CAMPBELL
COLLABORATION

Random and Mixed-effects Modeling

Overview

- Effect-size estimates
- Random-effects model
- Mixed model

Effect sizes

Suppose we have computed effect-size estimates from k studies, we will call them

$T_1, T_2, \ldots, T_k$

Call their variances (squares of their SE's)

$v_1, v_2, \ldots, v_k$

Call the population effect sizes

$\theta_1, \theta_2, \ldots, \theta_k$

When is the random-effects model appropriate?

If all population parameters are equal ($\theta_i = \theta$), we have the fixed-effects model

$$T_i = \theta + e_i \qquad \text{for } i = 1 \text{ to } k.$$

- T_i: Observed study outcome
- θ: Single population parameter
- e_i: Residual deviation due to sampling error

All studies are modeled as having the same effect θ.

When is the random-effects model appropriate?

Suppose that a test of homogeneity has indicated more between-studies variation than would be expected due simply to sampling error. Then we have the random-effects model

$$T_i = \theta_i + e_i \quad \text{for } i = 1 \text{ to } k.$$

- T_i: Observed study outcome
- θ_i: **Per-study** population parameter
- e_i: Residual deviation due to sampling error

Each study has its **own** population effect, θ_i.

When is the random-effects model appropriate?

This model can be written in one more form:

$$T_i = \theta_i + e_i \quad \text{for } i = 1 \text{ to } k.$$

$$T_i = \overbrace{\mu_\theta + u_i}^{\theta_i} + e_i \quad \text{for } i = 1 \text{ to } k,$$

where $\theta_i = \mu_\theta + u_i$

We have replaced the unique θ_i with an average effect μ_θ plus a component u_i representing between-studies variation.

When is the random-effects model appropriate?

Suppose that we have a predictor that explains some between-studies variation. Then we have the mixed-effects model

$$T_i = \beta X_i + u_i + e_i \quad \text{for } i = 1 \text{ to } k.$$

Observed study outcome (T_i)

Population parameter predicted by X (βX_i)

Residual deviation due to sampling and other error ($u_i + e_i$)

Each study has its **own** population effect, θ_i.

Random-effects Model

Random-effects analyses

$$T_i = \theta_i + e_i$$

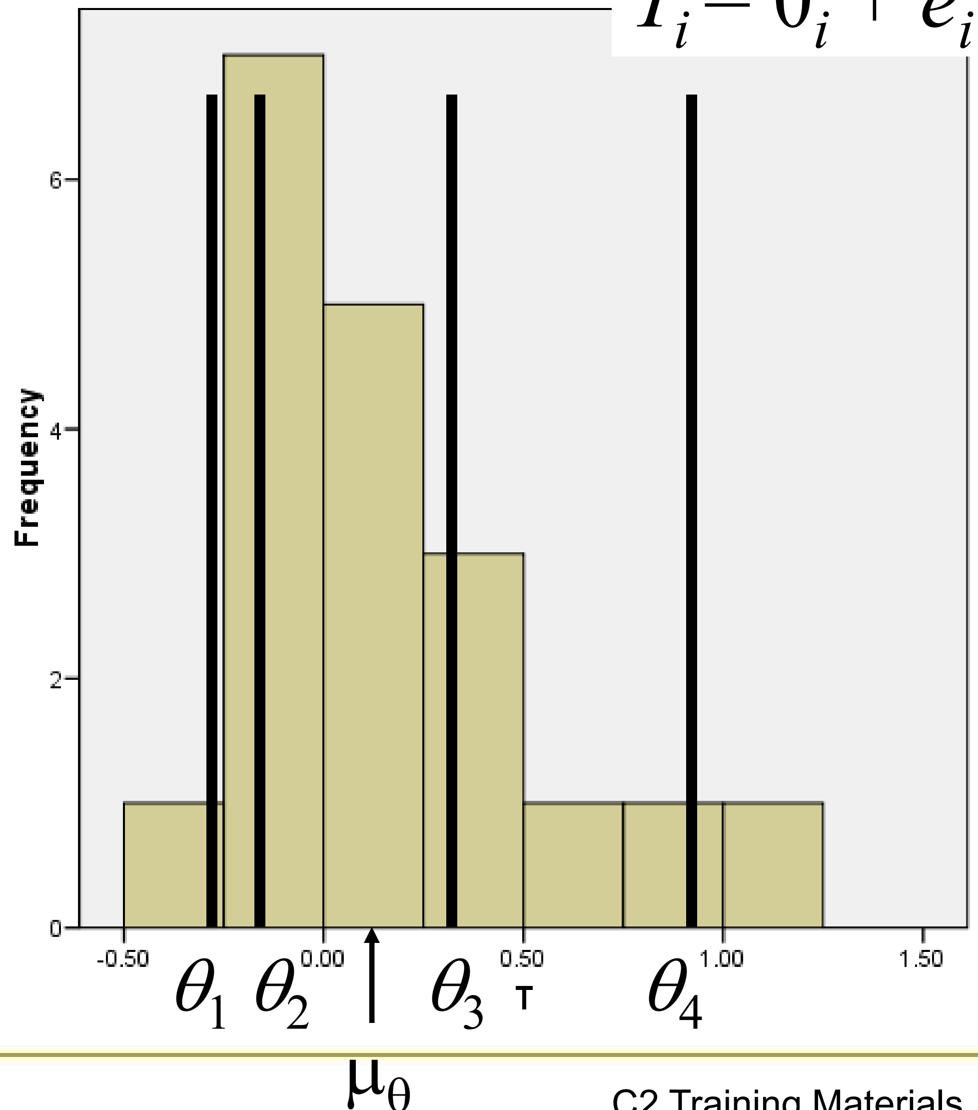

The bars represent four true θ_i values, and the histogram shows the observed effect sizes. μ_θ would be the mean of the θs.

