Educational and Psychological Measurement

http://epm.sagepub.com/

Published by:

$SAGE

http://www.sagepublications.com

>> Version of Record - Oct 21, 2005

What is This?

USING THE SPSS MIXED PROCEDURE TO FIT CROSS-SECTIONAL AND LONGITUDINAL MULTILEVEL MODELS

JAMES L. PEUGH AND CRAIG K. ENDERS
University of Nebraska

Beginning with Version 11, SPSS implemented the MIXED procedure, which is capable of performing many common hierarchical linear model analyses. The purpose of this article was to provide a tutorial for performing cross-sectional and longitudinal analyses using this popular software platform. In doing so, the authors borrowed heavily from Singer's overview of SAS PROC MIXED, duplicating her analyses using the SPSS MIXED procedure.

Keywords: *SPSS; multilevel modeling; hierarchical linear modeling; longitudinal analysis; software application*

Multilevel modeling techniques have become an increasingly popular method during the past two decades for analyzing nested data structures (Farmer, 2000; Hox & Kreft, 1994; Kreft, 1995). The number of statistical packages capable of performing multilevel analyses has grown in recent years, and researchers conducting these analyses now have a wealth of software options. For example, specialized software is widely available (for software reviews, see Kreft, de Leeuw, & van der Leeden, 1994; Zhou, Perkins, & Hui, 1999), and multilevel analyses can also be performed using popular statistical packages such as Statistical Analysis Software (SAS) and the Statistical Package for the Social Sciences (SPSS).

The SAS multilevel procedure, PROC MIXED, has been available for some time, and Singer (1998) provided an excellent tutorial for conducting multilevel analyses using SAS. Beginning with Version 11, SPSS imple-

mented a mixed-model procedure that is capable of performing a variety of multilevel analysis. Because SPSS enjoys widespread use among social science researchers (O'Connor, 1999), the purpose of this article was to illustrate the use of the SPSS mixed-model procedure for performing multilevel analyses of cross-sectional and longitudinal data. In doing so, we borrowed heavily from Singer (1998). For consistency with Singer's (1998) earlier work, we demonstrated the analysis of cross-sectional data using a subset of the high school and beyond data (e.g., see Raudenbush & Bryk, 2002) and used Willett's (1988) opposites-naming task data to illustrate the analysis of longitudinal growth models.

For comparison purposes, the corresponding SAS PROC MIXED syntax (see Singer, 1998) for all analyses is presented in Appendix A. The SPSS and SAS syntax commands provided the same parameter estimates and tests of significance. The corresponding SAS output was omitted in the interest of space but is available on request.

Data File Structure

Before proceeding, a brief description of multilevel database structures is warranted. In the cross-sectional context (e.g., students nested within schools), a separate data record is needed for each case, and an identification variable is needed for both the Level 1 (e.g., students) and Level 2 (e.g., schools) units. Some multilevel software programs (e.g., HLM 5) require separate data files for Level 1 and Level 2 units, but this is not the case with SPSS. Instead, Level 1 and Level 2 variables are included in the same file such that the values of Level 2 variables are identical for all cases nested within a particular Level 2 unit. For example, children within a particular school would share the same value for a Level 2 (i.e., school-level) variable that quantifies the mean socioeconomic status (SES) for their school.

The data file structure for a longitudinal growth model analysis may seem counterintuitive to researchers accustomed to repeated-measures ANOVA analyses. Typically, longitudinal data are arranged in a multivariate format in which each case occupies one record (i.e., row) and each repeated measurement appears as a separate variable (i.e., column). In the case of a growth model, SPSS requires a "stacked" data file in which each case has multiple records, one for each measurement occasion. In a stacked file, the repeated measurements are contained within a single variable, and the data from each measurement occasion occupy a separate record. Two additional features are required in a stacked longitudinal data file. First, consistent with the cross-sectional context, each individual (i.e., Level 2 unit) must be assigned a unique identification number because multiple records, or repeated measurements, are nested within individuals. Second, a variable that reflects the timing of the assessments must also be included because this variable will serve

as a Level 1 predictor variable. For example, if the repeated measurements were equally and identically spaced for all individuals, the time variable for a linear growth model might take on values of 0, 1, 2, . . . , $t - 1$, where t is the number of repeated measures. In a developmental context, the time variable might reflect chronological age at each measurement occasion. Singer and Willett (2003, pp. 17-23) provided a detailed discussion of longitudinal data formats, including examples of stacked data files.

The SPSS VARSTOCASES command can be used to convert longitudinal data files from a multivariate format to the stacked format required for growth model analyses. To illustrate, consider a hypothetical educational data set consisting of 10 variables: an identification variable (ID), gender (GENDER), verbal achievement scores taken at four measurement occasions (VERBAL1 through VERBAL4), and age in months at each measurement occasion (AGE1 through AGE4; see upper portion of Appendix B). The VARSTOCASES command would produce a restructured data set with only four variables (ID, GENDER, VERBAL, and AGE), the syntax for which is shown below.

```
VARSTOCASES
/MAKE VERBAL FROM VERBAL1 VERBAL2 VERBAL3 VERBAL4
/MAKE AGE FROM AGE1 AGE2 AGE3 AGE4
/INDEX = WAVE
/KEEP = ID GENDER SCHOOL.
```

The /MAKE subcommand restructures each set of repeated measurements into a single variable (VERBAL and AGE) by stacking their values into a single column, such that each measurement occasion appears on a separate record (e.g., VERBAL1 and AGE1 scores appear in the first record for the ith case, VERBAL2 and AGE2 in the second record, etc.). The /INDEX subcommand creates a new variable (named WAVE in this example) with integer values that correspond to the measurement occasion of each repeated measure (e.g., the records containing VERBAL1 scores would have WAVE values equal to 1, records with VERBAL2 scores would have WAVE values equal to 2, etc.). The /KEEP subcommand is used to specify the time-invariant variables (e.g., ID, GENDER) that are retained in the stacked file (see the lower portion of Appendix B). Consistent with the cross-sectional data structure, values of time-invariant Level 2 (i.e., person-level) variables are identical for all records nested within a particular Level 2 unit (e.g., a child's gender code would appear in multiple records).

