

Joop Hox

Multilevel Analysis

Techniques and Applications

Multilevel Analysis

Techniques and Applications

QUANTITATIVE METHODOLOGY SERIES
Methodology for Business and Management

George A. Marcoulides, Series Editor

The purpose of this series is to present methodological techniques to investigators and students from all functional areas of business, although individuals from other disciplines will also find the series useful. Each volume in the series will focus on a specific method (e.g., Data Envelopment Analysis, Factor Analysis, Multilevel Analysis, Structural Equation Modeling). The goal is to provide an understanding and working knowledge of each method with a minimum of mathematical derivations.

Proposals are invited from all interested authors. Each proposal should consist of the following: (i) a brief description of the volume's focus and intended market, (ii) a table of contents with an outline of each chapter, and (iii) a curriculum vita. Materials may be sent to Dr. George A. Marcoulides, Department of Management Science, California State University, Fullerton, CA 92634.

Marcoulides • Modern Methods for Business Research

Duncan/Duncan/Strycker/Al/Alpert • An Introduction to Latent Variable Growth Curve Modeling: Concepts, Issues, and Applications

Heck/Thomas • An Introduction to Multilevel Modeling Techniques

Marcoulides/Moustaki • Latent Variable and Latent Structure Models

Hox • Multilevel Analysis: Techniques and Applications

Multilevel Analysis

Techniques and Applications

Joop Hox

Utrecht University, the Netherlands



2002

LAWRENCE ERLBAUM ASSOCIATES, PUBLISHERS
Mahwah, New Jersey

London

The camera ready copy for the text of this book was supplied by the author.

Copyright © 2002 by Lawrence Erlbaum Associates, Inc.

All rights reserved. No part of this book may be reproduced in any form, by photostat, microform, retrieval system, or any other means, without the prior written permission of the publisher.

Lawrence Erlbaum Associates, Inc., Publishers
10 Industrial Avenue
Mahwah, New Jersey 07430

Cover Design by Kathryn Houghtaling Lacey

Library of Congress Cataloging-in-Publication Data

Hox, J.J.

Multilevel analysis : techniques and applications / by Joop Hox.

p. cm.

Includes bibliographical references and index.

ISBN 0-8058-3218-1 (alk. paper) - ISBN 0-8058-3219-X (alk. paper)

1. Social sciences-Statistical methods. 2. Analysis of variance.

3. Regression analysis.

I. Title.

HA29.H783 2002

001.4'22-dc21

2001057760

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

Contents

Preface	ix
1. Introduction to multilevel analysis	1
1.1. Why do we need special multilevel analysis techniques?	5
1.2. Multilevel theories	7
1.3. Models described in this book	8
2. The basic two-level regression model: introduction	11
2.1 Example	11
2.2 An extended example	15
2.3 Inspecting residuals	22
2.3.1 Examples of residuals plots	22
2.3.2 Examining slope variation: OLS and shrinkage estimators	27
2.4 Three- and more-level regression models	30
2.4.1 Multiple-level models	30
2.4.2 Intraclass-correlations in three-level models	31
2.5A note about notation and software	33
2.5.1 Notation	33
2.5.2 Software	34
3. Estimation and hypothesis testing in multilevel regression	37
3.1 Which estimation method?	37
3.1.1 Maximum Likelihood	37
3.1.2 Generalized Least Squares	39
3.1.3 Generalized Estimating Equations	40
3.1.4 Bootstrapping	40
3.1.5 Bayesian methods	41
3.2 Significance testing and confidence intervals	42
3.2.1 Testing regression coefficients and variance components	42
3.2.2 Comparing nested models	43
3.2.3 Comparing non-nested models	45
3.3 Contrasts and constraints	46