Mean =0.16
Std. Dev. =0.359
N =19

Random-effects analyses

In random-effects analyses, the goal is to estimate

- the mean population effect size μ_θ of the populations from which the observed studies are a sample (central tendency of the ❙ bars)

and

- between-study variation in effect sizes (in the θ_i or u_i), which has an impact on weights and the uncertainty of this mean. This is often called τ^2 or σ^2_θ . It is the variation in the ❙ bars.

The random-effects model

We will add the between-studies variance $\hat{\sigma}_{\theta}^{2}$ to each study's v_i and use weighted least squares (WLS) estimation with new random-effects weights.

The new variances for each study will be larger than the fixed-effects variances.

Because of this, sometimes means that were significant under the fixed-effects model may no longer be significant under the random model.

Random-effects analyses: Example

*teacher.sav [DataSet1] - SPSS Data Editor

2 : N_E 60

	ID	N_E	N_C	T
1	1	79	339	.03
2	2	60	198	.12
3	3	72	72	-.14
4	4	11	22	1.18
5	5	11	22	.26
6	6	129	348	-.06
7	7	110	636	-.02
8	8	26	99	-.32
9	9	75	74	.27
10	10	32	32	.80
11	11	22	22	.54
12	12	43	38	.18
13	13	24	24	-.02
14	14	19	32	.23
15	15	80	79	-.18
16	16	72	72	-.06
17	17	65	255	.30
18	18	233	224	.07
19	19	65	67	-.07
20				

Let us consider an example based on the teacher expectancy data from Raudenbush (1984).

The 19 effects compare IQ test scores for students randomly labeled as "bloomers" to other unlabeled students.

Random-effects analyses: Example

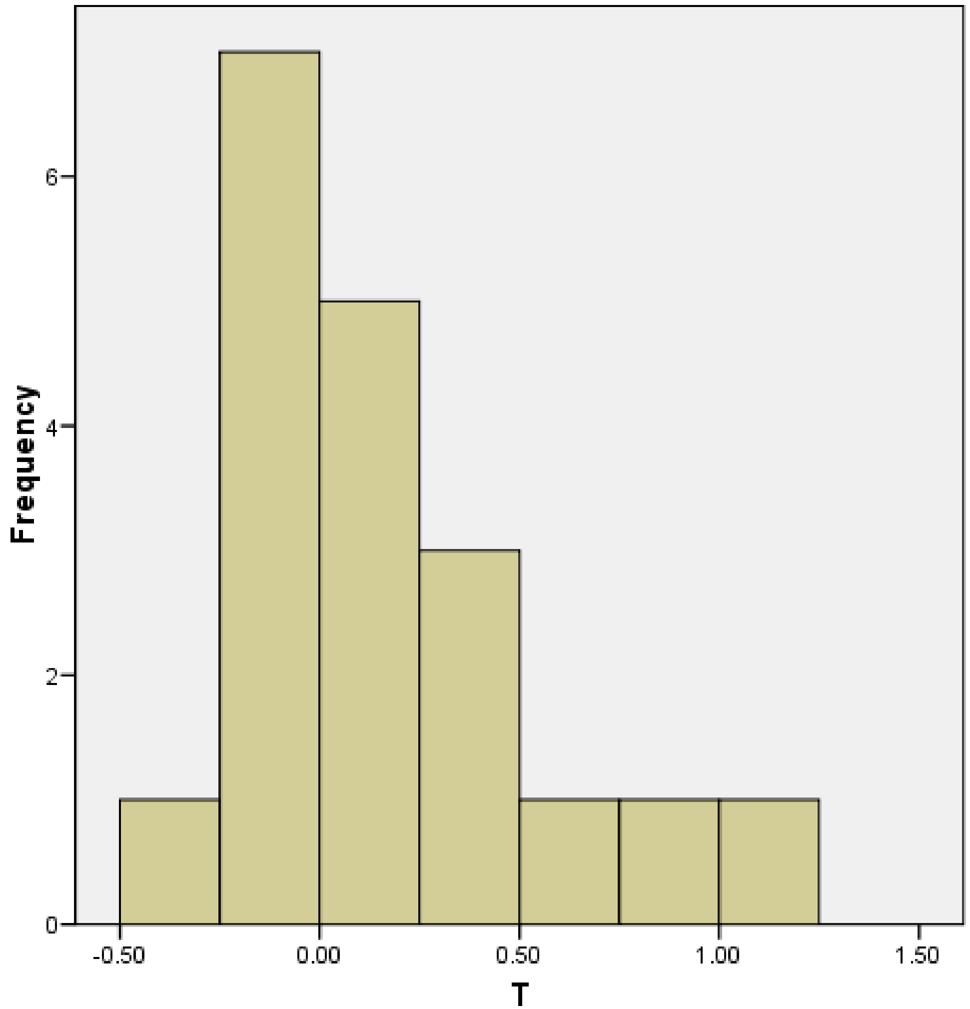

The histogram shows that effects vary from about -0.5 to above 1 standard deviation.

