

1	MODELS FOR CONTINUOUS OUTCOMES.....	7
1.1	MODELS BASED ON A SUBSET OF THE NESARC DATA.....	7
1.1.1	<i>The data.....</i>	7
1.1.1.1	Importing the data and defining variable types	8
1.1.1.2	Exploring the data	12
	Univariate graphs	13
1.1.2	<i>2-level random intercept model with 2 predictors.....</i>	15
1.1.2.1	The model	15
1.1.2.2	Setting up the analysis.....	17
1.1.2.3	Discussion of results.....	20
	Program information and syntax	20
	Data summary	21
	Descriptive statistics and starting values.....	22
	Fixed effects results	23
	Fit statistics	24
	Random effects results	25
1.1.2.4	Interpreting the results.....	26
	Model-based graphs	26
	ICCs and % variance explained	28
1.1.3	<i>A 2-level random intercept model with 4 predictors</i>	29
1.1.3.1	The model	29
1.1.3.2	Setting up the analysis.....	31
1.1.3.3	Discussion of results.....	32
	Fixed effects results	32
	Fit statistics	32
	Random effects results	33
1.1.3.4	Interpreting the results.....	33
	Estimated outcomes for different groups	33
	Fit statistics and % variation explained.....	35
1.2	MODELS BASED ON THE REISBY DATA	37
1.2.1	<i>The data.....</i>	37
1.2.1.1	Exploring the data	38
	Graphing the observed data.....	38
1.2.2	<i>A 2-level random intercept-and-slope model.....</i>	42
1.2.2.1	The model	42
1.2.2.2	Setting up the analysis.....	43
1.2.2.3	Discussion of results.....	45
	Descriptive statistics	45
1.2.2.4	Interpreting the results.....	46
	Fixed effects results	47
	Random effects results	47
	Fit statistics and ICC	48
1.2.3	<i>A 2-level random intercept-and-slope model with centered predictor</i>	51
1.2.3.1	Preparing the data.....	51
1.2.3.2	The model	53
1.2.3.3	Setting up the analysis.....	54
1.2.3.4	Discussion of results.....	55

1.2.3.5	Interpreting the results.....	56
	Comparison of models	56
1.2.4	<i>A random intercept-and-slope with a covariate and an interaction term.....</i>	57
1.2.4.1	The model	58
1.2.4.2	Setting up the analysis.....	59
1.2.4.3	Interpreting the results.....	61
	Fixed effects results	61
	Model comparison	62
1.2.5	<i>A random intercept-and-slope quadratic model.....</i>	64
1.2.5.1	The model	64
1.2.5.2	Preparing the data.....	64
1.2.5.3	Setting up the analysis.....	65
1.2.5.4	Interpreting the results.....	66
	Fixed effects results	66
1.2.5.5	Residuals	67
	Level 2 Bayes results	67
	Model-based graphs	70
1.2.6	<i>A 2-level random intercept-and-slope model with autocorrelated errors</i>	73
1.2.6.1	The non-stationary AR(1) model.....	74
	The analysis – step 1: starting values from a non-AR model	75
	The analysis – step 2: non-stationary AR model.....	77
1.2.6.2	Interpreting the output.....	78
	The starting values	78
	The maximum marginal likelihood (MML) estimates	79
	Correlation of the MML estimates	81
	Level-2 Bayes results	82
	Model comparison	83
1.3	MODELS BASED ON THE TVSFP DATA	84
1.3.1	<i>The data.....</i>	84
1.3.1.1	Exploring the data	85
	Univariate graphs	86
	Bivariate graphs	88
1.3.2	<i>A 2-level random intercept model using classroom as level-2 ID.....</i>	90
1.3.2.1	The model	90
1.3.2.2	Setting up the analysis.....	91
1.3.2.3	Discussion of results.....	93
	Program information and syntax	93
	Data summary	93
	Descriptive statistics and starting values.....	94
	Fixed effects estimates	96
	Random effect estimates	98
	End of the output.....	99
1.3.2.4	Interpreting the results.....	100
	Estimated outcomes for different groups	100
	Fit statistics and % variation explained.....	100
1.3.3	<i>2-level random intercept model by using school as level-2 ID.....</i>	100
1.3.3.1	The model	100
1.3.3.2	Setting up the analysis.....	101

1.3.3.3	Discussion of results.....	102
	Data summary	102
	Fixed effects estimates and descriptive statistics	102
	Random effect estimates and covariance/correlation matrices	103
1.3.4	<i>A 3-level random intercept model using class and school as IDs</i>	<i>104</i>
1.3.4.1	The model	104
1.3.4.2	Setting up the analysis.....	105
1.3.4.3	Discussion of results.....	106
	Data summary	106
	Fixed effects estimates	107
	Random effect estimates	108
1.3.4.4	Further discussion of the results	109
	Fit statistics and % variation explained	109
	Comparison of models	109
1.3.5	<i>A 3-level random intercept model including pre-THKS score.....</i>	<i>110</i>
1.3.5.1	The model	110
1.3.5.2	Setting up the analysis.....	111
1.3.5.3	Discussion of results.....	112
	Fixed effects estimates and descriptive statistics	112
	Random effect estimates	113
1.3.5.4	Interpreting the results.....	114
	Estimated outcomes for different groups	114
	ICCs and R square	114
	Model fit statistics and comparison of models	118
1.3.5.5	Residuals	119
	Level-2 Bayes results	119
	Level-3 Bayes results	121
1.4	NEW 3 LEVEL CONTINUOUS EXAMPLE USING A SUBSET OF SCHOENWALD DATA	124
1.4.1	<i>The data.....</i>	<i>124</i>
1.4.2	<i>Exploring the data</i>	<i>125</i>
1.4.2.1	Univariate graphs	125
	Histograms	125
	Transformation of variables	127
1.4.2.2	Bivariate graphs	131
	Exploratory graphs	131
	Transforming a variable	134
1.4.2.3	Exploratory graphs	135
	Box-and-whisker plots	137
	Bivariate bar charts	141
1.4.3	<i>Fitting a growth curve model to the data</i>	<i>142</i>
1.4.3.1	The model	142
1.4.3.2	Setting up the analysis.....	143
1.4.3.3	Discussion of results.....	145
	Fixed effects results	147
	Random effects results	147
1.4.3.4	Interpreting the results.....	148
	Estimated outcomes for different groups	148
	Model-based graphs	149

Fit statistics and ICC	152
<i>1.4.4 Fitting a random intercept model with 3 predictors and interaction term to the data</i>	<i>152</i>
1.4.4.1 The model	152
1.4.4.2 Setting up the analysis	153
1.4.4.3 Discussion of results.....	154
Fixed effects results	154
Random effects results	155
1.4.4.4 Interpretation of the results.....	156
Estimated outcomes for different groups	156
Equation modeling	158
Confidence intervals for random coefficients	161
Fit statistics	163
1.4.4.5 Residuals	164
Residual plots.....	164
<i>1.4.5 Fitting a random intercepts and slopes model.....</i>	<i>167</i>
1.4.5.1 The model	167
1.4.5.2 Setting up the analysis	168
1.4.5.3 Discussion of results.....	169
Fixed effects results	169
Random effects results	169
1.4.5.4 Interpreting the results.....	170
Fit statistics and ICC	170
<i>1.4.6 References</i>	<i>174</i>

FIGURE XXX.1: HISTOGRAM OF THE VARIABLE AGE_DEP	14
FIGURE XXX.2: PIE CHART OF THE ETHNICITY VARIABLE BLACK ERROR! BOOKMARK NOT DEFINED.	
FIGURE XXX.3: PLOT OF AGE_DEP VERSUS M_S_DEP FOR 2 GROUPS	27
FIGURE XXX.4: PLOT OF AGE_DEP VERSUS M_S_DEP FOR 2 GROUPS	28
TABLE XXX.1: EXPECTED AGE_DEP FOR VARIOUS GROUPS OF PATIENTS.....	34
TABLE XXX.2: COMPARISON OF RANDOM INTERCEPT MODELS FOR NESARC DATA.	35
FIGURE XXX.5: REACTION TRAJECTORIES OVER TIME FOR 66 PATIENTS.....	40
FIGURE XXX.6: REACTION TRAJECTORIES OVER TIME FOR PATIENTS WITH ENDOG=0.....	41
FIGURE XXX.7: REACTION TRAJECTORIES OVER TIME FOR PATIENTS WITH ENDOG=1.....	41
FIGURE XXX.8: SCORE TRENDS FOR INDIVIDUAL PATIENTS.....	42
TABLE XXX.3: ESTIMATES AND STANDARD ERRORS FOR TWO MODELS.....	56
FIGURE XXX.9: CHANGES IN COVARIANCE OVER TIME.....	57
TABLE XXX.4: PREDICTED HDRS VALUES FOR SELECTED PATIENTS	69
FIGURE XXX.10: PREDICTED HDRS FOR SELECTED PATIENTS.....	70
FIGURE XXX.11: PLOT OF LEVEL-1 RESIDUALS VS. PREDICTED VALUES	72
TABLE XXX.5: COMPARISON OF MODELS WITH AND WITHOUT AR(1) TERM.....	83
FIGURE XXX.12: BAR CHART OF POSTTHKS SCORES	87
FIGURE XXX.13: BOX-AND-WHISKER PLOTS OF POSTTHKS SCORES FOR DIFFERENT CC VALUES	89
TABLE XXX.6: COMPARISON OF OLS AND MIXED MODEL RESULTS	110
TABLE XXX.7: R^2 VALUES FOR A SET OF NESTED MODELS.....	117
TABLE XXX.8: COMPARISON OF OLS AND MIXED MODEL RESULTS	118
TABLE XXX.8 (CONTINUED).....	119
FIGURE XXX.14: 95% CONFIDENCE INTERVALS FOR LEVEL-2 BAYES ESTIMATES	123
FIGURE XXX.15: HISTOGRAM OF THE VARIABLE CBCTOT.....	127
FIGURE XXX.16: HISTOGRAM OF THE VARIABLE SQR_CBC	131

FIGURE XXX.17: RELATIONSHIP BETWEEN SQR_CBC AND VISIT FOR SELECTED PATIENTS.....	134
FIGURE XXX.18: RELATIONSHIP BETWEEN SQR_CBC AND SQ_VISIT FOR SELECTED PATIENTS.....	137
FIGURE XXX.19: BOX-AND-WHISKER PLOT OF SQR_CBC VS. VISIT	139
FIGURE XXX.20: BOX-AND-WHISKER PLOT OF CBCTOT VS. VISIT	139
FIGURE XXX.21: BOX-AND-WHISKER PLOT OF SQR_CBC VS. GENF	140
FIGURE XXX.22: BOX-AND-WHISKER PLOT OF CBCTOT VS. GENF	141
FIGURE XXX.23: BIVARIATE CHART OF VISIT VS GENF	142
FIGURE XXX.24: PLOT OF SQR_CBC VS. VISIT.....	151
FIGURE XXX.25: PLOT OF SQR_CBC VS. SQ_VISIT	151
TABLE XXX.9: EXPECTED SQUARE ROOT OF CBC SCORES.....	157
TABLE XXX.10: EXPECTED CBC SCORES IN ORIGINAL SCALE	158
FIGURE XXX.26: PLOT OF SQR_CBC VS. VISIT FOR GENDER GROUPS.....	159
FIGURE XXX.27: MODIFIED PLOT OF SQR_CBC VS. VISIT FOR GENDER GROUPS	160
FIGURE XXX.28: 95% CONFIDENCE INTERVALS FOR LEVEL-3 UNITS	163
FIGURE XXX.29: LEVEL-1 RESIDUAL PLOT BY GENDER GROUP	165
FIGURE XXX.30: MODIFIED LEVEL-1 RESIDUAL PLOT BY GENDER GROUP.....	167

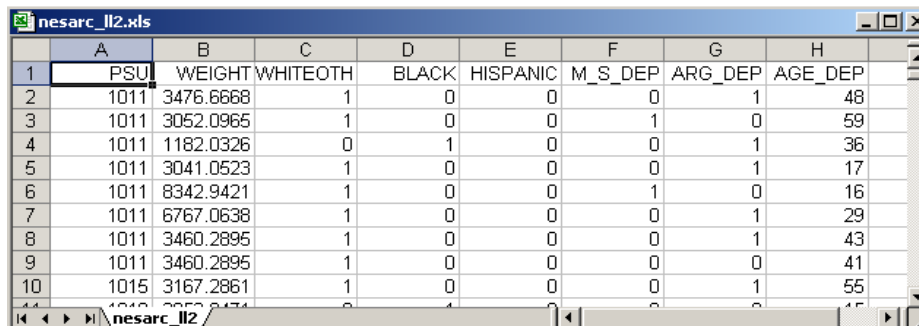
1 Models for continuous outcomes

1.1 Models based on a subset of the NESARC data

1.1.1 The data

The data set is from the National Epidemiologic Survey on Alcohol and Related Conditions (NESARC), a longitudinal survey with its first wave fielded in 2001-2002. The NESARC is a representative sample of the United States population, and 43,093 Americans participated in the first wave of the survey. The NESARC survey was conducted and sponsored by the National Institute on Alcohol Abuse and Alcoholism (NIAAA). Detailed information is available at <http://niaaa.census.gov/index.html>.

Section 4 of the NESARC data documentation describes data regarding major depression, family history of major depression and dysthymia. Together with the demographic information in Section 1, we produced the **nesarc_112.xls** data set as shown below. There are 2,339 dysthymia respondents in the survey. After listwise deletion, the sample size is 1,698.



	A	B	C	D	E	F	G	H
1	PSU	WEIGHT	WHITEOTH	BLACK	HISPANIC	M_S_DEP	ARG_DEP	AGE_DEP
2	1011	3476.6668	1	0	0	0	1	48
3	1011	3052.0965	1	0	0	1	0	59
4	1011	1182.0326	0	1	0	0	1	36
5	1011	3041.0523	1	0	0	0	1	17
6	1011	8342.9421	1	0	0	1	0	16
7	1011	6767.0638	1	0	0	0	1	29
8	1011	3460.2895	1	0	0	0	1	43
9	1011	3460.2895	1	0	0	0	0	41
10	1015	3167.2861	1	0	0	0	1	55

The variables of interest are:

- PSU is the Census 2000/2001 Supplementary Survey (C2SS) primary sampling unit (PSU).
- WEIGHT is the final weight, calculated as the product of the NESARC base weight and other individual weighting factors.
- WHITEOTH represents the white and other ethnicities, excluding African American and Hispanic. It is recoded from items S1Q1C, S1Q1D2, S1Q1D3 and S1Q1D5 in the NESARC source code (1 for white and other, 0 for African American and Hispanic).
- BLACK represents African Americans. It is recoded from items S1Q1C and S1Q1D3 in the NESARC source code (1 for African American, 0 for others).
- HISPANIC is an indicator for Hispanic. It is recoded from items S1Q1C, S1Q1D3 and S1Q1D5 (1 for Hispanic, 0 for others).
- M_S_DEP is recoded from item S4BQ10C. It is the response to the statement “Any of natural mother’s full sisters ever depressed,” with 1 for “Yes,” and 0 for “No.”
- ARG_DEP is recoded from item S4CQ43. It represents the response to the statement “Had arguments/friction with family, friends, people at work, or anyone else,” with 1 for “Yes,” 0 for “No.”
- AGE_DEP is a renamed version of item S4CQ7AR. It represents the age at onset of first episode of dysthymia.

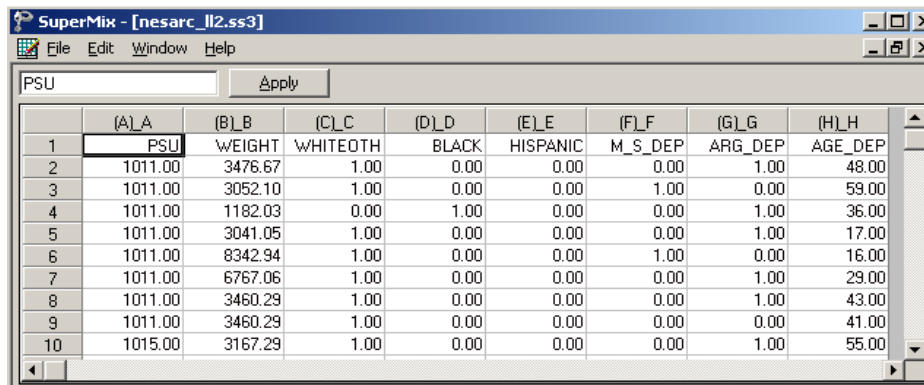
Inspection of the data shows that only about 2% of 43,093 respondents are of Asian and Pacific origin. Due to the skewness of the distribution of ethnicity, we recoded the variables representing ethnic origin. The resulting variable WHITEOTH represents this recoding of respondents as being either white or from other ethnic groups (blacks and Hispanics excluded).

1.1.1.1 Importing the data and defining variable types

The data set shown previously is available in the form of a spreadsheet file, named **nesarc_II2.xls**. This file contains a subset of the original NESARC data, *i.e.* data for the 1,698 respondents who reported some form of depression and for whom complete information on variables of interest was available.

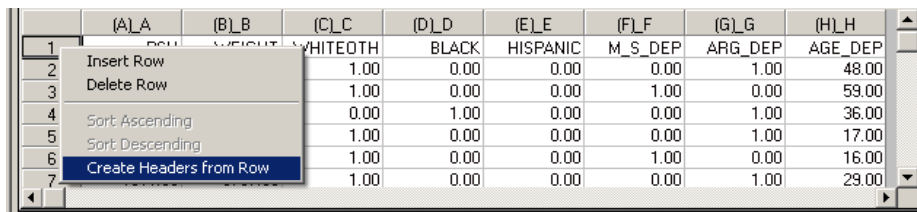
The first step is to create the SuperMix spreadsheet file (*.ss3) from the Excel file:

- Use the **Import Data File** option on the **File** menu to load the **Open** dialog box.
- Browse for the file **nesarc_II2.xls** in the **examples** folder of the SuperMix installation folder.
- Select the file and click on the **Open** button to open the following SuperMix spreadsheet window **nesarc_II2.ss3**.



	(A)_A	(B)_B	(C)_C	(D)_D	(E)_E	(F)_F	(G)_G	(H)_H
1	PSU	WEIGHT	WHITEOTH	BLACK	HISPANIC	M_S_DEP	ARG_DEP	AGE_DEP
2	1011.00	3476.67	1.00	0.00	0.00	0.00	1.00	48.00
3	1011.00	3052.10	1.00	0.00	0.00	1.00	0.00	59.00
4	1011.00	1182.03	0.00	1.00	0.00	0.00	1.00	36.00
5	1011.00	3041.05	1.00	0.00	0.00	0.00	1.00	17.00
6	1011.00	8342.94	1.00	0.00	0.00	1.00	0.00	16.00
7	1011.00	6767.06	1.00	0.00	0.00	0.00	1.00	29.00
8	1011.00	3460.29	1.00	0.00	0.00	0.00	1.00	43.00
9	1011.00	3460.29	1.00	0.00	0.00	0.00	0.00	41.00
10	1015.00	3167.29	1.00	0.00	0.00	0.00	1.00	55.00

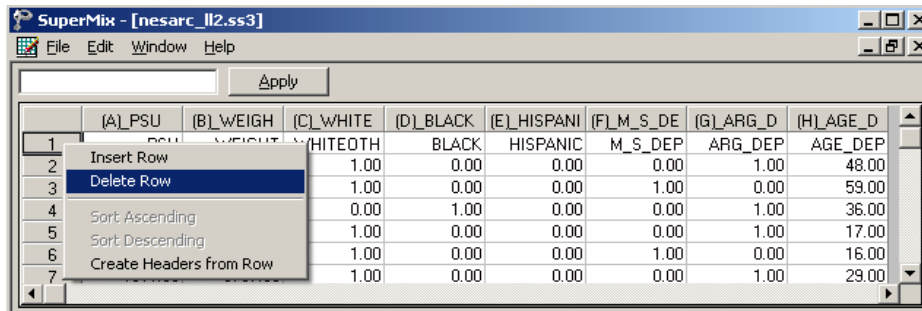
Note that row 1 of the spreadsheet contains the variable labels. To rename the column headers to reflect the variables contained in each column, first left-click on the row 1 tab to select the complete row and then right-click and select the **Create Headers from Row** option from the pop-up menu as shown below.



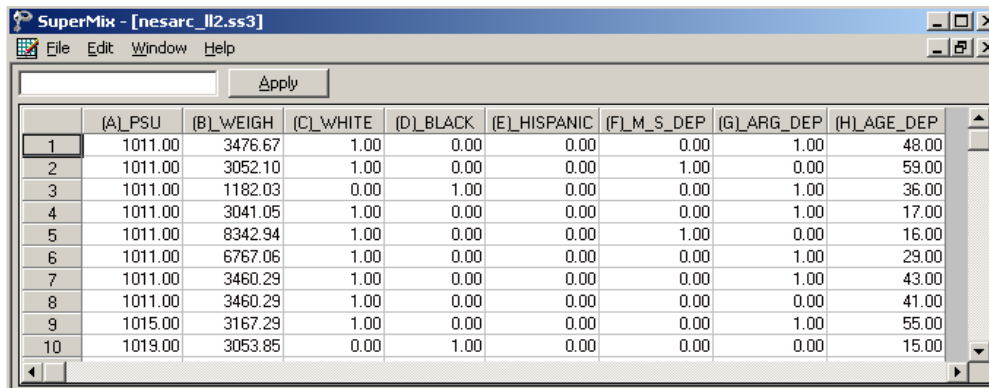
	(A)_A	(B)_B	(C)_C	(D)_D	(E)_E	(F)_F	(G)_G	(H)_H
1	PSU	WEIGHT	WHITEOTH	BLACK	HISPANIC	M_S_DEP	ARG_DEP	AGE_DEP
2	1011.00	3476.67	1.00	0.00	0.00	0.00	1.00	48.00
3	1011.00	3052.10	1.00	0.00	0.00	1.00	0.00	59.00
4	1011.00	1182.03	0.00	1.00	0.00	0.00	1.00	36.00
5	1011.00	3041.05	1.00	0.00	0.00	0.00	1.00	17.00
6	1011.00	8342.94	1.00	0.00	0.00	1.00	0.00	16.00
7	1011.00	6767.06	1.00	0.00	0.00	0.00	1.00	29.00

The spreadsheet headers now correspond to the names shown in the first row of the spreadsheet. This row is no longer needed, and is deleted by right-clicking anywhere

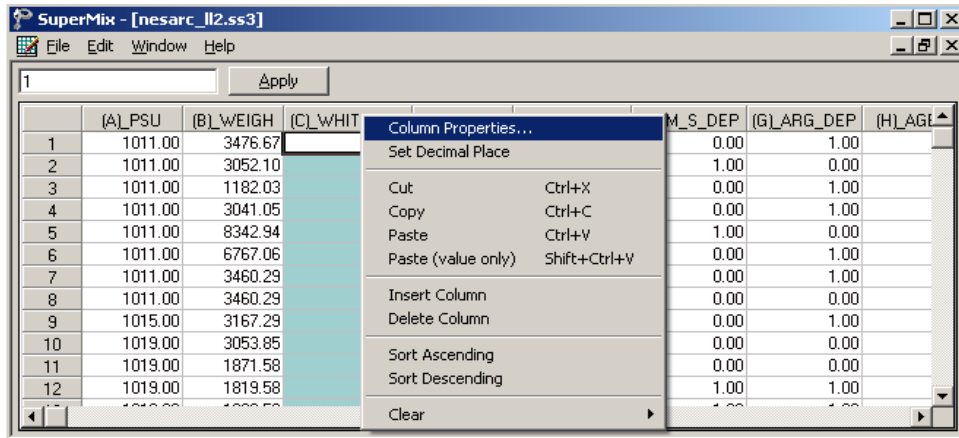
on row 1 and then selecting the **Delete Row** option from the pop-up menu as illustrated below



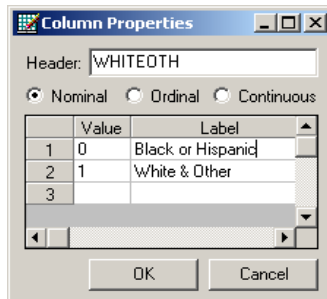
to produce the following SuperMix spreadsheet window.



Next, we define the variable types. Highlight WHITEOTH by clicking on the variable name, and then right click to open the following pop-up menu. Select the **Column Properties** option

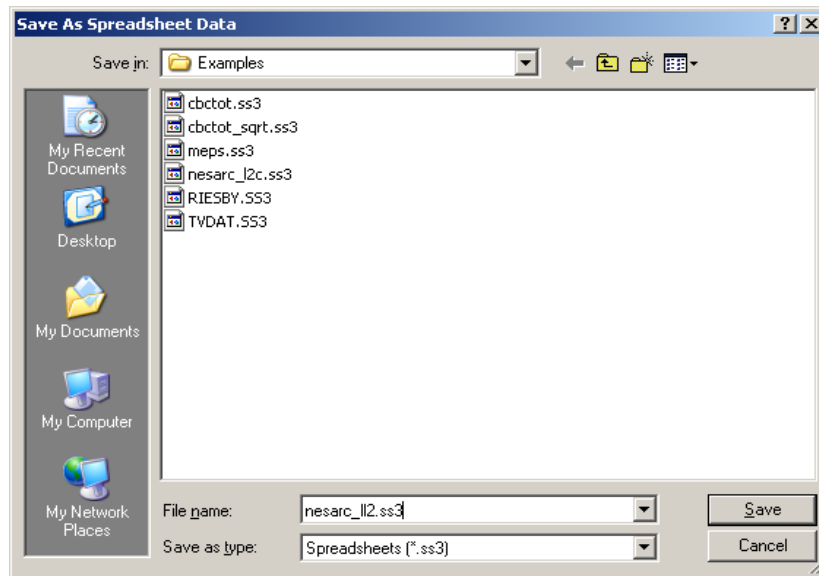


to open the **Column Properties** dialog box. Checking the **Nominal** radio button enables the user to define the labels. Input correct labels for the different categories as shown below.



Similarly define BLACK, HISPANIC, M_S_DEP and ARG_DEP as nominal variables and define AGE_DEP as continuous.

To save the **nesarc_ii2.ss3** spreadsheet, select the **Save As** option from the **File** menu to load the **Save As Spreadsheet Data** dialog box, and then enter the desired file name in the **File name** string field as shown below. Click on the **Save** button when done.



1.1.1.2 Exploring the data

Graphics are often a useful data-exploring technique through which the researcher may familiarize her- or himself with the data. Relationships and trends may be conveyed in an informal and simplified visual form via graphical displays. SuperMix offers both data-based and model-based graphs. Data-based graphing options are accessed via the **File, Data-based Graphs** option once a SuperMix data file (.ss3) is opened, and include **Exploratory, Univariate, Bivariate** and **Multivariate** graphs as shown on the pop-up menu below. Model-based graphs are available after the analysis has been performed, and will be discussed later in this section.

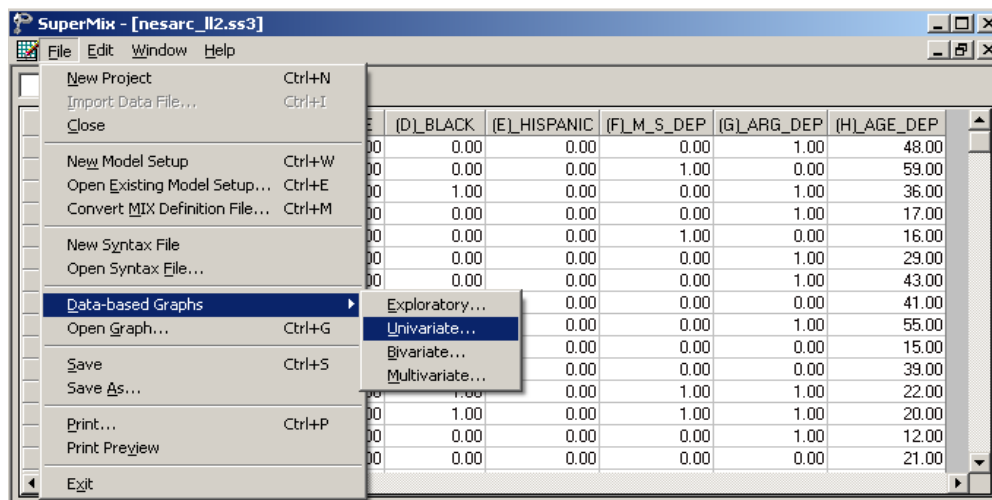
In the case of data-based graphs, we distinguish between three categories: univariate, bivariate, and multivariate graphs. Univariate graphs are particularly useful to obtain an overview of the characteristics of a variable. However, they do not necessarily offer the tools needed to explore longitudinal data as completely as one would wish. For that purpose, bivariate and multivariate data-based graphs are more appropriate.

Univariate graphs

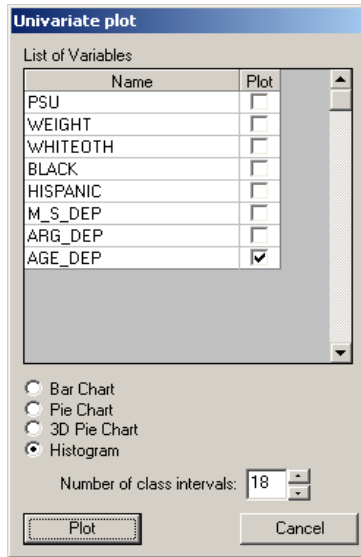
The pop-up menu below shows the data-based graphing options currently available in SuperMix. As a first step, we take a look at the distribution of age at onset of first depression episode (AGE_DEP), which is the potential dependent variable in this study.

Histograms

A histogram represents the frequency of cases per unit interval. It gives a good picture of the distribution of a variable. To create a histogram for AGE_DEP, select the **Univariate** option from the **Data-based Graphs** menu as shown below.



The **Univariate plot** dialog box appears. Select the variable AGE_DEP and indicate that a **Histogram** is to be graphed. The desired number of intervals shown on the histogram is controlled by the **Number of class intervals** field. It is specified as 18 in this case. Click the **Plot** button to display the histogram.



The histogram, as seen below, shows that the distribution of AGE_DEP is nearly symmetrical, and should satisfy the normality assumptions implicit in a multilevel model.

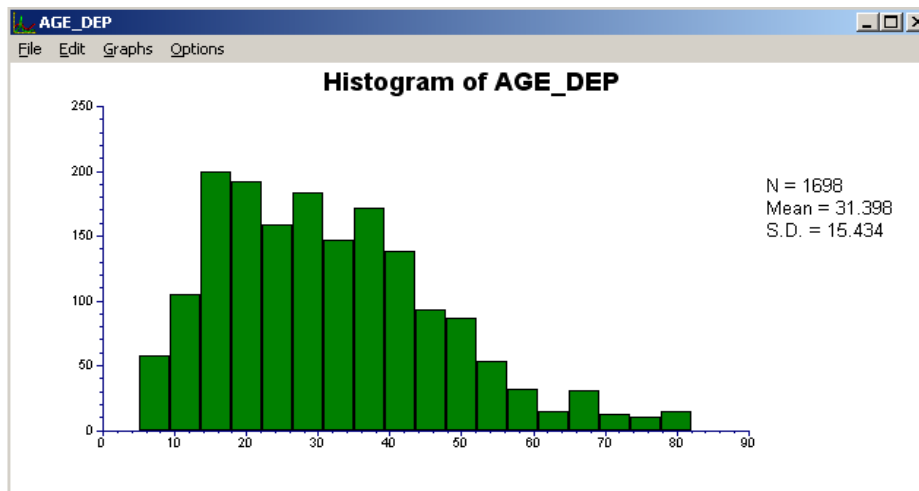


Figure XXX.1: Histogram of the variable AGE_DEP

1.1.2 2-level random intercept model with 2 predictors

1.1.2.1 The model

A two-level multilevel model consists of two submodels, one at each level of the hierarchy. A general two-level model for a continuous response variable y depending on a set of p predictors x_1, x_2, \dots, x_r can be written in the form

$$y_{ij} = \mathbf{x}'_{ij}\boldsymbol{\beta} + \mathbf{z}'_{ij}\mathbf{v}_i + \varepsilon_{ij}$$

where $i = 1, 2, \dots, N$ denotes the level-2 units, and $j = 1, 2, \dots, n_i$ the level-1 units. In this context, y_{ij} represents the response of individual j , nested within level-2 unit i . The model shown here consists of a fixed and a random part. The fixed part of the model is represented by the vector product $\mathbf{x}'_{ij}\boldsymbol{\beta}$, where \mathbf{x}'_{ij} is a typical row of the design matrix of the fixed part of the model with, as elements, a subset of the p predictors. The vector $\boldsymbol{\beta}$ contains the fixed, but unknown parameters to be estimated. $\mathbf{z}'_{ij}\mathbf{v}_i$ and ε_{ij} denote the random part of the model at levels 2 and 1 respectively. For example, \mathbf{z}'_{ij} represents a typical row of the design matrix of the random part at level 2, and \mathbf{v}_i the vector of random level-2 effects to be estimated. It is assumed that $v_{01}, v_{02}, \dots, v_{0N}$ are independently and identically distributed (i.i.d.) with mean vector $\mathbf{0}$ and covariance matrix $\Phi_{(v)}$. Similarly, the ε_{ij} are assumed i.i.d., with mean 0 and variance σ^2 .

The first model fitted to the NESARC data explores the relationship between AGE_DEP and the maternal-side depression and argument involvement, as represented by the variables M_S_DEP and ARG_DEP. The level-1 model is at a patient level, while the level-2 model is at a PSU level. The model can be expressed as

Level-1 model:

$$\text{AGE_DEP}_{ij} = b_{0i} + b_{1i} \times (\text{MS_DEP})_{ij} + b_{2i} \times (\text{ARG_DEP})_{ij} + e_{ij}$$

Level-2 model:

$$b_{0i} = \beta_0 + v_{0i}$$

$$b_{1i} = \beta_1$$

$$b_{2i} = \beta_2$$

where

$$\varepsilon_i \sim N(0, \sigma^2 \mathbf{I}_i)$$

$$\mathbf{v}_i \sim N(0, \boldsymbol{\Sigma}_i)$$

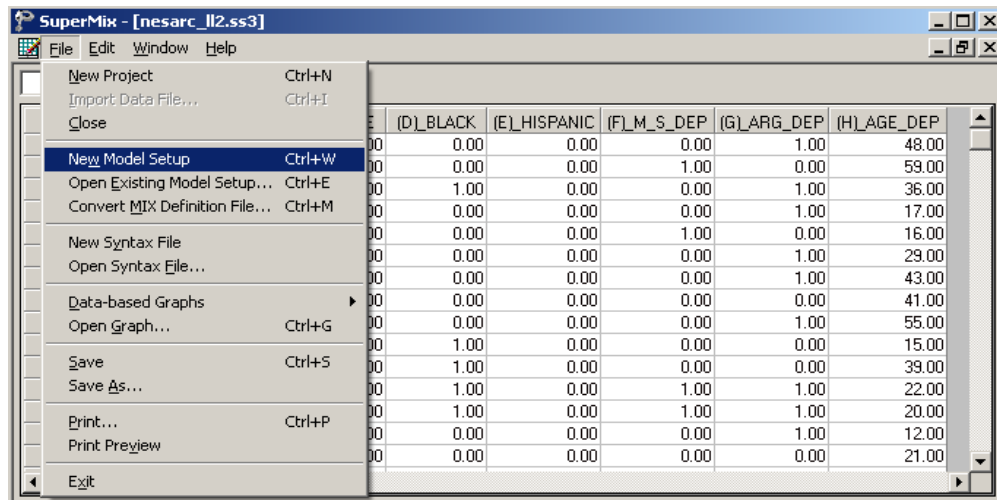
β_0 denotes the average expected age at onset of the first episode and β_1 denotes the coefficient of the predictor variable M_S_DEP (slope) in the fixed part of the model. Given that the variable M_S_DEP is an indicator variable, β_1 is in effect the expected change in age at onset for patients who reported maternal-side depression. Likewise, β_2 is in effect the expected change in age at onset for patients who reported arguments and stress. The random coefficients v_{i0} and e_{ij} denote the variation in the average expected AGE_DEP value between PSUs and between patients respectively.