4. Some Important Methodological and Statistical Issues	49
4.1 Analysis strategy	49
4.2 Centering and standardizing explanatory variables	54
4.3 Interpreting interactions	58
4.4 How much variance is explained?	63
5. Analyzing Longitudinal Data	73
5.1 Fixed and varying occasions	74
5.2 Example with fixed occasions	75
5.3 Example with varying occasions	86
5.4 Advantages of multilevel analysis for longitudinal data	93
5.5 Some statistical issues in longitudinal analysis	93
5.5.1 Investigating and analyzing patterns of change	93
5.5.2 Handling missing data and panel dropout	95
5.6 Complex covariance structures	96
5.7 Some software issues	102
6. The logistic model for dichotomous data and proportions	103
6.1 Generalized linear models	103
6.2 Multilevel generalized linear models	107
6.2.1 Estimating generalized multilevel models	108
6.3 Example: analyzing dichotomous data	110
6.4 Example: analyzing proportions	112
6.5 Analyzing event counts and other nonlinear models	122
7. Cross-classified multilevel models	123
7.1 Example of cross-classified data: pupils nested within (primary and secondary schools)	124
7.2 Example of cross-classified data: (sociometric ratings) in small groups	128
7.3 Statistical and computational issues	136
8. The multilevel approach to meta-analysis	139
8.1 Meta-analysis and multilevel modeling	139
8.2 The variance-known model	141
8.3 Example and comparison with classical meta-analysis	145
8.3.1 Classical meta-analysis	146
8.3.2 Multilevel meta-analysis	147
8.4 Correcting for artifacts	149
8.5 Statistical and software issues	153
Appendix	154

9. Multivariate multilevel regression models	157
9.1 The basic model	158
9.2 Example of multivariate multilevel analysis: multiple response variables	161
9.3 Example of multivariate multilevel analysis: measuring group characteristics	168
10. Sample sizes and power analysis in multilevel regression	173
10.1 Sample size and accuracy of estimates	173
10.1.1 Accuracy of fixed parameters and their standard errors	173
10.1.2 Accuracy of random parameters and their standard errors	174
10.1.3 Accuracy and sample size	174
10.1.4 Accuracy and sample size with proportions and dichotomous data	175
10.2 Estimating power in multilevel regression designs	177
10.2.1 Post hoc power analysis	179
10.2.2 A priori power analysis: general issues	180
10.2.3 A priori power analysis: the intercept variance and intraclass correlation	184
10.2.4 A priori power analysis: designing an intervention study	186
10.2.5 A general procedure for power analysis	194
11. Advanced methods for estimation and testing	197
11.1 The profile likelihood method	199
11.2 Robust standard errors	200
11.3 Bootstrapping estimates and standard errors	203
11.3.1 A simple example of bootstrapping	205
11.3.2 Bootstrapping multilevel regression models	207
11.3.3 An example of the bootstrap	208
11.4 Bayesian estimation methods	211
11.4.1 Simulating the posterior distribution	213
11.4.2 An example of Bayesian estimation using MlwiN: the estrone data	214
11.4.3 An example of Bayesian estimation using MlwiN: the popularity data	220
11.4.4 Some remarks on Bayesian estimation methods	223

12. Multilevel factor models	225
12.1 Decomposing multilevel variables	227
12.2 Muthén's pseudobalanced approach	228
12.2.1 An example of pseudobalanced multilevel factor analysis	232
12.2.2 Goodness of fit using the pseudobalanced approach	238
12.3 Direct estimation of the covariances at each level: the multivariate multilevel approach	242
12.4 Standardizing parameter estimates in multilevel structural equation modeling	248
13. Multilevel path models	251
13.1 Example of a multilevel path analysis	252
13.2 Statistical and software issues in multilevel factor and path models	259
14. Latent curve models	263
14.1 Example of latent curve modeling	266
14.2 A brief comparison of multilevel regression analysis and latent curve modeling of longitudinal data	273
Appendix: Data and stories	275
References	281
Author index	293
Subject index	299

Preface

*To err is human, to forgive divine;
but to include errors into your design is statistical.*

—Leslie Kish

This book is intended as an introduction to multilevel analysis for applied researchers. The term ‘multilevel’ refers to a hierarchical or nested data structure, usually people within organizational groups, but the nesting may also consist of repeated measures within people, or respondents within clusters as in cluster sampling. The expression *multilevel model* or *multilevel analysis* is used as a generic term for all models for nested data. This book presents two multilevel models: the multilevel regression model and a model for multilevel covariance structures.

Multilevel modeling used to be only for specialists. However, in the past decade, multilevel analysis software has become available that is both powerful and relatively accessible for applied researchers. As a result, there is a surge of interest in multilevel analysis, as evidenced by the appearance of several reviews and monographs, applications in different fields ranging from psychology and sociology to medicine, and a thriving Internet discussion list with more than 1400 subscribers.

