the covariance to 0. For this second model, we obtain:

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

Information Criteria[a]

-2 Restricted Log Likelihood	46645.169313
Akaike's Information Criterion (AIC)	46649.169313
Hurvich and Tsai's Criterion (AICC)	46649.170985
Bozdogan's Criterion (CAIC)	46664.928258
Schwarz's Bayesian Criterion (BIC)	46662.928258

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		37.034	.625	59.218	<.001	35.829	38.281
Intercept [subject = schoolid]	Variance	4.768	.655	7.278	<.001	3.642	6.241

a. Dependent Variable: mathach.

The difference between this -2 Restricted Log Likelihood, 46645.169, and the model obtained with a random slope and no covariance, 46640.663, is 4.506, which for a 1-df chi-square test (one-sided) with a critical value of 2.706, is significant. I should mention that I am wary of this particular test because the comparison model with no covariance is likely to be unreasonable and may have trouble with convergence or reflects an incorrect model.

R

Likelihood Ratio Tests

```
> library(lme4)
> 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
```

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
```

Partial Least Squares Profile Likelihood Approach for Random Effects

The profile likelihood confidence intervals (Bates et al., 2015) can be obtained with the confint() function to test any of the random effects. The method uses asymmetric CIs, so 95% level (default) can be used for any of them.

```
> confint(model1)
Computing profile confidence intervals ...
                  2.5 %     97.5 %
.sig01       1.91159331  2.5112989
.sig02      -1.00000000  0.3078580
.sig03       0.08487641  0.9666694
.sigma       5.96867353  6.1719565
(Intercept) 12.28852890 13.0406703
ses          2.15925014  2.6317003
```

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

The mixture (“chi-bar”) chi-square can be obtained with the rand() function has the same result and conclusion as the manual method I used with SPSS above.

```
> 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

I just conducted one comparison, the model with and without a random slope, both with the covariance constrained to be zero. To test the first model with no covariance between intercept and slope, under Other Settings, choose Estimation Settings and check the box Diagonalized Tau to set the correlation between intercept and slope to 0. REML estimation is fine because the two models will only differ in the random effect. (Not needed here, but for nested models involving any fixed effects, go to Estimation Settings, check the Full Maximum Likelihood button). Run that model allowing the slope to vary across schools (u_j is estimated). This part of the printout shows no covariance is estimated:

τ

INTRCPT1,β_0	4.85319	0.00000
SES,β_1	0.00000	0.42574

τ (as correlations)

INTRCPT1,β_0	1.000	0.000
SES,β_1	0.000	1.000

Statistics for current covariance components model

Deviance = 46638.825210
Number of estimated parameters = 3

Then run a second model but first go to Other Settings then Hypothesis Testing and then put in the value of the deviance (46638.825210) from the first model and the number of parameters from the first model (which was 3).

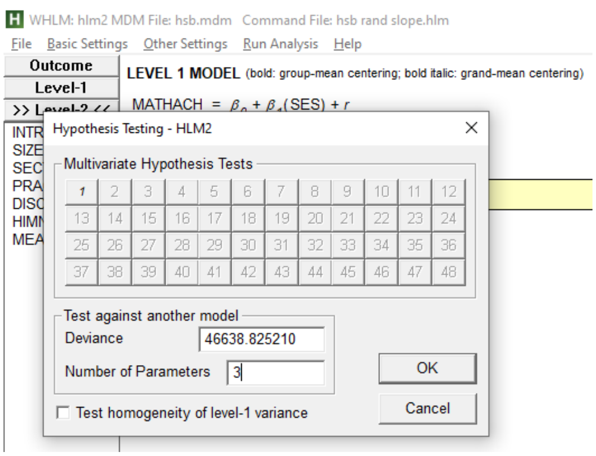

Hit ok, then rerun the model with the u_j not estimated. The nested test is printed at the bottom. At the bottom of the output, you will find the result of the likelihood ratio test.

Statistics for current covariance components model
Deviance = 46643.331427
Number of estimated parameters = 2

Variance-Covariance components test

χ^2 statistic = 4.50622
Degrees of freedom = 1
p-value = 0.032

The difference is 4.50622 and is significant. This value and conclusion match what was obtained for the same difference test manually computed from SPSS values.

References

Bates, D.,Maechler, M.M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software, 67*(1), 1-48.