The model can also be written in so-called mixed model notation, as shown below.

$$\text{AGE_DEP}_{ij} = \beta_0 + \beta_1 * \text{M_S_DEP}_{ij} + \beta_2 * \text{ARG_DEP}_{ij} + v_{i0} + e_{ij}$$

1.1.2.2 Setting up the analysis

Open the SuperMix spreadsheet **nesarc_II2.ss3** used during the exploratory analysis discussed previously. The next step is to describe the model to be fitted. We use the SuperMix interface to provide the model specifications. From the main menu bar, select the **File, New Model Setup** option.

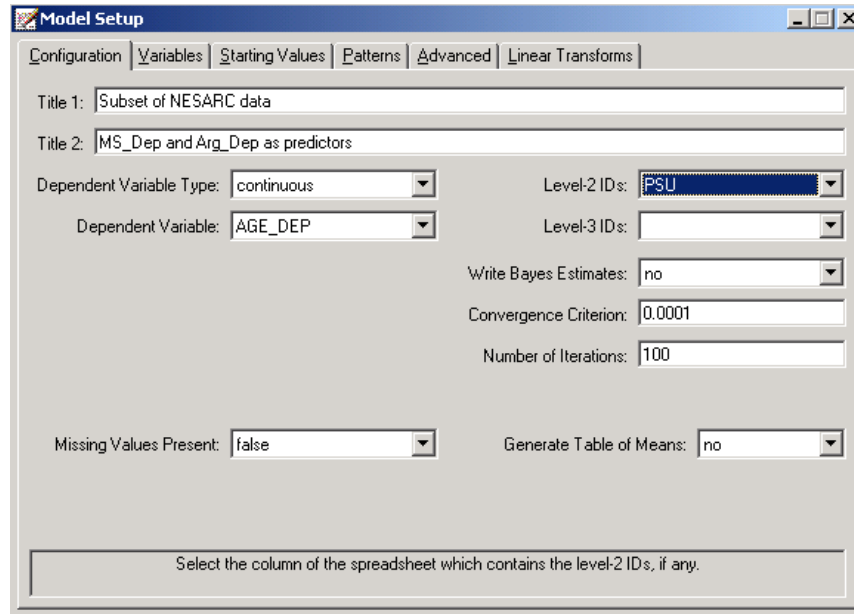


The **Model Setup** window that appears has six tabs. In this example, only the screens associated with the first two tabs are used. Information entered on these tabs are subsequently saved to a syntax file (*.mum) that can be retrieved later as needed.

The **Configuration** screen is the first tab on the **Model Setup** window. It enables the user to define the outcome variable, level-2 and level-3 IDs. Some other settings such as missing values, the convergence criterion, the number of iterations, etc. can be specified here. For all the available settings, please refer to chapter XXXX. To obtain the model we discussed, proceed as follows.

- Select the continuous outcome variable AGE_DEP from the **Dependent Variable** drop-down list box.
- Select PSU from **Level-2 ID** drop-down list box.
- Enter a title for the analysis in the **Title** text boxes (optional).

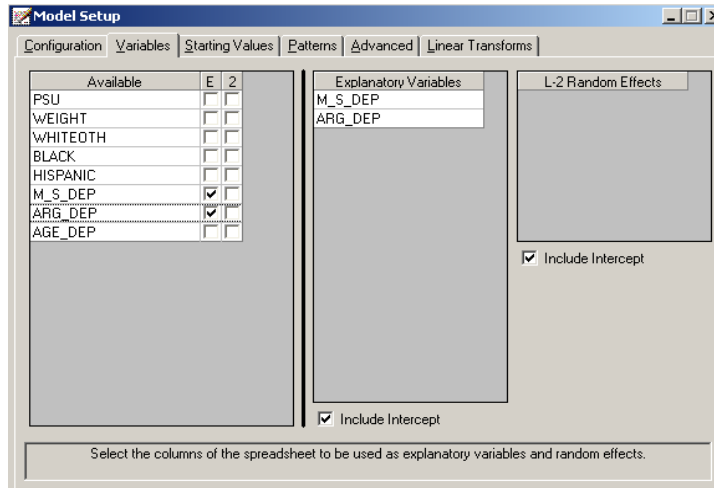
- Keep all the other settings on the **Configuration** screen at their default values. Proceed to the **Variables** screen by clicking on that tab.



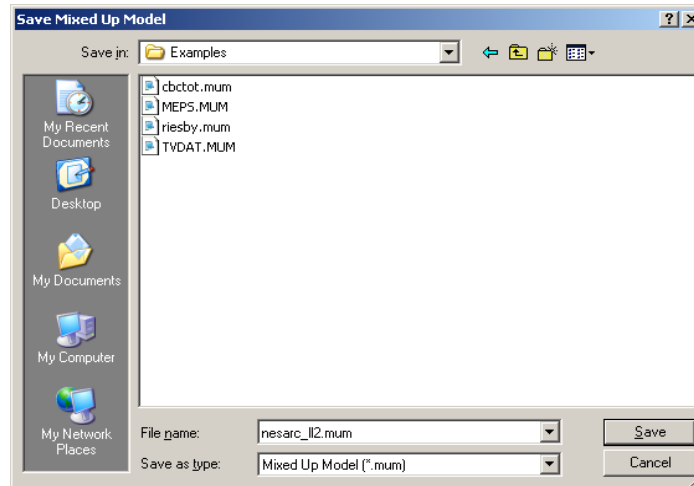
The **Variables** screen is used to specify the fixed and random effects to be included in the model. This screen shows the list of variables available for analysis and next to it two columns, with headings **E** (for explanatory variables) and **2** (for level-2 random effects).

- Select the explanatory (fixed) variables by checking the **E** check boxes next to the variables M_S_DEP and ARG_DEP in the **Available** grid at the left of the screen. Note that, as the variables are selected, they are listed in the **Explanatory Variables** grid.
- After selecting all the explanatory variables, the screen shown below is obtained.


Note that the **Include Intercept** check boxes in the **Explanatory Variables** grid and **L-2 Random Effects** grid are checked by default, indicating that an intercept term will automatically be included in the fixed and random parts of the model.



Before running the analysis, the model specifications have to be saved. Select the **File, Save As** option, provide a name (**nesarc_II2.mum**) for the model specification file, and save.



Run the analysis by selecting the **Run** option from the **Analysis** menu. The standard output file opens. It can also be viewed by selecting the **View Output** option from the same menu.

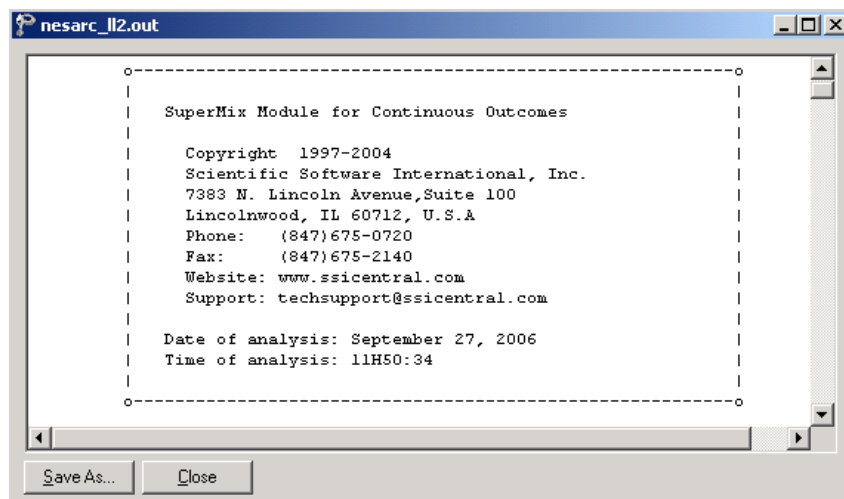


1.1.2.3 Discussion of results

Portions of the output file **nesarc_1l2.out** are shown below.

Program information and syntax

At the top of the output file, program information is given. It states the type, date and time of analysis, and provides contact information for technical support.



Program information is followed by model specifications. This section echoes the contents of the syntax file **nesarc_ll2.mum**. For more information on syntax and keywords, please see Section XXX.

```

Model specifications are as follows:

Options Converge=0.0001 Maxiter=100 Bayes=No;
Varnames= PSU WEIGHT WHITEOTH BLACK HISPANIC M_S_DEP ARG_DEP AGE_DEP intcept;
Title1=Subset of NESARC data;
Title2=MS_Dep and Arg_Dep as predictors;
DataFile=C:\Program Files\SuperMix\nesarc_ll2.dat;
Level2ID= PSU;
Dependent= AGE_DEP;
Predictors= intcept M_S_DEP ARG_DEP;
L1Random= intcept;
L2Random= intcept;
FixPatType=Free;
Cov2PatType=Correlated;
AutoCor=None;

Numbers of observations
  
```

Data summary

Numbers of observations

Level 2 observations =	371							
Level 1 observations =	1698							
N2 :	1	2	3	4	5	6	7	8
N1 :	8	1	5	5	4	1	4	1
N2 :	9	10	11	12	13	14	15	16
N1 :	3	2	3	2	26	15	3	2
N2 :	17	18	19	20	21	22	23	24
N1 :	6	4	2	2	8	1	8	6
N2 :	25	26	27	28	29	30	31	32

In the next section of the output file as shown above, a description of the hierarchical structure of the data is provided. Data from a total of 371 PSUs and 1,698 respondents were included at levels 2 and 1 of the model. In addition, a summary of the number of respondents nested within each PSU is provided. For

example, the PSU with N2:14 had 15 respondents. Note that N2:2 had only 1 observation, which means that the estimation for this PSU might not be reliable.

Descriptive statistics and starting values

The data summary is followed by descriptive statistics for all the variables included in the model. We note that the observed average age at the onset of depression is approximately 31 years.

Descriptive statistics for all variables

Variable	Minimum	Maximum	Mean	Stand. Dev.
Dependent				
AGE_DEP	5.00000	82.00000	31.39753	15.43436
Random-Effects				
intcept (2)	1.00000	1.00000	1.00000	0.00000
intcept (1)	1.00000	1.00000	1.00000	0.00000
Fixed Regressor(s)				
intcept	1.00000	1.00000	1.00000	0.00000
M_S_DEP	0.00000	1.00000	0.26443	0.44116
ARG_DEP	0.00000	1.00000	0.59600	0.49084

Descriptive statistics are followed by the starting values of the parameters that were used in the initial step of the iterative algorithm. These starting values are obtained by ordinary least squares (OLS) regression, which calculates the estimates by minimizing the sum of the squares of the residuals.

The starting values for the **fixed regressor(s)** are shown below. The **log likelihood** value and **number of free parameters** of the OLS regression are given in this part of the output.

Parameter starting values

Fixed regressor(s)

Variable	Estimate	Std. Err.	Z-value	p-value
intcept	37.46721	0.58425	64.12858	0.00000
M_S_DEP	-4.90811	0.81540	-6.01926	0.00000
ARC_DEP	-8.00650	0.73286	-10.92493	0.00000

Log Likelihood = -156505.9737
Number of free parameters = 5

Save As... Close

The starting values for the random effects are given next.

Variance/covariance components

Level 2	Estimate	Std. Err.	Z-value	p-value
intcept /intcept	-6.99789	0.10026	-69.79974	0.00000

Level 1	Estimate	Std. Err.	Z-value	p-value
intcept /intcept	217.71629	0.03861	5638.56252	0.00000

Save As... Close

Fixed effects results

The output describing the estimated **fixed effects** after convergence is shown next. The estimates are shown in the column with heading Estimate, and correspond to the coefficients β_0 , β_1 and β_2 in the model specification. From the z-values and associated exceedance probabilities, we see that all three estimates are highly significant.

TITLE1: Subset of NESARC data

Maximum likelihood estimates

Fixed regressor(s)


Variable	Estimate	Std. Err.	Z-value	p-value
intcept	37.47246	0.59754	62.71126	0.00000
M_S_DEP	-4.89876	0.81387	-6.01913	0.00000
ARG_DEP	-7.99211	0.73247	-10.91118	0.00000

Save As... Close

The estimated intercept is 37.472, which means that the average age of the first episode onset of the dysthymia respondents who do not have mother-side depression history and don't argue with others is around 37.4. The estimated coefficients associated with the mother-side history of depression (M_S_DEP) is -4.898 , which indicates that the respondents who have maternal-side depression history tend to get the first episode about five years earlier than those who do not (given the same response on ARG_DEP). The estimate for the indicator of argument involvement (ARG_DEP) shows that a respondent who has argument(s) with others is likely to have a first episode of depression about eight years earlier than a respondent who did not report arguing.

Fit statistics

In addition to the likelihood function value at convergence, a number of related statistical measures for assessing model adequacy are available. The most common of these are the likelihood ratio test and Akaike's and Schwarz's criteria. Both the Akaike information criterion (AIC) and the Schwarz Bayesian criterion (SBC) are functions of the number of estimated parameters, and therefore "penalize" models with large numbers of parameters. In the SuperMix output file, all three of these are reported. A chi-square scale factor, with which a chi-square value obtained from the difference between two deviance statistics should be multiplied to yield a corrected chi-square statistic in the case of a weighted analysis, may also be found in this section.



- The deviance is defined as $-2\ln L$. For a pair of nested models, the difference in $-2\ln L$ values has a χ^2 distribution, with degrees of freedom equal to the difference in number of parameters estimated in the models compared.
- The AIC was originally proposed for time-series models, but is also used in regression. It is defined as $-2\ln L + 2r$, where r denotes the number of parameters estimated in the model. The model with minimum AIC, in a set of nested models, will be the most parsimonious according to this criterion.
- The SBC is defined as $-2\ln L + r\log n$, where n denotes the number of units at the highest level of the hierarchy. A smaller value of this criterion would indicate the most parsimonious of the models being compared.

Random effects results

The output for the **random part** of the model follows, and is shown in the image below. In the case of a model with only a random intercept, there are two variances of interest: the variation in the random intercept over the patients, and the residual variation at level 1 over the measurements. There is no significant variation in the average estimated AGE_DEP at level 2 ($p=0.33$). This indicates that the expected average age at onset of depression does not differ significantly from PSU to PSU (the level-2 units). Significant differences between the patients (the level-1 units) are reported ($p=0.00$).

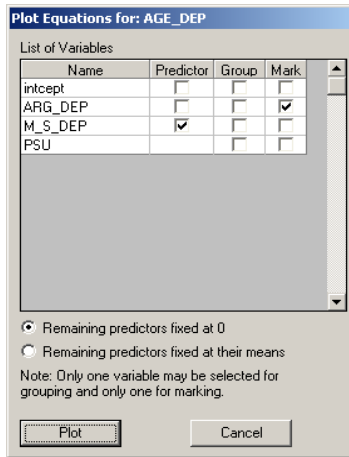
Variance/covariance components

Level 2	Estimate	Std. Err.	Z-value	p-value
intcept /intcept	2.78918	2.89726	0.96270	0.33570
Level 1	Estimate	Std. Err.	Z-value	p-value
intcept /intcept	213.07164	7.76003	27.45756	0.00000

1.1.2.4 Interpreting the results

Model-based graphs

Activate the **Model Setup** window by clicking on it. Using the **Plot Equations for: AGE_DEP** dialog box that appears when the **File, Model-based Graphs, Equations** option is selected, we can graphically depict the trend in expected age at onset of depression, taking the values of the predictors M_S_DEP and ARG_DEP into account. The dialog box below shows the selection of the predictor M_S_DEP. Marking of the plots by ARG_DEP is also requested. Two graphs will thus be displayed on the same set of axes: one for each value of the indicator variable ARG_DEP. By default, all variables present in the model, but not selected for inclusion in the graph, will be assumed to have a value of 0.



The graph below shows the result obtained when the **Plot** button is clicked after completion of the **Plot Equations for: AGE_DEP** dialog box as shown above. We note that patients who did not report arguing are expected to experience onset approximately 8 years later than patients reporting involvement in arguments.

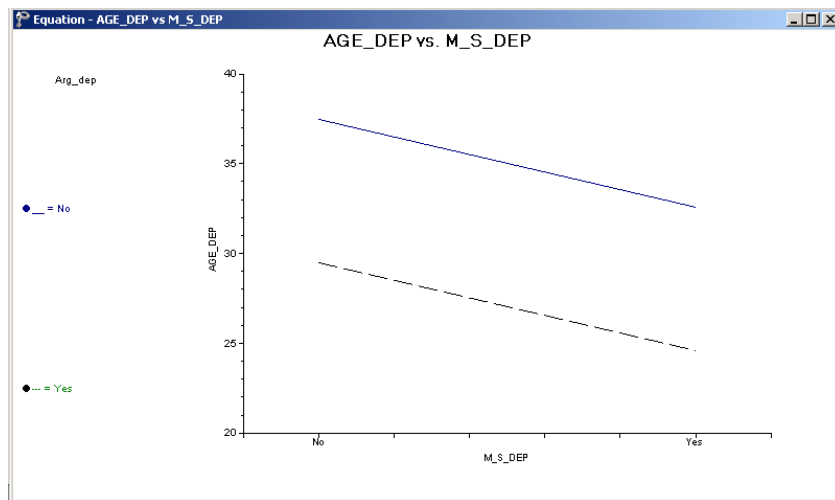


Figure XXX.3: Plot of AGE_DEP versus M_S_DEP for 2 groups

A similar plot for the predictor ARG_DEP is given next. This graph was obtained by swapping the positions of the M_S_DEP and ARG_DEP variables on the **Plot Equations for: AGE_DEP** dialog box. Note that patients with maternal-side depression had their first episode approximately 5 years earlier than patients with no history of maternal-side depression. The two graphs shown represent the graphic interpretation of the fixed effect estimates shown previously.

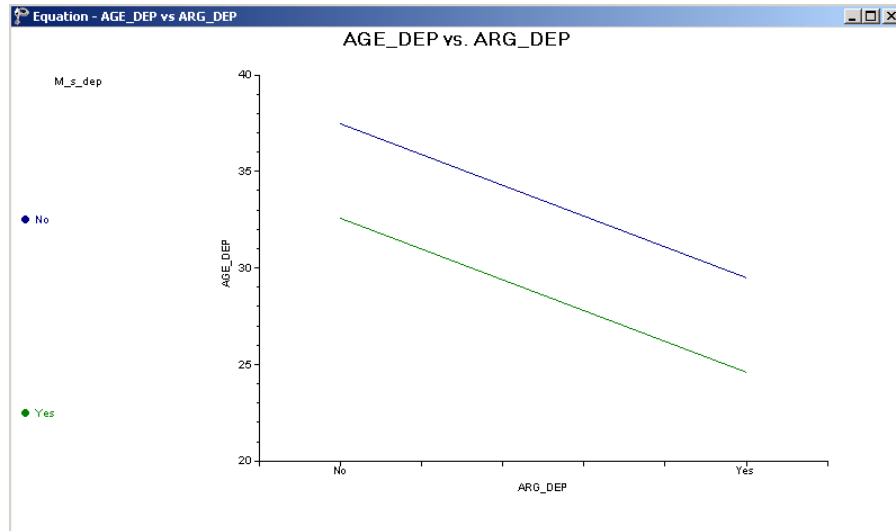


Figure XXX.4: Plot of AGE_DEP versus M_S_DEP for 2 groups

ICCs and % variance explained

By calculating the total variation in the age at onset as explained by the current model, we can obtain an estimate of the intraclass correlation coefficient. We first need to calculate the total variation in the outcome variable, which for this model is defined as $\widehat{\text{var}}(e_{ij}) + \widehat{\text{var}}(v_{i0})$.

The intraclass coefficient is then defined as

$$ICC = \frac{\widehat{\text{var}}(v_{i0})}{\widehat{\text{var}}(e_{ij}) + \widehat{\text{var}}(v_{i0})}$$

and represents the proportion of variation in age at onset that is between the groups (PSUs). An estimate of the percentage of variation in the outcome at a PSU level is obtained as

$$\frac{2.78918}{2.78918 + 213.07164} \times 100\% = 1.29\%$$

indicating that only 1.29% of the total variance is explained at PSU level; the rest of the variance remains at the respondent level.

1.1.3 A 2-level random intercept model with 4 predictors

1.1.3.1 The model

In the previous section, we modeled the outcome variable AGE_DEP as a function of M_S_DEP and ARG_DEP. The extended model discussed in this section takes the ethnicity of a respondent into consideration. The model fitted is expressed as follows:

$$\begin{aligned} \text{AGE_DEP}_{ij} = & \beta_0 + \beta_1 * \text{BLACK}_{ij} + \beta_2 * \text{HISPANIC}_{ij} \\ & + \beta_3 * \text{M_S_DEP}_{ij} + \beta_4 * \text{ARG_DEP}_{ij} + v_{i0} + e_{ij}. \end{aligned}$$

As before, β_0 denotes the average expected age at the onset of first episode, $\beta_1, \beta_2, \dots, \beta_4$ indicate the estimated coefficients associated with the fixed part of the model, and v_{i0} and e_{ij} represent the random part of the model.

Recall from Section XXX that ethnicity was represented by 3 indicator variables, namely WHITEOTH, BLACK and HISPANIC. In the model formulated above, only two of these variables have been included. This was done since the inclusion of all three indicators and the intercept term in the model would cause collinearity between the fixed effects. Any of the respondents will have a value of "1" on one of the three ethnicity indicators. If the values of the indicators are added together in a column-wise fashion, a column of 1s will result. The intercept variable is represented by just such a column of 1s in the program. If a linear combination of a subset of the columns of the design matrix is a constant multiple of another column, a condition referred to as multicollinearity is present and the model cannot be estimated properly.

Consider an example where three respondents, one from each of the three ethnic groups, are considered:

Patient	WHITEOTH	BLACK	HISPANIC	Sum of Ethnicity var.	Intercept
1	1	0	0	1	1
2	0	1	0	1	1
3	0	0	1	1	1

There are two ways in which the model can be formulated to avoid running into this problem. The first is to exclude the intercept and use only the three ethnicity indicators. Such a model, as shown below,

$$AGE_DEP_{ij} = \beta_0 * WHITEOTH_{ij} + \beta_1 * BLACK_{ij} + \beta_2 * HISPANIC_{ij} + \beta_3 * M_S_DEP_{ij} + \beta_4 * ARG_DEP_{ij} + v_{i0} + e_{ij}$$

would not offer an estimated coefficient of the average age at onset. Instead, the expected average age at onset for each of the three ethnic groups may be deduced from the estimated coefficients for WHITEOTH, BLACK and HISPANIC.

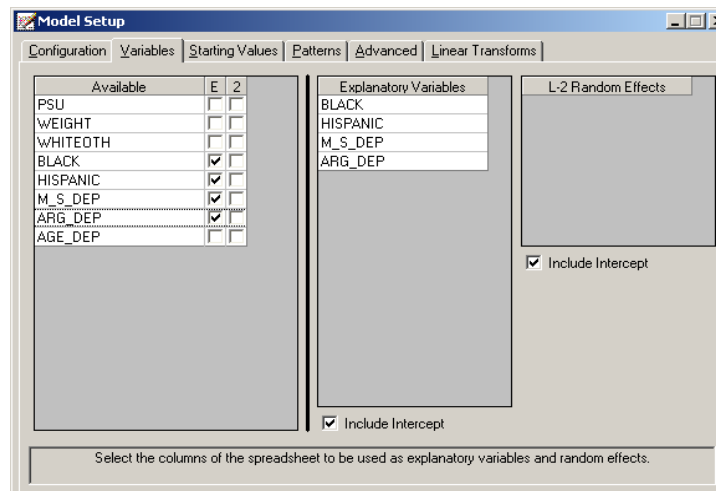
Alternatively, one can drop one of the ethnicity indicators from the model while retaining the intercept coefficient. This is what we have opted to do in the current example:

$$\text{AGE_DEP}_{ij} = \beta_0 + \beta_1 * \text{BLACK}_{ij} + \beta_2 * \text{HISPANIC}_{ij} + \beta_3 * \text{M_S_DEP}_{ij} + \beta_4 * \text{ARG_DEP}_{ij} + v_{i0} + e_{ij}$$

In the case of this formulation, the intercept coefficient represents the expected average age at onset for a patient with a value of zero on all the predictors. But if the indicators BLACK and HISPANIC assume a value of 0, it implies that the remaining ethnicity variable WHITEOTH must have a value of 1. As a result, the interpretation of the intercept coefficient would be the expected average onset age for a patient who is white or from some other ethnic origin (excluding African American and Hispanic). This ethnic group thus becomes the reference group in the current analysis. Any of the ethnic groups can be used as the reference group by simply adjusting the coding of the indicator variables; the only proviso being that the group of interest have sufficient data to serve as stable reference group.

1.1.3.2 Setting up the analysis

The SuperMix spreadsheet **nesarc_II2.ss3** and the model specification file **nesarc_II2.mum** discussed in the previous example are used a point of departure.



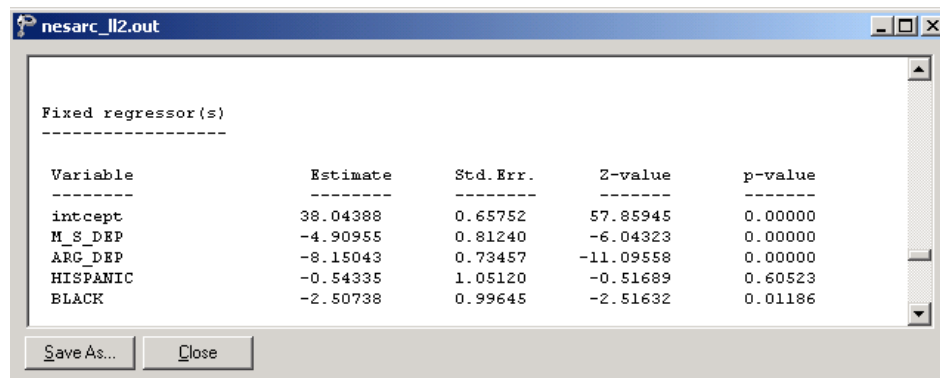
With the model specification file open, click on the **Variables** tab of the **Model Setup** window. Add the predictors **BLACK** and **HISPANIC** to the model by checking the boxes next to these variables in the **E** column, as shown above.

Save the modified model specification file, and select the **Run** option from the **Analysis** menu to perform the analysis.

1.1.3.3 Discussion of results

Fixed effects results

The maximum likelihood estimates of the coefficients in the fixed part of the model are shown below. Statistically the estimate for **HISPANIC** is not significant ($p=0.61$). Both estimates for **BLACK** and **HISPANIC** are negative, which indicates that African American and Hispanic respondents tend to have an earlier onset of the first episode when compare with patients from white and other ethnic groups.



Variable	Estimate	Std. Err.	Z-value	p-value
intcept	38.04388	0.65752	57.85945	0.00000
M_S_DEP	-4.90955	0.81240	-6.04323	0.00000
ARG_DEP	-8.15043	0.73457	-11.09558	0.00000
HISPANIC	-0.54335	1.05120	-0.51689	0.60523
BLACK	-2.50738	0.99645	-2.51632	0.01186

Fit statistics

Fit statistics for the current model are reported as shown below.


```

nesarc_ll2.out
Log Likelihood = -6968.6728
-2 Log Likelihood (Deviance) = 13937.3456
Akaike's Information Criterion = 13951.3456
Schwarz's Bayesian Criterion = 13978.7591
Number of free parameters = 7
Save As... Close

```

Random effects results

The output for the **random part** of the model is given next.

```

nesarc_ll2.out
Variance/covariance components
-----
Level 2
-----
intcept /intcept      Estimate   Std. Err.   Z-value   p-value
-----
                    3.31353    2.97617    1.11335   0.26556

Level 1
-----
intcept /intcept      Estimate   Std. Err.   Z-value   p-value
-----
                    211.80901   7.72711   27.41115   0.00000

Save As... Close

```

The random intercept effect at level 2 is not significant. As before, most of the variation in scores is found at a respondent level, with only about 2% of the variation remaining at the PSU level.

1.1.3.4 Interpreting the results

Estimated outcomes for different groups

The estimated outcome for any patient can be obtained using the formula

$$\widehat{\text{AGE_DEP}}_{ij} = \hat{\beta}_0 + \hat{\beta}_1 * \text{BLACK}_{ij} + \hat{\beta}_2 * \text{HISPANIC}_{ij} + \hat{\beta}_3 * \text{M_S_DEP}_{ij} + \hat{\beta}_4 * \text{ARG_DEP}_{ij} + v_{i0} + e_{ij}$$

For a white respondent, the expected AGE_DEP can be calculated as

$$\begin{aligned}\widehat{\text{AGE_DEP}}_{ij} &= \hat{\beta}_0 + \hat{\beta}_3 * \text{M_S_DEP}_{ij} + \hat{\beta}_4 * \text{ARG_DEP}_{ij} \\ &= 38.04388 - 4.90955 \times \text{M_S_DEP}_{ij} - 8.15043 \times \text{ARG_DEP}_{ij}.\end{aligned}$$

For African American respondents BLACK = 1, and thus the formula used to predict their AGE_DEP scores reduces to

$$\begin{aligned}\widehat{\text{AGE_DEP}}_{ij} &= \hat{\beta}_0 + \hat{\beta}_1 * \text{BLACK}_{ij} + \hat{\beta}_3 * \text{M_S_DEP}_{ij} + \hat{\beta}_4 * \text{ARG_DEP}_{ij} \\ &= 38.04388 - 2.50738 \times 1 - 4.90955 \times \text{M_S_DEP}_{ij} - 8.15043 \times \text{ARG_DEP}_{ij}.\end{aligned}$$

The formula for a patient of Hispanic origin can be derived in a similar way. In Table XXX.1, the same expected ages of the first episode onset for different groups are calculated based on the formulas above.

Table XXX.1: Expected AGE_DEP for various groups of patients

Origin	M_S_DEP = No ARG_DEP = No	M_S_DEP = Yes ARG_DEP = No	M_S_DEP = No ARG_DEP = Yes	M_S_DEP = Yes ARG_DEP = Yes
White & Other	38.04	33.13	29.89	24.98
African American	35.54	30.63	27.39	22.48
Hispanic	37.50	32.59	29.35	24.44

The results show that the respondent who has a history of maternal-side depression or gets involved into arguments generally has an earlier onset age for the first episode. For the respondents with the same M_S_DEP and ARG_DEP values, the average first episode onset ages of African American respondents are the lowest. We also conclude that a patient involved in arguments (ARG_DEP = 1) is likely to

have an earlier onset age of depression than a patient with maternal-side depression only (M_S_DEP = 1).

Fit statistics and % variation explained

The table below shows the fit indices for the previous and current models.

TABLE XXX.2: Comparison of random intercept models for NESARC data

Fit indices	Model with 2 indicators	Model with 4 indicators	Difference
Log Likelihood	-6971.8161	-6968.6728	
-2 Log Likelihood (Deviance)	13943.6321	13937.3456	6.2865
Akaike's Information Criterion	13953.6321	13951.3456	2.2865
Schwarz's Bayesian Criterion	13973.2131	13978.7591	-5.5460
Number of free parameters	5	7	

The difference in deviances can be used to assess the model fit. This method is valid for nested models. A nested model may be defined as any submodel of a given model that is based on the same number of observations. Given the difference in structure between the 2-level models these models cannot, however, be compared to each other.

The difference in the deviances follows a χ^2 distribution, where the degree of freedom is the difference of numbers of free parameters.

$$(-2 \ln_{model1}) - (-2 \ln_{model2}) \sim \chi^2(d.f.(-2 \ln_{model2}) - (-2 \ln_{model1}))$$

When the deviances of the two models are compared, a χ^2 -statistic of $13943.6321 - 13937.3456 = 6.2865$ with $7 - 5 = 2$ degrees of freedom is obtained. This indicates that the current model fits the data better. The AIC decreased from 13953.6321 to 13951.3456, and also favors the use of the 4-predictor model. The SBC, however, increased slightly, from 13973.2131 to 13978.7591, and thus favors the model

previously fitted as the more parsimonious. The definitions of these indices are given in the discussion of the output of the previous model. Note, however, that the changes in all three criteria are rather small.

The estimated percentages of variation in outcome at respondent level can be calculated using the variance components reported in the random effects part of the output file:

$$\frac{211.80901}{211.80901 + 3.31353} \times 100\% = 98.46\% .$$

Once the additional level-1 predictors are taken into account, there does not seem to be significant random variation in the outcome over the intercepts of the level-2 units. The estimated average onset age of the first episode does not vary significantly from PSU to PSU.

1.2 Models based on the Reisby data

1.2.1 The data

The data set is from a study described in Reisby *et. al.* (1977) that focused on the longitudinal relationship between imipramine (IMI) and desipramine (DMI) plasma levels and clinical response in 66 depressed inpatients (37 endogenous and 29 non-endogenous). Following a placebo period of 1 week, patients received 225 mg/day doses of imipramine for four weeks. In this study, subjects were rated with the Hamilton depression rating scale (HDRS) twice during the baseline placebo week (at the start and end of this week) as well as at the end of each of the four treatment weeks of the study. Plasma level measurements of both IMI and its metabolite DMI were made at the end of each week. The sex and age of each patient were recorded and a diagnosis of endogenous or non-endogenous depression was made for each patient.

Although the total number of subjects in this study was 66, the number of subjects with all measures at each of the weeks fluctuated: 61 at week 0 (start of placebo week), 63 at week 1 (end of placebo week), 65 at week 2 (end of first drug treatment week), 65 at week 3 (end of second drug treatment week), 63 at week 4 (end of third drug treatment week), and 58 at week 5 (end of fourth drug treatment week). The sample size is 375. Data for the first 10 observations of all the variables used in this section are shown below in the form of a SuperMix spreadsheet file, named **reisby.ss3**.



	(A)_Patient	(B)_HDRS	(C)_WEEK	(D)_WEEK	(E)_ENDO	(F)_WxEND
1	101.00	26.00	0.00	0.00	0.00	0.00
2	101.00	22.00	1.00	1.00	0.00	0.00
3	101.00	18.00	2.00	4.00	0.00	0.00
4	101.00	7.00	3.00	9.00	0.00	0.00
5	101.00	4.00	4.00	16.00	0.00	0.00
6	101.00	3.00	5.00	25.00	0.00	0.00
7	103.00	33.00	0.00	0.00	0.00	0.00
8	103.00	24.00	1.00	1.00	0.00	0.00
9	103.00	15.00	2.00	4.00	0.00	0.00
10	103.00	24.00	3.00	9.00	0.00	0.00

The variables of interest are:

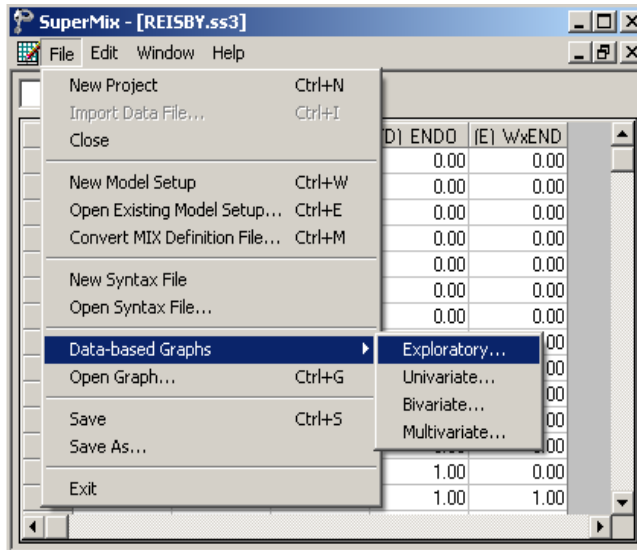
- Patient is the patient ID (66 patients in total).
- HDRS is the Hamilton depression rating scale.
- WEEK represents the week (0, 1, 2, 3, 4 or 5) at which a measurement was made.
- WEEKSQ represents the squared values of WEEK. The creation of this variable is illustrated in Section 1.2.5.2.
- ENDOG is a dummy variable for the type of depression a patient was diagnosed with (1 for endogenous depression and 0 for non-endogenous depression).
- WxENDOG represents the interaction between WEEK and ENDOG, and is the product of WEEK and ENDOG.

1.2.1.1 Exploring the data

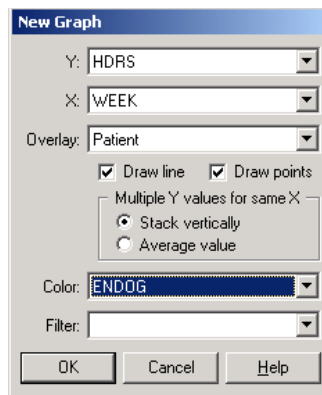
Graphing the observed data

In the previous example, we have shown a number of data-based graphs. Here, we use the **Exploratory** option of the **Data-Based Graphs** menu to explore the data in the **reisby.ss3** spreadsheet, stored in the **Continuous** subfolder.