Despite it being an introduction, the book includes a discussion of many extensions and special applications. As an introduction, it should be useable in courses in a variety of fields, such as psychology, education, sociology and business. The various extensions and special applications should make it useful to researchers who work in applied or theoretical research, and to methodologists who have to consult with these researchers. The basic models and examples are discussed in non-technical terms; the emphasis is on understanding the methodological and statistical issues involved in using these models. Some of the extensions and special applications contain discussions that are more technical, either because that is necessary for understanding what the model does, or as a helpful introduction to more advanced treatments in other texts. Thus, in addition to its role as an introduction, the book should be useful as a standard reference for a large variety of applications. It assumes that readers have a basic knowledge of social science statistics, including analysis of variance and multiple regression analysis. The section about multilevel structural equation models assumes a basic understanding of ordinary structural equation modeling.

I thank Peter van der Heijden, Herbert Hoijtink, Bernet Sekasanvu Kato, Edith de Leeuw, George Marcoulides, Mirjam Moerbeek, Ian Plewis, Ken Rowe, Godfried van den Wittenboer and Bill Yeaton for their comments on earlier drafts.

I thank my colleagues at the Department of Methodology and Statistics of the Faculty of Social Sciences at Utrecht University for providing me with many discussions and a generally stimulating research environment. My research has also benefited from the lively discussions by the denizens of the Internet *Multilevel Modeling* and the *Structural Equations Modeling (SEMNET)* discussion lists. I also want to mention the Dutch NOSMO research committee on Multilevel Research (MULOG) for their continued interest in multilevel modeling. Finally, draft chapters of this book have been used in a multilevel course at the Summer Institute of the Social Research Center of the University of Michigan. I thank the Summer Institute students for their critical comments. As always, any errors remaining in the book are entirely my own responsibility.

J.J. Hox

Amsterdam,
January, 2002

1

Introduction to Multilevel Analysis

Social research regularly involves problems that investigate the relationship between individual and society. The general concept is that individuals interact with the social contexts to which they belong, meaning that individual persons are influenced by the social groups or contexts to which they belong, and that the properties of those groups are in turn influenced by the individuals who make up that group. Generally, the individuals and the social groups are conceptualized as a hierarchical system of individuals and groups, with individuals and groups defined at separate levels of this hierarchical system. Naturally, such systems can be observed at different hierarchical levels, and variables may be defined at each level. This leads to research into the interaction between variables characterizing individuals and variables characterizing groups, a kind of research that is now often referred to as '*multilevel research*'.

In multilevel research, the data structure in the population is hierarchical, and the sample data are viewed as a multistage sample from this hierarchical population. Thus, in educational research, the population consists of schools and pupils within these schools, and the sampling procedure proceeds in two stages: first, we take a sample of schools, and next we take a sample of pupils within each school. Of course, in real research one may have a convenience sample at either level, or one may decide not to sample pupils but to study all available pupils in the sample of schools. Nevertheless, one should keep firmly in mind that the central statistical model in multilevel analysis is one of successive sampling from each level of a hierarchical population.

In this example, pupils are *nested* within schools. Other examples are cross-national studies where the individuals are nested within their national units, organizational research with individuals nested within departments within organizations, family research with family members within families and methodological research into interviewer effects with respondents nested within interviewers. Less obvious applications of multilevel models are longitudinal research and growth curve research, where a series of several distinct observations are viewed as nested within individuals, and meta-analysis where the subjects are nested within different studies. For simplicity, this book describes the multilevel models mostly in terms of individuals nested within groups, but note that the models apply to a much larger class of analysis problems.

In multilevel research, variables can be defined at any level of the hierarchy. Some of these variables may be measured directly at their ‘own’ natural level; for example, at the school level we may measure school size and denomination, and at the pupil level intelligence and school success. In addition, we may move variables from one level to another by aggregation or disaggregation. Aggregation means that the variables at a lower level are moved to a higher level, for instance, by assigning to the schools the school mean of the pupils' intelligence scores. Disaggregation means moving variables to a lower level, for instance by assigning to all pupils in the schools a variable that indicates the denomination of the school they belong to. Lazarsfeld and Menzel (1961) offer a typology to describe the relations between different types of variables, defined at different levels. A simplified scheme is presented below:

Level:	1		2		3		et cetera
Variable	global	⇒	analytical				
type:	relational	⇒	structural				
	contextual	⇐	global	⇒	analytical		
			relational	⇒	structural		
			contextual	⇐	global	⇒	
					relational	⇒	
					contextual	⇐	

The lowest level (level 1) in this scheme is usually defined by the individuals. However, this is not always the case. Galtung (1969), for instance, defines roles within individuals as the lowest level, and in longitudinal designs, one can define repeated measures within individuals as the lowest level (Goldstein, 1986, 1989).

