



Three-level models for continuous outcomes

Contents

3.1	3-LEVEL CONTINUOUS EXAMPLE USING A SUBSET OF THE SCHOENWALD DATA.....	1
3.1.1	<i>The data</i>	1
3.1.2	<i>Exploring the data</i>	2
3.1.3	<i>Fitting a growth curve model to the data</i>	15
3.1.4	<i>Fitting a random intercept model with 3 predictors and interaction term to the data</i>	23
3.1.5	<i>Fitting a random intercepts and slopes model</i>	34

3.1 3-level continuous example using a subset of the Schoenwald data

3.1.1 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.

3.1.2 Exploring the data

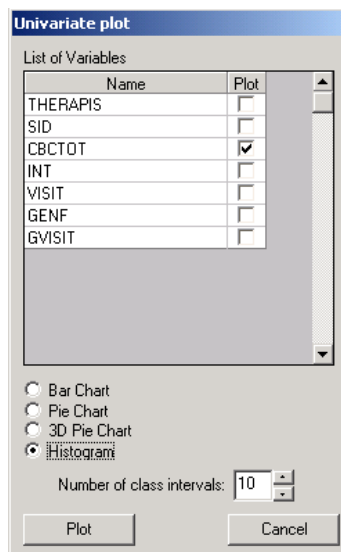
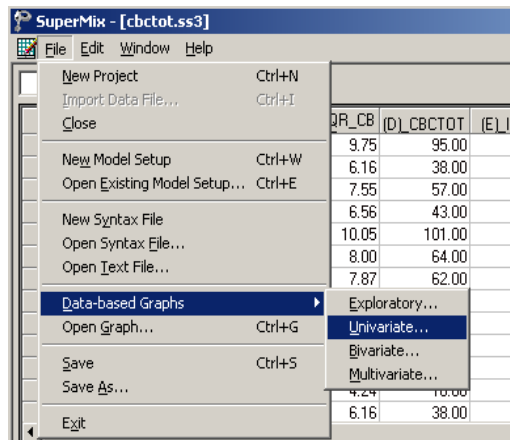
Relationships between variables, and trends over time in repeated measurement data, may be conveyed in an informal and simplified visual form via graphical displays. SuperMix offers both data-based and model-based graphs. Data-based graphing options are accessed via the **File, Data-based Graphs** option 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.

3.1.2.1 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.

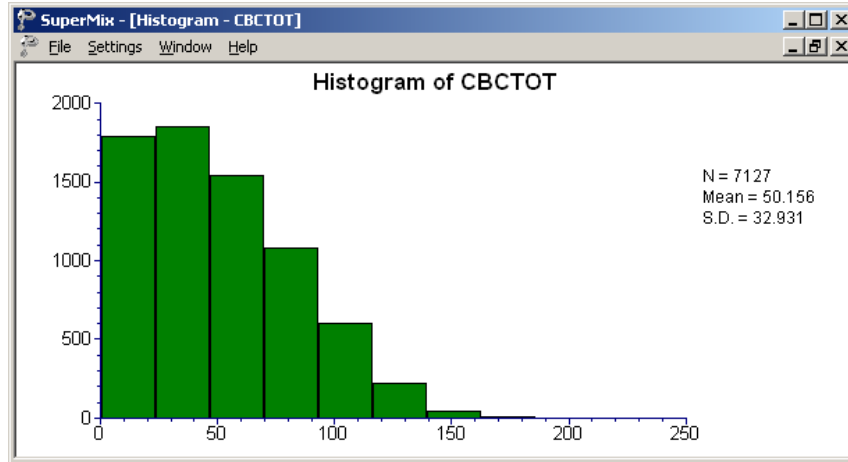
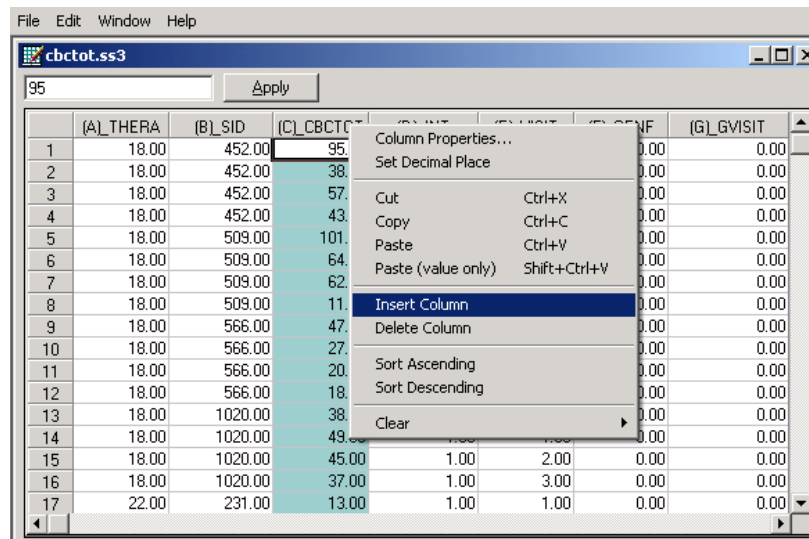


Figure 3.14: Histogram of the variable CBCTOT

The histogram above 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.

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.