Finally, the database conversion demonstration shown here with syntax commands can also be accomplished through the SPSS Windows point-and-click interface. Examples are not presented here due to page length concerns. However, interested readers can download a demonstration file that shows how to convert a database from a multivariate to a stacked format using the

SPSS Windows restructuring wizard interface. The Internet address for this demonstration file is provided in the Conclusion section of the article.

Centering

Centering plays a vital role in multilevel analyses. An overview of different centering options and their effect on model parameters is beyond the scope of this article, but interested readers are encouraged to consult Kreft, de Leew, and Aiken (1995) for an excellent overview of this topic. However, a brief demonstration of grand mean and group mean centering is warranted at this point. With grand mean centering, scores on a covariate are deviated around the sample mean. Grand mean centering is quite straightforward and can be implemented using the COMPUTE command. To illustrate, suppose that the sample mean for a particular covariate X was computed as $\bar{x} = 15$. The centered covariate scores, XCENT, would be computed using the following syntax.

COMPUTE XCENT = X – 15.

With group mean centering, scores on a covariate are deviated around the means of each Level 2 unit (e.g., student-level covariate scores are centered on school means). Group mean centering is somewhat more complicated and can be implemented using the SPSS AGGREGATE command. To illustrate, we draw on a subset of the high school and beyond variables used in the next section of this article. For clarity, all variable names are italicized in the following example. Specifically, it was of interest to group mean center a student-level predictor, SES. The AGGREGATE command can be used to compute the *SES* mean for each school and write these means to a new file that is subsequently merged with the student-level data file. To do so, it is necessary to specify a "break," or grouping, variable on the /BREAK subcommand line. In this example, the grouping variable is the Level 2 identification variable named *SCHOOL*. Next, a new variable, *MEANSES*, is computed using the MEAN function. The resulting school means are written to the output file specified after the OUTFILE keyword. The AGGREGATE syntax for this example is shown below.

```
AGGREGATE OUTFILE="c:\GRPMEANS.sav"
/BREAK=SCHOOL
/MEANSES=MEAN(SES).
```

The SPSS file containing the school means is subsequently merged with the source data file containing the student-level variables using the MATCH FILES command as follows.

```
SORT CASES BY SCHOOL.
MATCH FILES TABLE = "c:\GRPMEANS.sav"
/FILE=*
/BY SCHOOL.
```

Finally, the group mean centered covariate, *CSES*, is computed using the compute command, as shown below.

<div align="center">COMPUTE CSES = SES – MEANSES.</div>

Having demonstrated two common centering procedures, we now illustrate several multilevel analyses using the SPSS MIXED procedure.

Cross-Sectional Model Examples

Consistent with Singer (1998), we used a subset of the high school and beyond data (e.g., see Raudenbush & Bryk, 2002) to demonstrate cross-sectional multilevel analyses. These data consisted of 7,185 students from 160 different public and private schools. Two Level 1 (i.e., student-level) variables (math achievement and socioeconomic status) and two Level 2 (i.e., school-level) variables (the mean SES for a school and a dummy code differentiating public and private schools) were used in the subsequent analyses. Again, we borrowed heavily from Singer (1998), duplicating her analyses using the SPSS MIXED procedure. Although a number of different notational schemes are common in the multilevel literature (Ferron, 1997), we have adopted Raudenbush and Bryk's (2002) notation throughout this article. Because the SPSS syntax is provided for a number of analyses, all variable names have been italicized for clarity. Finally, although it is unconventional to do so, the parameter estimates from the example analyses are presented in the tabular format produced by SPSS, as we believe that this is consistent with the pedagogical nature of this article.

Unconditional Means Model

An unconditional means model is identical to a one-way ANOVA with random effects and can be used to compute the proportion of variability in outcome scores (e.g., math achievement) that exists between Level 2 units (e.g., schools). The random effects ANOVA model is given below.

$$\text{Level 1: } Y_{ij} = \beta_{0j} + r_{ij} \tag{1}$$

$$\text{Level 2: } \beta_{0j} = \gamma_{00} + u_{0j} \tag{2}$$

$$\text{Combined: } Y_{ij} = \gamma_{00} + u_{0j} + r_{ij}. \tag{3}$$

The Level 1 model describes the ith student's math achievement score as a function of the mean achievement score for the jth school (β_{0j}) plus a residual (r_{ij}) that reflects individual student differences around a given school mean. The Level 2 model describes mean math achievement scores for a given school (β_{0j}) as a function of the grand mean (γ_{00}) plus a school-specific devia- tion (u_{0j}). Substituting the Level 2 equation into the Level 1 equation gives the combined model in Equation 3. The three parameters estimated by the com- bined unconditional means model include one so-called fixed effect (i.e., the grand mean, γ_{00}) and variance estimates for the Level 1 and Level 2 residuals. The SPSS syntax for the unconditional means model is shown below.

```
MIXED MATHACH
/PRINT = SOLUTION TESTCOV
/METHOD = ML
/FIXED = INTERCEPT
/RANDOM = INTERCEPT | SUBJECT(SCHOOL).
```

For all analyses outlined in this article, the outcome variable (e.g., *MATHACH*) is listed immediately following the MIXED command and the /PRINT subcommand is used to control the printed output; SOLUTION prints fixed effect estimates and standard errors, and TESTCOV produces significance tests for the variance components. The /METHOD sub- command is used to specify one of two estimators, full or restricted maxi- mum likelihood (FML and REML, respectively). Although REML is the de- fault estimator, we used FML for all analyses presented in this article.