Mean =0.16
Std. Dev. =0.359
N =19

Random-effects analyses: Example

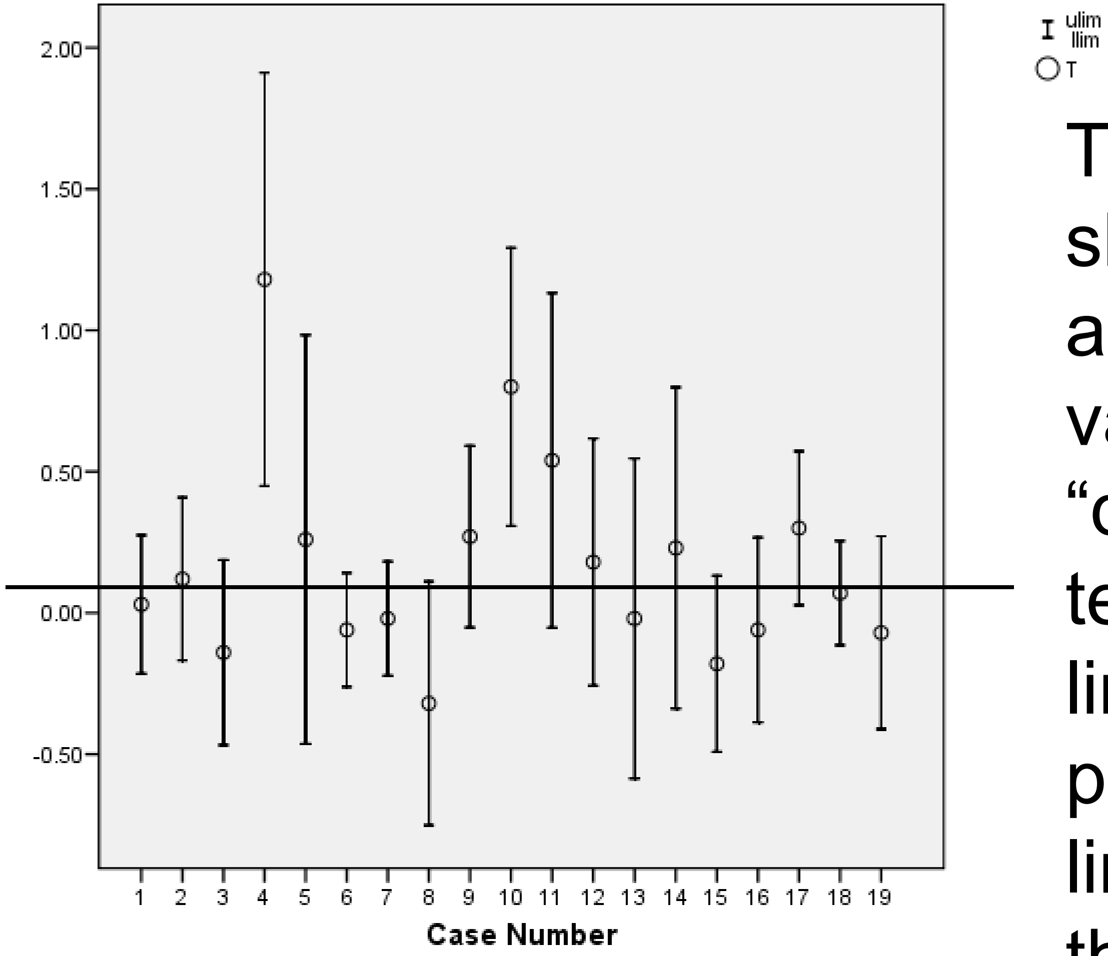

The 95% CI plot shows a fair amount of variation. The “quick and dirty” test of drawing a line across the plot shows no line can cross all the CIs.

Random-effects analyses: Example

The overall test of homogeneity, obtained from the SPSS GLM menu, shows significant between studies variation. This is a chi-square test with 18 df, and p = .007.

Tests of Between-Subjects Effects[b]

Dependent Variable: T

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	.000[a]	0	.	.	.
Intercept	2.735	1	2.735	1.374	.256
Error	35.825	18	1.990		
Total	38.561	19			
Corrected Total	35.825	18			

a. R Squared = .000 (Adjusted R Squared = .000)

b. Weighted Least Squares Regression - Weighted by w

Estimating the variance of $\theta_1, \theta_2, \ldots, \theta_k$

One estimator (SVAR on slide 18) is

$$\hat{\tau}^2 = \hat{\sigma}_\theta^2 = S_T^2 - \bar{v} = .129 - .0484 = .0806$$

where $S^2{}_T$ is the simple variance of the observed effects and $\bar{v}$ is the mean of the "known" fixed-effects variances v_i, $\bar{v} = \sum v_i / k$

$S^2{}_T$

Descriptive Statistics

	N	Minimum	Maximum	Mean	Std. Deviation	Variance
T	19	-.32	1.18	.1637	.35887	.129
V	19	.01	.14	.0484	.04063	.002
Valid N (listwise)	19					

$\bar{v}$

Estimating the variance of $\theta_1, \theta_2, \ldots, \theta_k$

Another estimator of τ^2 or σ^2_θ is

$$\hat{\tau}^2 = [Q - (k-1)]\,/c$$

where

$$c = \left(\sum_{i=1}^{k} w_i - \frac{\sum_{i=1}^{k} w_i^2}{\sum_{i=1}^{k} w_i} \right)$$

whenever the estimate is greater than 0, and 0 otherwise. This is called QVAR on slide 18.