	(A)_THERA	(B)_SID	(C)_C	(D)_CBCTO	(E)_INT	(F)_VISIT	(G)_GENF
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	0.00
3	18.00	452.00	7.55	57.00	1.00	2.00	0.00
4	18.00	452.00	6.56	43.00	1.00	3.00	0.00
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	0.00
7	18.00	509.00	7.87	62.00	1.00	2.00	0.00
8	18.00	509.00	3.32	11.00	1.00	3.00	0.00
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	0.00
11	18.00	566.00	4.47	20.00	1.00	2.00	0.00
12	18.00	566.00	4.24	18.00	1.00	3.00	0.00
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	0.00
15	18.00	1020.00	6.71	45.00	1.00	2.00	0.00
16	18.00	1020.00	6.08	37.00	1.00	3.00	0.00
17	22.00	231.00	3.61	13.00	1.00	1.00	0.00

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

(C)_C	(D)_CBCTO	(E)_INT	(F)_VISIT
6.71	45.00	1.00	2.00

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.

Column Properties

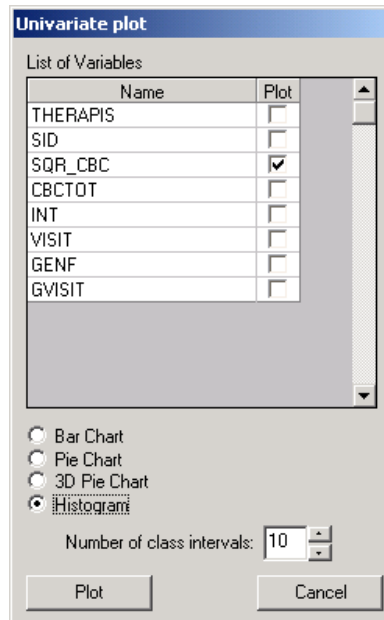
Header: SQR_CBC

Number of Decimal Places: 2

Continuous

OK Cancel

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.




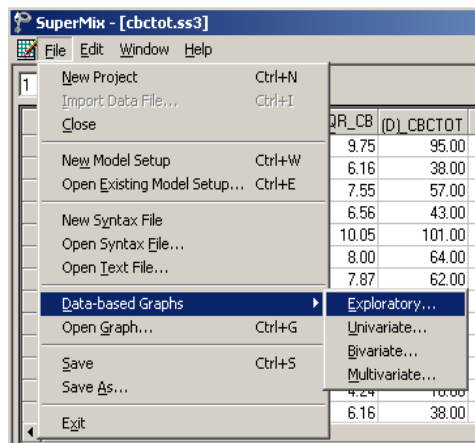
Figure 3.15: Histogram of the variable SQR_CBC

3.1.2.2 Bivariate graphs

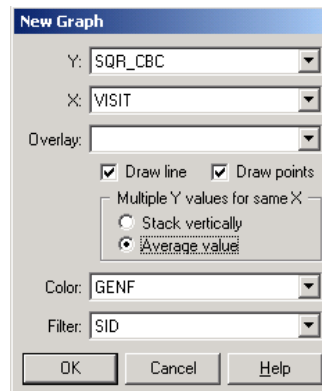
It is hoped that the total scores of patients would change over time, *i.e.*, with successive visits to their therapists. In addition, it is hypothesized that the gender of a patient may also have some relationship to the total score of a patient. Bivariate plots of possible relationships are a handy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



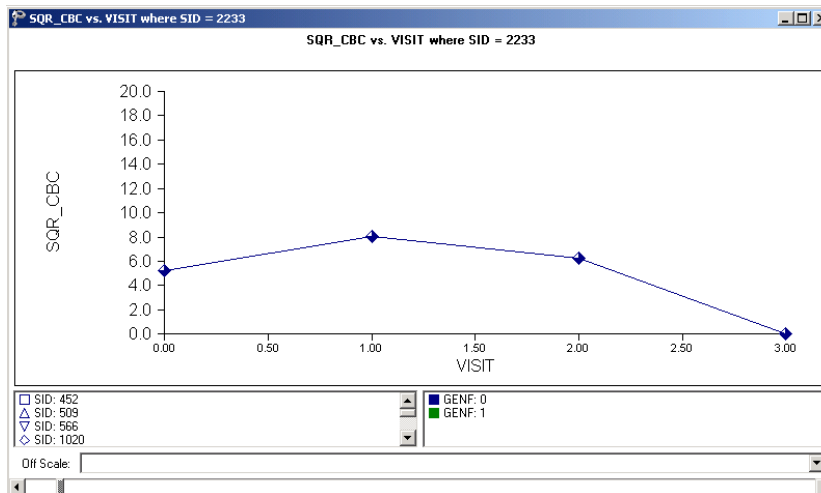


Figure 3.16: Relationship between SQR_CBC and VISIT for selected patients

Transforming a variable

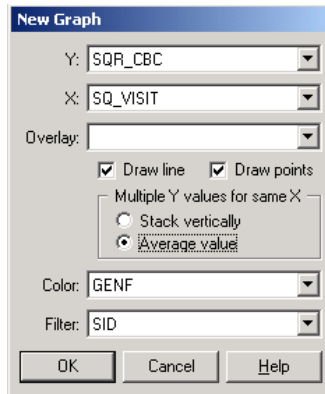
To examine the relationship between the outcome and the quadratic value of VISIT, a new variable has to be created. This is done in a similar way to adding the square root of the total score. First insert a column, then type the appropriate function into the formula box as shown below. Click the **Apply**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

3.1.2.3 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.