Start by opening the data file in the SuperMix spreadsheet. Then select the **Data-based Graphs, Exploratory** option on the **File** menu as shown below to activate the **New Graph** dialog box.



Specify HDRS as the dependent (vertical axis) variable by selecting it from the **Y** drop-down list box and WEEK as the independent (horizontal axis) variable by selecting it from the **X** drop-down list box. A graph on the same axis system is created for each patient by selecting the variable Patient from the **Overlay** drop-down list box. Furthermore, each graph is assigned a color by selecting ENDG from the **Color** drop-down list box to produce the following **New Graph** dialog box.



Click on the **OK** button to produce the following graph of the reaction trajectories over time for the 66 inpatients.


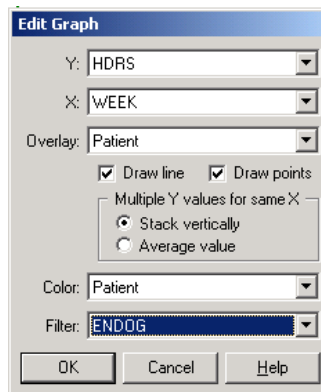


Figure XXX.5: Reaction trajectories over time for 66 patients

To modify the existing graphic display, select the **Edit Graph** option from the **Settings** menu to load the **Edit Graph** dialog box. To obtain different graphs for the two categories of the covariate ENDOG, select it from the **Filter** drop-down list box to produce the following **Edit Graph** dialog box.



Click on the **OK** button to open the following graphics window.




Figure XXX.6: Reaction trajectories over time for patients with ENDOG=0

At the bottom of the graphics window is a "slider" with left and right arrows. By clicking on the right arrow, one can obtain the next graphic shown below and by clicking on the left arrow, the graphic above.

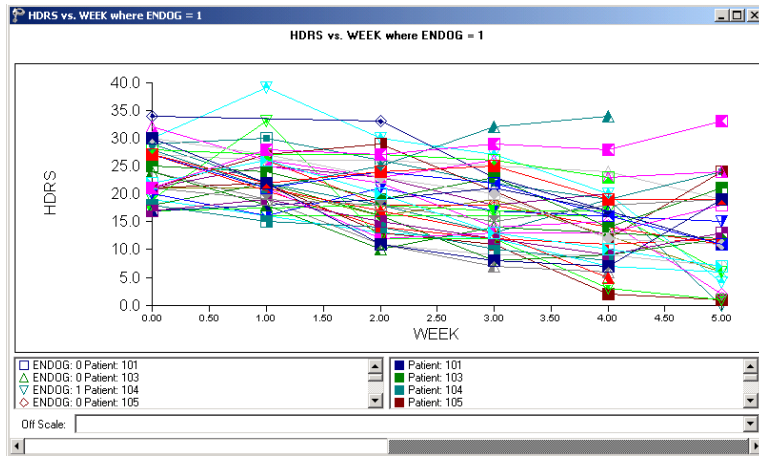


Figure XXX.7: Reaction trajectories over time for patients with ENDOG=1

The above graphs show a general, approximately linear decline over time and an increase in the variability of the HDRS scores across time for both types of depression.

1.2.2 A 2-level random intercept-and-slope model

From the graphical display obtained in the previous section, it seems as if the HDRS scores follow an approximately linear trend over time, decreasing over the course of the study. It is also apparent, however, that patients not only start out at different levels but also have differences in the slopes of the HDRS against WEEK lines. In this section, we explore a model that allows patients not only to have unique intercepts, but also unique slopes across time. In other words, we allow both intercept and WEEK (slope) to vary randomly over patients. The image below demonstrates the meaning of the random slope and random intercept in a hypothetical 2-level model.




Figure XXX.8: Score trends for individual patients

1.2.2.1 The model

The random intercept-and-slope model for the response variable HDRS may be expressed as

$$\text{HDRS}_{ij} = \beta_0 + \beta_1 \times (\text{WEEK})_{ij} + v_{i0} + v_{i1} (\text{WEEK})_{ij} + \varepsilon_{ij}$$

We can rewrite the model in the following way.

Level-1 model:

$$\text{HDRS}_{ij} = b_{0i} + b_{1i} \times (\text{WEEK})_{ij} + e_{ij}$$

Level-2 model:

$$b_{0i} = \beta_0 + v_{0i}$$

$$b_{1i} = \beta_1 + v_{1i}$$

where

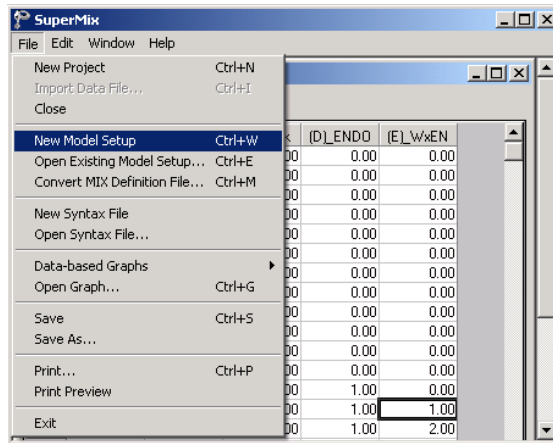
$$\varepsilon_i \sim N(0, \sigma^2 \mathbf{I}_i)$$

$$\mathbf{v}_i \sim N(\mathbf{0}, \Phi_{(v)})$$

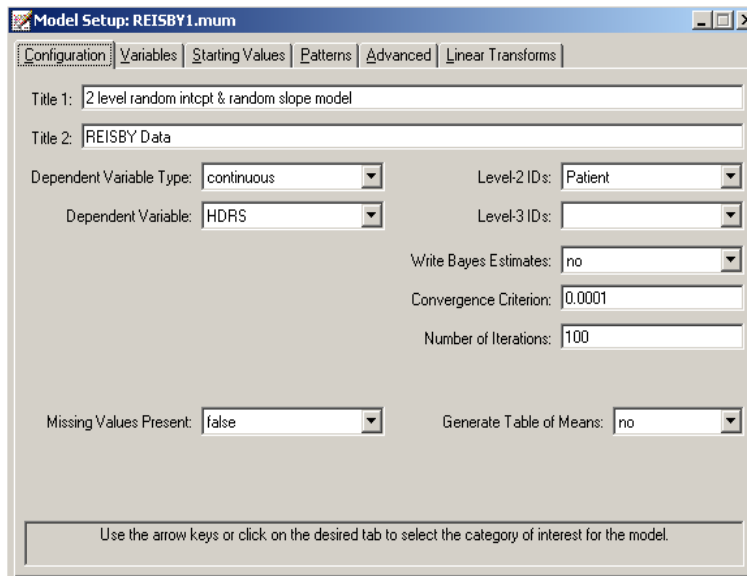
β_0 denotes the average expected depression rating scale value, β_1 denotes the coefficient of the predictor variable WEEK (slope) in the fixed part of the model, v_{1i} denotes the variation in the slopes over patients, and v_{0i} and e_{ij} denote the variation in the average expected HDRS value over patients and between patients respectively. Furthermore, $i = 1, 2, \dots, 66$ refers to the 66 patients; $j = 1, 2, \dots, n_i$ refers to the j^{th} observation for patient i . The maximum value for n_i is 6.

1.2.2.2 Setting up the analysis

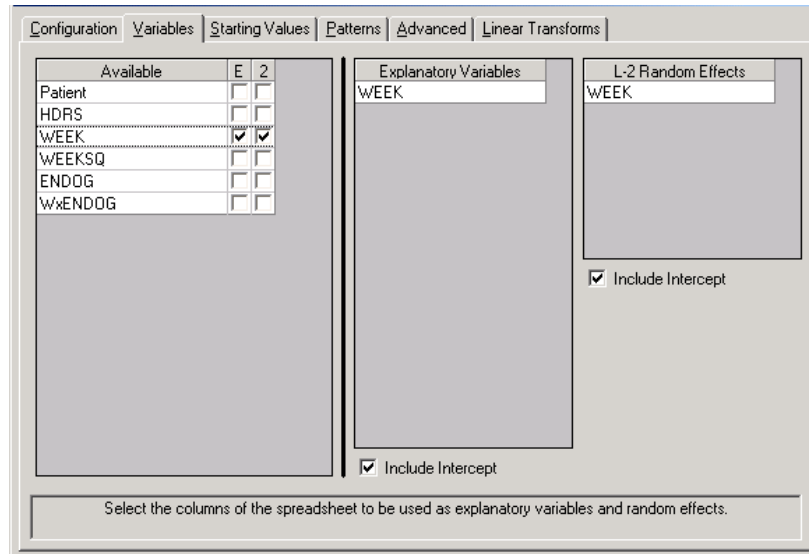
Start by opening the **reisby.ss3** file as a SuperMix spreadsheet. Next, select the **New Model Setup** option on the **File** menu as shown below to load the **Model Setup** window.



Starting with the **Configuration** screen, enter the (optional) title in the **Title 1** and **Title 2** text boxes respectively. The continuous outcome variable HDRS is selected from the **Dependent Variable** drop-down list box. The variable Patient, which defines the levels of the hierarchy, is selected as the Level-2 ID from the **Level-2 IDs** drop-down list box to produce the following **Configuration** screen.



Click the **Variables** tab to proceed to the **Variables** screen of the **Model Setup** window. The variable **Week** is specified as the covariate of the fixed part of the model by checking the **E** check box for **WEEK** in the **Available** grid. Mark the **2** check box for **Week** in the **Available** grid to specify the random slope at level 2 of the model. After completion, the **Variables** screen should look as shown below.



Before the analysis can be run, save the model specifications to **reisby1.mum**. Run the model to produce the output file **reisby1.out**.

1.2.2.3 Discussion of results

Descriptive statistics

The section of the output file shown below contains the descriptive statistics for all variables in the current model specification. If all patients' data were complete, the average for the time variable **WEEK** would have been exactly 2.5; the value of 2.48 indicates that the number of patients with information at each time point fluctuates somewhat.

Descriptive statistics for all variables

Variable	Minimum	Maximum	Mean	Stand. Dev.
Dependent				
HDRS	0.00000	39.00000	17.63733	7.19006
Random-Effects				
intcept (2)	1.00000	1.00000	1.00000	0.00000
WEEK (2)	0.00000	5.00000	2.48000	1.68320
intcept (1)	1.00000	1.00000	1.00000	0.00000
Fixed Regressor(s)				
intcept	1.00000	1.00000	1.00000	0.00000
WEEK	0.00000	5.00000	2.48000	1.68320

Save As... Close

1.2.2.4 Interpreting the results

The summary of the hierarchical structure of the data shows how the 375 measurements are nested within the 66 patients. It also indicates that the number of repeated measurements per patient varies from 4 to 6 observations. The convergence is attained in 5 iterations. The output file contains the final estimates of the fixed and random coefficients included in the model, along with some goodness of fit measures as shown below.

```

Maximum likelihood estimates
-----

Fixed regressor(s)
-----

Variable          Estimate      Std. Err.    Z-value      p-value
-----
intcept           23.57695     0.54555     43.21714     0.00000
WEEK              -2.37707     0.20865     -11.39280    0.00000

Log Likelihood          = -1109.5188
-2 Log Likelihood (Deviance) = 2219.0375
Akaike's Information Criterion = 2231.0375
Schwarz's Bayesian Criterion = 2244.1754
Number of free parameters = 6

Variance/covariance components
-----

Level 2
-----
Variable          Estimate      Std. Err.    Z-value      p-value
-----
intcept /intcept   12.62930     3.46653     3.64322     0.00027
WEEK /intcept     -1.42093     1.02595     -1.38500     0.16605
WEEK /WEEK         2.07899     0.50417     4.12363     0.00004

Level 1
-----
Variable          Estimate      Std. Err.    Z-value      p-value
-----
intcept /intcept   12.21663     1.10696     11.03615     0.00000

```

Fixed effects results

The results show a highly significant coefficient ($p < 0.00001$) for the time effect, as represented by the variable WEEK. At the beginning of the study, when WEEK = 0, the average expected HDRS score is 23.57695. For each subsequent week, a decrease of 2.37707 in average HDRS score is expected. At the end of the study period, the average expected HDRS score is $23.57695 - 5(2.37707) = 11.6916$.

Random effects results

With the exception of the WEEK-intcept covariance, all variance components are highly significant, as shown in the p -value column. From the output above we have $\widehat{\text{var}}(v_{i0}) = 12.62930$, $\widehat{\text{var}}(v_{i1}) = 2.07899$, $\widehat{\text{cov}}(v_{i0}, v_{i1}) = -1.42093$, and $\widehat{\text{var}}(e_{ij}) = 12.21663$. Typically, one would expect most of the variation in HDRS scores at the

measurement level, and thus would expect $\widehat{\text{var}}(e_{ij})$ to be larger than any of the other variances/covariances. With these data, however, there is more variation in the random intercepts over patients than in the measurements nested within patients. Due to this, it may be of interest to take a closer look at the variation in HDRS scores at the two levels of the hierarchy.

Fit statistics and ICC

In the case of a model with only a random intercept, there are two variances of interest: the variation in the random intercept over the patients (the level-2 units), and the residual variation at level 1, over the measurements. By calculating the total variation in the HDRS score explained by such a model, obtained as $\widehat{\text{var}}(e_{ij}) + \widehat{\text{var}}(v_{i0})$, we can obtain an estimate of the intraclass correlation coefficient.

The intraclass coefficient is defined as

$$ICC = \frac{\widehat{\text{var}}(v_{i0})}{\widehat{\text{var}}(e_{ij}) + \widehat{\text{var}}(v_{i0})}$$

and would, for a random intercept model for this data, represent the proportion of variation in HDRS scores between patients. The term intraclass correlation coefficient applies to random intercept models only; in more complicated models the focus is on explanation of variation in various coefficients.

In the current model, the situation is somewhat more complicated due to the inclusion of both random intercept and random slope. This implies a possible correlation between the level-2 random effects. When calculating an estimate of the total variation, the covariance(s) between random effects have to be taken into account in any attempt to estimate the proportion of variation in outcome at any level or for any random coefficient. In addition, the inclusion of a covariate such as ENDOG can affect the variance estimates.

The total variation in HDRS scores over patients is defined as

$$\text{Var}(\text{level } 2) = \text{var}(v_{i0}) + \text{var}(v_{i1})(\text{WEEK})_{ij}^2 + 2[\text{cov}(v_{i0}, v_{i1})](\text{WEEK})_{ij}$$

The total variation is a function of the value assumed by the predictor WEEK, which has a random slope. As such, the total variation at the beginning of the study is

$$\begin{aligned}\text{Var}(\text{level } 2) &= \text{var}(v_{i0}) + \text{var}(v_{i1})(0)^2 + 2[\text{cov}(v_{i0}, v_{i1})](0) \\ &= \text{var}(v_{i0})\end{aligned}$$

while at the end of the study we have

$$\begin{aligned}\text{Var}(\text{level } 2) &= \text{var}(v_{i0}) + \text{var}(v_{i1})(5)^2 + 2[\text{cov}(v_{i0}, v_{i1})](5) \\ &= \text{var}(v_{i0}) + 25 \text{var}(v_{i1}) + 10 \text{cov}(v_{i0}, v_{i1})\end{aligned}$$

An estimate of the total variation at this level can be obtained by using the estimates of the variances and covariance obtained under this model. By substituting $\widehat{\text{var}}(v_{i0})$, $\widehat{\text{var}}(v_{i1})$, and $\widehat{\text{cov}}(v_{i0}, v_{i1})$ into the equations above, we obtain the estimated variation in HDRS scores over patients at different points during the study period.

At the beginning of the study, the estimated total variation in HDRS scores over patients is simply the estimated variation in the random intercept, *i.e.*, $\widehat{\text{var}}(v_{i0}) = 12.62930$. At the end of the study, the total variation at level-2 is estimated as

$$\begin{aligned}\widehat{\text{Var}}(\text{level } 2) &= \widehat{\text{var}}(v_{i0}) + 25\widehat{\text{var}}(v_{i1}) + 10\widehat{\text{cov}}(v_{i0}, v_{i1}) \\ &= 12.62930 + 25(2.07899) + 10(-1.42093) \\ &= 50.39475.\end{aligned}$$

At the beginning of the study we obtain

$$\frac{\widehat{\text{var}}(\text{level 2})}{\widehat{\text{var}}(\text{level 2}) + \widehat{\text{var}}(\text{level 1})} = \frac{12.62930}{12.62930 + 12.21663} \\ = 0.5083$$

and thus conclude that 50.8% of the variation in HDRS scores at this time is over patients. At the end of the study, we find that

$$\frac{\widehat{\text{var}}(\text{level 2})}{\widehat{\text{var}}(\text{level 2}) + \widehat{\text{var}}(\text{level 1})} = \frac{50.39475}{50.39475 + 12.21663} \\ = 0.8049,$$

so that only 20% of the variation in HDRS scores are estimated to be at the measurement level, with 80% at the patient level. As mentioned before, the total variation in HDRS scores is a function of the time of measurement, as represented by the variable WEEK. The very different estimates of variation at a patient level show how the introduction of an important predictor, in this case at the measurement level, can have an impact on variance estimates at a different level of the hierarchy. By the end of the study period, the residual variation over measurements has been dramatically reduced, this being explained to a large extent by the inclusion of the time effect. Most of the remaining unexplained variation is at the patient level.

As a result of this finding and in the light of our original research question, whether the initial depression classification of a patient is also related to the HDRS scores over the time in which medication is administered, the model will be extended to include the covariate ENDOG. This dichotomous variable assumes a value of 1 when endogenous depression was diagnosed, and 0 if not. In addition, we will provide for a possible interaction between depression classification and measurement occasion by including the interaction term $W \times \text{ENDOG}$ in the model. While $W \times \text{ENDOG}$ can be viewed as a cross-level interaction, as WEEK is a measurement-level variable and

ENDOGENOUS a patient-level variable, the inclusion of the patient-level variable ENDOGENOUS may enable us to explain more of the remaining variation in the random intercepts and slopes at the patient level.

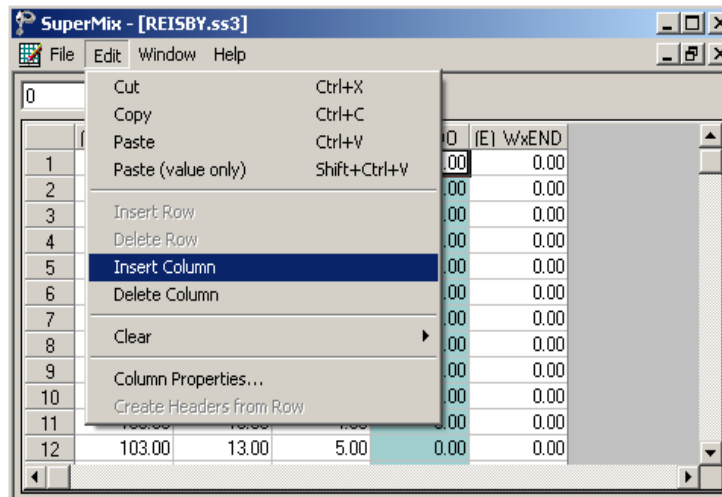
1.2.3 A 2-level random intercept-and-slope model with centered predictor

In the previous example, the time variable WEEK is coded from 0 to 5 and indicates the number of weekly follow-ups. The estimated average intercept of 23.577 obtained for this model represented the expected average HDRS score at the beginning of the study, i.e. WEEK = 0. An alternative formulation of the model that can be considered is one in which the estimated average intercept represents the expected average HDRS score midway through the study period. This linear transformation of the predictor variable WEEK, in which the grand mean of the variable is subtracted from each observed WEEK value, is referred to as grand mean centering. While the model based on the "raw" data and the model utilizing grand mean centered variables can be shown to be mathematically equivalent, the coefficients in these models have very different meanings.

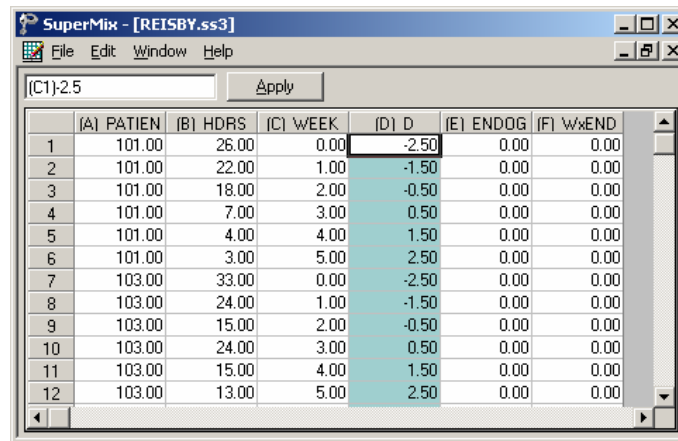
1.2.3.1 Preparing the data

Recall that the descriptive statistics in the previous model indicated a mean value over all level-1 observations of WEEK equal to 2.48. This is the true observed mean, compared to the value of 2.5 that would have been obtained if all patients had complete data over the course of the study. Here, we opt to use the value of 2.5 to center the WEEK variable.

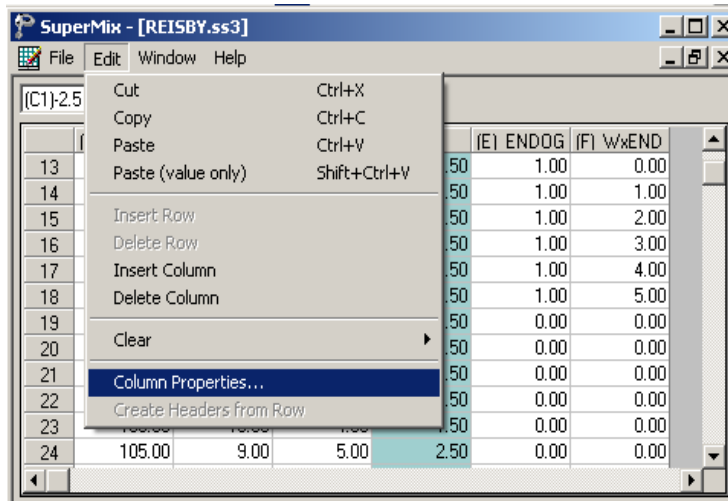
To grand mean center the predictor WEEK, proceed as follows. Open the **reisby.ss3** in the SuperMix spreadsheet, then highlight the column WEEK. Select the **Insert Column** option on the **Edit** menu as shown below to insert a blank column named D after WEEK.




Keep the column D highlighted, type the formula (C1)-2.5 in the string field of the top-left corner and click on the **Apply** button to produce the following screen.



Rename the newly created variable to WEEKC by first highlighting the column, then selecting the **Column Properties** option on the **Edit** menu as shown below.



Input the desired variable name, *e.g.* WEEKC, in the **Header** string field as shown below and click on the **OK** button. By default, all variables are assumed to be continuous.



Save the changes to **reisby.ss3** by selecting the **Save** option on the **File** menu.

1.2.3.2 The model

The revised random intercept-and-slope model for the response variable HDRS may be expressed as

$$\text{HDRS}_{ij} = \beta_0 + \beta_1 \times (\text{WEEKC})_{ij} + v_{i0} + v_{i1} (\text{WEEKC})_{ij} + \varepsilon_{ij}$$

or, alternatively, as

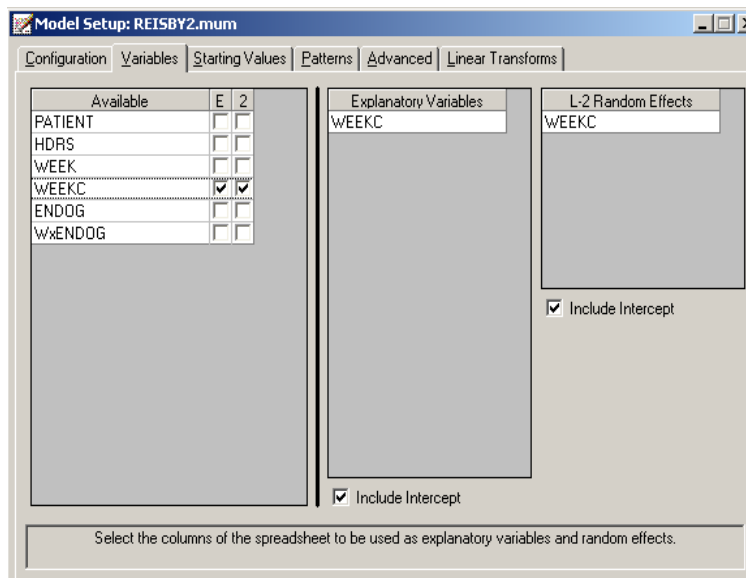
$$\text{HDRS}_{ij} = \beta_0 + \beta_1 \times \left[(\text{WEEK})_{ij} - \overline{\text{WEEK}} \right] + v_{i0} + v_{i1} \left[(\text{WEEK})_{ij} - \overline{\text{WEEK}} \right] + \varepsilon_{ij}$$

where $\overline{\text{WEEK}} = 2.5$.

1.2.3.3 Setting up the analysis

Open the previous model setup for **reisby1.mum**. Save the file as **reisby2.mum** by using the **Save As** option on the **File** menu. Change the title on the **Configuration** tab if desired.

Click on the **Variables** tab and select WEEKC both as **Explanatory Variable** and **L-2 Random Effects** instead of WEEK as shown below.



Save the changes to the file **reisby2.mum**. Select the **Run** option on the **Analysis** menu to produce the output file **reisby2.out**. Use the **Analysis, View Output** option to open the output file.

1.2.3.4 Discussion of results

The output file contains the final estimates of the fixed and random coefficients included in the model, along with some goodness of fit measures as given below. Note that the use of grand mean centering of the time variable has no effect on the fit statistics.

```

SuperMix - [REISBY2.out]
File Analysis Window Help

Maximum likelihood estimates
-----

Fixed regressor(s)
-----

Variable          Estimate      Std. Err.      Z-value      p-value
-----
intcept           17.63428     0.56031       31.47258    0.00000
WEEKC             -2.37707     0.20865      -11.39281    0.00000

Log Likelihood          = -1109.5188
-2 Log Likelihood (Deviance) = 2219.0375
Akaike's Information Criterion = 2231.0375
Schwarz's Bayesian Criterion = 2244.1754
Number of free parameters = 6

Variance/covariance components
-----

Level 2
-----
intcept /intcept      18.51833     3.61203     5.12685     0.00000
WEEKC /intcept        3.77654     1.05839     3.56821     0.00036
WEEKC /WEEKC          2.07899     0.50416     4.12364     0.00004

Level 1
-----
intcept /intcept      12.21663     1.10697    11.03614     0.00000
  
```

1.2.3.5 Interpreting the results

Comparison of models

Table XXX.3 contains the estimates and standard errors of the above two analyses. The coefficient for WEEKC is the same as for the uncentered variable WEEK. However, the variance of the random intercept ($\sigma_{v_0}^2$) and the covariance term $\sigma_{v_0v_1}$ have changed. The covariance between the intercept and the WEEKC slope is now significant.

TABLE XXX.3: Estimates and standard errors for two models

Coefficient	Level-2 model	
	WEEK = 0 ~ 5	WEEKC = -2.5 ~ 2.5
β_0	23.57695 (0.54555)	17.63428 (0.56031)
β_1	-2.37707 (0.20865)	-2.37707 (0.20865)
$\sigma_{v_0}^2$	12.6293 (3.46653)	18.51833 (3.61203)
$\sigma_{v_0v_1}$	-1.42093 (1.02595)	3.77654 (1.05839)
$\sigma_{v_1}^2$	2.07899 (0.50417)	2.07899 (0.50416)
σ_e^2	12.21663 (1.10697)	12.21663 (1.10697)
Deviance	2219.0375	2219.0375
AIC	2231.0375	2231.0375
SBC	2244.1754	2244.1754

Number of free parameters	6	6
---------------------------	---	---

As shown above, the estimates of the slope and its variance are the same. This is because the scale of WEEK was not changed; only its location changed. The estimated intercept decreased from 23.58 to 17.63, which corresponds to the average HDRS score at week 2.5 instead of week 0. Similarly, the $\sigma_{v_0}^2$ of intercept increased to 18.52, which shows the increase of the individual variance at week 2.5. The change of $\sigma_{v_0v_1}$ is interesting: not only the value changed, but also the sign. The covariance of the first analysis tells us that the higher the variance of intercept, the lower the variance of slope. Or say, at week 1, the HDRS score decreases at a faster rate for those patients who started with higher HDRS. However, at week 2.5, the patients with higher HDRS tend to improve less. When looking at the three HDRS versus WEEK plots for patient 604, 302 and 361 below, we can see why this could happen. The graphs show the change of $\sigma_{v_0v_1}$ from week 0 to week 2.5.




Figure XXX.9: Changes in covariance over time

1.2.4 A random intercept-and-slope with a covariate and an interaction term

The type of depression a patient was diagnosed with was recorded as part of the study and information on this patient characteristic is represented by the variable ENDOG, which assumes a value of 1 for patients with endogenous depression and 0 otherwise. Including this variable in the model allows us to explore the potential relationship between a patient's HDRS score and the type of depression the patient was diagnosed with. Moreover, it is possible that the trend in HDRS scores over the

study period may differ for the two ENDOG groups. Including an interaction term between the time of measurement and the type of depression in the model will allow us to evaluate this potential relationship as well.

1.2.4.1 The model

We now include ENDOG and WxENDOG in the level-1 model. ENDOG is a dummy variable representing the type of depression a patient was diagnosed with, and WxENDOG represents the interaction between WEEK and ENDOG. The model shows changes at both levels: at level 2, the covariate ENDOG is now included, while at level 1 the interaction between WEEK and ENDOG, which can potentially change from week to week, is added. The revised model for the response variable HDRS may be expressed as

Level-1 model:

$$\text{HDRS}_{ij} = b_{0i} + b_{1i} \times (\text{WEEK})_{ij} + b_{2i} \times (\text{WxENDOG})_{ij} + e_{ij}$$

Level-2 model:

$$b_{0i} = \beta_0 + \beta_3 \times (\text{ENDOG})_i + v_{0i}$$

$$b_{1i} = \beta_1 + v_{1i}$$

$$b_{2i} = \beta_2$$

or, in mixed model formulation, as

$$\begin{aligned} \text{HDRS}_{ij} = & \beta_0 + \beta_1 \times (\text{WEEK})_{ij} + \beta_2 \times (\text{WxENDOG})_{ij} + \beta_3 \times (\text{ENDOG})_i \\ & + v_{0i} + v_{1i} \times (\text{WEEK})_{ij} + e_{ij} \end{aligned}$$

where β_0 denotes the average HDRS level at week 0 for the non-endogenous depression patients (ENDOG=0), β_1 refers to the weekly improvement for the non-

endogenous group, β_2 indicates the expected change in HDRS score for a unit change in the value of the interaction term WxENDOG, and β_3 refers to the average expected change in HDRS level for endogenous patients. v_{0i} is the individual deviation from the average intercept. v_{1i} denotes the average deviation from the slope, or say, average improvement of the HDRS.

We can also write the model in terms of our original variables (WEEK and ENDOG) as:

Level-1 model:

$$\text{HDRS}_{ij} = b_{0i} + b_{1i} \times (\text{WEEK})_{ij} + b_{2i} \times (\text{WxENDOG})_{ij} + e_{ij}$$

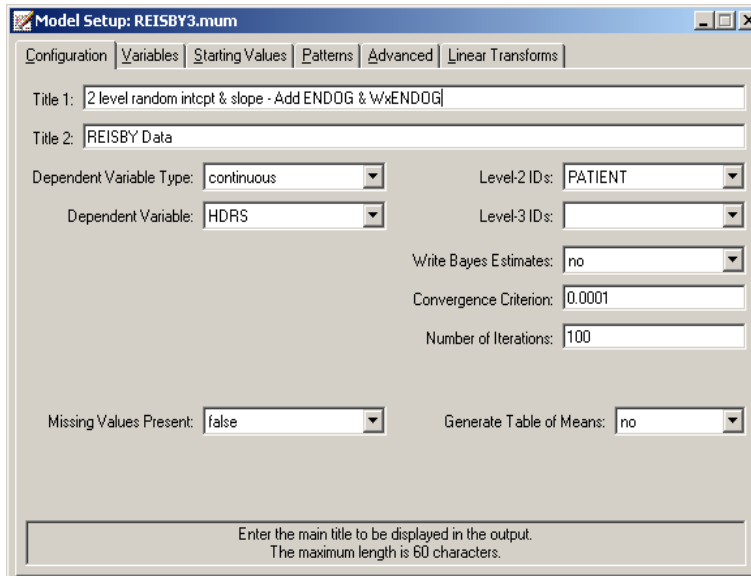
Level-2 model:

$$b_{0i} = \beta_0 + \beta_2 \times (\text{ENDOG})_i + v_{0i}$$

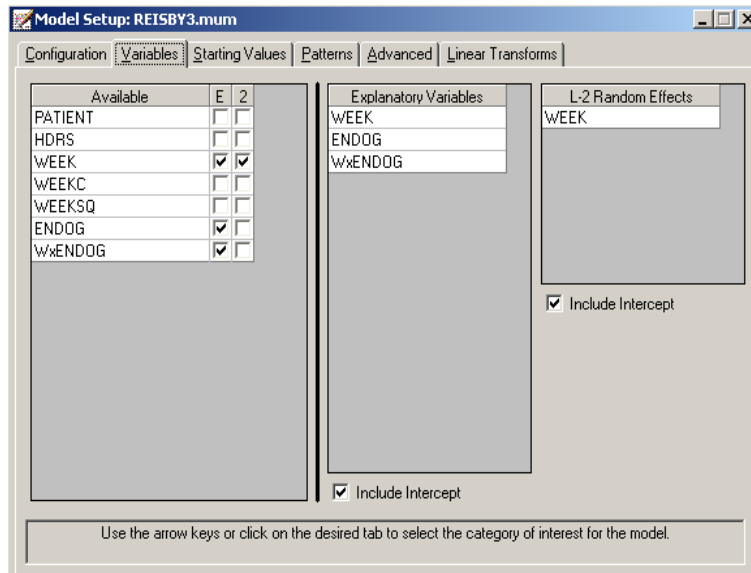
$$b_{1i} = \beta_1 + \beta_4 \times (\text{ENDOG})_i + v_{1i}$$

1.2.4.2 Setting up the analysis

To create the model specifications for this model, we start by opening **reisby.ss3** in a SuperMix spreadsheet window. Then we use the **Open Existing Model Setup** option on the **File** menu to load the **Model Setup** window for **reisby1.mum**. Save the file as **reisby3.mum** by using the **Save As** option on the **File** menu. Change the string in the **Title 1** text box on the **Configuration** screen to reflect the new model, thereby producing the following dialog box.



Next, click on the **Variables** tab to proceed to the **Variables** screen of the **Model Setup** window.



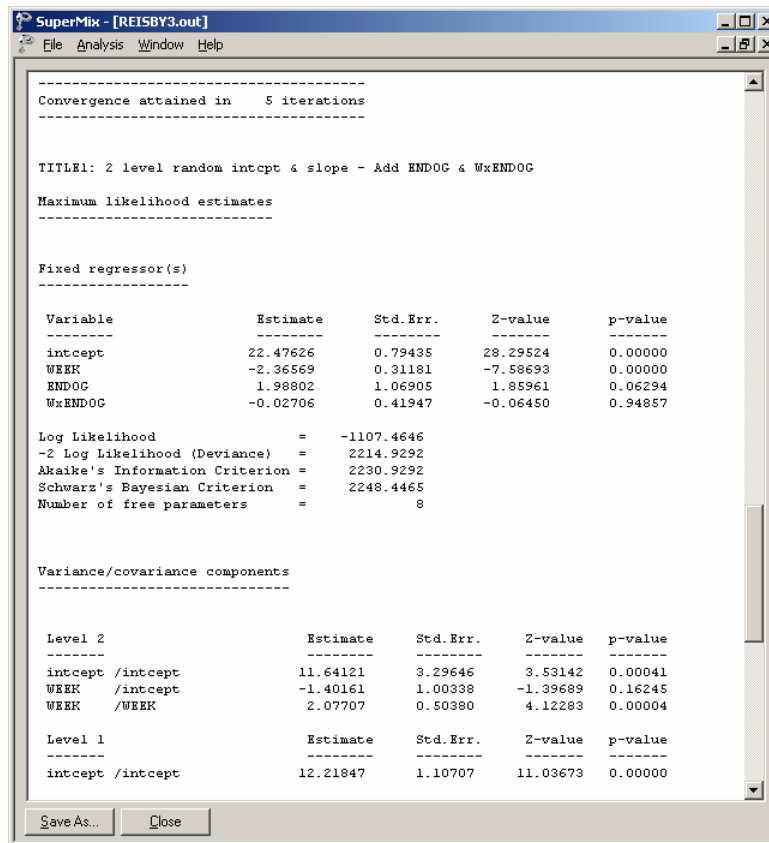
The two covariates are specified by checking the **E** check boxes for ENDOG and WxENDOG respectively in the **Available** grid respectively to produce the following **Variables** tab.