Random-effects analyses: Example

We can also get the Q test and variances from SAS. Here we see results of estimation of the simple random-effects variance for the data – the more conservative is .08:

OBS	K	Q	P	LL_T_DOT	T_DOT	UL_T_DOT	V_T_DOT	SE_T_DOT
1	19	35.8254	.0074284	-0.011168	0.060343	0.13185	.0013312	0.036485

Fixed-Effects Effect Size Analysis (Exercise 2), tchrexp.dat

OBS	MODVAR	MODSD	QVAR	QSD
1	0.022398	0.14966	0.025920	0.16100

$\hat{\sigma}_\theta^2$

OBS	_TYPE_	_FREQ_	K	SUMV	SUMT	SUMT2	SVAR	SSD
1	0	19	19	0.92009	3.11	2.8273	0.080366	0.28349

Random-effects analyses: Example

For the random-effects model, we need to compute an average of the differing effects.

We use a weighted mean – BUT we weight each data point by the inverse of its random-effects variance (i.e., $w^*_i = 1/[V(T_i) + \hat{\sigma}_\theta^2]$:

$$T.^* = \sum_{i=1}^{k} \frac{w_i^* T_i}{w_i^*} = \sum \frac{T_i /[V(T_i) + \hat{\sigma}_\theta^2]}{1/[V(T_i) + \hat{\sigma}_\theta^2]}$$

The random-effects model: Variances

The fixed effects variance is v and the random effects variance is vstar = $v_i + .08$.

Not only are the random effects variances larger, but they are also more equal in size.

	T	V	vstar
1	.03	.0156	.0966
2	.12	.0216	.1026
3	-.14	.0279	.1089
4	1.18	.1391	.2201
5	.26	.1362	.2172
6	-.06	.0106	.0916
7	-.02	.0106	.0916
8	-.32	.0484	.1294
9	.27	.0269	.1079
10	.80	.0630	.1440
11	.54	.0912	.1722
12	.18	.0497	.1307
13	-.02	.0835	.1645
14	.23	.0841	.1651
15	-.18	.0253	.1063
16	-.06	.0279	.1089
17	.30	.0193	.1003
18	.07	.0088	.0898
19	-.07	.0303	.1113
20			

Data View | Variable View

Random-effects analyses: Example

We use the random-effects variance to compute a new mean. It is somewhat larger than the fixed effects mean. The printed SE needs to be corrected, via $SE = SE_{printed}\sqrt{MSE}$.

Here $SE = .071 \times .89 = .064$ (vs. $SE_{fixed} = .036$)

Grand Mean[a]

Dependent Variable: T

		95% Confidence Interval	
Mean	Std. Error	Lower Bound	Upper Bound
.114	.071	-.035	.264

a. Weighted Least Squares Regression - Weighted by wstar

Random-effects analyses: Example

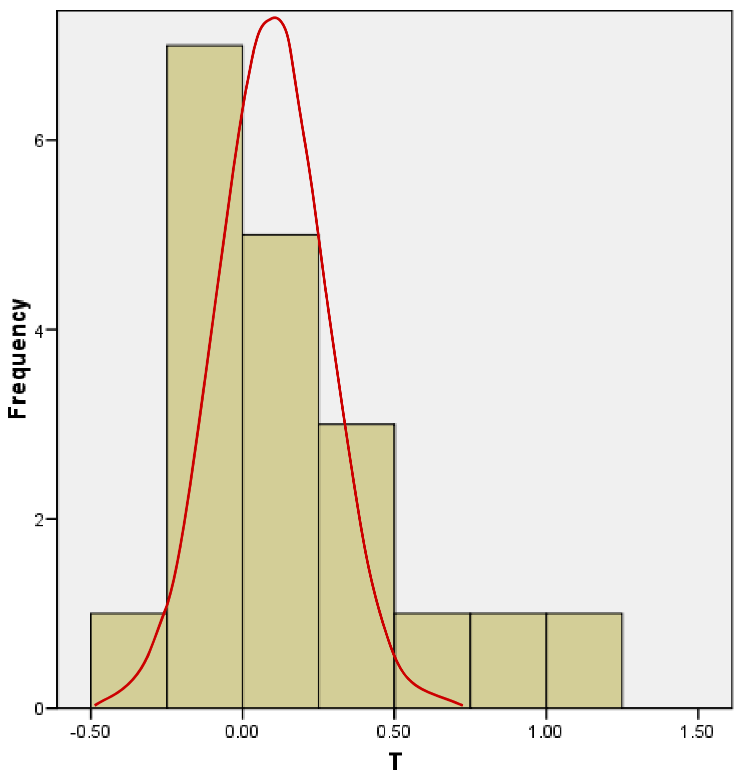

If we compute the mean $\pm$ 1.96 $\hat{\sigma}_\theta$ we can get a range of possible population θ_i values. Computing

0.114 $\pm$ 1.96 x 0.283, 95% of the θ_i values are estimated to be between -.44 and .67,assuming a normal distribution of θ_is.

Mixed-effects Model

When is the mixed-effects model appropriate?

Suppose we have attempted to fit a model (either a regression model or the ANOVA-like categorical model) and the model is significant but does not explain all variation. That is, we find that both Q_{Model} and Q_{Error} are significant.

We want to keep the predictors that are useful, but also to account for the remaining uncertainty. In this case, the mixed-effects model can be adopted.

The components in the mixed-effects model

This terminology is different from the use of the label "mixed model" in ANOVA where sometimes this term can refer to a model with both between-subjects and within-subjects terms.

Mixed-effects models in meta-analysis contain both *fixed* and *random* components.

The components in the mixed-effects model

For the ANOVA-like model, with study i from the jth group or category of studies, we may have

$$T_{ji} = \theta_{j\bullet} + u_{ji} + e_{ji} \qquad \text{for study } i \text{ in group } j$$

and for the regression model

$$T_i = \beta_0 + \beta_1 X_{1i} + \ldots + \beta_p X_{pi} + u_i + e_i$$

Those familiar with hierarchical modeling will recognize these as hierarchical linear models.