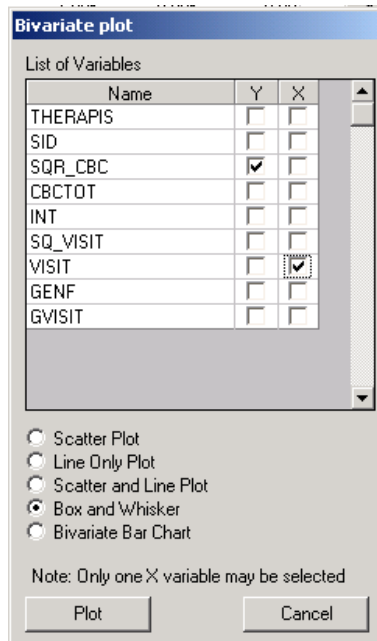


Figure 3.17: relationship between SQR_CBC and SQ_VISIT for selected patients

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 3.18: 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).

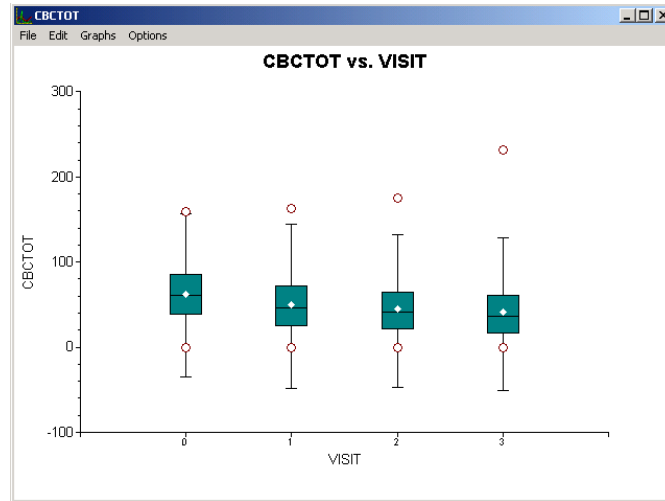


Figure 3.19: Box-and-whisker plot of CBCTOT vs. VISIT

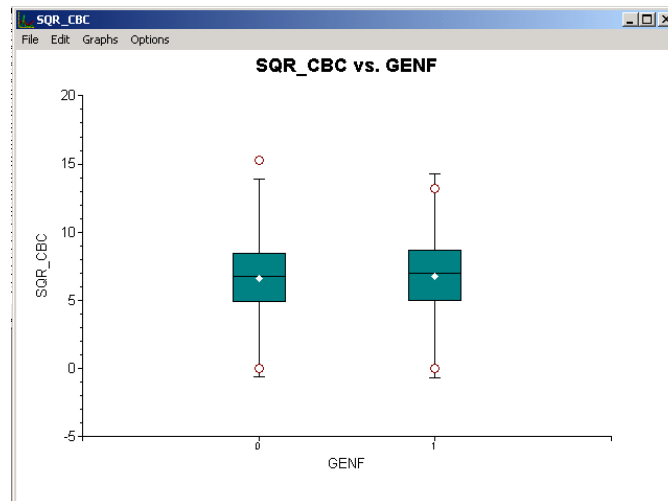


Figure 3.20: Box-and-whisker plot of SQR_CBC vs. GENF

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).

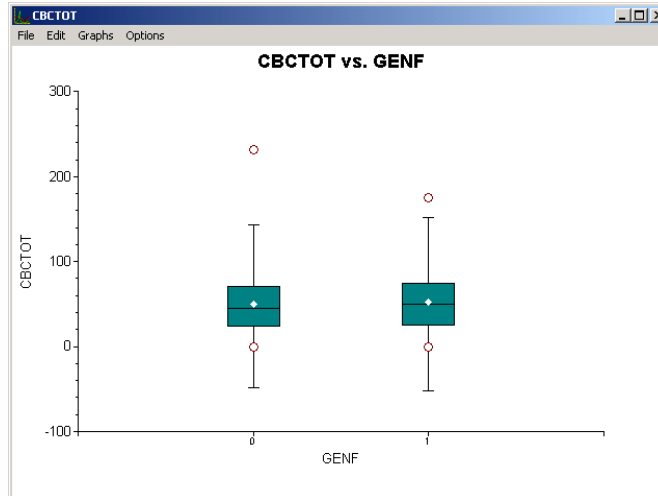
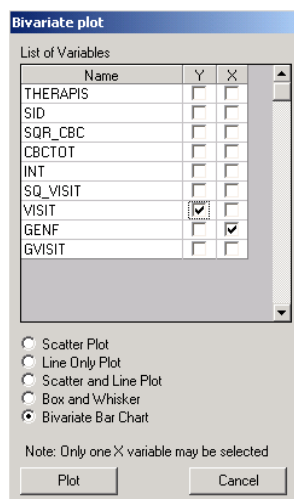


Figure 3.21: Box-and-whisker plot of CBCTOT vs. GENF

When Figure 3.20 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.

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.






Figure 3.22: Bivariate chart of VISIT vs GENF

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.