Save the changes to the file **reisby3.mum**. To fit the revised model to the data, select the **Run** option on the **Analysis** menu to produce the output file **reisby3.out**.

1.2.4.3 Interpreting the results

Fixed effects results

A portion of the output file **reisby3.out** is shown below. The interaction WxENDOG between the time variable WEEK and the depression classification variable ENDOG, is not significant. Given this, we can take a closer look at the estimated coefficients for the main effects WEEK and ENDOG respectively. Note, however, that the p -value for the ENDOG coefficient is larger than 0.05, and thus can only be considered significant at a 10% level of significance. The effect of time, on the other hand, is found to be highly significant. While the average HDRS score is predicted to decrease by -2.37 score scale units each week, patients classified as having endogenous depression (*i.e.*, ENDOG = 1) are predicted to have a HDRS score of 2 units higher at all occasions.



To obtain the predicted average HDRS scores, the estimates obtained from the output are used:

$$\hat{y} = \hat{\beta}_0 + \hat{\beta}_1(\text{WEEK}) + \hat{\beta}_2(\text{ENDOG}) + \hat{\beta}_3(\text{WxENDOG})$$

$$= 22.47626 - 2.36569(\text{WEEK}) + 1.98802(\text{ENDOG}) - 0.02706(\text{WxENDOG})$$

Model comparison

A question that arises from inspection of the results obtained thus far is whether the interaction term contributes overall to the explanation of the variation in the HDRS

scores. To test this, we can fit a model without the interaction term and use the deviance reported in the output to compare results for the model with interaction and the model without this term. The relevant output from an analysis without the interaction term is shown below. We note that the deviance obtained for the simpler model is almost identical to that of the model considered in this section. Based on this, we conclude that a model without the interaction $W \times \text{ENDOG}$ would fit the data as well as the one with the interaction term included.

SuperMix - [REISBY31.out]

File Analysis Window Help

Fixed regressor(s)

Variable	Estimate	Std.Err.	Z-value	p-value
intcept	22.49344	0.74839	30.05592	0.00000
WEEK	-2.38064	0.20859	-11.41317	0.00000
ENDOG	1.95650	0.95083	2.05769	0.03962

Log Likelihood = -1107.4667
-2 Log Likelihood (Deviance) = 2214.9334
Akaike's Information Criterion = 2228.9334
Schwarz's Bayesian Criterion = 2244.2610
Number of free parameters = 7

Save As... Close

In addition, we can test the hypothesis that the model with covariate (ENDOG) fits the data better than the random intercept and slope model considered previously. To test this hypothesis, we calculate the difference between the -2 log likelihood value obtained for the previous model and the -2 log likelihood value for the current model. It can be shown that this difference of $2219.04 - 2214.93 = 4.11$ has a χ^2 distribution with associated degrees of freedom equal to the difference in the number of parameters estimated in the two examples, *i.e.*, $8 - 7 = 1$ degrees of freedom. Since the p -value for this test statistic is less than 0.05, it is concluded that the random intercept-and-slope model with ENDOG as a covariate provides a better description of the data than the original random intercept-and-slope model. This finding is supported by the fact that the p -value for ENDOG when the interaction effect between WEEK and ENDOG is excluded equals 0.04.

1.2.5 A random intercept-and-slope quadratic model

1.2.5.1 The model

In this section we include an additional predictor and a random term to examine a possible quadratic response trend in HDRS scores over time. Keeping the level-2 model the same as before, the corresponding model for the response variable HDRS may be expressed as

Level-1 model:

$$\text{HDRS}_{ij} = b_{0i} + b_{1i} \times (\text{WEEK})_{ij} + b_{2i} \times (\text{WEEK}^2)_{ij} + e_{ij}$$

Level-2 model:

$$b_{0i} = \beta_0 + v_{0i}$$

$$b_{1i} = \beta_1 + v_{1i}$$

$$b_{2i} = \beta_2 + v_{2i}$$

1.2.5.2 Preparing the data

Create a new blank variable named WEEKSQ as shown in section XXXX. Highlight the column WEEKSQ, type the formula `SQUARE(C1)` where C = WEEK in the string field and click on the **Apply** button to produce the following screen. Save the change to **reisby.ss3**.

	(A) PATIEN	(B) HDRS	(C) WEEK	(D) WEEKC	(E) WEEKSQ	(F) ENDOG	(G) WxEN
1	101.00	26.00	0.00	-2.50	0.00	0.00	0.00
2	101.00	22.00	1.00	-1.50	1.00	0.00	0.00
3	101.00	18.00	2.00	-0.50	4.00	0.00	0.00
4	101.00	7.00	3.00	0.50	9.00	0.00	0.00
5	101.00	4.00	4.00	1.50	16.00	0.00	0.00
6	101.00	3.00	5.00	2.50	25.00	0.00	0.00
7	103.00	33.00	0.00	-2.50	0.00	0.00	0.00
8	103.00	24.00	1.00	-1.50	1.00	0.00	0.00
9	103.00	15.00	2.00	-0.50	4.00	0.00	0.00
10	103.00	24.00	3.00	0.50	9.00	0.00	0.00

1.2.5.3 Setting up the analysis

Again, we can modify the model setup file of **reisby1.mum** by first opening it, then saving the file as **reisby4.mum**. Change the title on the **Configuration** tab and request Bayes estimates by selecting the **means & (co)variances** option from the **Write Bayes Estimates** drop-down list.

Model Setup: REISBY4.mum

Configuration | Variables | Starting Values | Patterns | Advanced | Linear Transforms

Title 1: 2 level random intcpt & random slope model - quadratic trend

Title 2: REISBY Data

Dependent Variable Type: continuous | Level-2 ID: PATIENT

Dependent Variable: HDRS | Level-3 ID:

Write Bayes Estimates: means & (co)variances

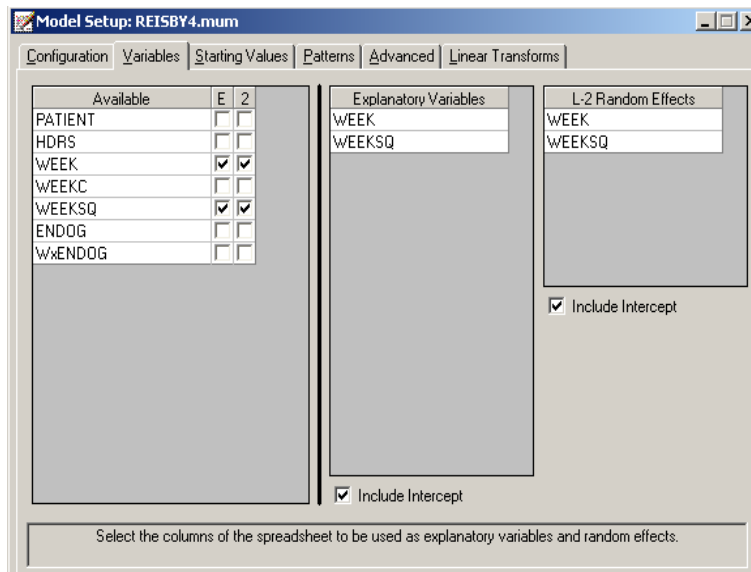
Convergence Criterion: 0.0001

Number of Iterations: 100

Missing Values Present: false | Generate Table of Means: no

Select between writing the Bayes estimates to an optional results file or suppressing them.

Next, click on the **Variables** tab to proceed to the **Variables** screen of the **Model Setup** window. The two covariates are specified by checking the **E** and **2** check boxes for WEEKSQ in the **Available** grid to produce the following **Variables** screen.



Save the changes to the file **reisby4.mum** and run the model.

1.2.5.4 Interpreting the results

A portion of the output file **reisby4.out** is shown below.

Fixed effects results

The level-1 estimate of the WEEKSQ coefficient is 0.05, which turns out not to be significant ($p = 0.56$). On the other hand, the WEEKSQ random effect is significant at a 5% level ($p = 0.04$). Comparing the present results with those reported in **reisby1.out**, we see that the deviance difference of $2219.04 - 2207.65 = 11.19$ with $10 - 7 = 3$ degrees of freedom, indicating an improved overall model fit at a 5% significance level. These results imply that, although the mean trend of HDRS scores over time is linear, some of the individuals' trajectories are quadratic.

SuperMix - [REISBY4.out]

File Analysis Window Help

Fixed regressor(s)

Variable	Estimate	Std.Err.	Z-value	p-value
intcept	23.76025	0.55206	43.03916	0.00000
WEEK	-2.63258	0.47900	-5.49603	0.00000
WEEKSQ	0.05148	0.08835	0.58272	0.56008

Log Likelihood = -1103.8239
-2 Log Likelihood (Deviance) = 2207.6479
Akaike's Information Criterion = 2227.6479
Schwarz's Bayesian Criterion = 2249.5444
Number of free parameters = 10

Variance/covariance components

Level 2	Estimate	Std.Err.	Z-value	p-value
intcept /intcept	10.44021	3.57924	2.91688	0.00354
WEEK /intcept	-0.91538	2.41817	-0.37854	0.70503
WEEK /WEEK	6.63806	2.74573	2.41759	0.01562
WEEKSQ /intcept	-0.11217	0.42143	-0.26617	0.79011
WEEKSQ /WEEK	-0.93648	0.48449	-1.93293	0.05324
WEEKSQ /WEEKSQ	0.19374	0.09391	2.06309	0.03910

Level 1	Estimate	Std.Err.	Z-value	p-value
intcept /intcept	10.51598	1.10143	9.54754	0.00000

Save As... Close

1.2.5.5 Residuals

Level 2 Bayes results

Up to this point, we have considered results averaged over all patients. We now turn our attention to the residual file **reisby4.ba2**, which offers the opportunity to take a closer look at the results by individual patient. After running the above model, select the **Analysis, View L-2 Bayes Results** option to open the image below. The contents of this file are displayed for the first 5 patients. Three lines of information are given for each patient, containing, in order of appearance,

- the number of the patient in the data set,
- the number of the empirical Bayes coefficient,

- the empirical Bayes estimate,
- the estimated variance of the Bayes coefficient, and
- the name of the associated coefficient as used in the model.

Patient ID	Week	Intercept	Slope	Quadratic Term	Label
101.00	1	1.4054	3.9425		intcept
101.00	2	-2.6506	2.6732		WEEK
101.00	3	0.99315E-01	0.10181		WEEKSQ
103.00	1	3.7472	3.9425		intcept
103.00	2	-0.80235	2.6732		WEEK
103.00	3	0.68413E-01	0.10181		WEEKSQ
104.00	1	2.2378	3.9425		intcept
104.00	2	-0.46145	2.6732		WEEK
104.00	3	-0.22836	0.10181		WEEKSQ
105.00	1	-2.7488	3.9425		intcept
105.00	2	-0.32003	2.6732		WEEK
105.00	3	0.11318	0.10181		WEEKSQ
106.00	1	-0.11679	3.9627		intcept
106.00	2	1.8882	2.7393		WEEK
106.00	3	-0.19335	0.13006		WEEKSQ

To obtain patient-specific predicted HRS scores, the empirical Bayes estimates for each patient have to be taken into account, as these estimates indicate the extent to which the random intercept or slope for that patient deviates from the intercept and slope over all patients. Patient-specific predicted HRS scores are calculated as

$$\widehat{y}_{ij} | \widehat{\beta} = 23.76025 - 2.63258 \times \text{WEEK}_{ij} + 0.05148 \times \text{WEEK}_{ij}^2 + v_{0i} + v_{1i} \times \text{WEEK}_{ij} + v_{2i} \times \text{WEEK}_{ij}^2$$

For the first patient shown in the residual file above, we have $\widehat{v}_{i0} = 1.4054$, $\widehat{v}_{i1} = -2.6506$ and $\widehat{v}_{i2} = 0.099315$. From this information, we can already tell that the intercept for the patient is higher than average, but that the WEEK slope for this patient is lower than average. The positive value of the quadratic term indicates that the decreasing rate slows down more quickly than average with an increase in time. The predicted HRS score for this patient (PATIENT = 101) is found to be

$$\widehat{y}_{ij} | \widehat{\beta} = 23.76025 - 2.63258 \times \text{WEEK}_{ij} + 0.05148 \times \text{WEEK}_{ij}^2 + 1.4054 - 2.6506 \times \text{WEEK}_{ij} + 0.099315 \times \text{WEEK}_{ij}^2$$

Substituting the WEEK with values 0, 1, ..., 5, we get the predicted HDRS scores for Patient 101, and similarly, for all the other patients. Table XXX.4 and the graphical display below give the predicted HDRS for the first 5 patients.

Table XXX.4: Predicted HDRS values for selected patients

	Patient 101	Patient 103	Patient 104	Patient 105	Patient 106	Population Avg.
Week 0	25.166	27.507	25.998	21.011	23.643	23.760
Week 1	20.033	24.192	22.727	18.224	22.757	21.179
Week 2	15.202	21.117	19.102	15.765	21.587	18.701
Week 3	10.673	18.282	15.124	13.636	20.133	16.326
Week 4	6.446	15.686	10.792	11.836	18.396	14.054
Week 5	2.520	13.330	6.106	10.365	16.375	11.884




Figure XXX.10: Predicted HDRS for selected patients

We find that Patient 101 had a higher initial HDRS score, but over time obtained a lower than average score. For Patient 103, a higher than average predicted HDRS score is obtained at each time point. In contrast, Patient 105 scored lower at each time point. The quadratic term doesn't affect much of the population average; however the effect is obvious for Patients 105 and 106.

Model-based graphs

Residual plot

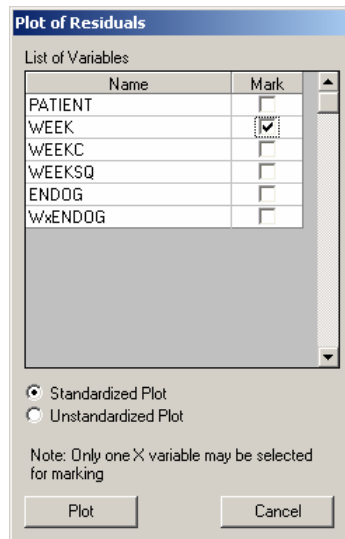
Level-1 residuals can also be obtained, either for a typical or specific patient, by using the empirical Bayes estimates. The residuals for a typical patient are obtained as

$$\begin{aligned} \text{Patient residual} &= \text{Observed HDRS score} - \hat{y} | \beta \\ &= \text{Observed HDRS score} - \left[23.76025 - 2.63258 \times \text{WEEK}_{ij} + 0.05148 \times \text{WEEK}_{ij}^2 \right] \end{aligned}$$

The residuals for a specific patient use the additional information given by the empirical Bayes residuals and have the form

$$\begin{aligned} \text{Patient-specific residual} &= \text{Observed HDRS score} - \hat{y} | \beta \\ &= \text{Observed HDRS score} - \\ &\left[23.76025 - 2.63258 \times \text{WEEK}_{ij} + 0.05148 \times \text{WEEK}_{ij}^2 + v_{0i} + v_{1i} \times \text{WEEK}_{ij} + v_{2i} \times \text{WEEK}_{ij}^2 \right] \end{aligned}$$

Select the **Residuals** option on the **File, Model-based Graphs** menu to activate the **Plot of Residuals** dialog box. Check the **Mark** check box for WEEK as shown below, then click on the **Plot** button.



The graph obtained, as shown below, shows that, in general, the range of the level-1 residuals is $(-5; 5)$.




Figure XXX.11: Plot of level-1 residuals vs. predicted values

Inspection of these residuals can be useful in examining the distributional assumptions for the level-1 data, in this case at the measurement level. For the current example, residuals for a typical patient have a mean of 0.000 with standard error of 2.66. Click on the middle of the graph to open an additional window that shows the detailed residual data for each observation.

Patient	Hdrs	Estimate	Error
101	26	25.16565	0.83435
101	22	20.033265	1.9667
101	18	15.20247	2.7975
101	7	10.673265	-3.6733
101	4	6.44565	-2.4457
101	3	2.519625	0.48038
103	33	27.50745	5.4925
103	24	24.192413	-0.19241
103	15	21.117162	-6.1172
103	24	18.281697	5.7183
103	15	15.686018	-0.68602
103	13	13.330125	-0.33012
104	29	25.99805	3.002
104	22	22.72714	-0.72714
104	18	19.10247	-1.1025
104	13	15.12404	-2.124
104	19	10.79185	8.2081
104	0	6.1059	-6.1059
105	22	21.01145	0.98855
105	12	18.2235	-6.2235
105	16	15.76487	0.23513
105	16	13.63556	2.3644
105	13	11.83557	1.1644
105	9	10.3649	-1.3649

We note that the estimate for Patient 101 at the beginning of the study was 25.166, and 2.520 at the end of the study. On both occasions, the residuals associated with these estimates were positive, indicating that the estimates are above estimated average.

1.2.6 A 2-level random intercept-and-slope model with autocorrelated errors

In the mixed models discussed so far, it was assumed that the level-1 errors are conditionally independent from each other. However, the errors could be correlated over time. Different types of correlated error structures are available in SuperMix: the first-order stationary autoregressive process, stationary AR(1), the first-order non-stationary autoregressive process, non-stationary AR(1), the first-order stationary moving average process, MA(1), the first-order stationary autoregressive moving average process, ARMA(1), and a general Toeplitz autocorrelation structure.

The stationary AR(1) and ARMA(1) use the stationary assumption, that is that the variance of errors is constant over time and that the covariance of errors from

differing times depends only on the time interval between these time points and not on the starting time point. The assumption of stationarity is relaxed in the other two types of models. In SuperMix, the maximum marginal likelihood (MML) solution at convergence is obtained by first using the EM algorithm and then Fisher scoring iterations.

1.2.6.1 The non-stationary AR(1) model

The model here is essentially the same as the one we had in section XXX, apart from the autocorrelated error term.

Level-1 model:

$$\text{HDRS}_{ij} = b_{0i} + b_{1i} \times \text{WEEK}_{ij} + b_{2i} \times \text{WxENDOG}_{ij} + e_{ij},$$

where

$$e_{ij} = \rho e_{i,j-1} + \varepsilon_{ij}$$

with ρ denoting the AR coefficient.

Level-2 model:

$$b_{0i} = \beta_0 + \beta_3 \times \text{ENDOG}_i + v_{0i}$$

$$b_{1i} = \beta_1 + v_{1i}$$

$$b_{2i} = \beta_2$$

We can rewrite the model as follows:

$$\begin{aligned} \text{HDRS}_{ij} &= \beta_0 + \beta_1 \times \text{WEEK}_{ij} + \beta_2 \times \text{ENDOG}_{ij} + \beta_3 \times \text{WEEK} \times \text{ENDOG}_{ij} \\ &\quad + u_{0i} + u_{1i} \times \text{WEEK}_{ij} + e_{ij} \\ &= \beta_0 + (\beta_1 + u_{1i}) \times \text{WEEK}_{ij} + \beta_2 \times \text{ENDOG}_{ij} + \beta_3 \times \text{WEEK} \times \text{ENDOG}_{ij} \\ &\quad + e_{ij} + u_{0i} \end{aligned}$$

The difference between the present and previous models lies in the assumption concerning the error term. Previously, we assumed that $\mathbf{e}_i = (e_{i1}, e_{i2}, \dots, e_{in_i})' \sim N(\mathbf{0}, \sigma^2 \mathbf{I}_i)$, where \mathbf{I}_i is an identity matrix of order $n_i \times n_i$. Now we assume that the errors are autocorrelated, and that $\mathbf{e}_i \sim N(\mathbf{0}, \sigma^2 \mathbf{\Omega}_i)$, where $\mathbf{\Omega}_i$ is the autocorrelation matrix.

The analysis – step 1: starting values from a non-AR model

Two types of iteration algorithms, EM and Fisher scoring, are used for fitting an autoregressive model. The EM solution proceeds by assigning starting values for the structural and population parameters. The Fisher scoring procedure utilizes the first derivatives and expected values of the second derivatives to obtain improved parameter estimates.

Although the Fisher scoring solution is a significant improvement in terms of speed of convergence over the EM solution, it can fail in the estimation of the covariance matrix of the random effects when these terms become very small. The most reliable way to minimize the chance of encountering convergence problems is first obtaining the starting values by running the model without autocorrelated errors, then substituting the starting values obtained prior to fitting the AR model.

Recall that in Section XXXX we fitted the model

Level-1 model:

$$\text{HDRS}_{ij} = b_{0i} + b_{1i} \times \text{WEEK}_{ij} + b_{2i} \times \text{WxENDOG}_{ij} + e_{ij}$$

Level-2 model:

$$b_{0i} = \beta_0 + \beta_3 \times \text{ENDOG}_i + v_{0i}$$

$$b_{1i} = \beta_1 + v_{1i}$$

$$b_{2i} = \beta_2$$

The estimates obtained for that model are repeated below.

```

-----
Convergence attained in 5 iterations
-----

TITLE: 2 level random intcpt & slope - Add ENDOG & WxENDOG

Maximum likelihood estimates
-----

Fixed regressor(s)
-----

Variable          Estimate      Std.Err.      Z-value      p-value
-----
intcpt            22.47626     0.79435       28.29524     0.00000
WEEK             -2.36569     0.31181       -7.58693     0.00000
ENDOG             1.98802     1.06905        1.85961     0.06294
WxENDOG          -0.02706     0.41947       -0.06450     0.94857

Log Likelihood          = -1107.4646
-2 Log Likelihood (Deviance) = 2214.9292
Akaike's Information Criterion = 2230.9292
Schwarz's Bayesian Criterion = 2248.4465
Number of free parameters = 8

Variance/covariance components
-----

Level 2
-----
Estimate      Std.Err.      Z-value      p-value
-----
intcpt /intcpt  11.64121     3.29646       3.53142     0.00041
WEEK /intcpt   -1.40161     1.00338       -1.39689     0.16245
WEEK /WEEK      2.07707     0.50380        4.12283     0.00004

Level 1
-----
Estimate      Std.Err.      Z-value      p-value
-----
intcpt /intcpt  12.21847     1.10707       11.03673     0.00000
  
```

The level-2 estimated variance of intercept and WEEK are 11.64121 and 2.07707 respectively. The estimated level-2 covariance is -1.40161. The estimated level-1 variance is 12.21847. These numbers will be used as the starting values in the non-stationary AR model to be fitted next.

The analysis – step 2: non-stationary AR model

We modify the model setup file, **reisby3.mum**, by first saving the file as **reisby_ar2.mum**. Change the title on the **Configuration** screen. Keep the settings of the **Variables** tab the same as before.

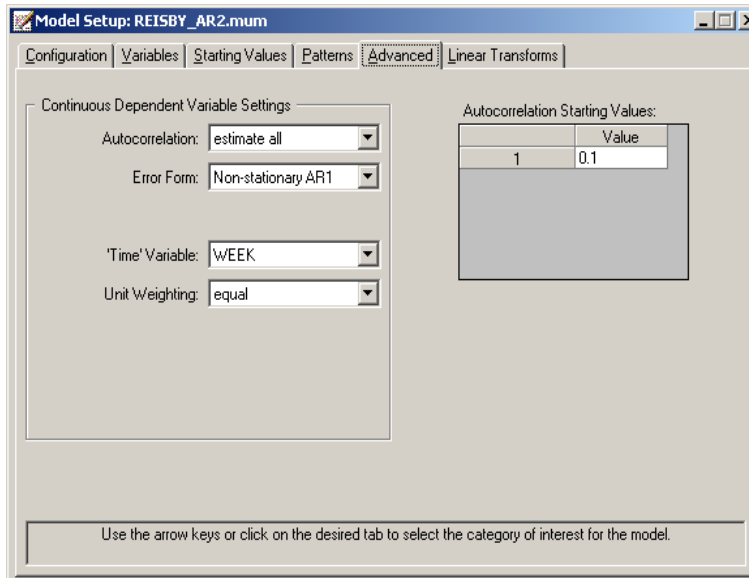
Click on the **Starting Values** tab. Select the **user-defined** option from the **Starting Values** drop down list to activate the grid fields for the starting values. Input the starting values we obtained from **reisby.out** to generate the following screen.

The screenshot shows the 'Starting Values' tab of a software interface. At the top, there are tabs for 'Configuration', 'Variables', 'Starting Values', 'Patterns', 'Advanced', and 'Linear Transforms'. Below the tabs, there is a dropdown menu for 'Starting Values' set to 'user-defined' and a text field for 'Starting Error Variance' with the value '12.218'. Below this, there are two tables: 'Explanatory Variables' and 'Level-2 (Co)variance Starting Values'. The 'Explanatory Variables' table has columns 'Variable' and 'Value' with rows for 'intcept', 'Week', 'ENDOG', and 'WxENDOG'. The 'Level-2 (Co)variance Starting Values' table has columns 'Variable' and 'Value' with rows for 'intcept variance', 'intcept, Week', and 'Week variance'. At the bottom of the window, there is a text box with the instruction: 'Enter the starting values for the variances and covariances of the level-2 random effects.'

Variable	Value
intcept	22.47626
Week	-2.36569
ENDOG	1.99802
WxENDOG	-0.02706

Variable	Value
intcept variance	11.641
intcept, Week	-1.402
Week variance	2.077

Click on the **Advanced** tab to proceed to the **Advanced** screen. First, select the **estimate all** option from the **Autocorrelation** drop down list; then select **Non-stationary AR1** as the **Error Form** and specify **WEEK** as the '**Time**' Variable. Input 0.1 in the **Autocorrelation Starting Values** grid field to get the **Advanced** screen as shown below.



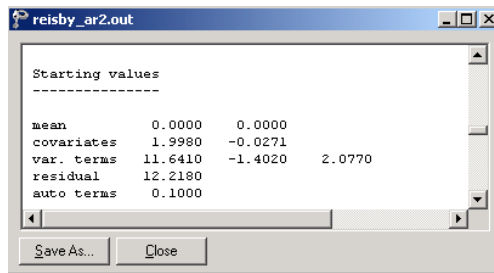
Save the changes to **reisby_ar2.mum** and run the model to produce the output file **reisby_ar2.out**.

1.2.6.2 Interpreting the output

The output for the AR model first shows the syntax information of the model setup. The number of observations, hierarchical structure of the 2-level model and descriptive statistics follow next.

The starting values


The starting values could either be user-defined or program generated. In our case the user-defined starting values are listed below.



The **Starting values** section in the output file corresponds with the starting values we specified in the **Starting Values** and **Advanced** screens. The **mean** row refers to the starting values for the fixed regressors, which are intercept and WEEK in this example. The **covariates** row contains the starting values for ENDOG and WxENDOG. The elements of the **var. terms** row are the starting values for the level-2 variance/covariance matrix. The **residual** value is the starting error variance. The **auto term(s)** is the autocorrelation starting value(s).

The maximum marginal likelihood (MML) estimates

The starting values section is followed by the **Final Results**. The maximum marginal likelihood (MML) solution at convergence is obtained by first using EM algorithm and then Fisher scoring iterations. The AIC, SBC and -2 log likelihood (deviance) are given right below the iteration information.



As shown above, the convergence is obtained after 10 EM and 10 Fisher iterations. The log likelihood value can be used to perform likelihood ratio tests.

For each model parameter, maximum marginal likelihood estimates, standard errors, z-values, and p -values are provided. These p -values are two-tailed, except for the variance parameters where one-tailed p -values are given.

Variable	Estimate	Stand. Error	Z	p-value
intcept	22.47646	0.78704	28.55812	0.00000
WEEK	-2.33888	0.30299	-7.71930	0.00000
ENDOG	1.85677	1.05917	1.75304	0.07960
WxENDOG	-0.01205	0.40784	-0.02954	0.97643
Random-effect variance & covariance term(s)				
intcept	6.05178	4.01616	1.50686	0.06592
covariance	-0.12213	1.06412	-0.11477	0.90863
WEEK	1.42305	0.52640	2.70335	0.00343
Residual variance				
	15.26661	1.86228	8.19781	0.00000
Autocorrelation term(s)				
	0.33247	0.11190	2.97108	0.00297
note: p-values are 2-tailed except for variances which are 1-tailed				

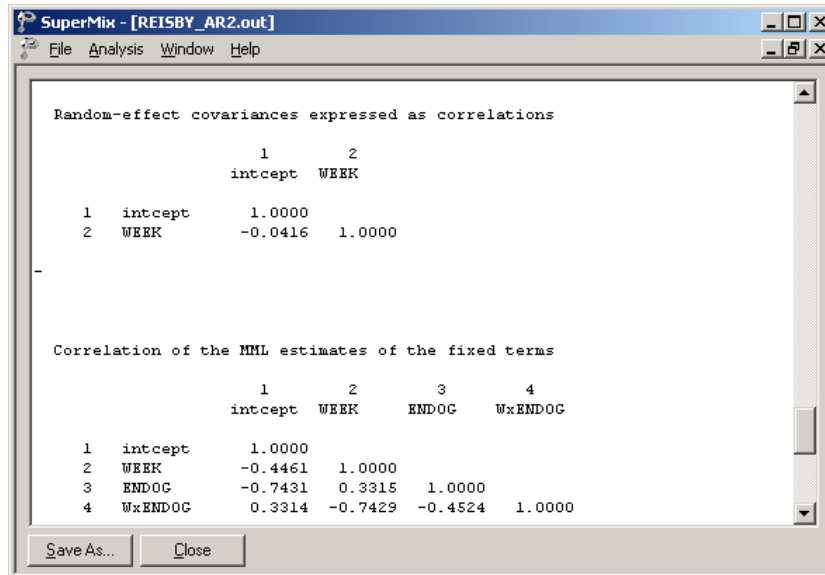
Considering the estimated fixed effects, the initial level of severity for non-endogenous patients is approximately 22.5 on the HDRS, while the endogenous patients start about 1.9 units higher. The difference in initial severity is almost significant ($p < 0.0790$). The reason that the intercept and endogenous effect reflect HDRS levels at week 0 is due to the coding of WEEK that was used, namely, 0 to 6. Using other codings of WEEK would change the meaning of these regression coefficients.

Both groups exhibit an overall weekly rate of improvement of roughly 2.3 units which is highly significant. In terms of the random-effect variance and covariance terms, there is a significant rate of improvement ($p < 0.00343$). The variation in

patients' initial severity is marginally significant at 0.066. However, the overall covariation between those two terms are significant at a 10% level ($p > 0.90864$).

Correlation of the MML estimates

Finally, correlation matrices are also provided for the estimates of all model parameters. It is important to realize that these correlation matrices are not correlations of the variables themselves, but correlations of the estimated model parameters. These matrices may be helpful in determining the degree to which collinearity is present in terms of the model parameters.



It is interesting to note that, when the correlations are rounded to two decimal places, equalities exist between the correlations:

$$\begin{aligned}
 r(\text{INTCEPT}, \text{ENDOG}) &= r(\text{WEEK}, \text{WxENDOG}) = -0.74 \\
 r(\text{INTCEPT}, \text{WEEK}) &= r(\text{ENDOG}, \text{WxENDOG}) = -0.45 \\
 r(\text{ENDOG}, \text{WEEK}) &= r(\text{INTCEPT}, \text{WxENDOG}) = 0.33
 \end{aligned}$$

From the fixed effect results we see that the WxENDOG effect was not significant ($p = 0.976$). It is reasonable to assume that, with the interaction term omitted from the model, the correlations between the intercept, ENDOG and WEEK coefficients will be close to those reported above.

Level-2 Bayes results

The residual file **reisby_ar2.ba2** offers the opportunity to take a closer look at the results by individual patient. Select the **Analysis, View L-2 Bayes Results** option to open the Bayes results as shown below. The contents of this file are displayed for the first 7 patients. Two lines of information are given for each patient, containing, in order of appearance,

- the number of the patient in the data set,
- the number of the empirical Bayes coefficient,
- the empirical Bayes estimate,
- the estimated variance of the Bayes coefficient, and
- the name of the associated coefficient as used in the model.

101	6	22.77116	3.59726	intcept
101	6	-3.72001	0.48751	WEEK
103	6	25.13207	3.59726	intcept
103	6	-2.40102	0.48751	WEEK
104	6	23.27914	3.59726	intcept
104	6	-3.61190	0.48751	WEEK
105	6	21.60314	3.59726	intcept
105	6	-2.43564	0.48751	WEEK
106	5	22.43105	3.59803	intcept
106	5	-1.58380	0.65885	WEEK
107	5	21.54122	3.60701	intcept
107	5	-2.88802	0.48829	WEEK
108	6	20.75563	3.59726	intcept
108	6	-3.16081	0.48751	WEEK

The user can obtain patient-specific predicted HDRS scores using the empirical Bayes estimates for each patient by using the method discussed earlier in section XXX.

Model comparison

In Table XXX.5, the estimates of the regular model without an autoregressive term and the non-stationary AR(1) are summarized. Note that the AIC and BIC values obtained from the AR(1) model were multiplied by -2.0 in order to facilitate comparison over the models.

TABLE XXX.5: Comparison of models with and without AR(1) term

	no AR term	Non-stationary AR(1)
intcept	22.47626 (0.79435)	22.47646 (0.78704)
WEEK	-2.36569 (0.31181)	-2.33888 (0.30299)
ENDOG	1.98802 (1.06905)	1.85677 (1.05917)
WxENDOG	-0.02706 (0.41947)	-0.01205 (0.40784)
Log Likelihood	-1107.4646	-1103.72
Akaike's Information Criterion	2214.9292	2217.44
Schwarz's Bayesian Criterion	2230.9292	2237.076
-2 Log Likelihood	2248.4465	2207.441
Number of free parameters	8	9

We notice that the estimates of both models are close to each other. The estimated variances of the non-stationary AR(1) model are smaller for all the parameters. The deviance is $2248.4465 - 2207.441 = 41.0055$ with 1 degree of freedom, which is highly significant. Thus, we conclude that in this example, the non-stationary AR(1) model fits the data better.

1.3 Models based on the TVSFP data

1.3.1 The data

The data set used here is from the Television School and Family Smoking Prevention and Cessation Project (TVSFP)(Flay *et. al.*, 1988). The study was designed to test independent and combined effects of a school-based social-resistance curriculum and a television-based program in terms of tobacco use and cessation. The data from the study included a total of 1,600 students from 135 classrooms drawn from 28 schools. Schools were randomized to one of four study conditions:

- a social-resistance classroom curriculum
- a media (television) intervention
- a social-resistance classroom curriculum combined with a mass-media intervention, and
- a no-treatment control group

A tobacco and health knowledge scale (THKS) was used in classifying subjects as knowledgeable or not. In its original form, the student's score was defined as the number of correct answers to seven items on tobacco and health knowledge.

While the structure of this study indicates a three-level hierarchical structure, the present application uses these data to fit a two-level model, with students nested within either classes or schools, in order to present an introduction to the analysis of ordinal outcomes. A 3-level model is presented in Section XXX.

Data for the first 10 students on most of the variables used in this section are shown below in the form of an SuperMix spreadsheet file, named **TVSFP.ss3**.

	A	B	C	D	E	F	G
1	SCHOOL	CLASS	POSTTHKS	PRETHKS	CC	TV	CCxTV
2	403	403101	3	2	1	0	0
3	403	403101	4	4	1	0	0
4	403	403101	3	4	1	0	0
5	403	403101	4	3	1	0	0
6	403	403101	4	3	1	0	0
7	403	403101	3	4	1	0	0
8	403	403101	2	2	1	0	0
9	403	403101	4	4	1	0	0
10	403	403101	5	5	1	0	0

The variables of interest are:

- SCHOOL indicates the school a student is from (28 schools in total).
- CLASS identifies the classroom (135 classrooms in total).
- POSTTHKS represents the post-intervention tobacco and health knowledge scale. It is treated as a continuous variable in the examples in this chapter. See Chapters XXX & XXX for examples where POSTTHKS is treated as a binary or ordinal outcome.
- PRETHKS indicates the pre-intervention THKS score.
- CC is a binary variable indicating whether a social-resistance classroom curriculum was introduced, where 0 indicates “no” and 1 “yes.”
- TV is an indicator variable for the use of media (television) intervention, with a “1” indicating the use of media intervention, and “0” the absence thereof.
- CCxTV was constructed by multiplying the variables TV and CC, and represents the CC by TV interaction.