The components in the mixed-effect model

In the ANOVA-like mixed model, we have

$$T_{ji} = \boxed{\theta_{j\bullet}} + \boxed{u_{ji}} + \boxed{e_{ji}}$$ for study i in group j

Fixed part — Random part(s) — Sampling error

For the regression model with p predictors

$$T_i = \boxed{\beta_0 + \beta_1 X_{1i} + \ldots + \beta_p X_{pi}} + \boxed{u_i} + \boxed{e_i}$$

Fixed part — Random part — Sampling error

Estimating the variance of the random part

We will need to estimate the variation of the u_i error terms in the mixed-effects model. We will denote this variance as $\hat{\sigma}^2_{\theta|X}$.

We don't want to use simple random-effects estimators because they ignore the fact we have useful predictors in our model. They often will be too large.

So we need similar variances, but which estimate only what is left unexplained by the model we choose.

Estimating mixed models

- It is possible to use standard statistical packages (such as SAS or SPSS) to estimate mixed models using a mullti-step process, but it can also be done in one pass with SPSS or SAS macros or by using more specialized software such as HLM or Stata.
- The metareg routine for Stata routinely estimates a mixed model, as does the V-known option of HLM.

Estimating mixed models

- In these slides we will walk through the step-by-step process using output from SAS and SPSS for the teacher expectancy data.
- An example of Stata metareg output for levels of political interest (Block & Becker, 2008) will also be shown, as will analyses for a dataset on the relationship of science teacher knowledge to student science achievement (Becker & Aloe, 2008).

The mixed-effects model: Regression model

We will begin with an ordinary least squares (OLS) regression. This means we will **not** use the weights typically used for meta-analysis.

We use whatever predictors we've found significant in a fixed-effects weighted analysis, and run an unweighted regression with them.

Then we will use the MSE from the OLS regression model to compute a variance that is similar to the simple methods-of-moments between-studies variance.

The mixed-effects model: Regression model and $\hat{\sigma}^2_{\theta|X}$

The method-of-moments random-effects variance is

$$\hat{\sigma}^2_{\theta} = S^2_T - \bar{v}$$

where $\bar{v} = \sum v_i / k$ is the mean fixed-effects variance. For the mixed model we compute the mixed-effects model variance

$$\hat{\sigma}^2_{\theta|X} = MSE_{\text{OLS}} - \bar{v}$$

where MSE_{OLS} is the mean squared error from the OLS (unweighted) regression.

The mixed-effects model: Regression model

If the regression predictor is useful, MSE_{OLS} will be smaller than $S_T^{\ 2}$ (this is true even though we are not using the proper weighting – also note that we do not use the slopes from the OLS regression for anything!).

Therefore, since $S_T^{\ 2} > MSE_{\text{OLS}}$

then $\hat{\sigma}_{\theta}^{2} > \hat{\sigma}_{\theta|X}^{2}$

So we expect the mixed-model variance to be lower than the simple random-effects variance.

The mixed-effects model: Regression model

We add the new variance $\hat{\sigma}^2_{\theta|X}$ to each study's v_i and use weighted least squares (WLS) regression with new mixed-model weights.

The mixed model variances for each study should also be larger than the fixed-effects variances but smaller than the random-effects values.

Because of this, predictors that were significant under the fixed-effects model may no longer be significant under the mixed model.

The mixed-effects model: Regression model

Thus, we have three possible weights:

Fixed $w_i = 1/v_i$

Random $w_i{}^* = 1/[v_i + \hat{\sigma}_\theta^2]$

Mixed $w_i^M = 1/[v_i + \hat{\sigma}_{\theta|X}^2]$

The mixed-model weights $w_i^M = 1/[v_i + \hat{\sigma}_{\theta|X}^2]$ are used instead of w_i or $w_i{}^*$ to estimate the mixed-effects model.

Also there is not one unique set of mixed model weights – since the value of $\hat{\sigma}_{\theta|X}^2$ depends on exactly what Xs are included in your model.

Example: Teacher expectancy data

We have two values of the simple random-effects variance for the data – the more conservative is .08:

OBS	K	Q	P	LL_T_DOT	T_DOT	UL_T_DOT	V_T_DOT	SE_T_DOT
1	19	35.8254	.0074284	-0.011168	0.060343	0.13185	.0013312	0.036485

Fixed-Effects Effect Size Analysis (Exercise 2), tchrexp.dat

OBS	MODVAR	MODSD	QVAR	QSD
1	0.022398	0.14966	0.025920	0.16100

$\hat{\sigma}_{\theta}^{2}$

OBS	_TYPE_	_FREQ_	K	SUMV	SUMT	SUMT2	SVAR	SSD
1	0	19	19	0.92009	3.11	2.8273	0.080366	0.28349

We will run the WLS and OLS regressions on the teacher expectancy data. The WLS regression model on X= weeks shows a slight amount of misfit.

Example: Teacher expectancy data plot

The relationship of "weeks" to the effects is somewhat nonlinear. The OLS "best fit" line that SPSS plots will not be correct unless all studies are equal in size.

This occurs because we need to use a weighted regression and the plotted line is unweighted.