3.1.3 Fitting a growth curve model to the data

3.1.3.1 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:

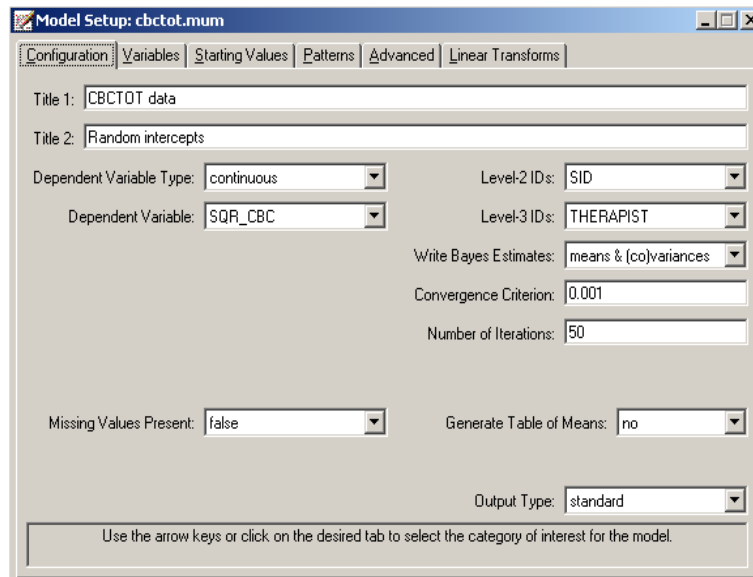
$$\text{SQR_CBC}_{ijk} = \beta_0 + \beta_1 * \text{VISIT}_{ijk} + \beta_2 * \text{SQ_VISIT}_{ijk} + v_{i0} + v_{ij0} + e_{ijk}$$

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.

3.1.3.2 Setting up the analysis

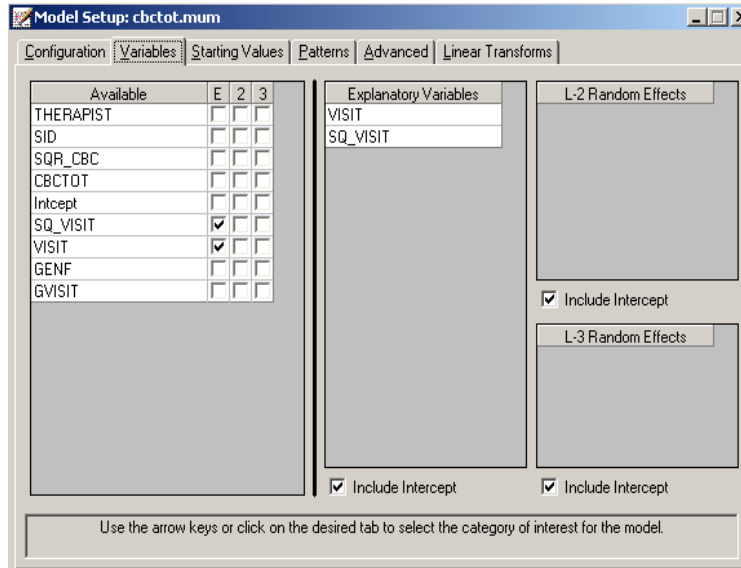
Open the SuperMix spreadsheet **cbtot.ss3**. The next step is to describe the model to be fitted. We use the SuperMix interface to provide the model specifications. From the main menu bar, select the **File, New Model Setup** option.

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 **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.

3.1.3.3 Discussion of results

Portions of the output file **cbtot.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

OK

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.

SuperMix - [cbctot.out]

File Analysis Window Help

Descriptive statistics for all variables

Variable		Minimum	Maximum	Mean	Stand. Dev.
Dependent					

SQR_CBC		0.0000	15.2320	6.6187	2.5199
Random-Effects					

intercept	(3)	1.0000	1.0000	1.0000	0.0000
intercept	(2)	1.0000	1.0000	1.0000	0.0000
intercept	(1)	1.0000	1.0000	1.0000	0.0000
Fixed Regressor(s)					

intercept		1.0000	1.0000	1.0000	0.0000
VISIT		0.0000	3.0000	1.4416	1.1124
SQ_VISIT		0.0000	9.0000	3.3153	3.4508

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 $\beta_0, \beta_1, \dots, \beta_3$ 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	7.59246	0.06797	111.69777	0.00000
VISIT	-0.94119	0.05679	-16.57386	0.00000
SQ_VISIT	0.13671	0.01829	7.47447	0.00000

Log Likelihood = -14946.6389
 -2 Log Likelihood (Deviance) = 29893.2778
 Akaike's Information Criterion = 29905.2778
 Schwarz's Bayesian Criterion = 29929.8797
 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.

Variance/covariance components

Level	Estimate	Std. Err.	Z-value	p-value
Level 3				
intcept /intcept	0.58201	0.10414	5.58882	0.00000
Level 2				
intcept /intcept	3.08098	0.13341	23.09426	0.00000
Level 1				
intcept /intcept	2.34084	0.04599	50.90214	0.00000

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).