The combined model from Equation 3 is specified using the /FIXED and /RANDOM subcommands. The /FIXED subcommand lists the variables associated with desired fixed effects. Because there were no predictors in the model, listing *INTERCEPT* on the /FIXED line will produce the grand mean (i.e., γ_{00}). It is important to note that *INTERCEPT* is an optional specification that will be estimated by default unless the intercept is specifically excluded by listing | NOINT on the /FIXED subcommand. The /RANDOM sub- command is used to specify the so-called random effects (i.e., variance com- ponents). In this example, there are two such effects, the Level 1 and Level 2 variance component estimates (i.e., σ^2 and τ_{00}). The variation in the school means around the grand mean is specified by listing *INTERCEPT* on the /RANDOM subcommand. Unlike the /FIXED command, random inter- cepts are not estimated by default, so this keyword must be explicitly men- tioned. It is not necessary to explicitly specify the Level 1 residual variance on the /RANDOM line because this component is estimated by default. How- ever, it is necessary to specify the Level 2 identification variable (e.g., *SCHOOL*) in parentheses following | SUBJECT.

Table 1 shows excerpts from the SPSS output for this analysis. As shown in Table 1, the math achievement grand mean was $\gamma_{00} = 12.64$, and variance

Table 1

Cross-Sectional Unconditional Means Model Parameter Estimates

							95% Confidence Interval	
Parameter	Estimate	SE	df	t	Wald Z	Significance	Lower Bound	Upper Bound
Estimates of fixed effects[a]								
INTERCEPT	12.63707	0.2436171	157.621	51.873		.000	12.1558946	13.1182450
Estimates of covariance parameters[a]								
Residual	39.14840	0.6606470			59.258	.000	37.8747347	40.4648957
INTERCEPT (subject variance)	8.5534642	1.0686329			8.004	.000	6.6957094	10.9266614

a. Dependent variable = *MATHACH*.

723

components suggested statistically significant variability at the between-school and within-school level, respectively ($\tau_{00} = 8.55$ and $\sigma^2 = 39.15$). The presence of such heterogeneity provides an impetus for adding Level 1 and Level 2 explanatory variables (i.e., covariates) to the model in subsequent analyses. As noted previously, the variance components can be used to calculate the proportion of variability that exists between Level 2 units. In this example, the intraclass correlation coefficient was computed as ($8.55/[8.55 + 39.15]$) = .18, which suggests that 18% of total math achievement score variability occurred between schools.

Adding a Level 2 Covariate

Because the unconditional means model indicated significant heterogeneity in math achievement means between schools, a Level 2 covariate, mean SES (*MEANSES*), was added to the Level 2 model; the Level 1 model remains the same as the previous example. The Level 2 model (and thus the combined model) now contains an additional fixed effect estimate, γ_{01}, which quantifies the influence of *MEANSES* on the school means. However, the influence of *MEANSES* on achievement is not allowed to randomly vary across schools and thus does not add to the number of estimated random effects. The expanded model is given below.

$$\text{Level 1: } Y_{ij} = \beta_{0j} + r_{ij} \tag{4}$$

$$\text{Level 2: } \beta_{0j} = \gamma_{00} + \gamma_{01}(MEANSES_j) + u_{0j} \tag{5}$$

$$\text{Combined: } Y_{ij} = \gamma_{00} + \gamma_{01}(MEANSES_j) + u_{0j} + r_{ij}. \tag{6}$$

And the SPSS syntax for this example is as follows:

```
MIXED MATHACH WITH MEANSES
/PRINT = SOLUTION TESTCOV
/METHOD = ML
/FIXED = INTERCEPT MEANSES
/RANDOM = INTERCEPT | SUBJECT(SCHOOL).
```

In this model, WITH was added to the MIXED command line and is followed by the variable names of any covariates (*MEANSES* in this case). *MEANSES* must also be specified on the /FIXED subcommand line as an added fixed effect to be estimated (i.e., γ_{01}). The remaining subcommands are the same as the previous analysis.

Excerpts from the SPSS output for the previous analysis are displayed Table 2. The expected math achievement mean for a school of average SES (*MEANSES* has a mean of 0) was $\gamma_{00} = 12.65$ and is labeled "intercept" on the

Table 2

Cross-Sectional Level 2 Covariate Model Parameter Estimates

| | | | | | | | 95% Confidence Interval | |
Parameter	Estimate	SE	df	t	Wald Z	Significance	Lower Bound	Upper Bound
Estimates of fixed effects[a]								
INTERCEPT	12.64974	0.1483178	155.693	85.288		.000	12.3567663	12.9427161
MEANSES	5.8629216	0.3591254	155.350	16.326		.000	5.1535225	6.5723207
Estimates of covariance parameters[a]								
Residual	39.15732	0.6608058			59.257	.000	37.8833501	40.4741335
INTERCEPT (subject variance)	2.5931856	0.3965863			6.539	.000	1.9215687	3.4995426

a. Dependent variable: *MATHACH*.

725

SPSS output. In addition, the influence of *MEANSES* on the school means was γ_{01} = 5.86 (labeled *MEANSES*) and was statistically significant. The residual variability in math achievement means, τ_{00} = 2.59, was notably smaller than the previous analysis (τ_{00} = 8.55) and was reduced by 70% (i.e., [8.55 – 2.59]/8.55 = .70) as a result of adding *MEANSES* as a covariate. The within-school variance remained virtually the same, however, because no Level 1 covariates were included in the model.

Adding a Level 1 Covariate

The previous models indicated that significant variability existed in math achievement scores at level 1, so a covariate can be added to the model to account for this variation as well. To illustrate the inclusion of a Level 1 covariate, a measure of student socioeconomic status (*CSES*) was added to the model. Note that this variable was group mean centered (i.e., each student's SES was deviated around their school mean SES) using SPSS syntax presented previously. The multilevel model for this analysis is shown below.

$$\text{Level 1: } Y_{ij} = \beta_{0j} + \beta_{1j}(CSES_{ij}) + r_{ij} \tag{7}$$

$$\text{Level 2: } \beta_{0j} = \gamma_{00} + u_{0j} \tag{8}$$

$$\text{Level 2: } \beta_{1j} = \gamma_{10} + u_{1j} \tag{9}$$

$$\text{Combined: } Y_{ij} = \gamma_{00} + \gamma_{10}(CSES_{ij}) + u^{0j} + u1j(CSES_{ij}) + r_{ij}. \tag{10}$$