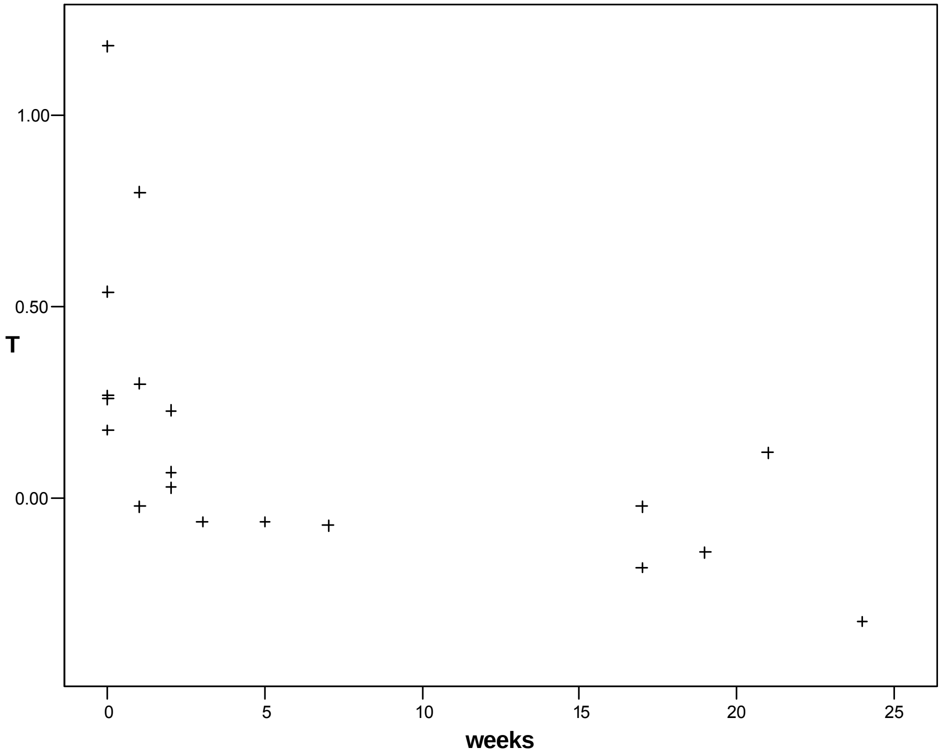

The fixed-effects model: Estimating the line

The fixed-effects equation is obtained by using the SPSS regression menu with “w” =1/v_i clicked into the box for the “WLS weight”.

Coefficients[a,b]

Model		Unstandardized Coefficients		Standardized Coefficients	t	Sig.
		B	Std. Error	Beta		
1	(Constant)	.158	.064		2.479	.024
	weeks	-.013	.006	-.477	-2.239	.039

a. Dependent Variable: T

b. Weighted Least Squares Regression - Weighted by w

The model is

$$d_i = 0.158 - 0.013 \text{ weeks}_i + e_i$$

The fixed-effects model: Plot of the line

Here's the regression line from the weighted analysis:

$$d_i = 0.158 - 0.013 \text{ weeks}_i + e_i$$

Clearly the points with the large effects at weeks = 0 are not well explained by this line.

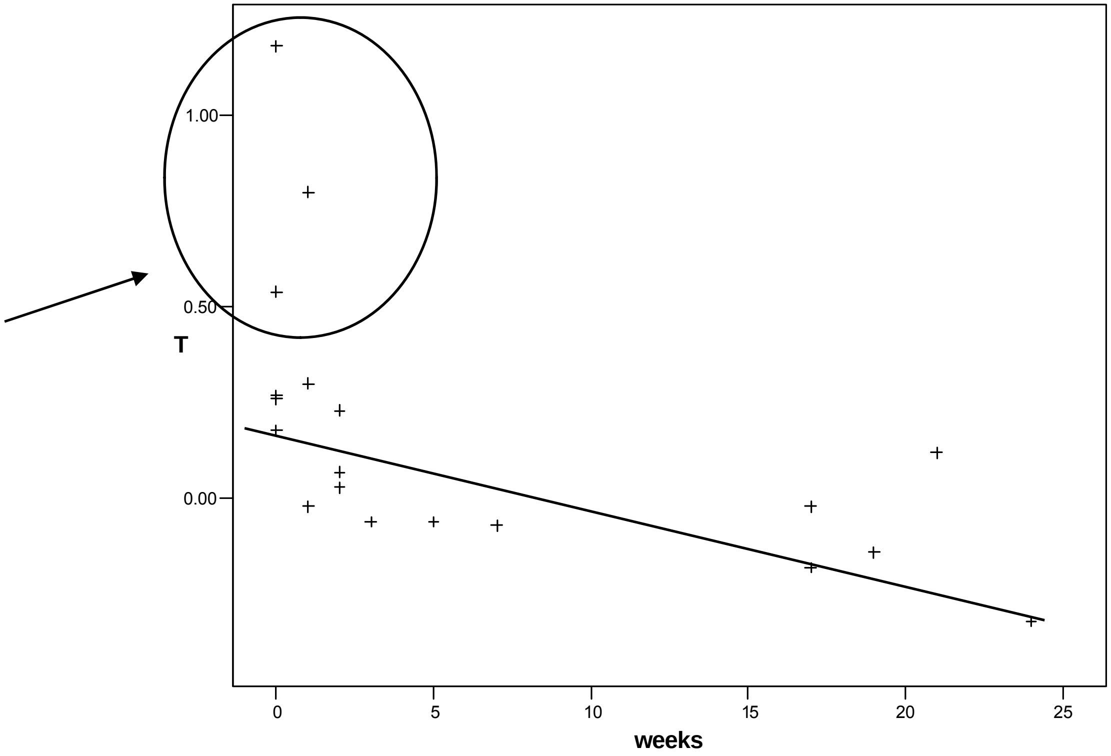

The fixed-effects model: Testing significance of X

Is the predictor "weeks" significant under the fixed effects model? If not, we do not need to proceed to a mixed model for this predictor. Here, it is significant:

$SS_{\text{Regression}} = Q_{\text{Model}}$ = 8.61 is significant with p = 1 df.