1.3.1.1 Exploring the data

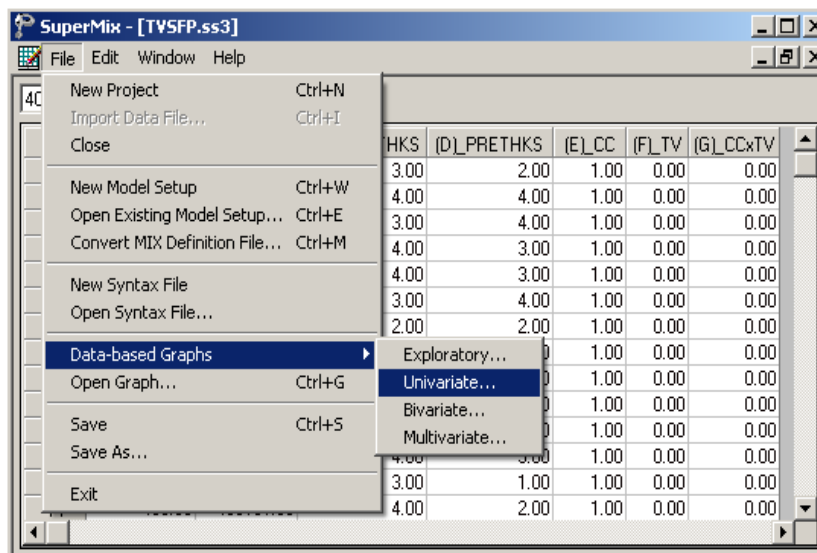
In this section, a univariate bar chart and a bivariate box-and-whisker plot are given. More information on other types of plots available are given in Chapter XXX.

Univariate graphs

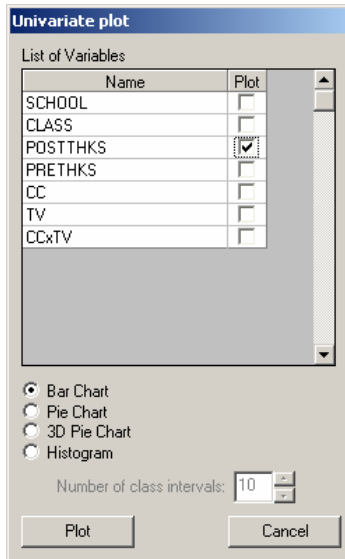
The pop-up menu below shows the data-based graphing options currently available in SuperMix. As a first step, we will take a closer look at the distribution of the total post-intervention scores (POSTTHKS), which is the potential dependent variable in this study. While scores such as these are not truly continuous variables, they are often treated as if they were.

Bar chart

To do so, select the **Univariate** option from the **Data-based Graphs** menu as shown below.



The **Univariate plot** dialog box appears. Select the variable POSTTHKS and indicate that a **Bar Chart** is to be graphed. Click the **Plot** button to display the bar chart.



The bell-shaped bar chart below shows that the variable POSTTHKS is approximately normally distributed. Note that histograms are usually used for the depiction of the distribution of a continuous variable.




Figure XXX.12: Bar chart of POSTTHKS scores

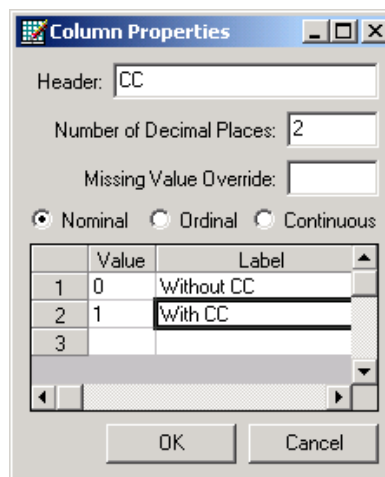
Bivariate graphs

It is hoped that the social-resistance classroom curriculum (CC), the television intervention (TV) and the CC and TV interaction combination (CCxTV) would affect the tobacco and health knowledge (POSTTHKS). Before we start with the model, we would like to show a box-and-whisker plot of POSTTHKS for each category of CC.

Box-and-whisker plots

A box-and-whisker plot is useful for depicting the locality, spread and skewness of variables in a data set and may be used to examine the distributions of continuous variables, such as for the different values of discrete valued predictors. This option is accessed via the **Data-based Graphs, Bivariate** option on the **File** menu.

To assign labels to the categories of CC, right-click on the CC column in the spreadsheet and select **Column Properties**. On the **Column Properties** dialog box, select the **Nominal** option and assign the appropriate labels.



The **Bivariate plot** dialog box is completed as shown below: select the outcome variable POSTTHKS as the **Y**-variable of interest, and the predictor CC to be plotted on the **X**-axis. Check the **Box and Whisker** option, and click **Plot**.


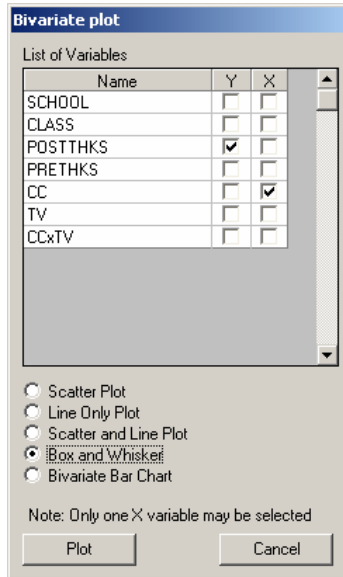


Figure XXX.13: Box-and-whisker plots of POSTTHKS scores for different CC values

The bottom line of a box represents the first quartile (q_1), the top line the third quartile (q_3), and the in-between line the median (me). The arithmetic mean is represented by a diamond. Here, the mean of POSTTHKS is lower in the group without the social-resistance classroom curriculum (CC). The box-and-whisker plot indicates a positive relationship between CC and POSTTHKS.

1.3.2 A 2-level random intercept model using classroom as level-2 ID

1.3.2.1 The model

The first model fitted to the data explores the cluster effects of each classroom on the outcome. The mixed model can be expressed as

$$\text{POSTTHKS}_{ij} = \beta_0 + \beta_1 \text{CC}_i + \beta_2 \text{TV}_i + \beta_3 (\text{CC}_i \times \text{TV}_i) + v_{0i} + \varepsilon_{ij},$$

where v_{0i} represents the classroom influence on POSTTHKS. To understand the model better, we can rewrite the model in the following way. The level-1 or within-cluster model is shown below.

Level-1 model: ($j = 1, \dots, n_i$)

$$\begin{aligned} \text{POSTTHKS}_{ij} &= b_{0i} + \varepsilon_{ij}, \\ \varepsilon_{ij} &\sim \text{NID}(0, \sigma^2) \end{aligned}$$

The level-1 model estimates POSTTHKS as a function of the intercept b_{0i} and error term ε_{ij} . Subscript i denotes the subscript for classroom, while subscript j refers to the student j . n_i is used to denote the number of students in each classroom. Because we have different numbers of students in different classrooms, n_i also varies. In this data set, $1 \leq n_i \leq 28$.

The level-2, or between-cluster, model describes the intercept b_{0i} as a function of cluster characteristics.

Level-2 model: ($i = 1, \dots, N$)

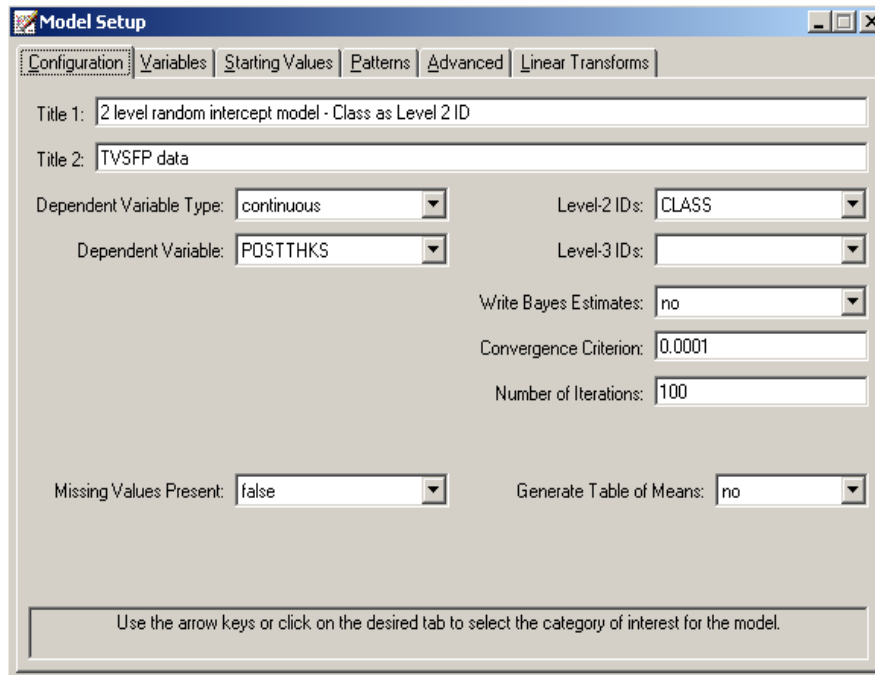
$$b_{0i} = \beta_0 + \beta_1 CC_i + \beta_2 TV_i + \beta_3 (CC_i \times TV_i) + v_{0i}$$
$$v_{0i} \sim NID(0, \sigma_v^2)$$

As shown above, the intercept b_{0i} is estimated as a function of the population average β_0 , the covariates CC_i , TV_i , and $CC_i \times TV_i$, and the classroom difference v_{0i} . The coefficient v_{0i} represents the amount that unit i deviates from the average β_0 , after controlling for the effects of the covariates included. The level-2 residual v_{0i} is assumed to follow $NID(0, \sigma_v^2)$ for all the is . If $v_{0i} = 0$ for all i , which implies $\sigma_v^2 = 0$, the model is the same as the ordinary regression model.

1.3.2.2 Setting up the analysis

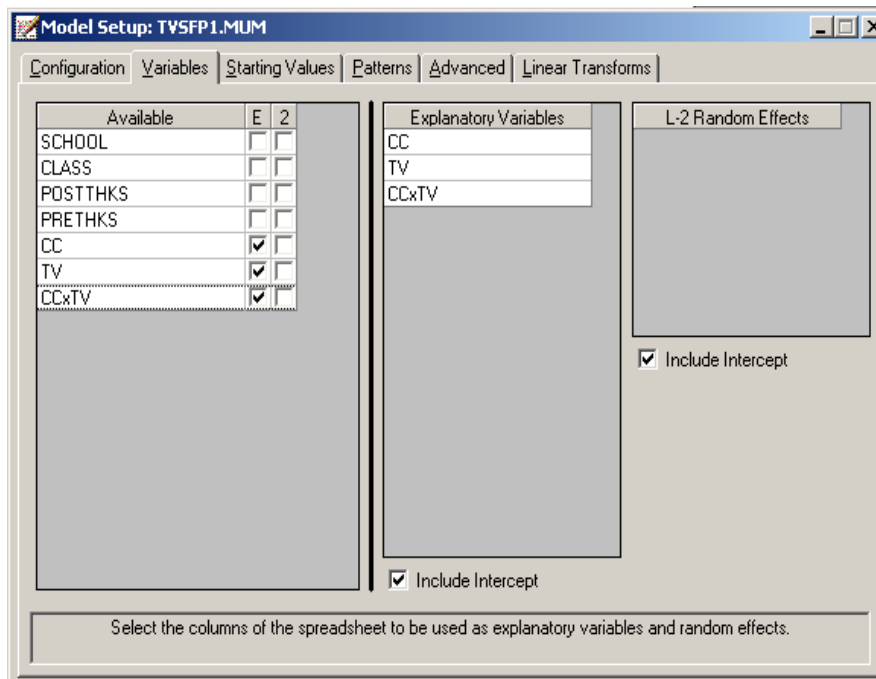
Open the SuperMix spreadsheet **TVSFP.ss3** used during the exploratory analysis discussed previously in this chapter. The next step is to describe the model to be fitted. We use the SuperMix interface to provide the model specifications. From the main menu bar, select the **File, New Model Setup** option.

Select the continuous outcome variable POSTTHKS from the **Dependent Variable** drop-down list box. Select the classroom number CLASS from the **Level-2 IDs** drop-down list box. Enter a title for the analysis in the **Title** text boxes. In this example, default settings for all other options associated with the **Configuration** screen are used.



Proceed to the **Variables** screen by clicking on that tab. The **Variables** screen is used to specify the fixed and random effects to be included in the model. Select the explanatory (fixed) variables using the **E** check boxes next to the variables names in the **Available** grid at the left of the screen. Note that, as the variables are selected, the selected variables are listed in the **Explanatory Variables** grid. After selecting all the explanatory variables, the screen shown below is obtained. The **Include Intercept** check box in the **Explanatory Variables** grid is checked by default, indicating that an intercept term will automatically be included in the fixed part of the model.

Next, specify the random effects at level 2 the hierarchy. In this example, we want to fit a model with random intercepts at level 2. By default, the **Include Intercept** check box in the **L-2 Random Effects** grid is checked. If this box is left checked, and no additional random effects are indicated using the **2** column in the **Available** grid to the left, the model fitted will be the random-intercepts-only model we intend to use. No further changes on this screen are necessary.



Before running the analysis, the model specifications have to be saved. Select the **File, Save As** option, and provide a name (**TVSFP1.mum**) for the model specification file. Run the analysis by selecting the **Run** option from the **Analysis** menu.

1.3.2.3 Discussion of results

Program information and syntax

Data summary

In the **numbers of observations** section, a summary of the hierarchical structure is provided.

As shown below, data from a total of 1600 students within 135 classrooms were included at levels 2 and 1 of the model. This corresponds to the study design described earlier. In addition, a summary of the number of students nested within

each classroom is provided. The classroom with $N2 = 6$, for example, had 26 students ($N1: 26$). By contrast, classroom 26 had only 1 student.

SuperMix - [TVSFP1.out]
File Analysis Window Help

Numbers of observations

Level 2 observations = 135
Level 1 observations = 1600

N2	:	1	2	3	4	5	6	7	8
N1	:	20	3	11	9	5	26	11	10
N2	:	9	10	11	12	13	14	15	16
N1	:	15	12	12	10	21	10	17	19
N2	:	17	18	19	20	21	22	23	24
N1	:	2	4	21	16	15	13	2	14
N2	:	25	26	27	28	29	30	31	32
N1	:	13	1	12	18	21	17	16	15
N2	:	33	34	35	36	37	38	39	40
...	:

Save As... Close

Descriptive statistics and starting values

Next, the **descriptive statistics for all variables** are given.

Descriptive statistics for all variables

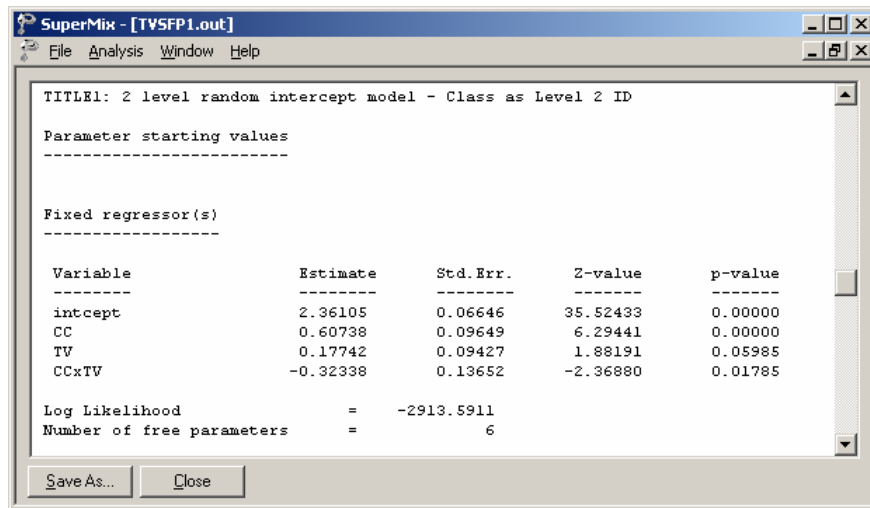
Variable	Minimum	Maximum	Mean	Stand. Dev.
Dependent				
POSTTHKS	0.00000	7.00000	2.66188	1.38293
Random-Effects				
intcept (2)	1.00000	1.00000	1.00000	0.00000
intcept (1)	1.00000	1.00000	1.00000	0.00000
Fixed Regressor(s)				
intcept	1.00000	1.00000	1.00000	0.00000
CC	0.00000	1.00000	0.47687	0.49962
TV	0.00000	1.00000	0.49938	0.50016
CCxTV	0.00000	1.00000	0.23938	0.42684

Buttons: Save As..., Close

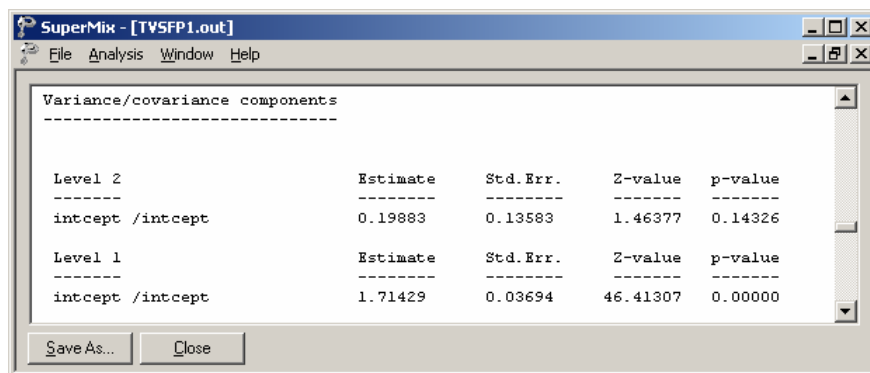
The minimum value, maximum value, mean and standard deviation are given for all the variables included in the model. For example, the mean POSTTHKS is 2.6618 with a standard deviation of 1.38293.

Starting values – OLS estimates

The starting values for the **fixed regressor(s)** are shown below. The **log likelihood** value and **number of free parameters** of the OLS regression are given in this part of the output.

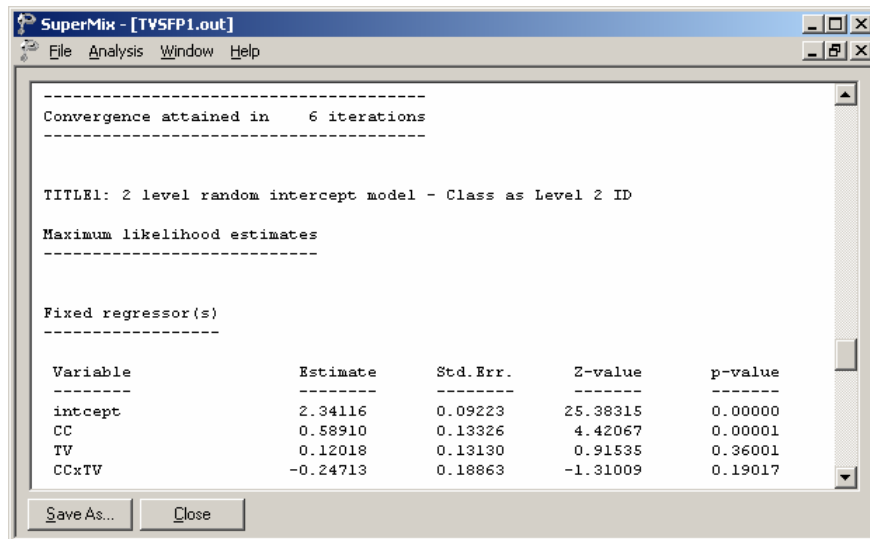


After the **number of free parameters**, the starting values of **variance/covariance components** are reported as shown.



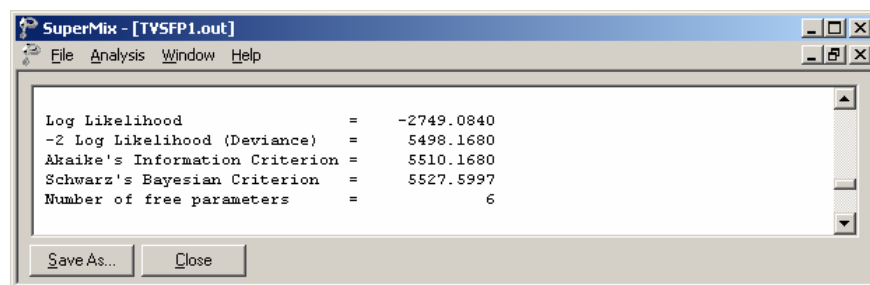
Fixed effects estimates

The number of iterations needed to obtain convergence is given after the starting values. The output describing the estimated **fixed regressor(s)** after convergence is shown next.



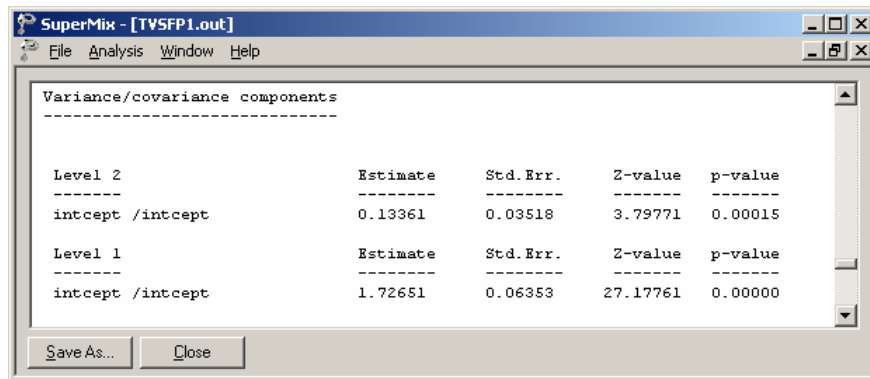
As shown above, the estimates for CC and TV are both positive. On average, a social-resistance classroom curriculum can improve the tobacco and health knowledge by 0.58910, and television intervention can increase the POSTTHKS score by 0.12018. However the estimate of CCxTV is negative, which implies that the students who had both CC and TV are expected to show a decrease of 0.24713 in their POSTTHKS score. The estimates associated with intercept and TV are highly significant, but estimates of the other two coefficients are not statistically significantly different from zero.

The estimates for the fixed regressors and model fit statistics are given next. For more information on these statistics, see Section XXX.



Random effect estimates

The estimates for the random part of the model are reported next. The variation in the average estimated intercept at level 2 is highly significant, which indicates that the classroom difference in intercepts does help to explain the variation in POSTTHKS scores.



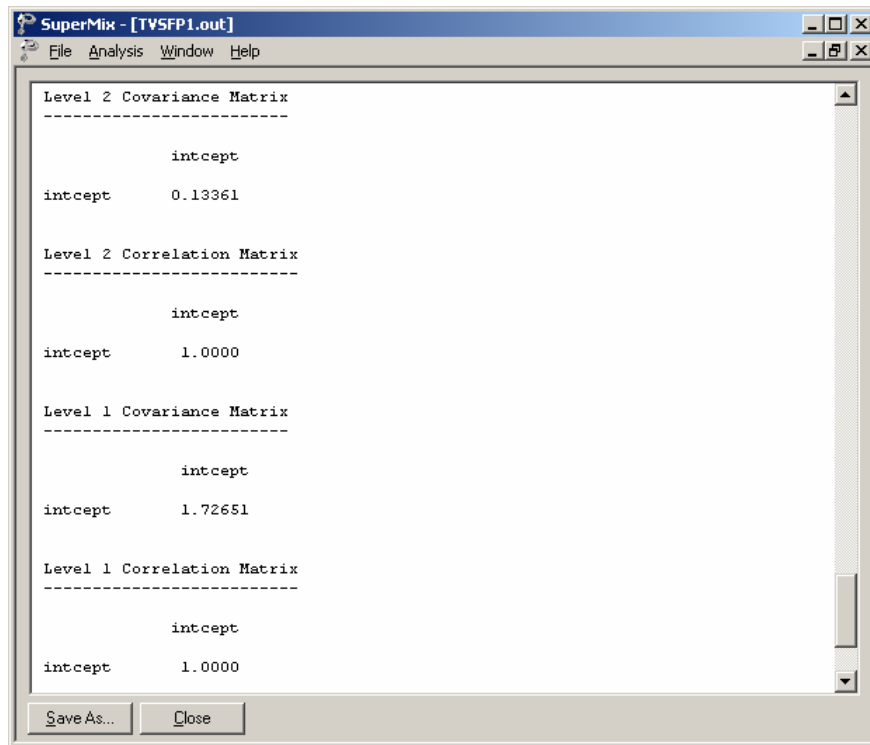
The screenshot shows a window titled "SuperMix - [TVSFP1.out]" with a menu bar (File, Analysis, Window, Help). The main content area displays the following text:

```
Variance/covariance components
-----
Level 2
-----
intcept /intcept      Estimate   Std.Err.   Z-value   p-value
                        0.13361   0.03518   3.79771   0.00015

Level 1
-----
intcept /intcept      Estimate   Std.Err.   Z-value   p-value
                        1.72651   0.06353   27.17761  0.00000
```

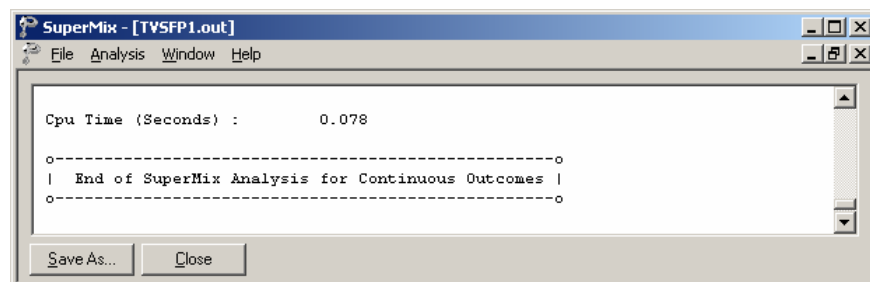
At the bottom of the window are two buttons: "Save As..." and "Close".

The covariance and correlation matrix of level-2 and level-1 random effects are given at the end of the output file. These values are the same as the estimates of variance/covariance components as shown above.



End of the output

After successfully running a SuperMix model, the following message is shown at the end of the output file to indicate the CPU time and the type of the outcome variable.



1.3.2.4 Interpreting the results

Estimated outcomes for different groups

For a student who participated in neither social-resistance classroom curriculum nor television intervention ($CC = 0$; $TV = 0$), the expected POSTTHKS is equal to just the intercept 2.36105. For a student who participated in both programs ($CC = 1$; $TV = 1$; $CC \times TV = 1$), the predicted POSTTHKS is calculated as follows:

$$\begin{aligned}\widehat{POSTTHKS}_{ij} &= \hat{\beta}_0 + \hat{\beta}_1 CC_i + \hat{\beta}_2 TV_i + \hat{\beta}_3 (CC_i \times TV_i) \\ &= 2.34116 + 0.5891 + 0.12018 - 0.24713 \\ &= 2.80331\end{aligned}$$

Fit statistics and % variation explained

An estimate of the percentage of variation in the outcome at classroom level is obtained as

$$\frac{0.13361}{0.13361 + 1.72651} \times 100\% = 7.18\%$$

indicating that about 7.18% of the total variance lies between the clusters/classrooms and that 92.82% of the variance remains at the student level.

1.3.3 2-level random intercept model by using school as level-2 ID

The model in the previous section shows that only about 7% of the total variation in outcome is at the classroom level. The question that arises is whether clustering within schools may provide a better explanation of the way in which post-intervention scores vary. In this section, the model is fitted using SCHOOL, rather than classroom, as the level-2 ID.

1.3.3.1 The model

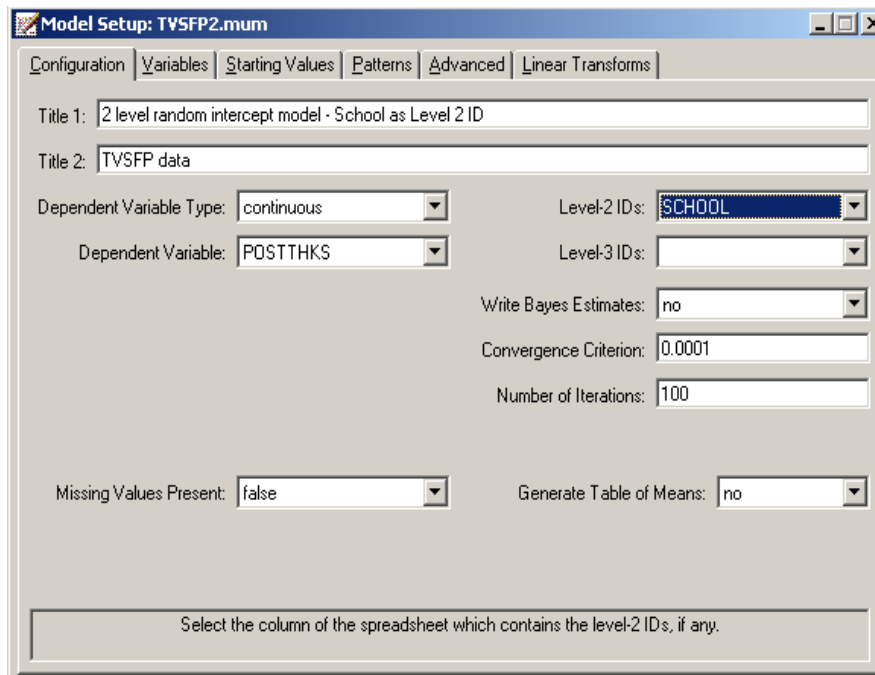
The mathematical equation of the model to be fitted is exactly the same as for the previous model.

$$\text{POSTTHKS}_{ij} = \beta_0 + \beta_1 \text{CC}_i + \beta_2 \text{TV}_i + \beta_3 (\text{CC}_i \times \text{TV}_i) + v_{0i} + \varepsilon_{ij},$$

The difference here is in the meaning of the subscript i . In the previous model, we used i to refer the classroom. However, the i s here refer to the schools.

1.3.3.2 Setting up the analysis

To create the model specifications for this model, we start by opening **TVSFP.ss3** in a SuperMix spreadsheet window. Then we use the **Open Existing Model Setup** option on the **File** menu to load the **Model Setup** window for **TVSFP1.mum**. Click on **File, Save as** to save the model setup in a new file, such as **TVSFP2.mum**. Next, change the string in the **Title 1** text box on the **Configuration** screen, and select **SCHOOL** as the **Level-2 ID** as shown below.



Keep all the other settings unchanged. Save the changes to the file **TVSFP2.mum** and select the **Run** option on the **Analysis** menu to produce the output file **TVSFP2.out**.

1.3.3.3 Discussion of results

Data summary

The **number of observations** section clearly shows that the data set contains 28 schools and each school has between 18 and 137 students as shown below.

Numbers of observations									

Level 2 observations =		28							
Level 1 observations =		1600							
N2	:	1	2	3	4	5	6	7	8
N1	:	23	25	26	70	31	42	52	55
N2	:	9	10	11	12	13	14	15	16
N1	:	39	33	52	65	27	80	33	18
N2	:	17	18	19	20	21	22	23	24
N1	:	34	38	67	73	70	74	82	114
N2	:	25	26	27	28				
N1	:	113	33	94	137				

Fixed effects estimates and descriptive statistics

The estimates for the fixed estimates are shown below. They are close to the estimates in the previous example, but not exactly the same. For example, the estimate for CC increased by 0.06326 ($0.65236 - 0.58910 = 0.06326$), and the estimate for the effect of television intervention is about 0.07811 higher when using school as the level-2 ID ($0.19829 - 0.12018 = 0.07811$). However the estimate of the interaction of CC and TV is about 0.17 lower.

SuperMix - [TVSFP2.out]

File Analysis Window Help

TITLE1: 2 level random intercept model - School as Level 2 ID

Maximum likelihood estimates

Fixed regressor(s)

Variable	Estimate	Std. Err.	Z-value	p-value
intcept	2.36132	0.12430	18.99641	0.00000
CC	0.65236	0.17828	3.65923	0.00025
TV	0.19829	0.17453	1.13611	0.25591
CCxTV	-0.41737	0.24951	-1.67280	0.09437

Save As... Close

Both the deviance and Akaike information criterion (AIC) are slightly higher than the previous model. The SBC is smaller.

SuperMix - [TVSFP2.out]

File Analysis Window Help

Log Likelihood	=	-2750.7190
-2 Log Likelihood (Deviance)	=	5501.4379
Akaike's Information Criterion	=	5513.4379
Schwarz's Bayesian Criterion	=	5521.4312
Number of free parameters	=	6

Save As... Close

Random effect estimates and covariance/correlation matrices

The estimates for the random part of the model are reported next.

The screenshot shows a window titled "SuperMix - [TVSFP2.out]" with a menu bar (File, Analysis, Window, Help). The main content area is titled "Variance/covariance components" and contains two tables. The first table is for "Level 2" and the second is for "Level 1". Both tables have columns for "Estimate", "Std. Err.", "Z-value", and "p-value".

Variance/covariance components				

Level 2	Estimate	Std. Err.	Z-value	p-value

intcept /intcept	0.07131	0.02868	2.48639	0.01290

Level 1	Estimate	Std. Err.	Z-value	p-value

intcept /intcept	1.78756	0.06374	28.04315	0.00000

The variation in the average estimated intercept at level 2 is highly significant, which indicates that the difference in school intercepts also explains the variation of POSTTHKS scores. Similarly, we can calculate that about 3.84% of the total variance can be explained by the school difference:

$$\frac{0.07131}{0.07131 + 1.78756} \times 100\% = 3.84\% .$$

1.3.4 A 3-level random intercept model using class and school as IDs

The previous two models show that both school and classroom contribute to the explanation of the total variation of the POSTTHKS scores. We now construct a three-level model that uses both CLASS and SCHOOL as level-2 and level-3 IDs.

1.3.4.1 The model

The level-1 and level-2 models are the same as the previous two models, as shown below.

Level-1 model ($k = 1, \dots, n_{ij}$)

$$\text{POSTTHKS}_{ijk} = b_{0ij} + \varepsilon_{ijk} ,$$

$$\varepsilon_{ijk} \sim NID(0, \sigma^2)$$

Level-2 model ($j = 1, \dots, n_i$)

$$b_{0ij} = b_{0i} + b_{1i}CC_{ij} + b_{2i}TV_{ij} + b_{3i}(CC_{ij} \times TV_{ij}) + v_{0ij}$$

$$v_{0ij} \sim NID(0, \sigma_{v(2)}^2)$$

Level-3 model ($i = 1, \dots, N$)

$$b_{0i} = \beta_0 + v_{0i}$$

$$b_{1i} = \beta_1$$

$$b_{2i} = \beta_2$$

$$b_{3i} = \beta_3$$

$$v_{0i} \sim NID(0, \sigma_{v(3)}^2)$$

In this mixed model the intercept b_{0ij} is estimated by a level-2 equation. It indicates that classroom j 's initial value is not only determined by the population average b_{0i} , but also by the classroom difference v_{0ij} . The level-2-intercept b_{0ij} is estimated by a level-3 equation which takes the school difference v_{0i} into consideration, where i denotes the school ID.