The Level 1 model now expresses student *i*'s math achievement as a function of his or her school mean (β_{0j}), an influence due to *CSES* that is unique to school *j* (β_{1j}) and a residual (r_{ij}). There are now two Level 2 equations, one for the random intercepts (β_{0j}) and another for the random slopes (β_{1j}) that quantify the influence of *CSES* within each school. Both Level 2 equations express the Level 1 regression coefficients using a grand mean and school-specific residual. The combined model includes two fixed effects estimates and a complex error term that results in four variance components (three elements in the Level 2 covariance matrix and the Level 1 residual variance). The SPSS syntax for this analysis is shown below.

```
MIXED MATHACH WITH CSES
/PRINT = SOLUTION TESTCOV
/METHOD = ML
/FIXED = INTERCEPT CSES
/RANDOM = INTERCEPT CSES | SUBJECT(SCHOOL) COVTYPE(UN).
```

As before, the covariate *CSES* is listed after WITH on the MIXED command line. The fixed and random components for *CSES* (i.e., the mean influence of

CSES and the variation of the *CSES* coefficient across schools, respectively) were added to the /FIXED and /RANDOM subcommand lines, as shown above. The keyword COVTYPE is inserted on the /RANDOM subcommand line to specify the form of the Level 2 covariance matrix, τ. In this analysis, UN denotes an unstructured covariance matrix, meaning that the intercept variance, *CSES* slope variance, and the covariance between the intercept and slope values are estimated.

Excerpts from the SPSS output for the previous analysis are displayed in Table 3. The expected math achievement score for a student whose SES is at his or her school's mean was $\gamma_{00} = 12.65$ and is labeled "intercept." A statistically significant and positive *CSES* fixed effect estimate ($\gamma_{10} = 2.19$) indicates that math achievement scores increased as student SES increased. With respect to the random effects estimates, the Level 1 residual variance component was $\sigma^2 = 36.70$; including *CSES* as a Level 1 covariate reduced the Level 1 variance by approximately 6% (i.e., $[39.15 - 36.70]/39.15 = .06$). The Level 2 variance components are contained in a covariance matrix, where the diagonal elements (1, 1) and (2, 2) represent the intercept and slope variance, respectively, and the off-diagonal element (2, 1) represents the covariance between the random intercepts and slopes. These variance components quantify the variation in math achievement means across schools ($\tau_{00} = 8.62$), variation in the influence of student SES on math achievement scores across schools ($\tau_{11} = .68$), and the covariance between the intercepts and slopes ($\tau_{01} = .05$), which was not statistically significant. The relationship between the Level 2 covariance matrix, τ, and output from the SPSS MIXED procedure is shown below.

$$\tau = \begin{bmatrix} \tau_{00} & \\ \tau_{01} & \tau_{11} \end{bmatrix} = \begin{bmatrix} UN(1,1) & \\ UN(2,1) & UN(2,2) \end{bmatrix} = \begin{bmatrix} 8.62 & \\ 0.05 & 0.68 \end{bmatrix}$$

Adding Level 1 and Level 2 Covariates

In the final cross-sectional example, individual- (i.e., Level 1) and school-level (i.e., Level 2) covariates were simultaneously added to the hierarchical model. In addition to *CSES* and *MEANSES*, an additional Level 2 covariate, *SECTOR* (public = 0, Catholic = 1), was added to the model. The model for this analysis is shown below.

$$\text{Level 1: } Y_{ij} = \beta_{0j} + \beta_{1j}(CSES_{ij}) + r_{ij} \tag{11}$$

$$\text{Level 2: } \beta_{0j} = \gamma_{00} + \gamma_{01}(MEANSES_j) + \gamma_{02}(SECTOR_j) + u_{0j} \tag{12}$$

$$\text{Level 2: } \beta_{1j} = \gamma_{10} + \gamma_{11}(MEANSES_j) + \gamma_{12}(SECTOR_j) + u_{1j} \tag{13}$$

Table 3

Cross-Sectional Level 1 Covariate Model Parameter Estimates

Parameter	Estimate	SE	df	t	Wald Z	Significance	95% Confidence Interval	
							Lower Bound	Upper Bound
Estimates of fixed effects[a]								
INTERCEPT	12.64943	0.2437371	157.723	51.898		.000	12.1680207	13.1308401
CSES	2.1931474	0.1278475	156.141	17.154		.000	1.9406135	2.4456812
Estimates of covariance parameters[a]								
Residual	36.70004	0.6257356			58.651	.000	35.4938870	37.9471820
INTERCEPT + CSES (subject + *SCHOOL*)								
UN (1,1)	8.6210646	1.0694522			8.061	.000	6.7603385	10.9939399
UN (2,1)	0.0504005	0.4025569			0.125	.900	−0.7385966	0.8393976
UN (2,2)	0.6782500	0.2779774			2.440	.015	0.3037590	1.5144345

a. Dependent variable = *MATHACH*.

728

Combined: $Y_{ij} = \gamma_{00} + \gamma_{01}(MEANSES_j) + \gamma_{02}(SECTOR_j) + \gamma_{10}(CSES_{ij}) +$
$\gamma_{11}(MEANSES_j * CSES_{ij}) + \gamma_{12}(SECTOR_j * CSES_{ij}) + u_{0j} + u_{1j}(CSES_{ij}) + r_{ij}.$ (14)

The Level 1 model is unchanged from the previous analysis. The Level 2 model now includes the covariates *MEANSES* and *SECTOR*, the effects of which were not allowed to randomly vary across schools. The combined model (obtained by substituting Level 2 equations into the Level 1 equation) contains six fixed effects estimates and four variance components. The previous model would be specified using the syntax below.

```
MIXED MATHACH WITH MEANSES SECTOR CSES
/PRINT = SOLUTION TESTCOV
/METHOD = ML
/FIXED = INTERCEPT MEANSES SECTOR CSES MEANSES *
    CSES SECTOR * CSES
/RANDOM = INTERCEPT CSES | SUBJECT(SCHOOL) COVTYPE(UN).
```

It is worth noting that the SPSS MIXED procedure does not differentiate between Level 1 and Level 2 covariates; both Level 1 and Level 2 covariates are listed on the MIXED command line following the keyword WITH. In this analysis, the covariate main effects were listed on the /FIXED subcommand line along with the two cross-level interaction terms from the combined model. All other subcommand specifications were identical to the previous analysis, so no further discussion of random effect specification is presented.