ANOVA[b,c]

Model		Sum of Squares	df	Mean Square	F	Sig.
1	Regression	8.161	1	8.161	5.015	.039[a]
	Residual	27.664	17	1.627		
	Total	35.825	18			

a. Predictors: (Constant), weeks

b. Dependent Variable: T

c. Weighted Least Squares Regression - Weighted by w

The fixed-effects model: Testing specification

We'd like the Q_{Error} to be small and not significant.

$SS_{\text{Residual}} = Q_{\text{Error}}$ = 27.66 is just barely significant with p = 17 df. We may want to fit the mixed model because we know some studies with weeks = 0 do not fit well.

ANOVA[b,c]

Model		Sum of Squares	df	Mean Square	F	Sig.
1	Regression	8.161	1	8.161	5.015	.039[a]
	Residual	27.664	17	1.627		
	Total	35.825	18			

a. Predictors: (Constant), weeks

b. Dependent Variable: T

c. Weighted Least Squares Regression - Weighted by w

The mixed-effects model: Variance estimation

We get the mixed-model variance. MSE_{OLS} is .094. Also $\bar{v}$ is the same at .0484.

ANOVA[b]

Model		Sum of Squares	df	Mean Square	F	Sig.
1	Regression	.728	1	.728	7.786	.013[a]
	Residual	1.590	17	.094		
	Total	2.318	18			

a. Predictors: (Constant), weeks

b. Dependent Variable: T

MSE_{OLS}

$$\hat{\sigma}^2_{\theta|X} = MSE_{OLS} - \bar{v} = .094 - .0484 = .0456$$

Recall the random-effects variance was larger, at .08

The mixed-effects model: Significance testing

We add $\hat{\sigma}^2_{\theta|X}$ = .0456 to each study's variance and run a weighted mixed-model regression. The model is still significant but Q_{Error} has dropped since we are accounting for $\hat{\sigma}^2_{\theta|X}$.

Q_{Model} = 5.56 with df = 1, Q_{Error} is NS.

ANOVA[b,c]

Model		Sum of Squares	df	Mean Square	F	Sig.
1	Regression	5.564	1	5.564	7.031	.017[a]
	Residual	13.452	17	.791		
	Total	19.015	18			

a. Predictors: (Constant), weeks

b. Dependent Variable: T

c. Weighted Least Squares Regression - Weighted by wm

The mixed-effects model: Estimating the line

Recall that the FE model was

$$T_i = 0.158 - 0.013\ weeks_i + e_i$$

The mixed model is quite similar

$$T_i = 0.238 - 0.019\ weeks_i + u_i + e_i$$

Coefficients[a,b]

Model		Unstandardized Coefficients B	Unstandardized Coefficients Std. Error	Standardized Coefficients Beta	t	Sig.
1	(Constant)	.238	.078		3.057	.007
	weeks	-.019	.007	-.541	-2.652	.017

a. Dependent Variable: T

b. Weighted Least Squares Regression - Weighted by wm

Output from a mixed model in Stata

This model from Block and Becker (2008) predicts the proportion of people expressing any level of political interest (called any_x) from the year a survey was conducted.

There appear to be increasing levels of interest.

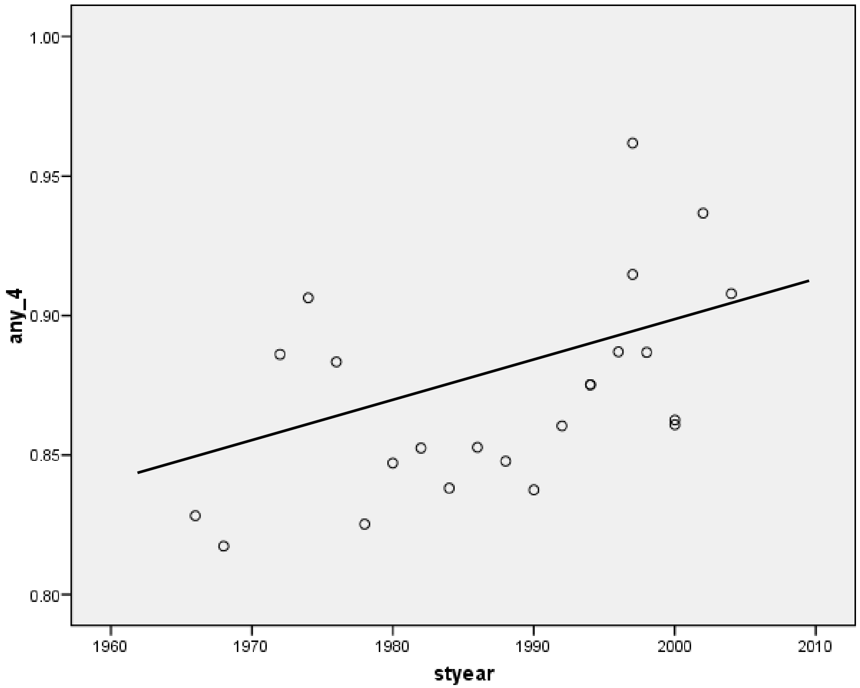

Output from a mixed model in Stata

Stata produces this output from its routine called meta-reg

$\hat{\sigma}^2_{\theta|X}$

```
. metareg   any_4 yr40, wsse(seany4) bsest(mm)
Meta-regression                                   Number of studies =        25
Fit of model without heterogeneity (tau2=0):      Q (23 df)         =   301.155
                                                  Prob > Q          =     0.000
Proportion of variation due to heterogeneity      I-squared         =     0.924
Moment-based estimate of between-study variance:  tau2              =    0.0010
------------------------------------------------------------------------------
       any_4 |      Coef.   Std. Err.       t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
        yr40 |   .0016399   .0006019     2.72   0.012      .0003948     .002885
       _cons |   .7976401   .0296878    26.87   0.000      .7362262     .859054
------------------------------------------------------------------------------
```