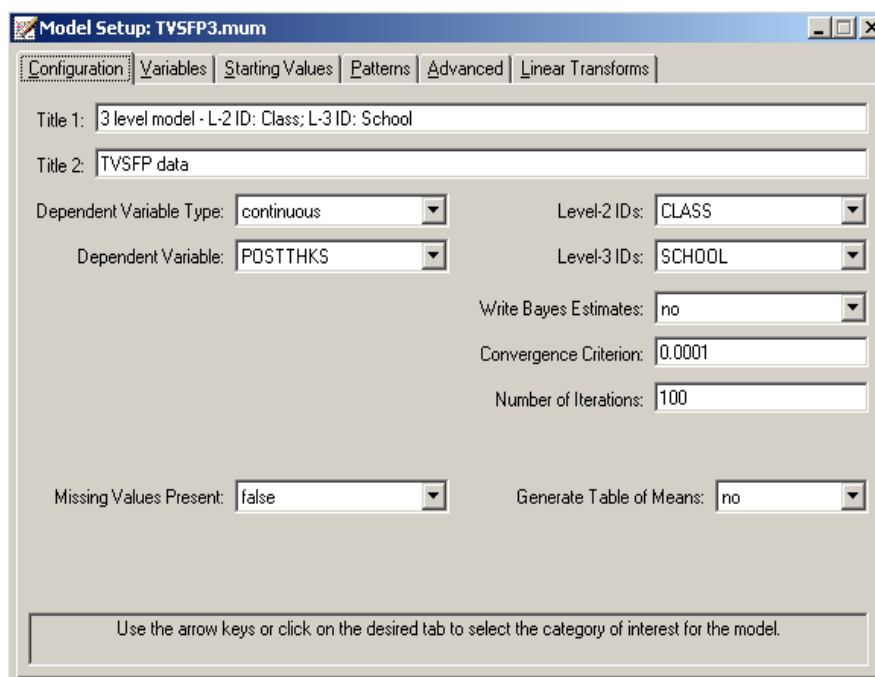
The above model can also be written in the following format.

$$POSTTHKS_{ijk} = \beta_0 + \beta_1 CC_{ij} + \beta_2 TV_{ij} + \beta_3 (CC_{ij} \times TV_{ij}) + v_{0ij} + v_{0i} + \varepsilon_{ijk}.$$

1.3.4.2 Setting up the analysis

We modify our model setup saved to the syntax file **TVSFP1.mum** by first using the **Open Existing Model Setup** option on the **File** menu of the **TVSFP.ss3** window to retrieve the syntax file. Then click on **File, Save as** to save the model setup in a new

file, such as **TVSFP3.mum**. Next, we change the string in the **Title 1** text box on the **Configuration** screen, and select **SCHOOL** as the **Level-3 ID** as shown below. We now have both level-2 and level-3 IDs selected.



Keep all the other settings unchanged. Save the changes to the file **TVSFP3.mum** and select the **Run** option on the **Analysis** menu to produce the output file **TVSFP3.out**.

1.3.4.3 Discussion of results

Data summary

The **number of observations** section clearly shows the hierarchical structure of the data. The data contains 1600 students from 135 classes nested in 28 schools. In school number 20 (LEVEL 3: 20), the data of 73 students (N1: 73) from 7 (N2: 7) classes are present in this data set.

SuperMix - [TVSFP3.out]

File Analysis Window Help

Numbers of observations

Level 3 observations = 28
Level 2 observations = 135
Level 1 observations = 1600

LEVEL3 :	1	2	3	4	5	6	7	8
N2 :	2	3	1	6	2	4	3	6
N1 :	23	25	26	70	31	42	52	55
LEVEL3 :	9	10	11	12	13	14	15	16
N2 :	2	2	3	3	4	4	4	2
N1 :	39	33	52	65	27	80	33	18
LEVEL3 :	17	18	19	20	21	22	23	24
N2 :	6	5	5	7	11	7	4	8
N1 :	34	38	67	73	70	74	82	114
LEVEL3 :	25	26	27	28				
N2 :	7	4	7	13				
N1 :	113	33	94	137				

Save As... Close

Fixed effects estimates

As shown below, the estimates are not markedly different from the estimates of the previous two models.

SuperMix - [TVSFP3.out]

File Analysis Window Help

TITLE1: 3 level model - L-2 ID: Class_ L-3 ID: School

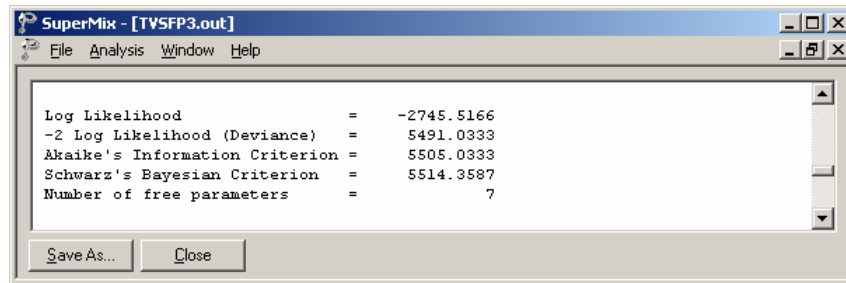
Maximum likelihood estimates

Fixed regressor(s)

Variable	Estimate	Std. Err.	Z-value	p-value
intcept	2.35537	0.12784	18.42442	0.00000
CC	0.61496	0.18243	3.37095	0.00075
TV	0.17156	0.17858	0.96068	0.33671
CCxTV	-0.35078	0.25465	-1.37750	0.16836

Save As... Close

Both the deviance and Akaike information criterion (AIC) are slightly higher than the previous model. The SBC is smaller.



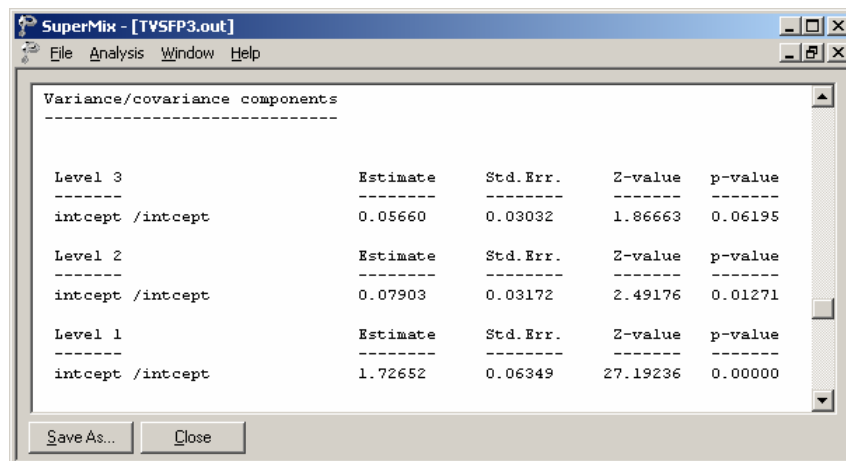
The screenshot shows a window titled "SuperMix - [TVSFP3.out]" with a menu bar (File, Analysis, Window, Help). The main area displays the following statistics:

```
Log Likelihood = -2745.5166
-2 Log Likelihood (Deviance) = 5491.0333
Akaike's Information Criterion = 5505.0333
Schwarz's Bayesian Criterion = 5514.3587
Number of free parameters = 7
```

Buttons for "Save As..." and "Close" are visible at the bottom.

Random effect estimates

The estimates for the random part of the model are reported next.



The screenshot shows a window titled "SuperMix - [TVSFP3.out]" with a menu bar (File, Analysis, Window, Help). The main area displays "Variance/covariance components" with the following table:

Level	Estimate	Std. Err.	Z-value	p-value
Level 3				
intcept /intcept	0.05660	0.03032	1.86663	0.06195
Level 2				
intcept /intcept	0.07903	0.03172	2.49176	0.01271
Level 1				
intcept /intcept	1.72652	0.06349	27.19236	0.00000

Buttons for "Save As..." and "Close" are visible at the bottom.

The estimated level-2 random effect is highly significant ($p = 0.08$), but the level-3 is not ($p = 0.06$).

1.3.4.4 Interpreting the results

Fit statistics and % variation explained

The variation of POSTTHKS scores can be explained by individual differences, classroom differences and school differences.

For schools,

$$\frac{0.05660}{0.05660 + 0.07903 + 1.72652} \times 100\% = 3.04\%$$

while for classrooms

$$\frac{0.07903}{0.05660 + 0.07903 + 1.72652} \times 100\% = 4.24\%$$

As calculated above, the school difference contributes 3.04% to the explanation of the total variance in the outcome, and classroom difference contributes 4.24%. The rest, 92.72% of the variation, is explained by the student differences.

Comparison of models

In Table XXX.6 the estimates of the previous three models and OLS in this chapter are summarized. The three-level estimates all lie between the corresponding two level-2 estimates.

Table XXX.6: Comparison of OLS and mixed model results

	OLS Estimates	Mixed Model		
		L-2 model		L-3 model
		L-2 ID: CLASS	L-2 ID: SCHOOL	L-2 ID: CLASS L-3 ID: SCHOOL
intcept	2.361 (0.066)	2.341 (0.092)	2.361 (0.124)	2.355 (0.128)
CC	0.607 (0.096)	0.589 (0.133)	0.652 (0.178)	0.615 (0.182)
TV	0.177 (0.094)	0.120 (0.131)	0.198 (0.175)	0.172 (0.179)
CCxTV	-0.323 (0.137)	-0.247 (0.189)	-0.417 (0.250)	-0.351 (0.255)
Deviance		5498.168	5501.438	5491.033
AIC		5510.168	5513.438	5505.033
SBC		5527.600	5521.431	5514.359
Number of free parameters		6	6	7

1.3.5 A 3-level random intercept model including pre-THKS score

The PRETHKS variable indicates the observed score before implementation of intervention. It might have an impact on the POSTTHKS scores. In this section, a three-level model including the PRETHKS as predictor is fitted.

1.3.5.1 The model

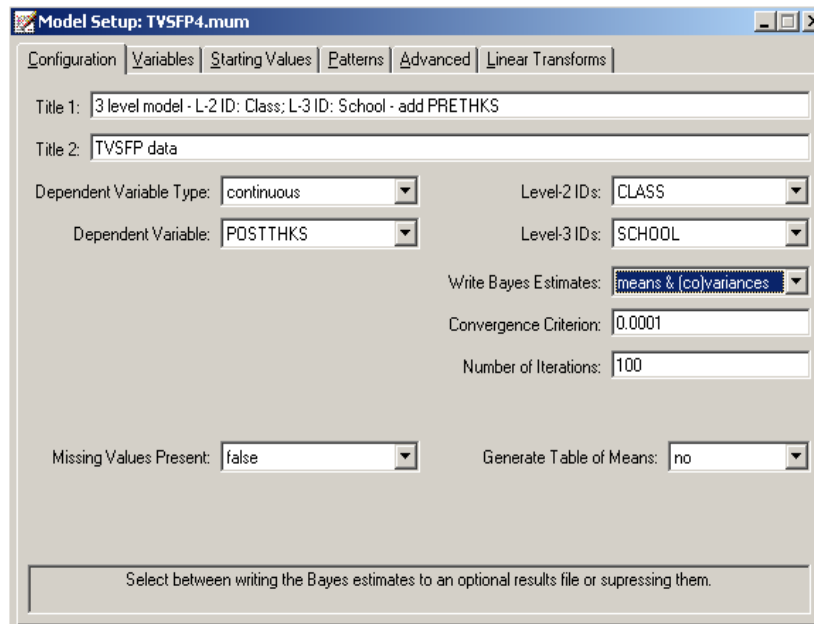
The only difference between this model and the previous one is the addition of the variable PRETHKS as a covariate:

$$\text{POSTHKS}_{ijk} = \beta_0 + \beta_1 \text{CC}_{ij} + \beta_2 \text{TV}_{ij} + \beta_3 (\text{CC}_{ij} \times \text{TV}_{ij}) + \beta_4 (\text{PRETHKS}_{ijk}) + v_{0ij} + v_{0i} + \varepsilon_{ijk}.$$

From the subscripts associated with the coefficients, we note that while CC, TV and CCxTV were measured at a classroom level, the pre-intervention score PRETHKS is measured on the individual level. Such a variable may also be referred to as a level-1 predictor, while CC, TV and CCxTV may be called level-2 predictors, covariates, or mediating effects.

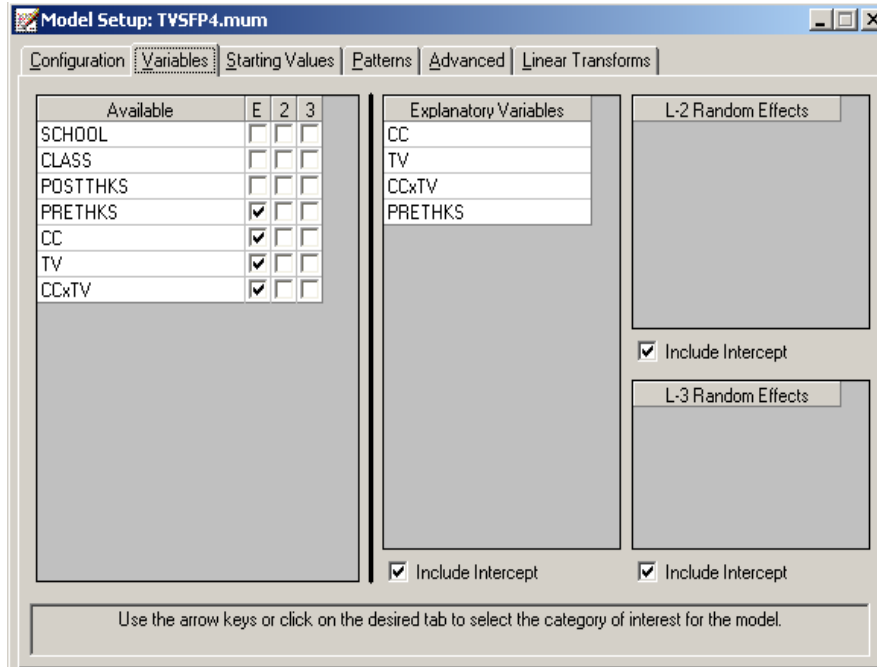
1.3.5.2 Setting up the analysis

The easiest way to set up this model is to modify the model setup in the syntax file **TVSFP3.mum** by first using the **Open Existing Model Setup** option on the **File** menu. Then click on **File, Save as** to save the model setup in a new file, such as **TVSFP4.mum**.



Next, we change the string in the **Title 1** text box on the **Configuration** screen. Notice that we would like to request Bayes estimates as part of the program output. To do so, select **means & (co)variances** option from the **Write Bayes Estimates** drop down list as shown above.

Click on the **Variables** tab and select PRETHKS as an additional **Explanatory Variable** by checking the corresponding **E** check box.



Save the changes to the file **TVSFP4.mum** and select the **Run** option on the **Analysis** menu to produce the output file **TVSFP4.out**.

1.3.5.3 Discussion of results

Fixed effects estimates and descriptive statistics

As shown below, the estimated coefficient of PRETHKS is highly significant. The estimate of the intercept coefficient decreased because part of the variation in the intercept can now be explained by PRETHKS.

SuperMix - [TVSFP4.out]

File Analysis Window Help

TITLE1: 3 level model - L-2 ID: Class_ L-3 ID: School - add PRETHKS

Maximum likelihood estimates

Fixed regressor(s)

Variable	Estimate	Std. Err.	Z-value	p-value
intcpt	1.69700	0.11666	14.54688	0.00000
CC	0.63919	0.14722	4.34189	0.00001
TV	0.17811	0.14365	1.23986	0.21503
CCxTV	-0.32042	0.20551	-1.55910	0.11897
PRETHKS	0.30720	0.02584	11.88788	0.00000

Save As... Close

The fit statistics are given below. A comparison of these two three-level examples will be given in the next section.

SuperMix - [TVSFP4.out]

File Analysis Window Help

Log Likelihood	=	-2678.6793
-2 Log Likelihood (Deviance)	=	5357.3586
Akaike's Information Criterion	=	5373.3586
Schwarz's Bayesian Criterion	=	5384.0163
Number of free parameters	=	8

Save As... Close

Random effect estimates

The third-level random intercept estimate is not significant at a 5% level of significance, which implies that after taking PRETHKS into account, the school differences are not significant.

Variance/covariance components				
Level	Estimate	Std. Err.	Z-value	p-value
Level 3	0.02575	0.01971	1.30665	0.19133
Level 2	0.06358	0.02767	2.29768	0.02158
Level 1	1.60201	0.05889	27.20544	0.00000

1.3.5.4 Interpreting the results

Estimated outcomes for different groups

For example, if a typical student who only participated in television intervention had a PRETHKS score of 2 (CC = 0; TV = 1; CCxTV = 0), the expected POSTTHKS score is calculated as follows:

$$\begin{aligned}
 \widehat{\text{POSTTHKS}}_{ijk} &= \hat{\beta}_{00} + \hat{\beta}_{02} \text{TV}_{ij} + \hat{\beta}_{04} (\text{PRETHKS}_{ijk}) \\
 &= 1.697 + 0.17811 + 2 \times 0.3072 \\
 &= 2.48951.
 \end{aligned}$$

ICCs and R square

ICCs

The so-called ICC (interclass correlation) measures the proportion of variation in the outcome variable between units at the different levels. It is occasionally referred to as the cluster effect, and is defined as the ratio of the between-cluster variance to the total variance. From the output for the random effects, we have

Level-1: $\widehat{\text{error var}} = 1.6020$

Level-2: $\widehat{\text{class var}} = 0.0636$

Level-3: $\widehat{\text{school var}} = 0.0258$.

Based on this information, we can calculate the ICCs as shown below.

Similarity of students within the same school:

$$ICC = \frac{0.0258}{1.6020 + 0.0636 + 0.0258} = 0.0153$$

Similarity of students within the same classrooms (and schools):

$$ICC = \frac{0.0636 + 0.0258}{1.6020 + 0.0636 + 0.0258} = 0.0529$$

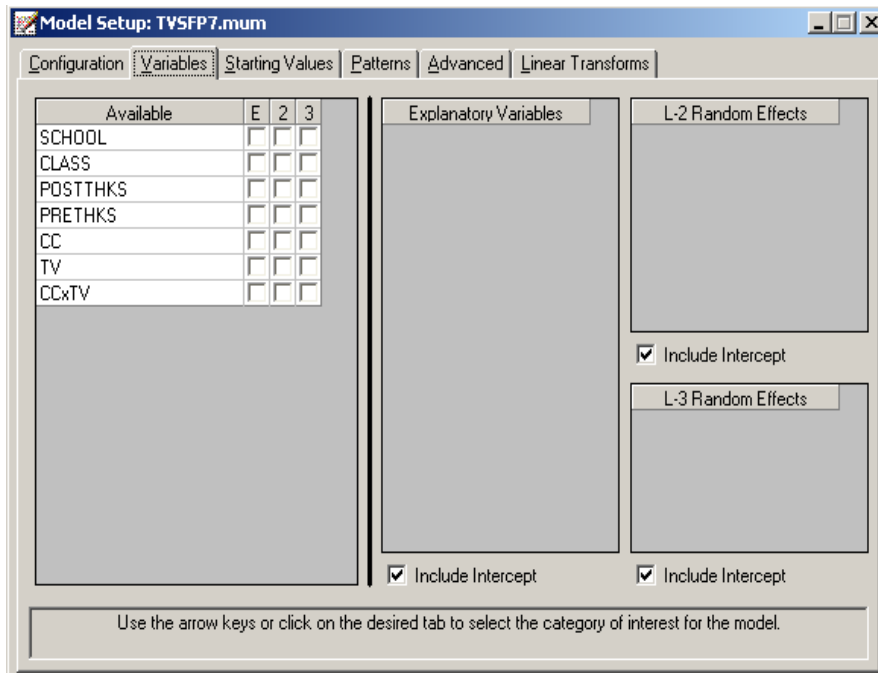
Similarity of classes within the same school:

$$ICC = \frac{0.0258}{0.0636 + 0.0258} = 0.289$$

R square

Another way to evaluate the explanation of variation in the outcome is to compute a statistic analogous to the familiar R^2 used in multiple linear regression. In a multilevel model, however, there is an R^2 for each variance component. Use of these statistics is not without problems, however, because the R^2 may at times have negative values, and in other cases the addition of explanatory variables can lead to an increase rather than a decrease in variance components. The more complex a hierarchical model is, the more likely is the occurrence of anomalies in variance-explained statistics.

To calculate the R^2 s for different levels of the level-3 model, we first need to get the variances for the null model, which is a 3-level model with no covariates. Open the file **TVSFP4.mum**, click on the **Variables** tab, and uncheck the check boxes of the selected **Explanatory Variables** as shown below.



Save the setup as **TVSFP7.mum** and run the model to get the following output of the variance/covariance component.

Variance/covariance components				
Level	Estimate	Std. Err.	Z-value	p-value

Level 3				

intcept /intcept	0.11032	0.04573	2.41251	0.01584

Level 2				

intcept /intcept	0.08481	0.03281	2.58504	0.00974

Level 1				

intcept /intcept	1.72367	0.06341	27.18391	0.00000

The R^2 s are calculated as

$$R_1^2 = 1 - \frac{\hat{\sigma}_p^2}{\hat{\sigma}_0^2} \quad R_2^2 = 1 - \frac{\hat{\sigma}_{v(2)p}^2}{\hat{\sigma}_{v(2)0}^2} \quad R_3^2 = 1 - \frac{\hat{\sigma}_{v(3)p}^2}{\hat{\sigma}_{v(3)0}^2}$$

where subscript 0 refers to a model with no covariates (i.e., the null model, **TVSFP7.out**) and subscript p refers to a model with p covariates (i.e., the full model, **TVSFP4.out**). The R^2 s for different levels are given in Table XXX.7.

Table XXX.7: R^2 values for a set of nested models

level	variance	null	full	R^2
1 (students)	$\hat{\sigma}^2$	1.724	1.602	.071
2 (classrooms)	$\hat{\sigma}_{v(2)}^2$.085	.064	.247
3 (schools)	$\hat{\sigma}_{v(3)}^2$.110	.026	.764

In the current example, only the intercept coefficient is allowed to vary randomly over classrooms and schools, thus making the calculation of the R^2 relatively

straightforward. In the case of models with random slopes, the calculation of R^2 statistics becomes more difficult. For an extensive discussion of the rationale and calculation of R^2 statistics, the user is referred to Snijders & Bosker (2000, pp. 99 – 109).

Model fit statistics and comparison of models

Now, we consider two level-2 models using the same covariates but different level-2 IDs: one uses CLASS as level-2 ID, the other uses SCHOOL. The models' setups are given in TVSFP5.mum and TVSFP6.mum. The comparison of estimates is summarized in the following table.

Table XXX.8: Comparison of OLS and mixed model results

	OLS Estimates	Mixed Model		
		L-2 model		L-3 model
		L-2 ID: CLASS	L-2 ID: SCHOOL	L-2 ID: CLASS L-3 ID: SCHOOL
intcept	1.6613 (0.0844)	1.6776 (0.0988)	1.6952 (0.1145)	1.6970 (0.1167)
CC	0.6406 (0.0921)	0.6330 (0.1186)	0.6601 (0.1440)	0.6392 (0.1472)
TV	0.1987 (0.0900)	0.1597 (0.1167)	0.2024 (0.1401)	0.1781 (0.1437)
CCxTV	-0.3216 (0.1303)	-0.2747 (0.1678)	-0.3697 (0.2011)	-0.3204 (0.2055)
PRETHKS	0.3252 (0.0259)	0.3116 (0.0258)	0.3103 (0.0259)	0.3072 (0.0258)

Table XXX.8 (continued)


error variance	1.6030 (0.0589)	1.6523 (0.0589)	1.6020 (0.0589)
class variance	0.0870 (0.0277)		0.0636 (0.0277)
school variance		0.0372 (0.0184)	0.0258 (0.0197)
Deviance	5359.9641	5366.0133	5357.3586
AIC	5373.9641	5380.0133	5373.3586
SBC	5394.3010	5389.3387	5384.0163
Number of free parameters	7	7	8

When comparing the deviances, AIC and SBC of the level-3 model with the level-2 models, we conclude that the three-level model has a better fit to the data.

1.3.5.5 Residuals

Level-2 Bayes results

Returning to the **TVSFP4.mum** output, click on the **Analysis** menu of the output window or the model set up window, and note that **View Level-2 Bayes Results** is now activated. Select the option to open the level-2 Bayes results.



Note that the default extension for the level-2 Bayes estimates is **.ba2**. Part of the file is shown below.

Level-3 ID	Level-2 ID	Students	Empirical Bayes Coefficients	Empirical Bayes Estimate	Estimated Variance	Coefficient Name
401.00	401101.00	13	1	-0.13230	0.40075E-01	intcept
401.00	401101.00	1	1	-0.30231E-01	0.61136E-01	intcept
401.00	401101.00	12	1	0.10878	0.41392E-01	intcept
401.00	401102.00	18	1	0.13415	0.33783E-01	intcept
402.00	402101.00	21	1	0.40196E-01	0.30746E-01	intcept
402.00	402102.00	17	1	0.91264E-01	0.34802E-01	intcept
405.00	405102.00	16	1	-0.78120E-01	0.35947E-01	intcept
405.00	405103.00	15	1	0.11404	0.37454E-01	intcept
405.00	405103.00	16	1	-0.23619	0.36287E-01	intcept
407.00	407101.00	21	1	0.58924E-01	0.31114E-01	intcept
407.00	407102.00	21	1	0.38397	0.31261E-01	intcept
407.00	407103.00	27	1	0.19960	0.26267E-01	intcept
408.00	408104.00	17	1	-0.20588	0.35282E-01	intcept

The representations of these seven columns are given in order below.


- Column 1: level-3 ID, which is school in our example.
- Column 2: level-2 ID, which refers to classroom.
- Column 3: number of the observations within level-2 ID, number of students within each classroom.
- Column 4: the number of the empirical Bayes coefficients.
- Column 5: the empirical Bayes estimate.
- Column 6: the estimated variance of the Bayes coefficient.
- Column 7: the name of the associated coefficient as used in the model.

Classroom 407102 has the largest Bayes estimate with a value of 0.38397. When considering the class difference, the predicted POSTTHKS score for a student in this specific class who only participated in television intervention with a PRETHKS score of 2 (CC = 0; TV = 1; CCxTV = 0) is calculated as follows.

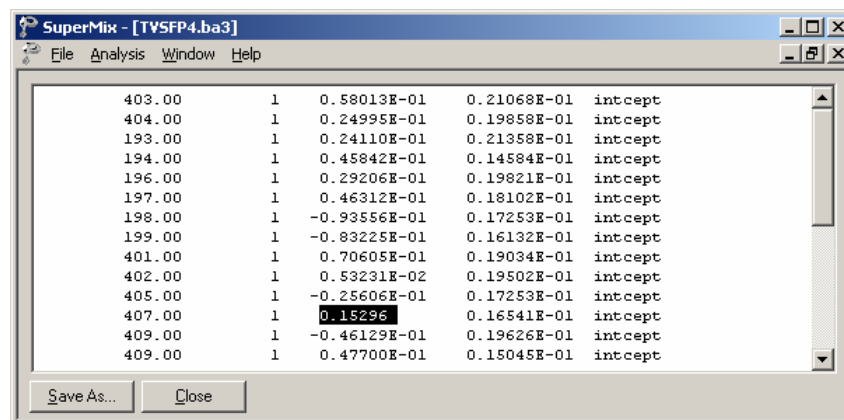
$$\begin{aligned}
 \widehat{\text{POSTTHKS}}_{ijk} &= \hat{\beta}_0 + \hat{\beta}_2 \text{TV}_{ij} + \hat{\beta}_4 (\text{PRETHKS}_{ijk}) + \hat{u}_{0i} \\
 &= 1.697 + 0.17811 + 2 \times 0.3072 + 0.38397 \\
 &= 2.87348.
 \end{aligned}$$

Level-3 Bayes results

Similarly, the level-3 Bayes results can be viewed by clicking on the **Analysis, View Level-3 Bayes Results**.



Part of the **TVSFP.ba3** is shown below.

A screenshot of the SuperMix software window displaying a table of Level-3 Bayes results. The window title is 'SuperMix - [TVSFP4.ba3]'. The table has five columns: a numerical value, a binary indicator (1), two scientific notation values, and a label 'intcept'. The value '0.15296' is highlighted in the row for school number 407.00. Buttons for 'Save As...' and 'Close' are visible at the bottom of the window.

403.00	1	0.58013E-01	0.21068E-01	intcept
404.00	1	0.24995E-01	0.19858E-01	intcept
193.00	1	0.24110E-01	0.21358E-01	intcept
194.00	1	0.45842E-01	0.14584E-01	intcept
196.00	1	0.29206E-01	0.19821E-01	intcept
197.00	1	0.46312E-01	0.18102E-01	intcept
198.00	1	-0.93556E-01	0.17253E-01	intcept
199.00	1	-0.83225E-01	0.16132E-01	intcept
401.00	1	0.70605E-01	0.19034E-01	intcept
402.00	1	0.53231E-02	0.19502E-01	intcept
405.00	1	-0.25606E-01	0.17253E-01	intcept
407.00	1	0.15296	0.16541E-01	intcept
409.00	1	-0.46129E-01	0.19626E-01	intcept
409.00	1	0.47700E-01	0.15045E-01	intcept

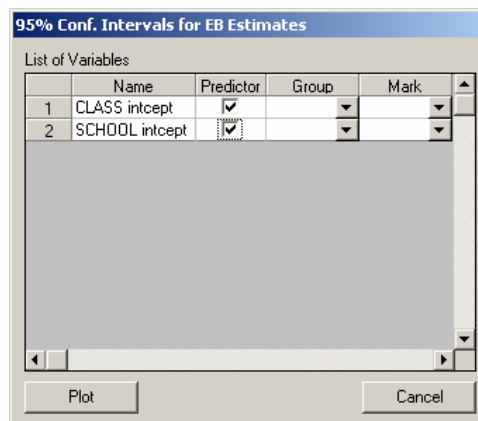
The same classroom (ID = 407102) discussed above is nested in school number 407. Now, considering both the class and school differences, the estimated POSTTHKS for a student from this classroom who only participated in television intervention with a pre-intervention score of 2 (CC = 0; TV = 1; CCxTV = 0) is calculated as follows.

$$\begin{aligned}\widehat{\text{POSTTHKS}}_{ijk} &= \hat{\beta}_0 + \hat{\beta}_2 \text{TV}_{ij} + \hat{\beta}_4 (\text{PRETHKS}_{ijk}) + \hat{v}_{0ij} + \hat{v}_{0i} \\ &= 1.697 + 0.17811 + 2 \times 0.3072 + 0.38397 + 0.15296 \\ &= 3.02644.\end{aligned}$$

Confidence intervals for random coefficients

The **Confidence Intervals** option on the **File, Model-based Graphs** menu provides the option to display confidence intervals for the empirical Bayes estimates of the random effects specified in a given model. This option is now used to examine the confidence intervals of the random intercepts for the schools, which represent the highest level of the hierarchy in the current example.

Select the **Confidence Intervals** option on the **File, Model-based Graphs** menu to activate the **95% Conf. Intervals for EB estimates** dialog box. Two graphs of the confidence intervals for the empirical Bayes estimates of the intercepts at the classroom level and school level are obtained by selecting CLASS intcept and SCHOOL intcept in the **Predictor** column before clicking **Plot**.



The graph obtained, as shown below, shows that, in general, the range of the confidence intervals for the level-3 empirical Bayes estimates of the intercepts is $(-0.2; 0.2)$, and the range for level-2 is about $(-0.4; 0.4)$.

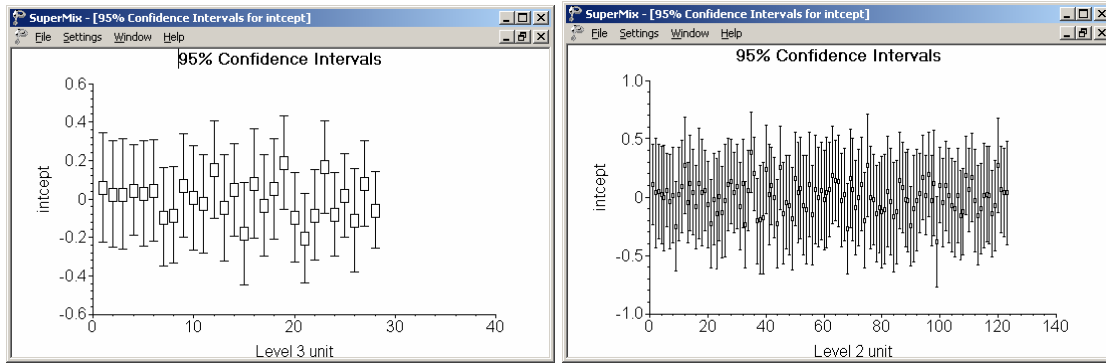


Figure XXX.14: 95% confidence intervals for level-2 Bayes estimates

The deviations from the estimated population intercept over schools are also apparent. Each confidence interval is obtained using the formula

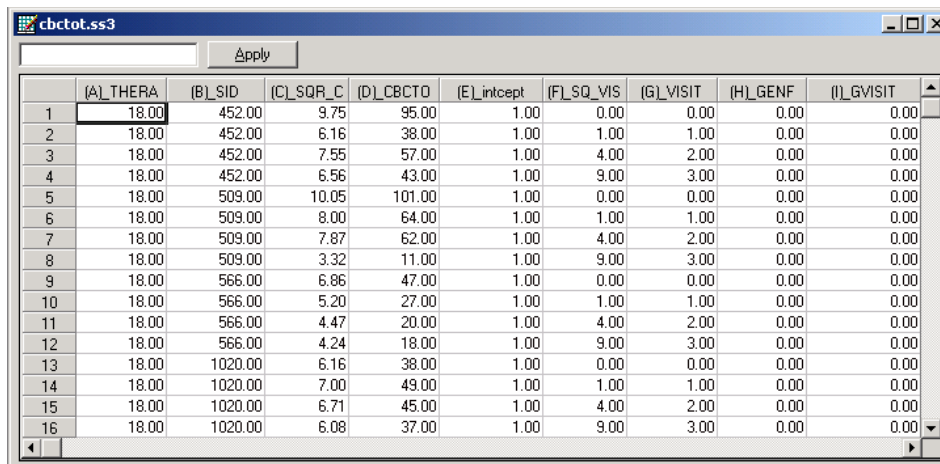
$$\text{Empirical Bayes residual} \pm 1.96\sqrt{\text{var}(\text{Empirical Bayes residual})}.$$

1.4 New 3 level continuous example using a subset of Schoenwald data

1.4.1 The data

The data set for this example is taken from a study described in Schoenwald & Henggeler (2005). Children in the study were assigned to therapists and followed across time. In this study, respondents were rated with the Child Behavioral Checklist (Achenbach, 1991) at four occasions. The gender of each respondent was also recorded.

Although the total number of patients in this study was 1,951, the number of patients treated by any single therapist ranged between 1 and 19. A total of 7,127 measurements were made for all patients over the course of the study. Data for the observations of all the variables for the first four patients treated by therapist number 18 are shown below in the form of a SuperMix spreadsheet file, named **cbtot.ss3**.



	(A)_THERA	(B)_SID	(C)_SQR_C	(D)_CBCTO	(E)_intcept	(F)_SQ_VIS	(G)_VISIT	(H)_GENF	(I)_GVISIT
1	18.00	452.00	9.75	95.00	1.00	0.00	0.00	0.00	0.00
2	18.00	452.00	6.16	38.00	1.00	1.00	1.00	0.00	0.00
3	18.00	452.00	7.55	57.00	1.00	4.00	2.00	0.00	0.00
4	18.00	452.00	6.56	43.00	1.00	9.00	3.00	0.00	0.00
5	18.00	509.00	10.05	101.00	1.00	0.00	0.00	0.00	0.00
6	18.00	509.00	8.00	64.00	1.00	1.00	1.00	0.00	0.00
7	18.00	509.00	7.87	62.00	1.00	4.00	2.00	0.00	0.00
8	18.00	509.00	3.32	11.00	1.00	9.00	3.00	0.00	0.00
9	18.00	566.00	6.86	47.00	1.00	0.00	0.00	0.00	0.00
10	18.00	566.00	5.20	27.00	1.00	1.00	1.00	0.00	0.00
11	18.00	566.00	4.47	20.00	1.00	4.00	2.00	0.00	0.00
12	18.00	566.00	4.24	18.00	1.00	9.00	3.00	0.00	0.00
13	18.00	1020.00	6.16	38.00	1.00	0.00	0.00	0.00	0.00
14	18.00	1020.00	7.00	49.00	1.00	1.00	1.00	0.00	0.00
15	18.00	1020.00	6.71	45.00	1.00	4.00	2.00	0.00	0.00
16	18.00	1020.00	6.08	37.00	1.00	9.00	3.00	0.00	0.00

The variables of interest are:

- THERAPIS is the therapist ID (446 in total).
- SID is the patient ID (1951 in total).