Excerpts from the SPSS output for the previous analysis are displayed Table 4. As seen in the table, the intercept (i.e., the expected math achievement score for public school students with average student- and school-level SES) was $\gamma_{00} = 12.11$. The influence of *MEANSES* on the intercepts ($\gamma_{01} = 5.34$) was statistically significant, as was the effect of *SECTOR* ($\gamma_{02} = 1.22$), which indicated that Catholic school students had higher math achievement scores than public school students did, on average. The two significant cross-level interaction fixed effects were $\gamma_{11} = 1.04$ (*MEANSES_j * CSES_{ij}*) and $\gamma_{12} = -1.64$ (*SECTOR_j * CSES_{ij}*), the significance tests for which indicated that *MEANSES* and *SECTOR* explained a significant portion of the variation in *CSES* regression slopes across schools. With respect to the random effects, the nonsignificant slope residual variance ($\tau_{22} = .07$) suggested that this variance component (and the covariance component by extension) could be constrained to zero by removing *CSES* from the /RANDOM subcommand line. Both the Level 1 residual variance ($\sigma^2 = 36.72$) and the Level 2 intercept residual variance ($\tau_{22} = 2.31$) were statistically significant, however.

The cross-sectional multilevel analyses demonstrated here via syntax commands can also be accomplished through the SPSS Windows point-and-click interface. These examples are also not presented here due to page length concerns. However, the downloadable file mentioned earlier also contains an

Table 4

Cross-Sectional Level 1 and Level 2 Covariate Model Parameter Estimates

Parameter	Estimate	SE	df	t	Wald Z	Significance	95% Confidence Interval	
							Lower Bound	Upper Bound
Estimates of fixed effects								
INTERCEPT	12.11360	0.1969132	162.972	61.517		.000	11.7247725	12.5024329
MEANSES	5.3379502	0.3656776	153.728	14.597		.000	4.6155484	6.0603520
SECTOR	1.2169438	0.3033674	152.323	4.011		.000	0.6175930	1.8162947
CSES	2.9393565	0.1534947	140.203	19.150		.000	2.6358931	3.2428200
MEANSES + CSES	1.0423705	0.2960358	162.756	3.521		.001	0.4578043	1.6269367
*SECTOR * CSES*	−1.64386	0.2373583	145.075	−6.926		.000	−2.1129911	−1.1747369
Estimates of covariance parameters[a]								
Residual	36.72115	0.6261133			58.649	.000	35.5142727	37.9690485
INTERCEPT + CSES (subject = *SCHOOL*)								
UN (1,1)	2.3188948	0.3610217			6.423	.000	1.7090709	3.1463138
UN (2,1)	0.1881010	0.1985144			0.948	.343	−0.2009800	0.5771821
UN (2,2)	0.0652362	0.2071072			0.315	.753	0.0001295	32.8717526

a. Dependent variable = *MATHACH*.

730

example of how to conduct a multilevel cross-sectional analysis with Level 1 and Level 2 covariates via the SPSS Windows Mixed Linear Models interface. The Internet address for this file is presented in the Conclusion section of the article.

Longitudinal Model Examples

Consistent with Singer (1998), we demonstrate growth modeling analyses using Willett's (1988) simulated opposites-naming task data. These data consisted of four repeated measures collected on four consecutive days for 35 cases. Because the assessments were equally and identically spaced for all individuals, the *TIME* variable (the Level 1 covariate) took on integer values between 0 (at the initial assessment) and 3 (at the final assessment). Centering *TIME* in this manner allows the intercept to be interpreted as the estimated initial status (i.e., the expected value of the outcome variable when *TIME* = 0).

Unconditional Linear Growth Model

The multilevel formulation of the unconditional linear growth model is shown below.

$$\text{Level 1: } Y_{ij} = \pi_{0i} + \pi_{1i}(TIME_{ij}) + r_{ij} \tag{15}$$

$$\text{Level 2: } \pi_{0i} = \beta_{00} + u_{0i} \tag{16}$$

$$\text{Level 2: } \pi_{1i} = \beta_{10} + u_{1i} \tag{17}$$

$$\text{Combined: } Y_{ij} = \beta_{00} + \beta_{10}(TIME_{ij}) + u_{0i} + u_{1i}(TIME_{ij}) + r_{ij}. \tag{18}$$

The Level 1 model describes the growth trajectory for the repeated measures. Specifically, individual *i*'s opposites-naming score at assessment *j* is a function of the intercept (i.e., initial status, π_{0i}), the slope (i.e., the growth rate, π_{1i}), and a time-specific residual term (r_{ij}) that captures the deviation between an individual's observed data points and their estimated linear trajectory. In turn, Equation 16 describes the individual intercepts as a function of the mean initial status (β_{00}) plus an individual deviation (u_{0i}) from this mean. Similarly, individual growth rates are expressed as a function of the mean growth rate (β_{10}) and a residual (u_{1i}). As before, the combined model is obtained by substituting the Level 2 equations into the Level 1 equation. The SPSS syntax to perform an unconditional linear growth model analysis is given below.

```
MIXED Y WITH TIME
/PRINT = SOLUTION TESTCOV
```

    /METHOD = ML
    /FIXED = *INTERCEPT TIME*
    /RANDOM *INTERCEPT TIME* | SUBJECT(*ID*) COVTYPE(UN).

The syntax for this particular growth model is identical to that of the cross-sectional analyses presented previously. Random intercepts are specified by listing *INTERCEPT* on the /FIXED and /RANDOM subcommand lines. In a similar vein, the Level 1 predictor, *TIME*, was added to the /FIXED and /RANDOM subcommand lines, as slopes were allowed to randomly vary across individuals. As with the previous analyses, the Level 2 identification variable, *ID*, is given in parentheses following the keyword SUBJECT. Finally, COVTYPE(UN) produces estimates of the Level 2 covariance matrix, τ, which consists of three components (the intercept variance, slope variance, and the covariance between intercepts and slopes).