The mixed model is

$$p_i = 0.798 + 0.0016\, yr40_i + u_i + e_i$$

with $yr40 = \text{year - } 1940$

The mixed-effects model: ANOVA-like categorical model

The ANOVA model is more complicated than the regression model because there are several ways we can have a mixed model.

For instance, it may be that only one category or group has studies that are heterogeneous.

Another possibility is that all groups are heterogeneous, but to different degrees. A final case is that all are heterogeneous to the same degree.

The mixed-effects model: *RSSVAR* in the ANOVA-like model

The simplest analysis may be to compute the RE variance <u>within</u> each category or group (e.g., using the $\hat{\sigma}_{\theta}^{2}$ formulas– let's say we get $\hat{\sigma}_{\theta\ j}^{2}$ for group j) and add a different value to each study's v_{ij} depending on which group the study is in.

This approach allows all groups to vary (or not) and assigns different variances for the u_i values in each subset.

The mixed-effects model: $\hat{\sigma}^2_{\theta|X}$ in the ANOVA-like model

A weakness with estimating the within group variance for each separate group of studies is that some groups may be very small and $\hat{\sigma}^2_{\theta}$ and $\hat{\sigma}^2_{\theta|X}$ are poorly estimated with few studies.

If we are willing to assume that all groups are heterogeneous to the same degree, we can use a parallel approach to the regression mixed model.

The mixed-effects model: $\hat{\sigma}^2_{\theta|X}$ in the ANOVA-like model

We run an unweighted ANOVA to get a MSW (analogous to the MSE_{OLS}) and compute the $\hat{\sigma}^2_{\theta|X}$ in the way described above.

This approach has the weakness that it may make some studies have larger (or smaller) variances than are really appropriate, if the groups are not equally heterogeneous.

Example: Teacher expectancy data

We can use "weeks" as a categorical variable if we categorize studies as having 0, 1, 2 or 3+ weeks of exposure. We call it "weekcat".

Although the ANOVA model does not misfit for the TE data, let's consider how the analyses would work. If we choose to assume equal heterogeneity, we'd run an unweighted ANOVA and get the *MSW*, which is .077.

ANOVA

T

	Sum of Squares	df	Mean Square	F	Sig.
Between Groups	1.164	3	.388	5.039	.013
Within Groups	1.155	15	.077		
Total	2.318	18			

Example: Teacher expectancy data

We compute the common within groups variance as we did above:

$$\hat{\sigma}^2_{\theta|X} = .077 - .0484 = .0286$$

The value .0286 would be added to every study's variance and we would run a weighted ANOVA using the new weights = 1/[v_i + .0286].

Example: Teacher expectancy data

Alternately we may wish to examine each group of studies, to see if some groups are more heterogeneous than others.

From our initial plot of the weeks – effect size relationship we may recall that effects with no weeks of exposure (weeks = 0) seemed more variable.

If we find different heterogeneity in each group, we'd add different variances $\hat{\sigma}^2_{\theta|X\ j}$ to each study's v_i in group j.

Example: Teacher expectancy data

Here are values of S_T^2 and $\bar{v}$ for each *weekcat* group. $\hat{\sigma}^2_{\theta\ j}$ values are shown at the right – two values are negative and would be set to 0. For studies with *weekcat* = 2 or 3 we can use $w_i = 1/v_i$, but we'd use $w_i = 1/[v_i + .116]$ for *weekcat*= 1, etc. If some groups are more variable this approach may be best.

$\bar{v}$ S_T^2

Descriptive Statistics

weekcat		N	Mean	Std. Deviation	Variance
0	T	5	.4860	.41107	.169
	V	5	.0886	.05035	.003
	Valid N (listwise)	5			
1	T	3	.3600	.41328	.171
	V	3	.0553	.03279	.001
	Valid N (listwise)	3			
2	T	3	.1100	.10583	.011
	V	3	.0362	.04163	.002
	Valid N (listwise)	3			
3	T	8	-.0913	.12800	.016
	V	8	.0253	.01205	.000
	Valid N (listwise)	8			

Wkcat	$\hat{\sigma}^2_{\theta\ j}$
0	**.169 - .089 = .080**
1	**.171 - .055 = .116**
2	**.011 - .036 = -.025 = 0**
3	**.016 - .025 = -.009 = 0**

C2 Training Materials Team

Thanks are due to the following institutions and individuals

- Funder: Norwegian Knowledge Centre for the Health Sciences
- Materials contributors: Betsy Becker, Harris Cooper, Larry Hedges, Mark Lipsey, Therese Pigott, Hannah Rothstein, Will Shadish, Jeff Valentine, David Wilson
- Materials developers: Ariel Aloe, Betsy Becker, Sunny Kim, Jeff Valentine, Meng-Jia Wu, Ying Zhang
- Training co-coordinators: Betsy Becker and Therese Pigott