- CBTOT is the total score of the Child Behavior Checklist.
- INT is a column of ones, representing an (optional) intercept.
- VISIT represents the visit number (0, 1, 2, or 3) at which a measurement was made.
- GENF is an indicator variable for gender, and assumes the value 0 for males and 1 for females.
- GVISIT represents the interaction between GENF and VISIT, and is the product of GENF and VISIT.

1.4.2 Exploring the data

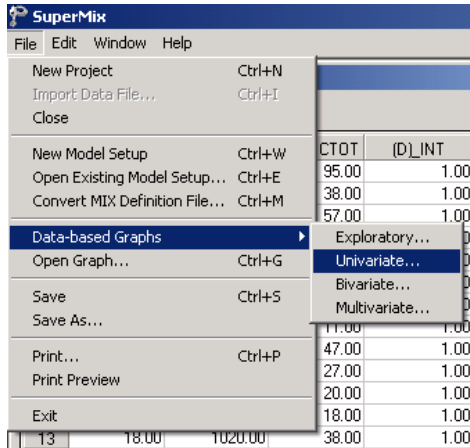
Relationships between variables, and trends over time in repeated measurement data, may be conveyed in an informal and simplified visual form via graphical displays. SuperMix offers both data-based and model-based graphs. Data-based graphing options are accessed via the **File, Data-based Graphs** option once a SuperMix data file (.ss3) is opened, while model-based graphs are available after the analysis has been performed, and will be discussed later in this section.

In the case of data-based graphs, we distinguish between three categories: univariate, bivariate, and multivariate graphs. Univariate graphs are particularly useful to obtain an overview of the characteristics of a single variable. In the sections to follow, we use data-based graphs to take a closer look at some of the variables in these data.

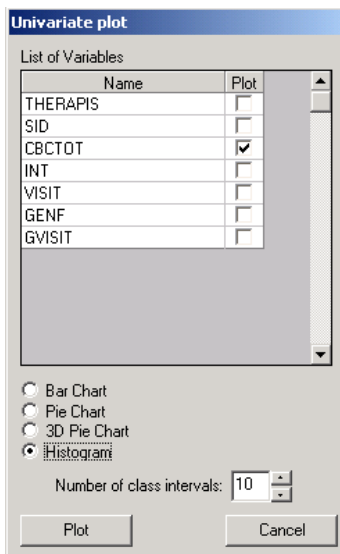
1.4.2.1 Univariate graphs

Histograms

As a first step, we take a look at the distribution of the total score on the Child Behavior Checklist (CBTOT) which is the potential dependent variable in this study. While scores such as these are not truly continuous variables, they are often treated as if they were. However, like personal income, the distribution of a score often is skewed. As a first step, we will take a closer look at the distribution of the intended outcome variable CBTOT. To do so, select the **Univariate** option from the **Data-based Graphs** menu as shown below.



The **Univariate plot** dialog box appears. Select the variable CBTOT and indicate that a **Histogram** is to be graphed. Note that the number of class intervals shown on the histogram is controlled by the **Number of class intervals** field, which is left at the default value of 10 in this case. Click the **Plot** button to display the histogram.



The histogram below shows that the distribution of total scores (CBCTOT) on the Child Behavior Checklist (CBC) is markedly asymmetrical. Given the normality assumptions used in fitting a 3-level linear multilevel model, it would be inappropriate to use CBCTOT in its current state. A transformation of this variable is required before it would make a suitable outcome variable for the intended analysis.


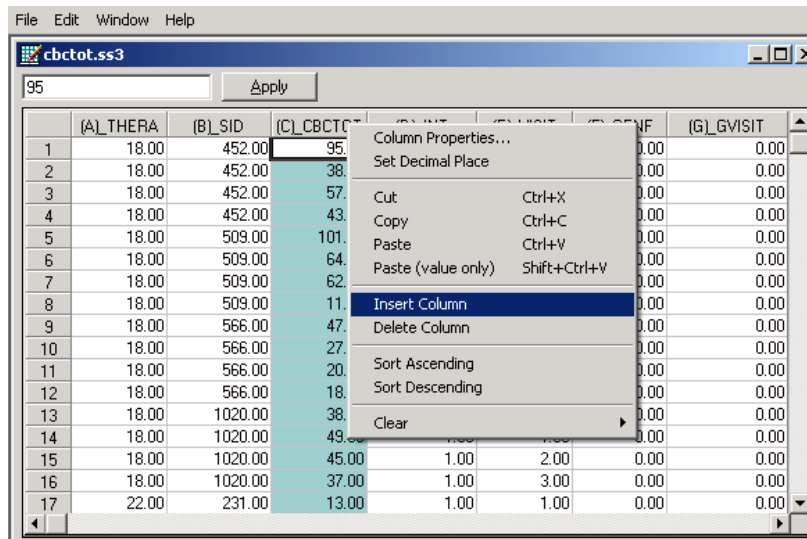


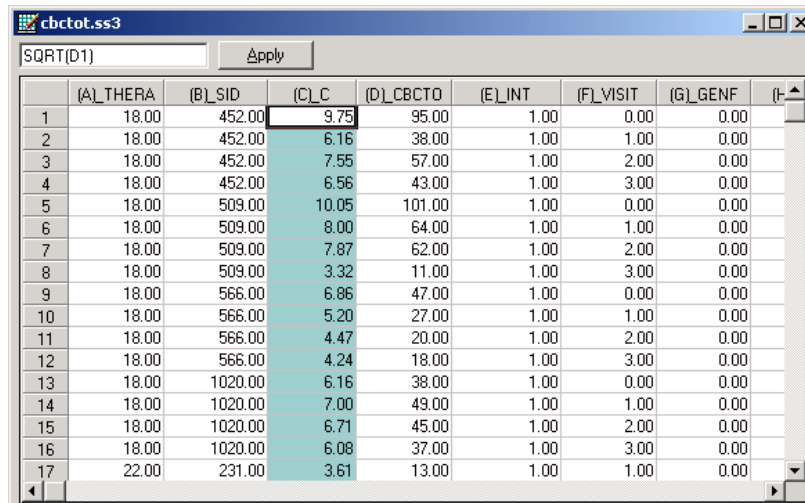
Figure XXX.15: Histogram of the variable CBCTOT

Transformation of variables

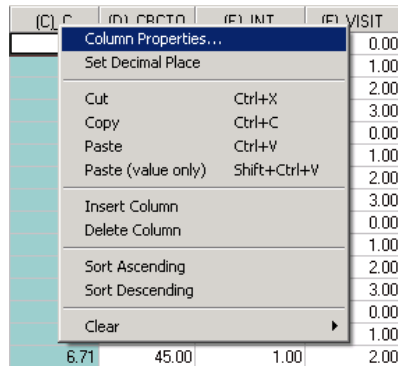
Common transformations used in the case of skewed variables include the natural logarithm of the variable in question, or the square root of the variable. We opt to explore the possibility of using the square root of the total score as outcome. To do so, a new variable containing the square root of the current total scores has to be created in the SuperMix spreadsheet. Right-click on the column with CBCTOT as heading, and select the **Insert Column** option from the pop-up menu that appears.




Select the new column and input the function **SQRT(D1)** in the formula box as shown below. Click the **Apply** button. Each value of the new variable is the square root value of the corresponding value of the variable CBCTOT as shown below.



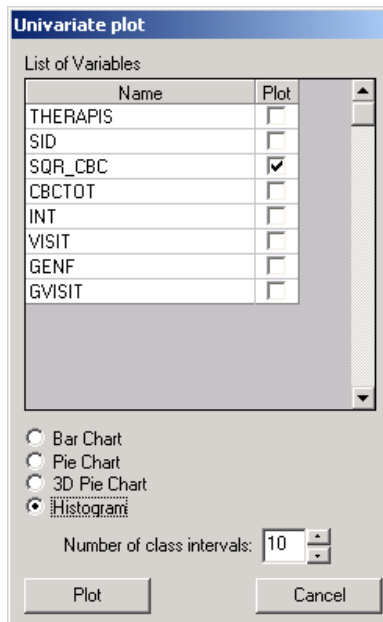
To rename the new variable, right-click again on the column header and select the **Column Properties** option.



Complete the **Header** field in the **Column Properties** dialog box as shown below. Also indicate that this is a continuous variable by selecting the **Continuous** option before clicking the **OK** button.



Check the distribution of the square root of the total score on the Child Behavior Checklist (SQR_CBC) by selecting the **Univariate** option from the **Data-based Graphs** menu to activate the **Univariate plot** dialog box. After selecting SQR_CBC by checking the appropriate box in the **Plot** column, select the **Histogram** option as before, and click **Plot**.



The histogram for the variable SQR_CBC is appreciably more symmetric than was the case for the original variable CBCTOT, as evident from the histogram shown below.





Figure XXX.16: Histogram of the variable SQR_CBC

1.4.2.2 Bivariate graphs

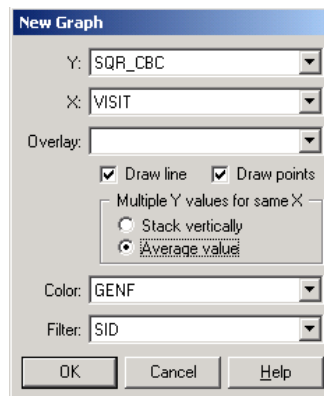
It is hoped that the total scores of patients would change over time, *i.e.*, with successive visits to their therapists. In addition, it is hypothesized that the gender of a patient may also have some relationship to the total score of a patient. Bivariate plots of possible relationships are a handy tool for the exploration of possible relationships.

Exploratory graphs

To explore the relationship between the time since the start of therapy and the square root of the total score, select the **Data-based Graphs, Exploratory** option from the **File** menu to activate the **New Graph** dialog box.




Select the outcome variable SQR_CBC as the **Y**-variable and VISIT as the **X** variable. Add the variable representing gender, GENF, from the **Color** field. Doing so will lead to the graphs of the gender groups to be displayed in different colors (blue and green being the default colors for two groups). Select the patient ID, as denoted by the variable SID, as the **Filter** variable to obtain individual graphs for patients. Click **OK** after completing the fields on this dialog box.



Graphs for patients with SIDs equal to 973, 790, and 2233 are shown below. These are but three of the 1951 graphs created via the graphing specification described above. Plotting symbols for each patient are shown at the bottom left of the

graphing window, and the legend for gender groups to the right. The slider at the bottom of the window is used to move from one graph to another.

For the first patient, with SID equal to 973, a roughly linear decrease in the outcome is observed as the visit number increases. This is not the case for patient 790, where an almost parabolic curve is observed, or for patient 2233 where an inverted parabola seems to be the most obvious line to fit. It can be concluded from these graphs that the relationship between SQR_CBC and VISIT differs from patient to patient, and moreover that it may not be strictly linear. The possible inclusion of a quadratic function of the time of measurement, i.e. VISIT, should be explored. No definite trend is immediately apparent for gender groups within the wide variety of curves plotted, but the possibility of an interaction between the gender and the number of the visit cannot be excluded.







Figure XXX.17: Relationship between SQR_CBC and VISIT for selected patients

Transforming a variable

To examine the relationship between the outcome and the quadratic value of VISIT, a new variable has to be created. This is done in a similar way to adding the square root of the total score. First insert a column, then type the appropriate function into the formula box as shown below. Click the **Apply** button. Each value of the new

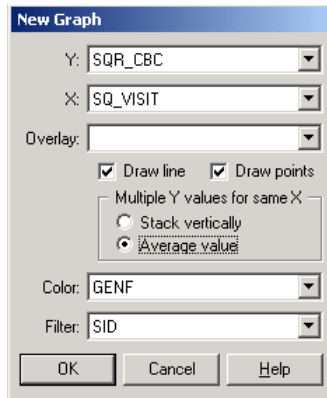
variable is the squared value of the corresponding value of the variable VISIT as shown below.

	(A)_THERA	(B)_SID	(C)_SQR_C	(D)_CBCTO	(E)_INT	(F)_F	(G)_VISIT	(H)_
1	18.00	452.00	9.75	95.00	1.00	0.00	0.00	
2	18.00	452.00	6.16	38.00	1.00	1.00	1.00	
3	18.00	452.00	7.55	57.00	1.00	4.00	2.00	
4	18.00	452.00	6.56	43.00	1.00	9.00	3.00	
5	18.00	509.00	10.05	101.00	1.00	0.00	0.00	
6	18.00	509.00	8.00	64.00	1.00	1.00	1.00	
7	18.00	509.00	7.87	62.00	1.00	4.00	2.00	
8	18.00	509.00	3.32	11.00	1.00	9.00	3.00	
9	18.00	566.00	6.86	47.00	1.00	0.00	0.00	
10	18.00	566.00	5.20	27.00	1.00	1.00	1.00	
11	18.00	566.00	4.47	20.00	1.00	4.00	2.00	
12	18.00	566.00	4.24	18.00	1.00	9.00	3.00	
13	18.00	1020.00	6.16	38.00	1.00	0.00	0.00	
14	18.00	1020.00	7.00	49.00	1.00	1.00	1.00	
15	18.00	1020.00	6.71	45.00	1.00	4.00	2.00	
16	18.00	1020.00	6.08	37.00	1.00	9.00	3.00	
17	22.00	231.00	3.61	13.00	1.00	1.00	1.00	


Right-click on the header of the newly inserted column to activate the **Column Properties** dialog box and enter a variable name such as SQ_VISIT into the **Header** field. Click **OK** to return to the spreadsheet.

1.4.2.3 Exploratory graphs

Remake the bivariate graphs shown previously for SQR_CBC and VISIT, using the squared value of VISIT (SQ_VISIT) instead. The completed **New Graph** dialog box, accessed via the **Data-based Graphs, Exploratory** option, is shown below. Click **OK** to display the graphs for individual patients.



Very little change in the shape of the plots is observed in the graphs obtained for the three patients considered earlier. To follow up on the possibility of a nonlinear relationship between the outcome and the visit number, both of the variables VISIT and SQ_VISIT will be included in the first model fitted, where the relationship of each with the outcome can be evaluated in the presence of the other.



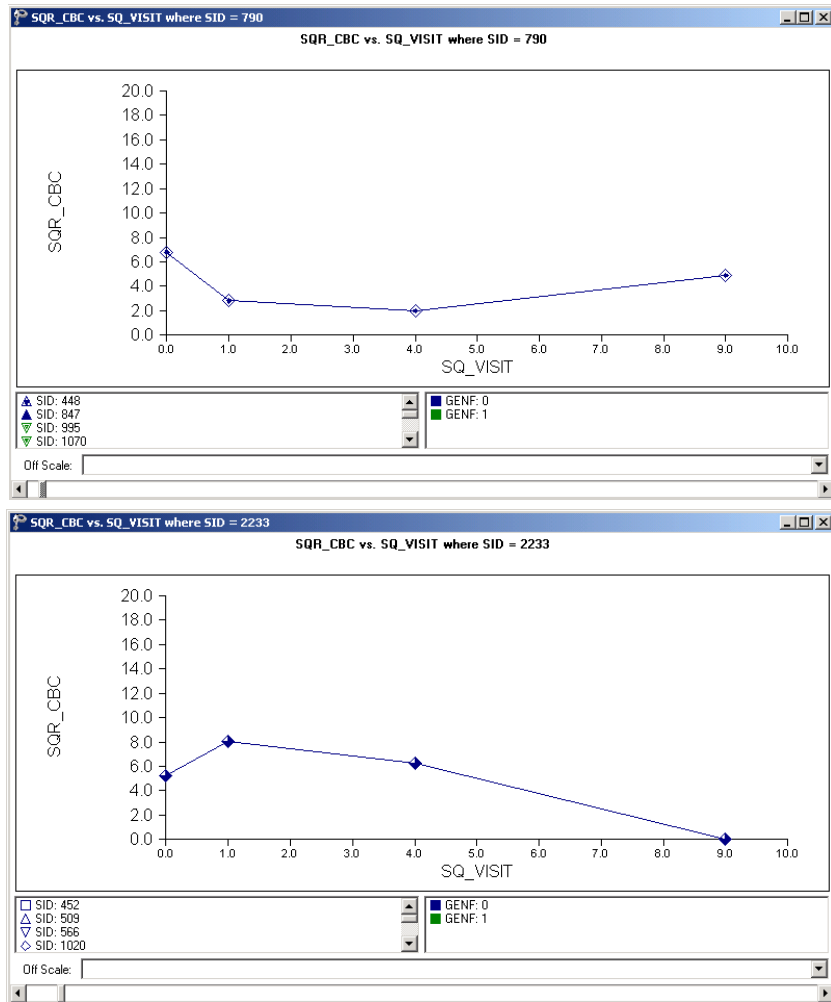


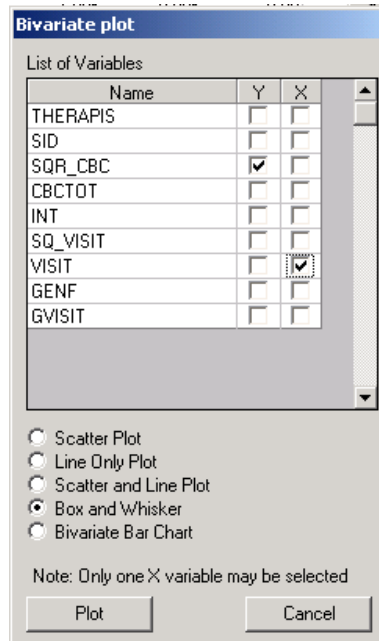
Figure XXX.18: relationship between SQR_CBC and SQ_VISIT for selected patients

Box-and-whisker plots

Another bivariate plot of interest is a box-and-whisker plot, which may be used to examine the distributions of continuous variables such as for the different values of discrete valued predictors. This option, accessed via the **Data-based Graphs, Bivariate** option on the **File** menu, is now used to take a closer look at the

distribution of the transformed outcome variable at different visits, and for the two gender groups.

The **Bivariate plot** dialog box is completed as follows: select the outcome variable SQR_CBC as the Y-variable of interest, and the predictor VISIT to be plotted on the X-axis. Check the **Box and Whisker** option, and click **Plot**.



In the plot shown below, the box-and-whisker plots for the square root of the CBCTOT scores are shown at each of the measurement occasions. Recall that the bottom line of a box represents the first quartile (q_1), the top line the third quartile (q_3), and the in-between line the median (me). Here, the arithmetic mean is represented by a diamond. A decrease in the mean HDRS rating is observed over the course of the study. In addition, the larger distances between the extremes of the boxes at the later measurement occasions indicate more variability in the transformed CBCTOT scores towards the end of the study.




Figure XXX.19: Box-and-whisker plot of SQR_CBC vs. VISIT

When a similar plot is made for the original total score as represented by the variable CBCTOT, it is clear that the distributions of the transformed scores, though still exhibiting more variability at later visits, are closer to normal for the transformed variable (figure below).

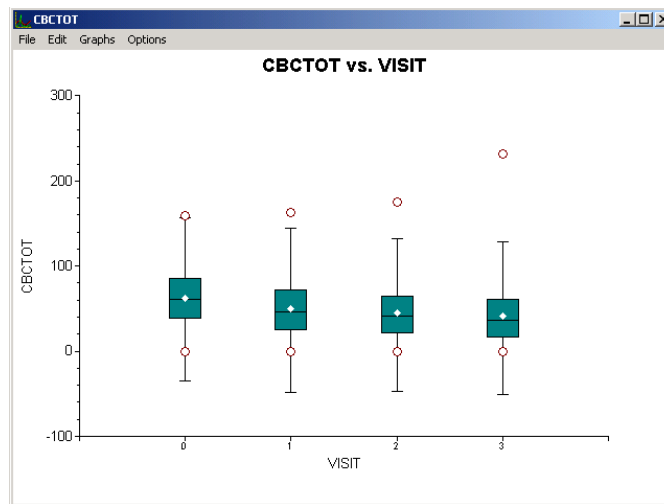


Figure XXX.20: Box-and-whisker plot of CBCTOT vs. VISIT

A box-and-whisker plot of the transformed scores for the two gender groups can easily be obtained. Simply close the graph window shown above, deselect VISIT as the **X**-variable and select the indicator of gender GENF instead. Click **Plot** to obtain the box-and-whisker plot shown below. A slightly larger range of scores is observed for males (GENF = 0) than for females (GENF = 1).




Figure XXX.21: Box-and-whisker plot of SQR_CBC vs. GENF

When this plot is compared to a similar one for the untransformed outcome variable CBCTOT, the same tendency towards a less normal distribution is observed, particularly with respect to the total scores of male patients.

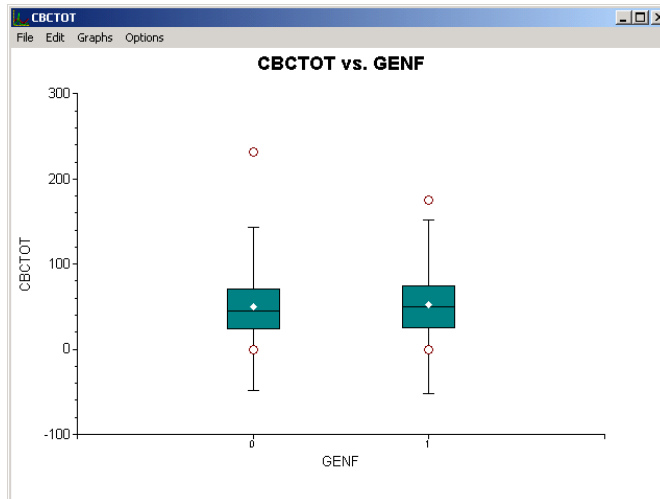
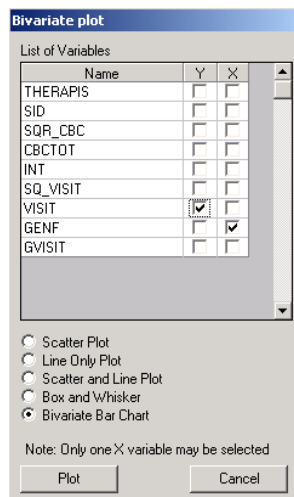


Figure XXX.22: Box-and-whisker plot of CBCTOT vs. GENF

Bivariate bar charts

Another bivariate plot that may provide insight is a plot of gender by the number of visits. The **Bivariate** option on the **File, Data-based Graphs** menu is again used to access the **Bivariate plot** dialog box. Select VISIT as the Y-variable and GENF as the X-variable, and request a bivariate bar chart prior to clicking the **Plot** button.



The bar chart for VISIT vs. GENF shows not only that more males than females are present in the data, but also that roughly equal numbers of observations/scores are available for the two groups at each of the visits. The pattern in terms of the number of observations available at each visit is the same for the two gender groups.




Figure XXX.23: Bivariate chart of VISIT vs GENF

1.4.3 Fitting a growth curve model to the data

1.4.3.1 The model

The first model fitted to the data explores the relationship between SQR_CBC and the visit number, as represented by the variables VISIT and SQ_VISIT:

$$SQR_CBC_{ijk} = \beta_0 + \beta_1 * VISIT_{ijk} + \beta_2 * SQ_VISIT_{ijk} + v_{i0} + v_{ij0} + e_{ijk}$$

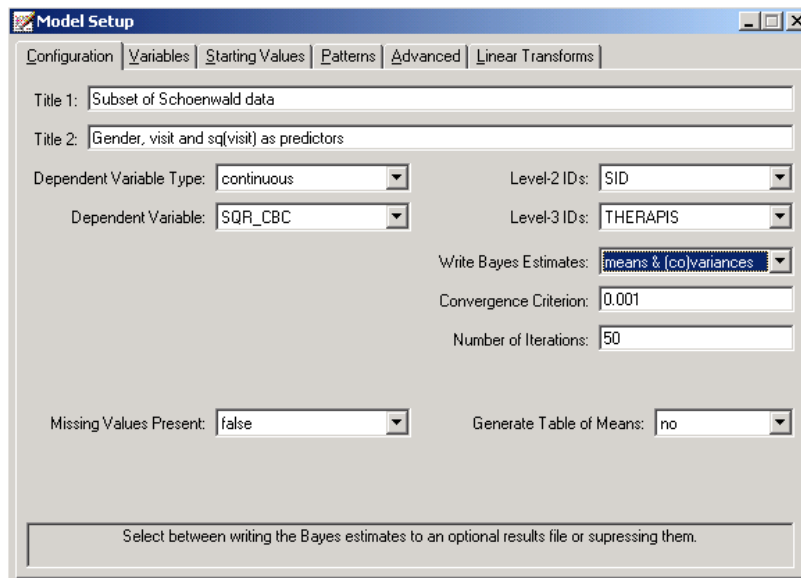
In this model, β_0 denotes the average expected total score, and β_1 and β_2 indicate the estimated coefficients associated with the fixed part of the model which contains the predictor variables VISIT and SQ_VISIT. The random part of the model is

represented by v_{i0} , v_{ij0} and e_{ijk} , which denote the variation in average total score over therapists, between patients (or, in other words, over patients nested within therapists) and between measurements at the lowest level of the hierarchy.

1.4.3.2 Setting up the analysis

Open the SuperMix spreadsheet **cbtot.ss3**. The next step is to describe the model to be fitted. We use the SuperMix interface to provide the model specifications. From the main menu bar, select the **File, New Model Setup** option.

Select the continuous outcome variable SQR_CBC from the **Dependent Variable** drop-down list box on the **Configuration** tab. The therapist number THERAPIS and respondent identification code SID used to define the levels of the hierarchy are specified as **Level-3 ID** and **Level-2 ID** respectively by selecting them from the **Level-3 IDs** and **Level-2 IDs** drop-down list boxes. Enter a title for the analysis in the **Title** text boxes. Select the **means & (co)variances** option from the **Write Bayes estimates** drop-down list box to request the writing of residuals to an external file. In this example, default settings for all other options associated with the **Configuration** screen are used. Proceed to the **Variables** screen by clicking on that tab.



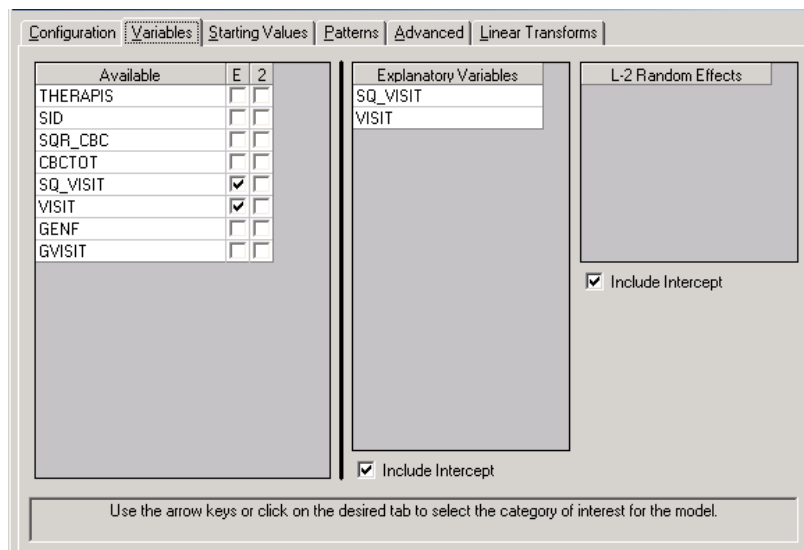
The screenshot shows the 'Model Setup' dialog box with the following settings:

- Configuration tab selected.
- Title 1: Subset of Schoenwald data
- Title 2: Gender, visit and sq[visit] as predictors
- Dependent Variable Type: continuous
- Dependent Variable: SQR_CBC
- Level-2 ID: SID
- Level-3 ID: THERAPIS
- Write Bayes Estimates: means & (co)variances
- Convergence Criterion: 0.001
- Number of Iterations: 50
- Missing Values Present: false
- Generate Table of Means: no

Select between writing the Bayes estimates to an optional results file or supressing them.

The **Variables** screen is used to specify the fixed and random effects to be included in the model. Select the explanatory (fixed) variables using the **E** check boxes next to the variables names in the **Available** grid at the left of the screen. Note that, as the variables are selected, the selected variables are listed in the **Explanatory Variables** grid. After selecting all the explanatory variables, the screen shown below is obtained. The **Include Intercept** check box in the **Explanatory Variables** grid is checked by default, indicating that an intercept term will automatically be included in the fixed part of the model.

Next, specify the random effects at levels 2 and 3 of the hierarchy. In this example, we want to fit a model with random intercepts at levels 2 and 3. By default, the **Include Intercept** check boxes in both the **L-2 Random Effects** and **L-3 Random effects** grids are checked. If these boxes are left checked, and no additional random effects are indicated using the **2** column in the **Available** grid to the left, the model fitted will be the random-intercepts-only model we intend to use. No further changes on this screen are necessary.

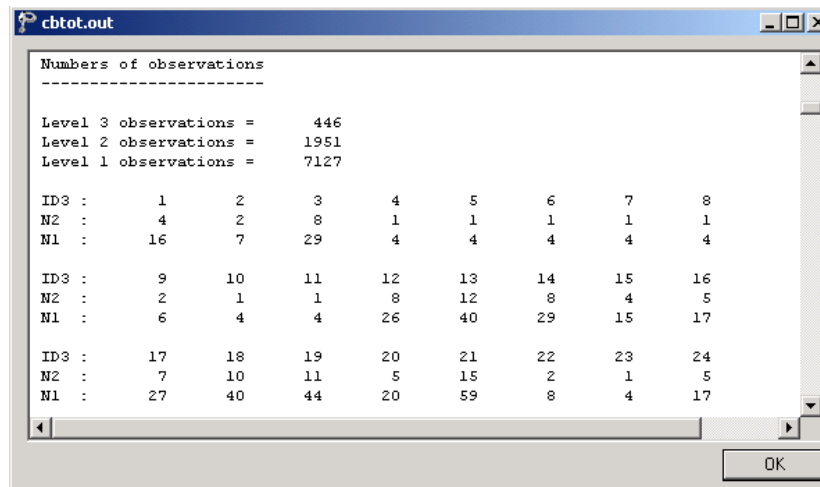


Before running the analysis, the model specifications have to be saved. Select the **File, Save As** option, and provide a name (**cbctot.mum**) for the model specification file. Run the analysis by selecting the **Run** option from the **Analysis** menu.

1.4.3.3 Discussion of results

Portions of the output file **cbctot.out** are shown below.

In the first section of the output file, a description of the hierarchical structure is provided. Data from a total of 446 therapists and 1951 patients at 7127 measurement occasions were included at levels 3, 2 and 1 of the model. This corresponds to the study design described earlier. In addition, a summary of the number of patients and measurements nested within each therapist is provided. The therapist with ID3 = 21, for example, had 15 patients (N2: 15). These patients were measured at 59 occasions. By contrast, therapist 23 had only 1 patient, for whom 4 measurements were made.



Numbers of observations

Level 3 observations = 446
Level 2 observations = 1951
Level 1 observations = 7127

ID3 :	1	2	3	4	5	6	7	8
N2 :	4	2	8	1	1	1	1	1
N1 :	16	7	29	4	4	4	4	4
ID3 :	9	10	11	12	13	14	15	16
N2 :	2	1	1	8	12	8	4	5
N1 :	6	4	4	26	40	29	15	17
ID3 :	17	18	19	20	21	22	23	24
N2 :	7	10	11	5	15	2	1	5
N1 :	27	40	44	20	59	8	4	17

The data summary is followed by descriptive statistics for all the variables included in the model. The mean of 6.61867 reported for the outcome SQR_CBC translates to a total score of 43.806 on the Child Behavior Checklist.

cbtot.out

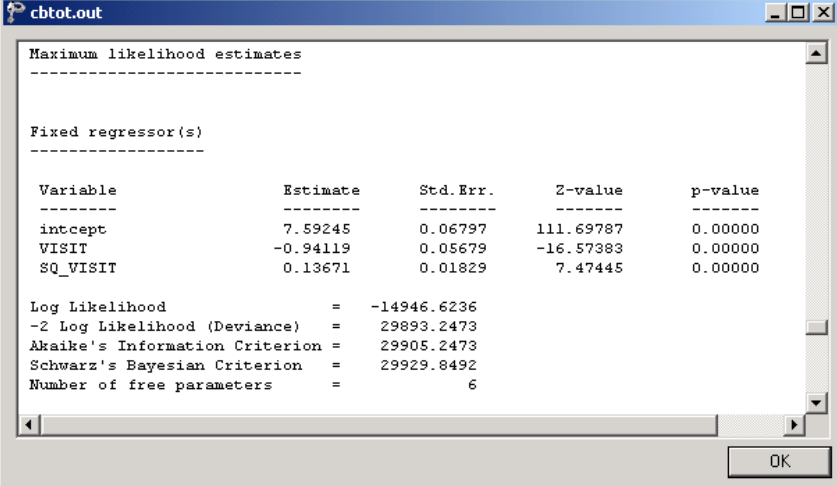
Descriptive statistics for all variables

Variable	Minimum	Maximum	Mean	Stand. Dev.
Dependent Variable				
SQR_CBC	0.00000	15.23155	6.61867	2.51988
Random-Effect Variable(s)				
intcept	1.00000	1.00000	1.00000	0.00000
intcept	1.00000	1.00000	1.00000	0.00000
Fixed Regressor(s)				
VISIT	0.00000	3.00000	1.44156	1.11237
SQ_VISIT	0.00000	9.00000	3.31528	3.45076

OK

Fixed effects results

The output describing the estimated **fixed effects** after convergence is shown next. The estimates are shown in the column with heading Estimate, and correspond to the coefficients $\beta_0, \beta_1, \dots, \beta_3$ in the model specification. From the z-values and associated exceedance probabilities, we see that the coefficients associated with both the time of measurement (VISIT) and squared value of the time of measurement (SQ_VISIT) are highly significant. The significance of the estimate associated with SQ_VISIT supports the tentative conclusion made during the exploratory analysis that the relationship between score and visit number cannot adequately be described by a linear relationship. While the average CBC score is expected to decrease with 0.94119 units between two successive visits, a smaller increase in score of 0.13671 is associated with the squared value of the time of measurement.



```
Maximum likelihood estimates
-----

Fixed regressor(s)
-----

Variable          Estimate      Std. Err.      Z-value        p-value
-----
intcept           7.59245      0.06797       111.69787     0.00000
VISIT             -0.94119     0.05679       -16.57383     0.00000
SQ_VISIT          0.13671     0.01829        7.47445     0.00000

Log Likelihood = -14946.6236
-2 Log Likelihood (Deviance) = 29893.2473
Akaike's Information Criterion = 29905.2473
Schwarz's Bayesian Criterion = 29929.8492
Number of free parameters = 6
```

Random effects results

The output for the **random part** of the model follows, and is shown in the image below. There is significant variation in the average estimated total health expenditure at all levels, with the most variation over the patients (level-2), and the least variation over therapists (level-3).

Variance/covariance components				
Level	Estimate	Std. Err.	Z-value	p-value
Level 3				
intcept /intcept	0.58201	0.10414	5.58881	0.00000
Level 2				
intcept /intcept	3.08097	0.13341	23.09426	0.00000
Level 1				
intcept /intcept	2.34083	0.04599	50.90214	0.00000

1.4.3.4 Interpreting the results

Estimated outcomes for different groups

A typical patient at the start of the study is expected to have a transformed CBC score of

$$\begin{aligned} \widehat{\text{SQR_CBC}}_{ij0} &= \hat{\beta}_0 + \hat{\beta}_1 * \text{VISIT}_{ij0} + \hat{\beta}_2 * \text{SQ_VISIT}_{ij0} \\ &= 7.59245 - 0.94119(0) + 0.13671(0^2) \\ &= 7.59245, \end{aligned}$$

that is, the estimated intercept. Similar equations for expected transformed scores at subsequent measurements (visits) are obtained in the same way:

$$\begin{aligned} \text{VISIT 1: } \widehat{\text{SQR_CBC}}_{ij1} &= 7.59245 - 0.94119(1) + 0.13671(1^2) \\ &= 7.59245 - 0.94119 + 0.13671 \\ &= 6.78797 \end{aligned}$$

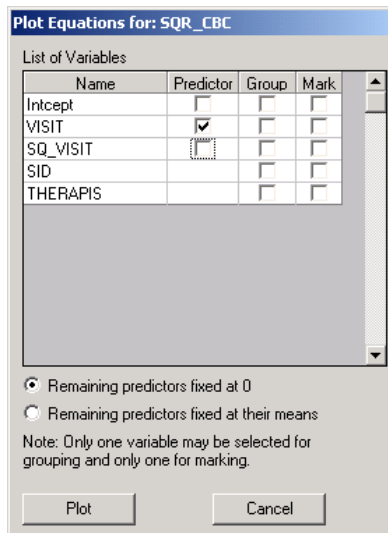
$$\begin{aligned} \text{VISIT 2: } \widehat{\text{SQR_CBC}}_{ij2} &= 7.59245 - 0.94119(2) + 0.13671(2^2) = \\ &= 7.59245 - 1.88238 + 0.54684 \\ &= 6.25691 \end{aligned}$$

$$\begin{aligned} \text{VISIT 3: } \widehat{\text{SQR_CBC}}_{ij3} &= 7.59245 - 0.94119(3) + 0.13671(3^2) \\ &= 7.59245 - 2.82357 + 1.23039 \\ &= 5.99927 \end{aligned}$$

The effect of the positive estimate for SQ_VISIT in slowing down the expected decrease in CBC scores over successive measurement occasions is clear from the equations above: without this estimate, the expected CBC scores at visits 1, 2, and 3 would have been 6.65126, 5.71007, and 4.76888 respectively. In terms of the actual rather than the square root of the CBC scores, the expected scores at the 4 measurement occasions under the current model are 57.6453, 46.0765, 39.1489, and 35.9914 respectively.