Excerpts from the SPSS output for the previous analysis are displayed Table 5. As seen in the table, the mean opposites-naming score at the initial measurement was $\beta_{00} = 164.37$, and the mean growth rate per day was $\beta_{10} = 26.96$. Furthermore, all variance components were significant, and the significant negative covariance ($\tau_{21} = -172.77$) between the random slopes and intercepts indicated that participants with lower initial scores experienced higher rates of growth (i.e., individuals became more similar over time). Expressed as a Pearson correlation, the relationship between the random intercepts and slopes was $r = -.45$ (i.e., $\tau_{21} / \sqrt{\tau_{11}\tau_{22}} = -172.77 / \sqrt{116.34(127.71)} = -.45$).

Nested models tests provide an alterative method for obtaining significance tests of variance components (Singer & Willett, 2003). A nested, or reduced, model is specified by constraining one or more parameters (in this case, variance estimates) of a full model to zero. For example, constraining the intercept variance to zero by removing *INTERCEPT* from the /RANDOM subcommand line will produce a nested model. In a similar vein, constraining slope variability to zero by removing *TIME* from the /RANDOM subcommand line will produce a second nested model. The differences in deviance (i.e., –2 log likelihood) statistics between a full and nested model is distributed as a chi-square statistic with degrees of freedom equal to the difference between the number of parameters in the two models. To illustrate, the likelihood ratio chi-square statistic for testing whether the intercept variance is greater than zero is obtained by subtracting the deviance from the full model from that of the reduced model, $\chi^2 = 1365.6 - 1275.4 = 90.2$. Because the reduced model has two fewer parameters (the intercept variance and covariance between the slopes and intercepts), the degrees of freedom value is 2, and the reduced model fits significantly worse than the full model ($p < .001$). The slope variance could be tested for significance in a similar fashion.

Table 5
Longitudinal Level 1 Covariate Model Parameter Estimates

							95% Confidence Interval	
Parameter	Estimate	SE	df	t	Wald Z	Significance	Lower Bound	Upper Bound
Estimates of fixed effects[a]								
INTERCEPT	164.3743	6.0308030	35	27.256		.000	152.1311047	176.6174668
TIME	26.96000	2.1354279	35	12.625		.000	22.6248509	31.2951491
Estimates of covariance parameters[a]								
Residual	159.4771	26.95656			5.916	.000	114.5035958	222.1149381
INTERCEPT + TIME (subject = *ID*)								
UN (1,1)	1161.336	304.8826			3.809	.000	694.2155641	1942.771804
UN (2,1)	–172.767	85.21003			–2.028	.043	–339.7756212	–5.7584360
UN (2,2)	127.7064	38.53118			3.314	.001	70.6953398	230.6930647

a. Dependent variable = *Y.*

733

Conditional Linear Growth Model

Because the unconditional growth model analysis indicated that individuals differed in their initial status and growth rates, a grand mean centered Level 2 (i.e., person-level) covariate was added to the model. Centering the Level 2 covariate at the grand mean facilitated the interpretation of the Level 2 intercepts, as β_{00} and β_{10} are now interpreted as the expected initial status and growth rates, respectively, for individuals who were at the mean of the covariate (Kreft et al., 1995). Because no predictors were added to the Level 1 (i.e., repeated measures) portion of the model, that portion of the model remains the same. The conditional growth model is shown below.

$$\text{Level 1: } Y_{ij} = \pi_{0i} + \pi_{1i}(TIME_{ij}) + r_{ij} \tag{19}$$

$$\text{Level 2: } \pi_{0i} = \beta_{00} + \beta_{01}(X_i - \overline{X}) + u_{0i} \tag{20}$$

$$\text{Level 2: } \pi_{1i} = \beta_{10} + \beta_{11}(X_i - \overline{X}) + u_{1i} \tag{21}$$

$$\text{Combined: } Y_{ij} = \beta_{00} + \beta_{01}(X_i - \overline{X}) + \beta_{10}(TIME_{ij})$$
$$+ \beta_{11}(X_i - \overline{X} * TIME_{ij}) + u_{0i} + u_{1i}(TIME_{ij}) + r_{ij}. \tag{22}$$

Only minor changes in the SPSS syntax are needed to incorporate Level 2 covariates, as shown below.

```
MIXED Y WITH TIME CCOVAR
/PRINT = SOLUTION TESTCOV
/METHOD = ML
/FIXED = INTERCEPT TIME CCOVAR TIME*CCOVAR
/RANDOM INTERCEPT TIME | SUBJECT(ID) COVTYPE(UN).
```

As before, the covariate *CCOVAR* was added to the MIXED command line following the keyword WITH, and the four fixed effects from Equation 22 are listed on the /FIXED subcommand. Both the initial status and slope residual variances are specified on the /RANDOM subcommand line (*INTERCEPT* and *TIME*, respectively), but *CCOVAR* was omitted from this subcommand because the influence of *CCOVAR* on the intercepts and slopes was not allowed to randomly vary across individuals.

Selected results from the conditional linear growth model analysis are given in Table 6. The regression of the intercepts on *CCOVAR* ($\beta_{01} = -.11$) was not significantly different from zero, and the residual variance in the intercepts was not dramatically lower than the unconditional variance estimated in the previous model (1159.38 vs. 1161.34, respectively). However, the cross-level interaction term, *TIME*CCOVAR* ($\beta_{11} = .43$), was positive and significantly different from zero, which indicated that higher values of *CCOVAR* were associated with accelerated growth. Note also that the addi-

Table 6
Longitudinal Level 2 Covariate Model Parameter Estimates

							95% Confidence Interval	
Parameter	Estimate	SE	df	t	Wald Z	Significance	Lower Bound	Upper Bound
Estimates of fixed effects[a]								
INTERCEPT	164.3743	6.0261702	35	27.277		.000	152.1405094	176.6080613
TIME	26.96000	1.9360748	35	13.925		.000	23.0295604	30.8904421
CCOVAR	−0.1135527	0.4893997	35	−0.232		.818	−1.1070870	0.8799816
*TIME * CCOVAR*	0.4328577	0.1572333	35	2.753		.009	0.1136572	0.7520583
Estimates of covariance parameters[a]								
Residual	159.4771	26.95656			5.916	.000	114.5035958	222.1149381
INTERCEPT + TIME (subject = *ID*)								
UN (1,1)	1159.381	304.4162			3.809	.000	692.9920737	1939.654771
UN (2,1)	165.315	78.28048			−2.112	.035	−318.7415105	−11.8876821
UN (2,2)	99.29807	31.82128			3.120	.002	52.9862111	186.0881633

a. Dependent variable = *Y*.