3.1.3.4 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

$$\begin{aligned} \text{SQR_CBC}_{ij0}^{\wedge} &= \hat{\beta}_0 + \hat{\beta}_1 * \text{VISIT}_{ij0} + \hat{\beta}_2 * \text{SQ_VISIT}_{ij0} \\ &= 7.59246 - 0.94119(0) + 0.13671(0^2) \\ &= 7.59246, \end{aligned}$$

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

$$\begin{aligned} \text{VISIT 1: SQR_CBC}_{ij1}^{\wedge} &= 7.59245 - 0.94119(1) + 0.13671(1^2) \\ &= 7.59246 - 0.94119 + 0.13671 \\ &= 6.78798 \end{aligned}$$

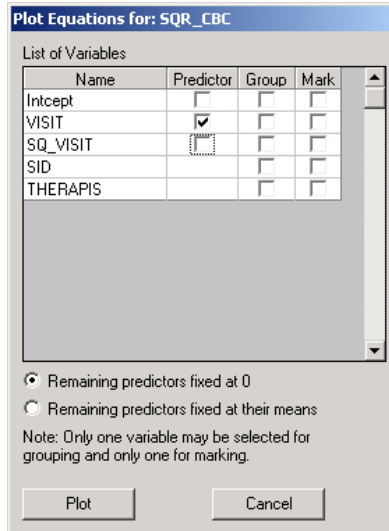
$$\begin{aligned} \text{VISIT 2: SQR_CBC}_{ij2}^{\wedge} &= 7.59245 - 0.94119(2) + 0.13671(2^2) \\ &= 7.59246 - 1.88238 + 0.54684 \\ &= 6.25692 \end{aligned}$$

$$\begin{aligned} \text{VISIT 3: SQR_CBC}_{ij3}^{\wedge} &= 7.59245 - 0.94119(3) + 0.13671(3^2) \\ &= 7.59245 - 2.82357 + 1.23039 \\ &= 5.99928 \end{aligned}$$

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.




Figure 3.23: Plot of SQR_CBC vs. VISIT

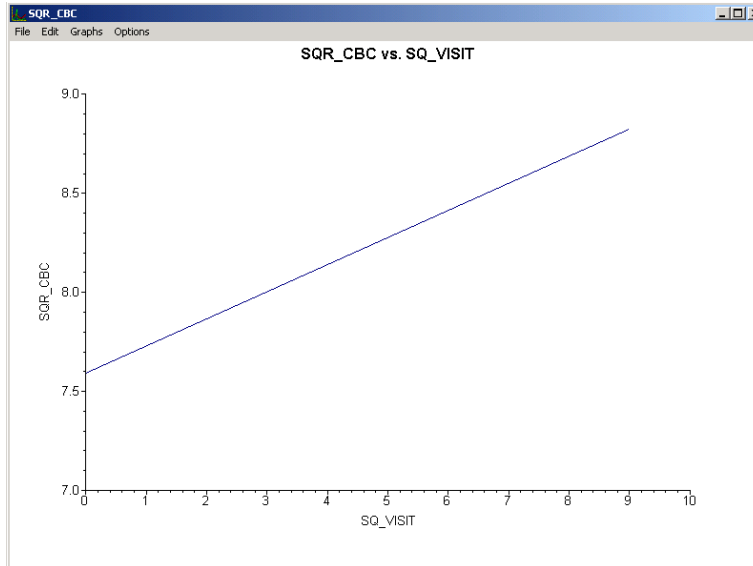


Figure 3.24: Plot of SQR_CBC vs. SQ_VISIT

A similar plot for the predictor SQ_VISIT is given in Figure 3.24. 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.

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.

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

3.1.4.1 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.

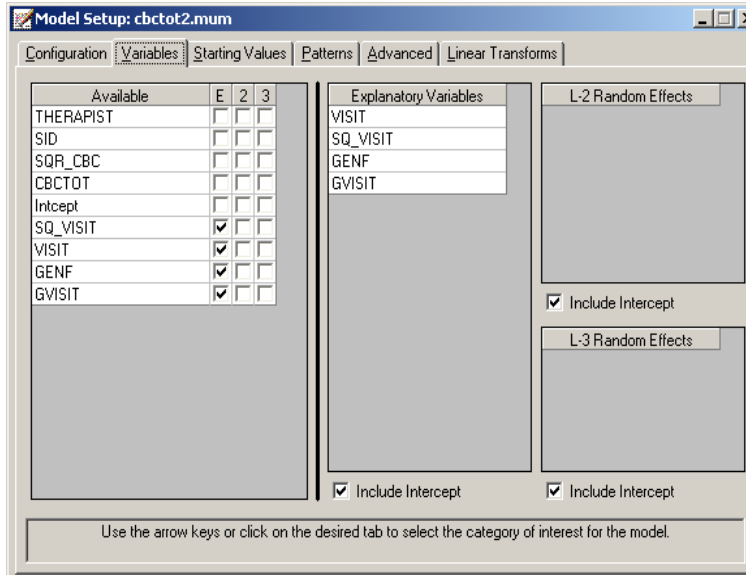
The model considered in this section uses the participant's gender, visit number, squared value of the visit number, and the interaction between visit number and gender (represented by the variable GVISIT in the data spreadsheet) to predict the square root of the total score on the Child Behavior Checklist. 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.

3.1.4.2 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.