Model-based graphs

Using the **Plot Equations for: SQR_CBC** dialog box that appears when the **File, Model-based Graphs, Equations** option is selected, we can graphically depict the trend in expected average squared score for the predictors VISIT and SQ_VISIT. The dialog box below shows the selection of the predictor VISIT, and in the graph requested, SQ_VISIT will be fixed at a value of 0.



The graph below shows the result obtained when the **Plot** button is clicked after completion of the **Plot Equations for: SQR_CBC** dialog box as shown above. A similar plot for the predictor **SQ_VISIT** is given directly after. Note that, in the second graph, the increase in expected score seems larger than implied by the estimate of 0.13671. This is due in part to the difference in the ranges of the two predictors, as reflected in the tick marks on the X-axes of the graphs.




Figure XXX.24: Plot of SQR_CBC vs. VISIT




Figure XXX.25: Plot of SQR_CBC vs. SQ_VISIT

Fit statistics and ICC

From the output for the **random part** of the model it is clear that there is significant variation in the average estimated total health expenditure at all levels, with the most variation over the patients (level-2), and the least variation over therapists (level-3).

An estimate of the percentage of variation in the outcome at a patient level, for example, is obtained as

$$\frac{3.08097}{0.58201 + 3.08097 + 2.34083} \times 100\% = 51.32\%$$

indicating that 51.32% of the total variance in scores is at the patient level. In contrast,

$$\frac{0.58201}{0.58201 + 3.08097 + 2.34083} \times 100\% = 9.69\%$$

is at the therapist level, with the remainder over measurements nested within patients.

1.4.4 Fitting a random intercept model with 3 predictors and interaction term to the data

1.4.4.1 The model

From the exploratory analysis, we are aware of a possibly nonlinear relationship between the transformed outcome variable SQR_CBC and the visit number, as represented by the variables VISIT and SQ_VISIT. Differences in the distributions of the transformed scores of the two gender groups also lead us to suspect that the outcome may depend to some extent on the gender of the patient. The possibility of an interaction between the time elapsed since the start of the study, as represented by VISIT and SQ_VISIT, and the gender of a patient cannot be ruled out.

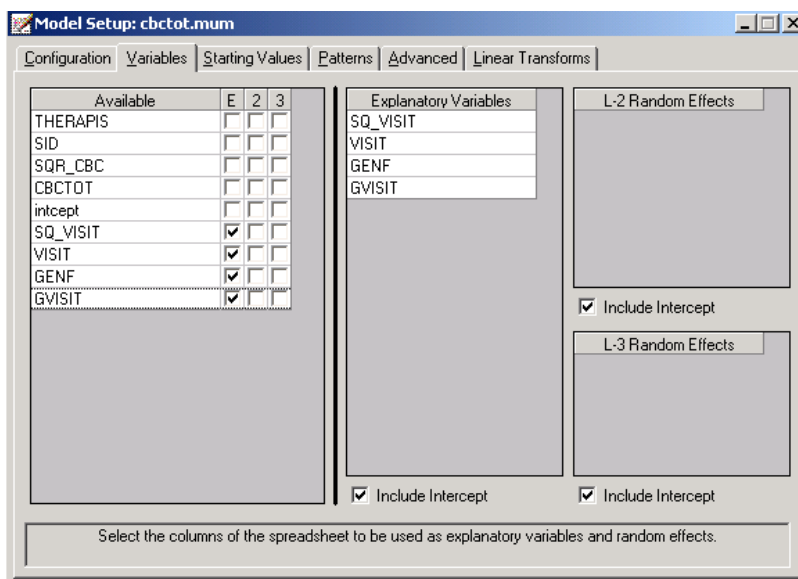
The model considered in this section uses the participant's gender, visit number, squared value of the visit number, and the interaction between visit number and gender (represented by the variable GVISIT in the data spreadsheet) to predict the square root of the total score on the Child Behavior Checklist. This second order growth curve with gender and the interaction term as covariates can be expressed as follows:

$$\text{SQR_CBC}_{ijk} = \beta_0 + \beta_1 * \text{GENF}_{ij} + \beta_2 * \text{VISIT}_{ijk} + \beta_3 * \text{SQ_VISIT}_{ijk} + \beta_4 * \text{GENF}_{ij} * \text{VISIT}_{ijk} + v_{i0} + v_{ij0} + e_{ijk}$$

As before, β_0 denotes the average expected total score, $\beta_1, \beta_2, \dots, \beta_4$ indicate the estimated coefficients associated with the fixed part of the model, and v_{i0}, v_{ij0} and e_{ijk} represent the random part of the model.

1.4.4.2 Setting up the analysis

The SuperMix spreadsheet **cbtot.ss3** and the model specification file **cbtot.mum** discussed in the previous example are used a point of departure. With the model specification file open, click on the **Variables** tab of the **Model Setup** window. Add the predictors GENF and GVISIT to the model by checking the boxes next to these variables in the **E** column, as shown below.



Save the modified model specification file, then select the **Run** option from the **Analysis** menu to perform the analysis.

1.4.4.3 Discussion of results

Fixed effects results

The maximum likelihood estimates of the coefficients in the fixed part of the model are shown below. The statistical significance of all the effects confirm our suspicion that the CBC scores measured over time not only depend on the time of measurement and the squared value thereof, but also on the gender of the patient. A significant interaction between gender and the time of measurement is also observed. Recall that for male patients GENF was coded 0, and for females GENF was assigned a value of 1. The positive estimate of 0.28977 for the effect of gender indicates that males tended to have a lower score on average than females: the expected average male score is 0.28977 units lower on the transformed CBC scores than for females. This effect is offset by the negative estimate of the interaction effect. For males, the interaction term GVISIT assumes the value 0, but for females GVISIT is equal to 0, 1, 2, and 3 respectively at the 4 measurement occasions. The transformed score of a female patient is thus expected to be 0.10034 units lower at

the second visit than the score of a male patient, or a female patient at the beginning of the study.

Maximum likelihood estimates

Fixed regressor(s)

Variable	Estimate	Std. Err.	Z-value	p-value
intcept	7.49255	0.07763	96.51258	0.00000
VISIT	-0.90730	0.05795	-15.65543	0.00000
SQ_VISIT	0.13703	0.01827	7.49868	0.00000
GENF	0.28977	0.10960	2.64397	0.00819
GVISIT	-0.10034	0.03485	-2.87888	0.00399

Log Likelihood = -14941.3366
 -2 Log Likelihood (Deviance) = 29882.6732
 Akaike's Information Criterion = 29898.6732
 Schwarz's Bayesian Criterion = 29931.4757
 Number of free parameters = 8

OK

Random effects results

The output for the **random part** of the model is given next.

Variance/covariance components

Level	Estimate	Std. Err.	Z-value	p-value
Level 3				
intcept /intcept	0.57846	0.10383	5.57142	0.00000
Level 2				
intcept /intcept	3.08224	0.13340	23.10575	0.00000
Level 1				
intcept /intcept	2.33647	0.04590	50.90206	0.00000

OK

As before, most variation in scores is found at a patient level, and the least variation at the therapist level. The estimated percentages of variation in outcome at patient and therapist level are

$$\frac{3.08224}{0.57846 + 3.08224 + 2.33647} \times 100\% = 51.39\%$$

and

$$\frac{0.57846}{0.57846 + 3.08224 + 2.33647} \times 100\% = 9.65\%$$

respectively. When compared to the similar percentages for the growth curve model, changes observed are negligible. The addition of the variables GENF and GVISIT did not contribute significantly to the explanation of remaining variation in the outcome at the various levels of the model.

1.4.4.4 Interpretation of the results

Estimated outcomes for different groups

For a typical patient, the expected CBC score can be calculated as

$$\begin{aligned} \widehat{\text{SQR_CBC}}_{ijk} &= \hat{\beta}_0 + \hat{\beta}_1 * \text{GENF}_{ij} + \hat{\beta}_2 * \text{VISIT}_{ijk} + \hat{\beta}_3 * \text{SQ_VISIT}_{ijk} + \hat{\beta}_4 * \text{GVISIT}_{ijk} \\ &= 7.49255 + 0.28977 * \text{GENF}_{ij} - 0.90730 * \text{VISIT}_{ijk} \\ &\quad + 0.13703 * \text{SQ_VISIT}_{ijk} - 0.10034 * \text{GVISIT}_{ijk} . \end{aligned}$$

For male patients GENF = 0, and thus the formula used to predict their CBC scores reduces to

$$\begin{aligned} \widehat{SQR_CBC}_{ijk} &= 7.49255 + 0.28977 * (0) - 0.90730 * VISIT_{ijk} \\ &\quad + 0.13703 * SQ_VISIT_{ijk} - 0.10034 * (0) \\ &= 7.49255 - 0.90730 * VISIT_{ijk} + 0.13703 * SQ_VISIT_{ijk}. \end{aligned}$$

For female patients GENF = 1, and thus the formula used to predict their CBC scores can be expressed as

$$\begin{aligned} \widehat{SQR_CBC}_{ijk} &= 7.49255 + 0.28977 * (1) - 0.90730 * VISIT_{ijk} \\ &\quad + 0.13703 * SQ_VISIT_{ijk} - 0.10034 * GVISIT_{ijk} \\ &= 7.78232 - 0.90730 * VISIT_{ijk} + 0.13703 * SQ_VISIT_{ijk} - 0.10034 * GVISIT_{ijk}. \end{aligned}$$

Table XXX.9 below shows the expected square roots of CBC scores for the various groups formed by the gender groups and interaction term at all measurement occasions. In Table XXX.10, the same expected scores are given in the scale of the original total score on the Child Behavioral Checklist. The initial impression, based on the positive coefficient of GENF, that females had higher expected CBC scores than males, seems to hold at the onset of the study. However, these tables show that, after the effects of the other variables are also taken into account, females are likely to have a slightly lower score than males at the end of the study period.

Table XXX.9: Expected square root of CBC scores

Gender	Visit			
	0	1	2	3
Male	7.4926	6.7223	6.2261	6.0039
Female	7.7823	6.9117	6.3152	5.9927

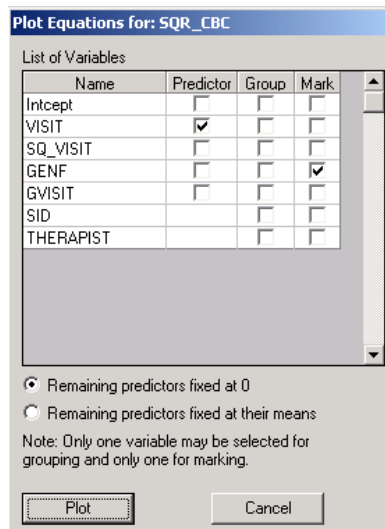
Table XXX.10: Expected CBC scores in original scale

Gender	Visit			
	0	1	2	3
Male	56.1383	45.1890	38.7639	36.0471
Female	60.5645	47.7717	39.8813	35.9121

The results in these tables can also be depicted graphically using the **File, Model-based Graphs** menu. This menu offers three options, namely equation modeling, residual plots and confidence intervals for random effects.

Equation modeling

To plot the trends in CBC scores for gender groups over successive visits, make sure the **Model Setup** window is activated by clicking on it before select the **Equations** option from the **File, Model-based Graphs** menu.



This activates the **Plot Equations for: SQR_CBC** dialog box. Select VISIT as the predictor, and request marking by gender as shown in the image below. Note that, by default, remaining predictors are fixed at 0. Click **Plot** to display the graphing window.




Figure XXX.26: Plot of SQR_CBC vs. VISIT for gender groups

By default, graphs using a two-category marking variable such as GENF will be displayed using blue and green to indicate the categories. To make the distinction between the groups of interest more clear, and create a graph that can be included in a report or paper to be printed in black and white, the plotting symbols can be changed. Here, we opt to change the line for female patients to a black, dotted line. Double-click on the line for this group in the legend box shown at the top right of the graph to activate the **Plot Parameters** dialog box. Next, click the **Line Attributes** button to load the **Line Parameters** dialog box.

Change the line style to dotted using the **Style** drop-down list box, and select black from the **Color** drop-down list box. Click **OK** on both the **Line Parameters** and **Plot Parameters** dialog boxes to obtain the final graph shown below.

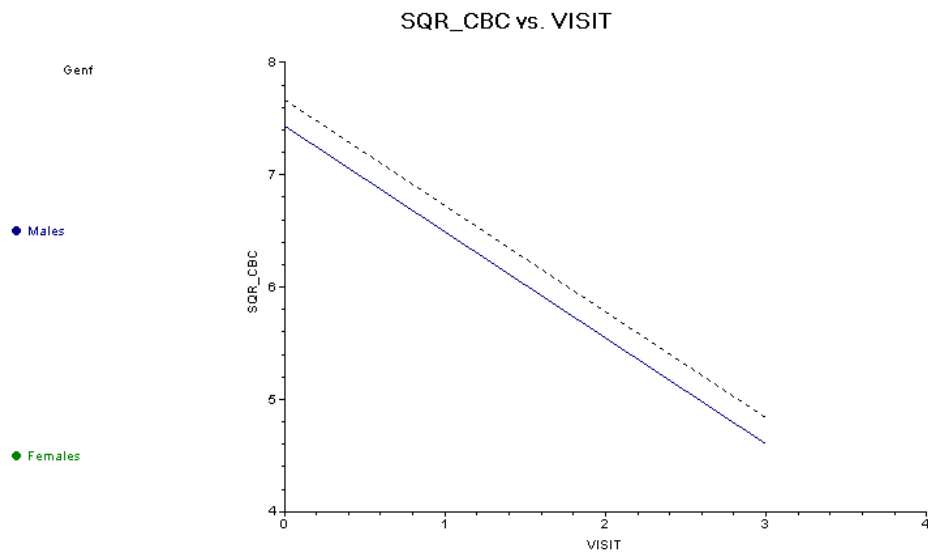
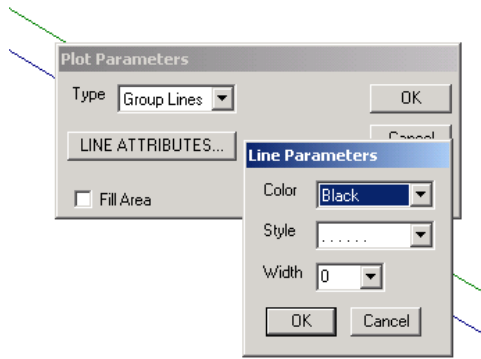


Figure XXX.27: Modified plot of SQR_CBC vs. VISIT for gender groups

The predicted decrease in CBC score echoes the results of the maximum likelihood estimation of the fixed effects, where a negative coefficient of -0.9073 was reported for the predictor VISIT. While the predicted intercept for males at the beginning of the study is approximately 7.5 as indicated in the graph at the top-left of the graphing window, the predicted intercept for the same group has decreased to approximately 4.75 by the final visit. This is lower than reported in Table XXX.10, where calculations showed an expected CBC score of 6.00 for males by the final

visit. The reason for this difference can be found in the formula used to produce the graph: recall that all remaining predictors were fixed to a value of 0. Whereas the result for males at the final visit shown in Table XXX.10 was based on the calculation

$$\begin{aligned}\widehat{\text{SQR_CBC}}_{ijk} &= 7.49255 + 0.28977 * (0) - 0.90730 * \text{VISIT}_{ijk} \\ &\quad + 0.13703 * \text{SQ_VISIT}_{ijk} - 0.10034 * (0) \\ &= 7.49255 - 0.90730 * \text{VISIT}_{ijk} + 0.13703 * \text{SQ_VISIT}_{ijk},\end{aligned}$$

the line shown for this group in the graph above is based on the formula

$$\widehat{\text{SQR_CBC}}_{ijk} = 7.49255 + 0.28977 * (0) - 0.90730 * \text{VISIT}_{ijk}.$$

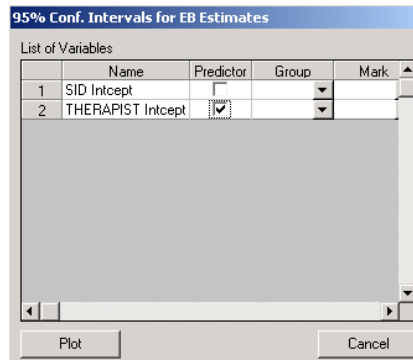
As a result, the predicted outcome shown in the graph for males at the end of the study will be $(0.13703)(9) = 1.2333$ units lower than reported in Table XXX.10. This difference underlines the fact that care should be taken when selecting the treatment for remaining predictors in the model. In this case, both SQ_VISIT and GVISIT can assume the value of 0, and thus using the remaining predictors fixed at zero option is permissible. In cases where predictors cannot assume a value of zero, the better choice would be to fix remaining predictors at their mean values instead when completing the **Plot Equations for:** dialog box.

Confidence intervals for random coefficients

The **Confidence Intervals** option on the **File, Model-based Graphs** menu provides the option to display confidence intervals for the empirical Bayes estimates of the random effects specified in a given model. This option is now used to examine the confidence intervals of the random intercepts for the therapists, who represent the highest level of the hierarchy in the current example.

Select the **Confidence Intervals** option on the **File, Model-based Graphs** menu to activate the **95% Conf. Intervals for EB estimates** dialog box. A simple graph of the

confidence intervals for the empirical Bayes estimates of the intercepts at the therapist level is obtained by selecting THERAPIST Intcept in the **Predictor** column before clicking **Plot**. Note that it is also possible to select both a grouping and marking variable to be used in the graph.



The graph obtained, as seen below, shows that, in general, the range of the confidence intervals for the level-3 empirical Bayes estimates of the intercepts is $(-2; 2)$. The deviations from the estimated population intercept over therapists are also apparent. Each confidence interval is obtained using

$$\text{Empirical Bayes residuals} \pm 1.96\sqrt{\text{var}(\text{Empirical Bayes residuals})}.$$

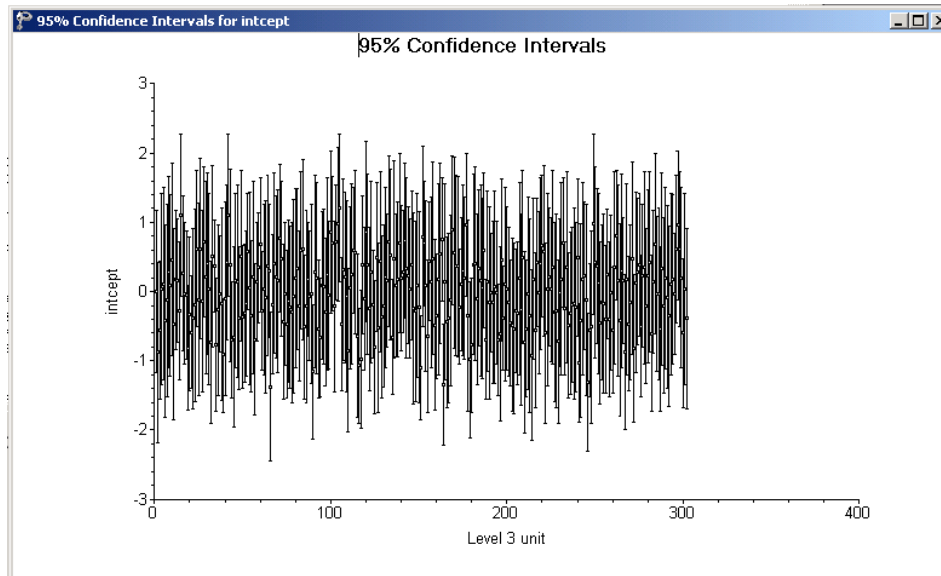


Figure XXX.28: 95% confidence intervals for level-3 units

Fit statistics

Recall that for the growth curve model the following indices were obtained:

- Log Likelihood: -14946.6236
- -2 Log Likelihood (Deviance): 29893.2473
- Akaike's Information Criterion: 29905.2473
- Schwarz's Bayesian Criterion: 29929.8492
- Number of free parameters: 6

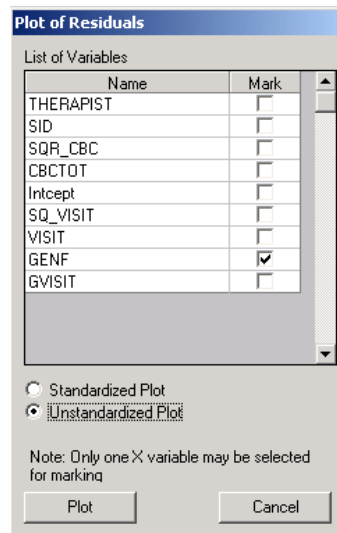
When the deviances of the two models are compared, a χ^2 -statistic of $29893.2473 - 29882.6732 = 10.57$ with $8 - 6 = 2$ degrees of freedom is obtained. This indicates that the current model fits the data better than the growth curve model. The AIC decreased from 29905.2473 to 29898.6732, and also favors the use of the current model. The SBC, however, increased slightly, from 29929.8492 to 29931.4757, and thus favors the growth curve model previously fitted as the more parsimonious. Note, however, that the changes in all three criteria are rather small.

1.4.4.5 Residuals

Residual plots

The **Residuals** option on the **File, Model-based Graphs** menu is used to examine the residuals obtained for a fitted model. It is useful for examining the fit of the model, and also as a check for possible distributional assumption violations. As residuals are defined as the difference between the observed and predicted outcomes, trends in residuals, for example over the course of a study in a longitudinal data set, may indicate that an important predictor was not included in the model fitted to the data.

Select the **Residuals** option on the **File, Model-based Graphs** menu to activate the **Plot of Residuals** dialog box. To simultaneously check for any differences in residuals for the gender groups, select GENF as marking variable. Opt to create an unstandardized plot of the residuals by selecting the **Unstandardized Plot** option rather than the default **Standardized Plot** option. Click **Plot**.



The graph below shows the residuals for the gender groups in the default colors of blue and green. To make the distinction between the groups more clear, click on the plotting symbol for the female group in the legend box.


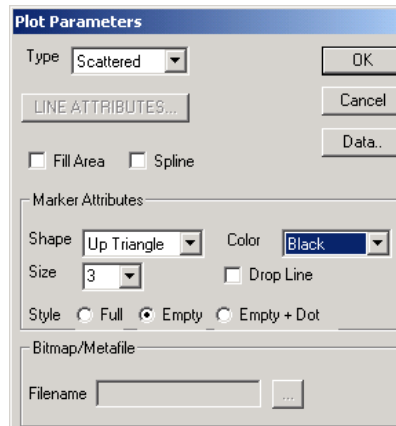
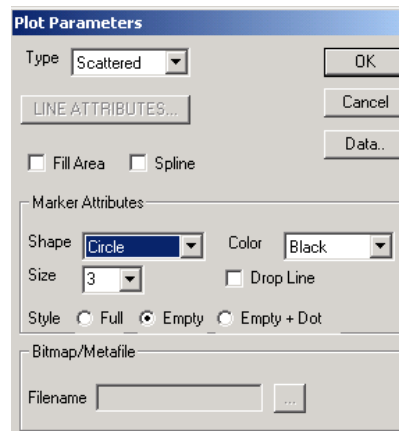


Figure XXX.29: Level-1 residual plot by gender group

The **Plot Parameters** dialog box appears. Change the **Shape** of the symbol to "Up Triangle," adjust the **Size** to 3 and change the display **Color** to black as shown below. Click **OK** when done.



Click on the symbol for the male group next, and change the parameters for this group to those shown in the dialog box below. Click **OK** to return to the graphing window.



The final plot is shown below. The residuals are clustered reasonably symmetrically around the 0 tick mark on the Y-axis, and no gender pattern can be discerned for the larger residuals. A single residual, for a male respondent, has a value larger than 10. This potential outlier can be identified using the **Data** option on the **Plot Parameters** dialog box (see above).




Figure XXX.30: Modified level-1 residual plot by gender group

1.4.5 Fitting a random intercepts and slopes model

1.4.5.1 The model

The graphs obtained during the exploratory analysis of the CBC data showed that the change in total CBC score over the course of the study differed from patient to patient. Because of this, the models fitted in Sections XXX.2 and XXX.3 allowed for the intercepts to vary randomly at both patient and therapist level. In effect, we assumed that each patient may have a different starting point. These models indicated a statistically significant relationship between the observed CBC score and the measurement occasion. To test whether there is significant variation in the way individual patients' scores change over the study period, a model in which both intercepts and slopes of the predictor VISIT are allowed to vary randomly can be used.

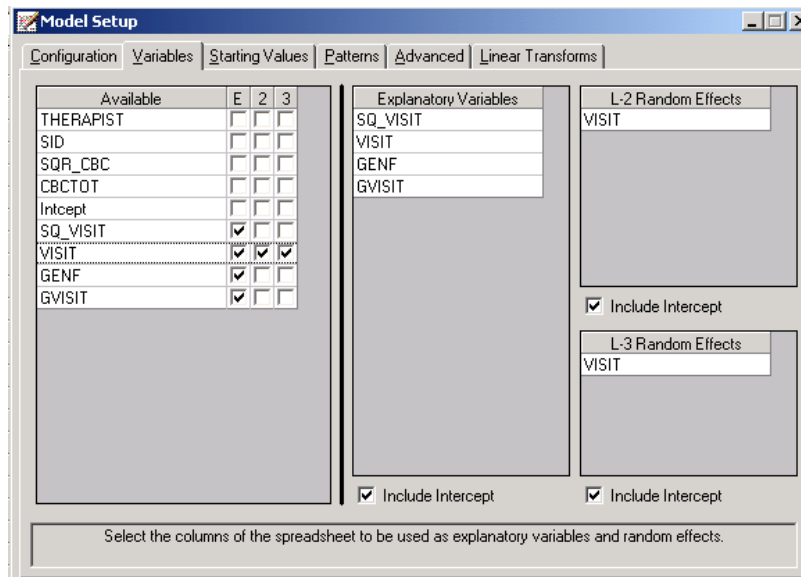
The model can be formulated as

$$\begin{aligned} \text{SQR_CBC}_{ijk} = & \beta_0 + \beta_1 * \text{GENF}_{ij} + \beta_2 * \text{VISIT}_{ijk} + \beta_3 * \text{SQ_VISIT}_{ijk} + \beta_4 * \text{GENF}_{ij} * \text{VISIT}_{ijk} \\ & + (v_{i0} + v_{i1} * \text{VISIT}_{ijk}) + (v_{ij0} + v_{ij1} * \text{VISIT}_{ijk}) + e_{ijk} \end{aligned}$$

At level 2, two random coefficients are now included: v_{ij0} represents the random intercept and v_{ij1} the random coefficient for the slope of the predictor VISIT. The random coefficients v_{i0} and v_{i1} serve the same purpose at level 3 (the therapist level) of the model.

1.4.5.2 Setting up the analysis

Again, we use the SuperMix spreadsheet **cbtot.ss3** and the model specification file **cbtot.mum** discussed in the previous example as the starting point. With the model specification file open, click on the **Variables** tab of the **Model Setup** window. Add random coefficients for the predictor VISIT to levels 2 and 3 of the model by checking the boxes next to these variables in the **2** and **3** column, as shown below. Save the changes to the model specification file, using the **File, Save** option to overwrite the previous specification file or the **File, Save as** option to assign a new filename. Click **Run** on the **Analysis** menu to perform the analysis.

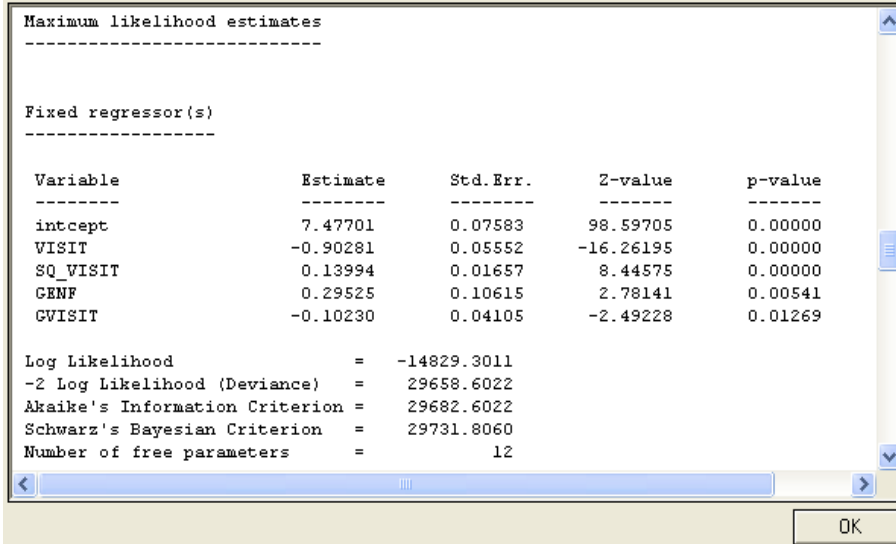


1.4.5.3 Discussion of results

Partial output is given below. We focus on those parts of the output that differ from the output obtained for the previous analysis, and conclude with a discussion of the additional output files containing the empirical Bayes residuals.

Fixed effects results

The inclusion of random VISIT slopes at levels 2 and 3 of the hierarchy has very little impact the estimated fixed coefficients. Results for the fixed part are shown below.



```
Maximum likelihood estimates
-----

Fixed regressor(s)
-----

Variable           Estimate      Std. Err.      Z-value      p-value
-----
intcept            7.47701      0.07583       98.59705    0.00000
VISIT              -0.90281     0.05552      -16.26195    0.00000
SQ_VISIT           0.13994     0.01657       8.44575     0.00000
GENF               0.29525     0.10615       2.78141     0.00541
CVISIT            -0.10230     0.04105      -2.49228     0.01269

Log Likelihood      = -14829.3011
-2 Log Likelihood (Deviance) = 29658.6022
Akaike's Information Criterion = 29682.6022
Schwarz's Bayesian Criterion = 29731.8060
Number of free parameters = 12
```

Random effects results

Turning to the estimated coefficients in the random part of the model, we note a change in the between measurement (level-1) variation, which has decreased from 2.33647 to 1.88939. This illustrates that the addition of a random coefficient at any level of a model can affect the random effect(s) at another level.

Variance/covariance components				

Level 3	Estimate	Std. Err.	Z-value	p-value
-----	-----	-----	-----	-----
intcept /intcept	0.58870	0.11432	5.14940	0.00000
VISIT /intcept	-0.03108	0.02873	-1.08181	0.27934
VISIT /VISIT	0.03869	0.01290	2.99923	0.00271
Level 2	Estimate	Std. Err.	Z-value	p-value
-----	-----	-----	-----	-----
intcept /intcept	3.04361	0.16072	18.93748	0.00000
VISIT /intcept	-0.11485	0.04929	-2.32999	0.01981
VISIT /VISIT	0.23823	0.02550	9.34216	0.00000
Level 1	Estimate	Std. Err.	Z-value	p-value
-----	-----	-----	-----	-----
intcept /intcept	1.88939	0.04582	41.23789	0.00000

At levels 2 and 3 we find evidence of significant variation in the VISIT slopes. While not of the same magnitude as the intercept variation, it is clear that it is more realistic to allow the slopes to vary from patient to patient, and from therapist to therapist, than to assume that the VISIT slope can be described adequately by a common fixed effect as was done in the previous model.

1.4.5.4 Interpreting the results

Fit statistics and ICC

Model fit

When the measures of fit are compared to those of the random-intercepts-only model, it becomes clear that the current model fits the data better. Recall that for the random intercepts model the following fit measures were obtained:

- Log likelihood: -14941.3366
- -2 log Likelihood (Deviance): 29882.6732
- Akaike's Information Criterion: 29898.6732
- Schwarz's Bayesian Criterion: 29931.4757
- Number of free parameters: 8

While four more parameters had to be estimated in the random intercepts and slopes model, the deviance decreased significantly. The χ^2 -statistic for comparing these models is $29882.6732 - 29658.6022 = 223.9185$, with 4 degrees of freedom. The improved fit of the current model is also clear from the other fit measures: both the AIC and the SBC clearly favor the current model, and have decreased substantially from those reported for the random-intercepts-only model.

Percentage variation explained

To take a closer look at the amount of variation explained at the levels of the hierarchy, the total variation at each level has to be calculated. At level 3, we have three variance/covariance components to take into account. Recall that the model included two random effects, namely v_{i0} and $v_{i1}(VISIT_{ijk})$. The total variation at this level follows as

$$\begin{aligned}
 Var(level - 3) &= \text{var}(v_{i0} + v_{i1}(VISIT_{ijk})) \\
 &= \text{var}(v_{i0}) + \text{var}(v_{i1}(VISIT_{ijk})) + \text{cov}(v_{i0}, v_{i1}(VISIT_{ijk})) \\
 &= \text{var}(v_{i0}) + (VISIT_{ijk})^2 \text{var}(v_{i1}) + 2(VISIT_{ijk}) \text{cov}(v_{i0}, v_{i1}) \\
 &= 0.58870 + 0.03869(VISIT_{ijk})^2 - 2(0.03108)(VISIT_{ijk}) \\
 &= 0.58870 + 0.03869(VISIT_{ijk})^2 - 0.06216(VISIT_{ijk})
 \end{aligned}$$

At level 2, the total variation can be expressed in similar fashion as

$$\begin{aligned}
 Var(level - 2) &= \text{var}(u_{ij0} + u_{ij1}(VISIT_{ijk})) \\
 &= \text{var}(u_{ij0}) + (VISIT_{ijk})^2 \text{var}(u_{ij1}) + 2(VISIT_{ijk}) \text{cov}(u_{ij0}, u_{ij1}) \\
 &= 3.04361 + 0.23823(VISIT_{ijk})^2 - 0.2297(VISIT_{ijk})
 \end{aligned}$$

The total variation in the model is

$$\begin{aligned} \text{Total Var} &= \text{var}(\text{level} - 1) + \text{var}(\text{level} - 2) + \text{var}(\text{level} - 3) \\ &= 5.5217 + 0.27692(\text{VISIT}_{ijk})^2 - 0.29186(\text{VISIT}_{ijk}) \end{aligned}$$

The variation at the higher levels and, consequently, the total variation are a function of the measurement occasion, as represented by the predictor VISIT. For example, at the start of the study we find that the total variation is equal to 5.5217, with 0.58870 at level 3 and 3.04361 at level 2. This indicates that at the time of the first visit,

$$\frac{0.58870}{5.5217} \times 100 = 10.66\%$$

of the total variation explained by this model is at a therapist level. By the end of the study, VISIT assumes a value of 3, and thus the total variation is equal to 7.1384. The total variation at the therapist level at the last measurement occasion is 0.75043, and thus the percentage of variation at therapist level at the end of the study is

$$\frac{0.75043}{7.1384} \times 100 = 10.51\%.$$

At a patient level, the corresponding percentages of variation at the first and last visit are

$$\frac{3.04361}{5.5217} \times 100 = 55.12\%.$$

and

$$\frac{4.49858}{7.1384} \times 100 = 63.02\%$$

respectively. While the total variation explained at a therapist level declines over visits, there is an increase of approximately the same size in the total variation explained at a patient level over visits. The variation over patients is consistently much higher than over therapists or over measurements nested within patients.

1.4.6 References

Achenbach, T.M. (1991). Manual for the Child Behavioral Checklist and 1991 profile. Burlington: University of Vermont Department of Psychiatry.

THIS IS INCOMPLETE STILL.