735

tion of *CCOVAR* decreased the slope variation (99.30 vs. 127.71 in the unconditional model), but the residual slope variance was still statistically significant.

Testing Alternative Covariance Structures

Considered together, the three residuals in the combined model equation $(u_{0i} + u_{1i}[TIME_{ij}] + r_{ij})$ form a term known as the composite residual, or complex error term. Unlike traditional regression, which assumes homoscedastic and uncorrelated residuals, the composite residual from the linear growth model allows for heteroscedasticity and autocorrelation, a fact that is readily apparent in the combined model notation; the $u_{1i}(TIME_{ij})$ term allows the magnitude of the error term to vary across time, and autocorrelation exists because u_{0i} and u_{1i} are constant components in the composite residual at each measurement occasion (i.e., they lack a j subscript). As described by Singer and Willett (2003), the traditional linear growth model imposes a specific covariance structure, such that the magnitude of the composite residual forms a U-shaped function across time. Alternative error structures may provide a better fit to the data and may increase the power of the fixed effects significance tests (Goldstein, Healy, & Rasbash, 1994; Lesaffre, Todem, Verbeke, & Kenward, 2000; Van Leeuwen, 1997). Incorrectly adopting the default error structure does not appear to bias fixed effects estimates in many cases, however (Ferron, Dailey, & Yi, 2002; Singer, 1998; Singer & Willett, 2003; Wolfinger, 1993).

To illustrate, consider the simplified linear growth model shown below.

$$\text{Level 1: } Y_{ij} = \pi_{0j} + \pi_{1j}(TIME_{ij}) + r_{ij} \tag{23}$$

$$\text{Level 2: } \pi_{0j} = \beta_{00} \tag{24}$$

$$\text{Level 2: } \pi_{1j} = \beta_{10} \tag{25}$$

$$\text{Combined: } Y_{ij} = \beta_{00} + \beta_{10}(TIME_{ij}) + r_{ij}. \tag{26}$$

In this model, a single average growth trajectory is assumed to describe all participants, and deviations between the observed data points and this mean trajectory are captured by a single residual, r_{ij}. Alternative covariance structures can be imposed on these residuals, and the fit of these models can be compared to that of the traditional model using fit indices such as the Bayesian Information Criterion (Schwartz, 1978) and the Akaike Information Criterion (Akaike, 1973); in most cases, models with alternative error structures are not nested within the traditional growth model, so likelihood ratio tests cannot be used. For example, an unstructured error covariance

matrix is the most general (and least parsimonious) covariance structure because all elements of the residual covariance matrix are free to vary. This model will always produce the best fit (i.e., the highest log likelihood) but does so at the expense of many estimated parameters; $t(t + 1)/2$ parameters are used to describe the error covariance structure in this model, where t is the number of repeated measurements.

The most recent release of SPSS (Version 13) offers 17 different error covariance structures. In the interest of space, we demonstrate the specification of a first-order autoregressive covariance matrix, but only minor changes in the syntax are required to test additional structures. Further details regarding alternative error structures can be found in Singer (1998) or Singer and Willett (2003). The SPSS syntax for implementing a first-order autoregressive residual covariance matrix is shown below.

```
MIXED Y WITH TIME CCOVAR
/PRINT = SOLUTION TESTCOV
/METHOD = ML
/FIXED = TIME CCOVAR TIME*CCOVAR
/REPEATED = WAVE | SUBJECT(ID) COVTYPE(AR1).
```

The SPSS syntax is identical to that in the previous example, except that /RANDOM has been replaced by the /REPEATED subcommand. The /REPEATED subcommand also requires an additional variable, which we named *WAVE* in this example. The variable specified in the /REPEATED subcommand (e.g., *WAVE*) is a categorical variable that assigns an integer value to each repeated measurement (e.g., 1, 2, . . . , t, where t is the number of repeated measures). Note that the variable required by the /REPEATED subcommand is automatically generated by the VARSTOCASES command described above. As before, | SUBJECT is used to specify the Level 2 identification variable, and the residual covariance structure of interest is specified by placing the appropriate keyword in parentheses (e.g., AR1 denotes a first-order autoregressive error structure). Because the SPSS syntax used to test other error covariance structures is identical to that shown above, no further discussion of this topic is warranted. The SPSS output for this analysis is omitted in the interest of space but is available on request.

The longitudinal multilevel analyses demonstrated here via syntax commands can also be accomplished through the SPSS Windows point-and-click interface, but these also are not presented here due to page length concerns. The downloadable file mentioned earlier also contains an example of how to conduct a multilevel longitudinal analysis to test the error covariance structure via the SPSS Windows Mixed Linear Models interface. The Internet address for this file is presented in the Conclusion section of the article.