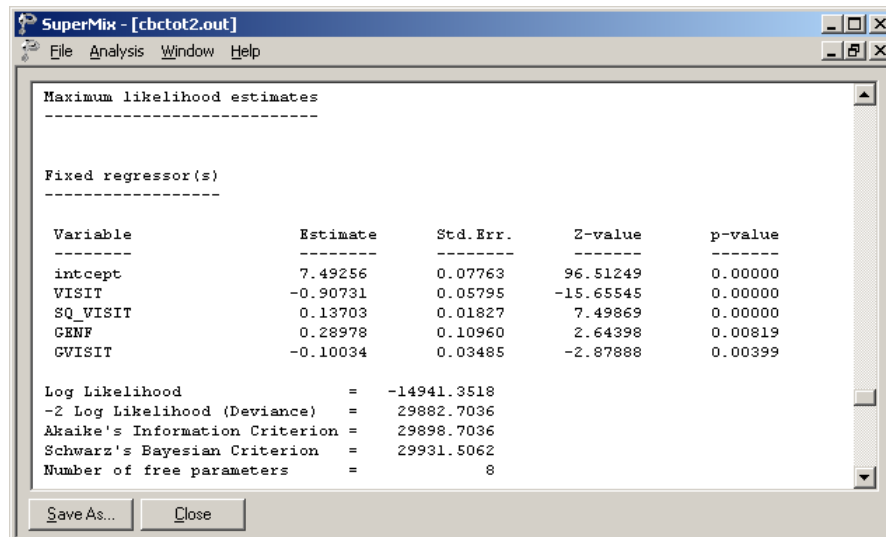


Save the modified model specification file as **cbtot2.mum**, then select the **Run** option from the **Analysis** menu to perform the analysis.

3.1.4.3 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.

Random effects results

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

The screenshot shows a window titled 'SuperMix - [cbctot2.out]' with a menu bar (File, Analysis, Window, Help). The main content area displays 'Variance/covariance components' with a table of results for three levels. The table has columns for Level, Estimate, Std. Err., Z-value, and p-value. The data is as follows:

Level	Estimate	Std. Err.	Z-value	p-value
Level 3 intcept /intcept	0.57846	0.10383	5.57142	0.00000
Level 2 intcept /intcept	3.08225	0.13340	23.10575	0.00000
Level 1 intcept /intcept	2.33648	0.04590	50.90206	0.00000

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

$$\frac{3.08225}{0.57846 + 3.08225 + 2.33648} \times 100\% = 51.39\%$$

and

$$\frac{0.57846}{0.57846 + 3.08225 + 2.33648} \times 100\% = 9.65\%$$

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.

3.1.4.4 Interpretation of the results

Estimated outcomes for different groups

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

$$\begin{aligned} \widehat{SQR_CBC}_{ijk} &= \hat{\beta}_0 + \hat{\beta}_1 * GENF_{ij} + \hat{\beta}_2 * VISIT_{ijk} + \hat{\beta}_3 * SQ_VISIT_{ijk} + \hat{\beta}_4 * GVISIT_{ijk} \\ &= 7.49255 + 0.28977 * GENF_{ij} - 0.90730 * VISIT_{ijk} \\ &\quad + 0.13703 * SQ_VISIT_{ijk} - 0.10034 * GVISIT_{ijk}. \end{aligned}$$

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 3.9 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 3.10, 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 3.9: 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 3.10: Expected CBC scores in original scale

Gender	Visit			
	0	1	2	3
Male	56.1385	45.1890	38.7638	36.0468
Female	60.5648	47.7719	39.8812	35.9120

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 Figure 3.25. Note that, by default, remaining predictors are fixed at 0. Click **Plot** to display the graphing window.


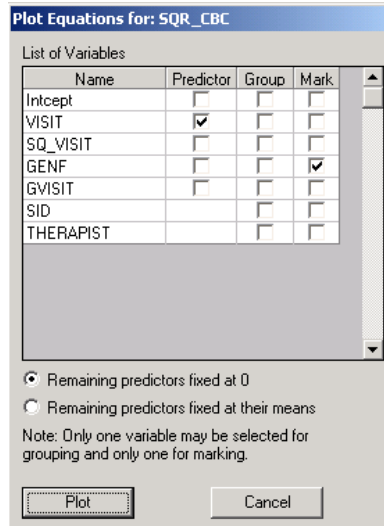


Figure 3.25: Plot of SQR_CBC vs. VISIT for gender groups

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.


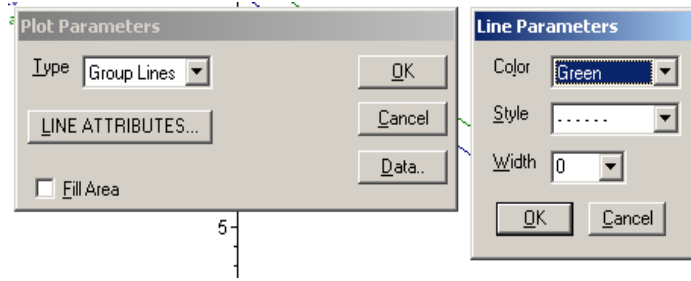


Figure 3.26: Modified plot of SQR_CBC vs. VISIT for gender groups

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 3.10, 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 3.10 was based on the calculation

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

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

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

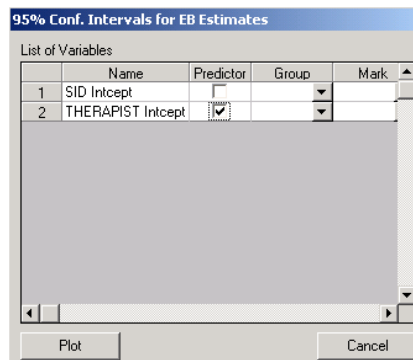
As a result, the predicted outcome shown in the graph for males at the end of the study will be $(0.13703)(9) = 1.2333$ units lower than reported in Table 3.10. This difference underlines the fact that care should be taken when selecting the treatment for remaining predictors in the model. In this case, 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.

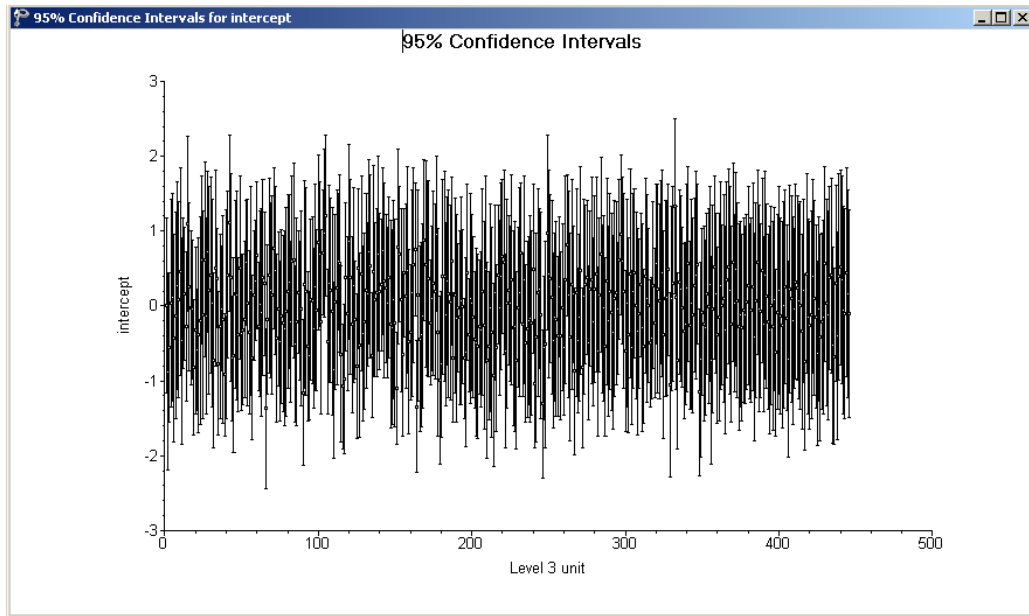


Figure 3.27: 95% confidence intervals for level-3 units

Each confidence interval is obtained using

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

Fit statistics

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

- Log Likelihood: -14946.6389
- -2 Log Likelihood (Deviance): 29893.2778
- Akaike's Information Criterion: 29905.2778
- Schwarz's Bayesian Criterion: 29929.8797
- Number of free parameters: 6

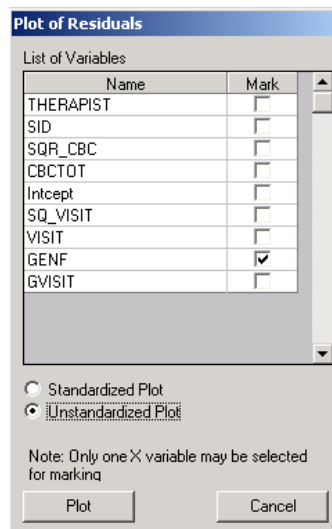
When the deviances of the two models are compared, a χ^2 -statistic of $29893.2778 - 29882.7036 = 10.57$ with $8 - 6 = 2$ degrees of freedom is obtained. This indicates that the current model fits the data better than the growth curve model. The AIC decreased from 29905.2778 to 29898.7036 , and also favors the use of the current model. The SBC, however, increased slightly, from 29929.8797 to 29931.5062 , 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.

3.1.4.5 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.


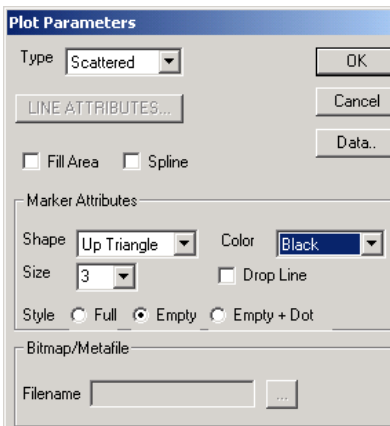
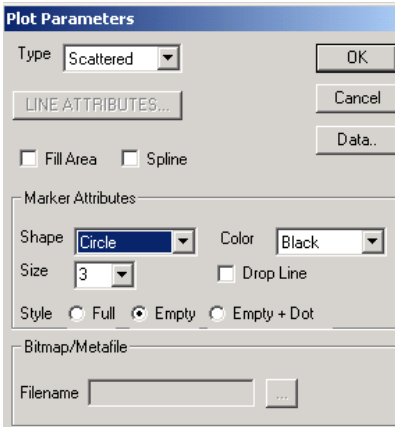


Figure 3.28: Level-1 residual plot by gender group

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).