Summary and Conclusion

Beginning with Version 11, SPSS implemented the MIXED procedure, which is capable of performing many common hierarchical linear model analyses. The purpose of this article was to provide a tutorial for performing cross-sectional and longitudinal analyses using this popular software platform. In doing so, we borrowed heavily from Singer's (1998) overview of SAS PROC MIXED, duplicating her analyses using the SPSS MIXED procedure. As mentioned previously, all operations (i.e., database restructuring) and analyses (i.e., multilevel cross-sectional and longitudinal) demonstrated here via syntax commands can also be performed using the SPSS Windows point-and-click interface. We have made available to interested readers a downloadable file, titled "SPSS MIXED Point&Click.doc" at the following Internet address: http://www.asu.edu/clas/psych/people/faculty/enders.htm. This file uses "screenshots" to demonstrate, in a step-by-step fashion, how to use the point-and-click interface to accomplish the database restructuring and multilevel analyses examples presented in this article. Researchers now have a wealth of options for performing multilevel analyses, and it is hoped that this article facilitates the work of researchers who wish to use this popular statistical package to perform their analyses.

Appendix A
SAS PROC MIXED Syntax Commands

Cross-Sectional Analyses

```
*Unconditional Means Model;
    proc mixed method = ml noclprint covtest noitprint
    class school;
    model mathach = /solution;
    random intercept /sub=school;

*Level-2 predictor = meanses;
    proc mixed method = ml noclprint covtest noitprint
    class school;
    model mathach = meanses/solution ddfm=bw;
    random intercept /sub=school;
```

*Level-1 predictor = Centered meanses (cses);
    proc mixed method = ml noclprint covtest noitprint
    class school;
    model mathach = cses/solution ddfm=bw notest;
    random intercept cses/sub=school type=un;

*Level-1 and Level-2 predictors;
    proc mixed method = ml noclprint covtest noitprint
    class school;
    model mathach = meanses sector cses meanses*cses
    sector*cses/solution ddfm=bw notest;
    random intercept cses/sub=school type=un;

Longitudinal Analyses

*Unconditional Linear Growth;
    proc mixed method = ml noclprint covtest noitprint
    class id;
    model y = time/solution ddfm=bw notest;
    random intercept time /subject=id type=un;

*Centered Covariate;
    proc mixed method = ml noclprint covtest noitprint
    class id;
    model y = time ccovar time*ccovar/s ddfm=bw notest;
    random intercept time /subject=id type=un;

Error Covariance Structure Test

*Autoregressive error covariance;
    proc mixed method = ml noclprint covtest noitprint
    class id wave;
    model y = time ccovar time*ccovar/s ddfm=bw notest;
    random intercept time/subject=id type=un;
    repeated wave/type=ar(1) subject=id;

Appendix B
Data Structure Conversion

ID	GENDER	AGE1	AGE2	AGE3	AGE4	VERBAL1	VERBAL2	VERBAL3	VERBAL4	SCHOOL
1	1	60	90	117	144	3	6	4	5	1
2	0	63	99	120	138	2	2	3	1	1

ID	GENDER	SCHOOL	WAVE	VERBAL	AGE
1	1	1	1	3	60
1	1	1	2	6	90
1	1	1	3	4	117
1	1	1	4	5	144
2	0	1	1	2	63
2	0	1	2	2	99
2	0	1	3	3	120
2	0	1	4	1	138

740

References

Akaike, H. (1973). Information theory as an extension of the maximum likelihood principle. In B. N. Petrov & F. Csaki (Eds.), *Second international symposium on information theory* (pp. 267-288). Budapest, Hungary: Akademiai Kiado.

Farmer, G. L. (2000). Use of multilevel covariance structure analysis to evaluate the multilevel nature of theoretical constructs. *Social Work Research*, *24*, 180-191.

Ferron, J. (1997). Moving between hierarchical model notations. *Journal of Educational and Behavioral Statistics*, *22*, 119-123.

Ferron, J., Dailey, R., & Yi, Q. (2002). Effects of misspecifying the first-level error structure in two-level models of change. *Multivariate Behavioral Research*, *37*, 379-403.

Goldstein, H., Healy, M. J. R., & Rasbash, J. (1994). Multilevel time series models with applications to repeated measures data. *Statistics in Medicine*, *13*, 1643-1655.

Hox, J. J., & Kreft, I. G. G. (Eds.). (1994). Multilevel analysis methods [Special issue]. *Sociological Methods & Research*, *22*(3).

Kreft, I. G. G. (Ed.). (1995). Hierarchical linear models: Problems and prospects [Special issue]. *Journal of Educational and Behavioral Statistics*, *20*(2).

Kreft, I. G. G., de Leew, J., & Aiken, L. S. (1995). The effect of different forms of centering in hierarchical linear models. *Multivariate Behavioral Research*, *30*, 1-21.

Kreft, I. G. G., de Leew, J., & van der Leeden, R. (1994). Review of five multilevel analysis programs: BMDP-5V, GENMOD, HLM, ML3, VARCL. *American Statistician*, *48*, 324-335.

Lesaffre, E., Todem, D., Verbeke, G., & Kenward, M. (2000). Flexible modeling of the covariance matrix in a linear random effects model. *Biometrical Journal*, *42*, 807-822.

O'Connor, B. P. (1999). Simple and flexible SAS and SPSS programs for analyzing lag-sequential categorical data. *Behavior Research Methods, Instruments & Computers*, *31*, 718-726.

Raudenbush, S. W., & Bryk, A. S. (2002). *Hierarchical linear models: Applications and data analysis methods* (2nd ed.). Newbury Park, CA: Sage.

Schwartz, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, *6*, 461-464.

Singer, J. D. (1998). Using SAS PROC MIXED to fit multilevel models, hierarchical models, and individual growth models. *Journal of Educational and Behavioral Statistics*, *24*, 323-355.

Singer, J. D., & Willett, J. B. (2003). *Applied longitudinal data analysis*. New York: Oxford University Press.

Van Leeuwen, D. M. (1997). A note on the covariance structure in a linear model. *American Statistician*, *51*, 140-144.

Willett, J. B. (1988). Questions and answers in the measurement of change. In E. Rothkopf (Ed.), *Review of research in education (1988-1989)*. Washington, DC: American Educational Research Association.

Wolfinger, R. D. (1993). Covariance structure selection in general mixed models. *Communications in Statistics–Simulations*, *22*, 1079-1106.

Zhou, X. H., Perkins, A. J., & Hui, S. L. (1999). Comparisons of software packages for generalized linear multilevel models. *American Statistician*, *53*, 282-290.