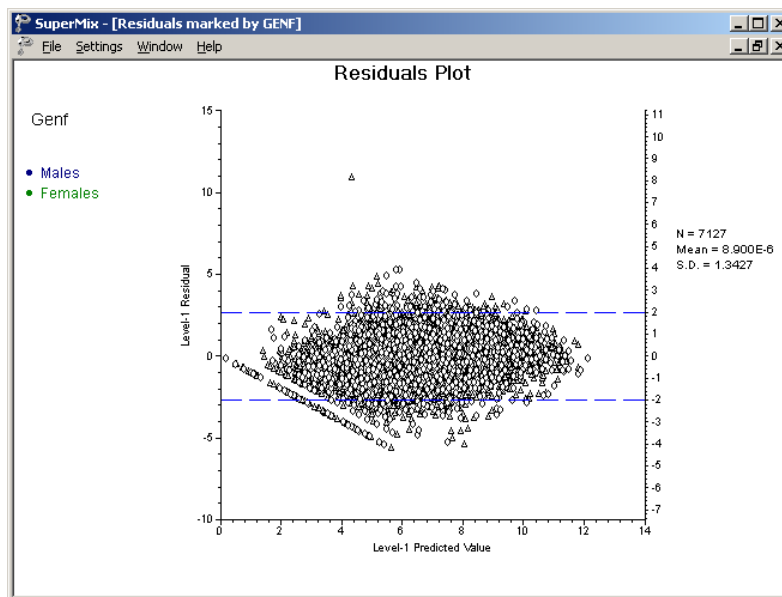


Figure 3.29: Modified level-1 residual plot by gender group

3.1.5 Fitting a random intercepts and slopes model

3.1.5.1 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 3.4.3 and 3.4.4 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.

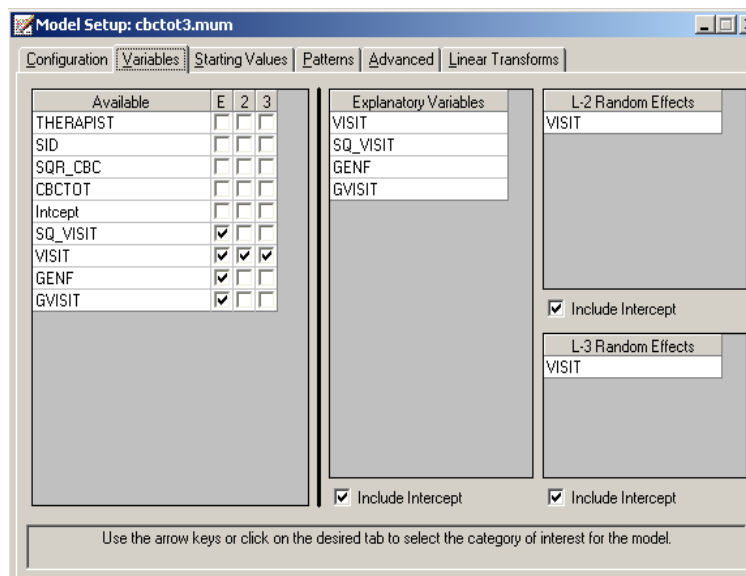
The model can be formulated as

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

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.

3.1.5.2 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.

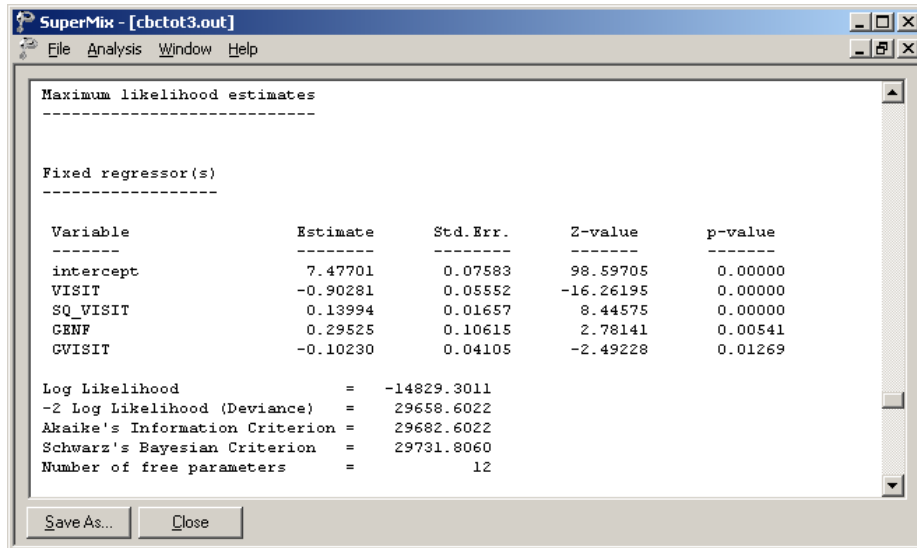


3.1.5.3 Discussion of results

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.



The screenshot shows a window titled "SuperMix - [cbctot3.out]" with a menu bar (File, Analysis, Window, Help). The main text area displays the following output:

```
Maximum likelihood estimates
-----

Fixed regressor(s)
-----

Variable          Estimate      Std. Err.    Z-value      p-value
-----
intercept         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
CENF             0.29525     0.10615     2.78141     0.00541
CWISIT          -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
```

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

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 2.33648 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.

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.

Variance/covariance components					

Level 3		Estimate	Std. Err.	Z-value	p-value

intercept	/intercept	0.58870	0.11432	5.14940	0.00000
VISIT	/intercept	-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

intercept	/intercept	3.04361	0.16072	18.93748	0.00000
VISIT	/intercept	-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

intercept	/intercept	1.88939	0.04582	41.23789	0.00000

3.1.5.4 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: -14941.3518
- -2 log Likelihood (Deviance): 29882.7036
- Akaike's Information Criterion: 29898.7036
- Schwarz's Bayesian Criterion: 29931.5062
- 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 $29882.7036 - 29658.6022 = 223.9185$, 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

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

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

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

The total variation in the model is

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

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.