At each level in the hierarchy, we may have several types of variables. *Global* variables are variables that refer only to the level at which they are defined, without reference to other units or levels. A pupil's intelligence or gender would be a global variable at the pupil level. School size would be a global variable at the school level. *Relational* variables also belong to one single level, but they describe the relationships of a unit to the other units at the same level. Many sociometric indices, such as indices of popularity or the reciprocity of relationships, are relational variables. *Analytical* and *structural* variables are measured by referring to the sub-units at a lower level. Analytical variables are constructed from variables at a lower level, for example, in defining the school variable ‘mean intelligence’ as the mean intelligence of the pupils in that school. Using the mean of a lower-level variable as an explanatory variable at a

higher level is a customary procedure in multilevel analysis. Other functions may also be valuable. For instance, using the standard deviation of a lower-level variable as an explanatory variable at a higher level could be used to test hypotheses about the effect of group heterogeneity on the outcome variable. Structural variables refer to the distribution of relational variables at the lower level; many social network indices are of this type. It is clear that constructing an analytical or structural variable from the lower-level data involves *aggregation* (which is indicated in the scheme by \Rightarrow); data on lower-level units are aggregated into data on a smaller number of higher-level units. *Contextual* variables, on the other hand, refer to the super-units; all units at the lower level receive the value of a variable for the super-unit to which they belong at the higher level. For instance, we can assign to all pupils in a school the school size, or the mean intelligence, as a pupil level variable. This is called *disaggregation* (indicated in the scheme by \Leftarrow); data on higher-level units are disaggregated into data on a larger number of lower-level units. The resulting variable is called a *contextual* variable, because it refers to the higher-level context of the units we are investigating.

In order to analyze multilevel models, it is not important to assign each variable to its proper place in the scheme given above. The benefit of the scheme is conceptual; it makes clear to which level a measurement properly belongs. Historically, multilevel problems led to analysis approaches that moved all variables by aggregation or disaggregation to one single level of interest followed by an ordinary multiple regression, analysis of variance, or some other 'standard' analysis method. However, analyzing variables from different levels at one single common level is inadequate, because it leads to two distinct problems.

The first problem is statistical. If data are aggregated, the result is that different data values from many sub-units are combined into fewer values for fewer higher-level units. As a result, much information is lost, and the statistical analysis loses power. On the other hand, if data are disaggregated, the result is that a few data values from a small number of super-units are 'blown up' into many more values for a much larger number of sub-units. Ordinary statistical tests treat all these disaggregated data values as independent information from the much larger sample of sub-units. The proper sample size for these variables is of course the number of higher-level units. Using the larger number of disaggregated cases for the sample size leads to significance tests that reject the null-hypothesis far more often than the nominal alpha level suggests. In other words: investigators come up with many 'significant' results that are totally spurious.

The second problem encountered is conceptual. If the analyst is not very careful in the interpretation of the results, s/he may commit the fallacy of the wrong level, which consists of analyzing the data at one level, and formulating conclusions at another level. Probably the best-known fallacy is the *ecological fallacy*, which is interpreting aggregated data at the individual level. It is also known as the 'Robinson effect' after Robinson (1950). Robinson presents aggregated data describing the

relationship between the percentage of blacks and the illiteracy level in nine geographic regions in 1930. The *ecological correlation*, that is, the correlation between the aggregated variables at the region level is 0.95. In contrast, the individual-level correlation between these global variables is 0.20. Robinson concludes that in practice an ecological correlation is almost certainly not equal to its corresponding individual-level correlation. For a statistical explanation why this happens, see Robinson (1950) or Kreft and de Leeuw (1987). This problem occurs also the other way around. Formulating inferences at a higher level based on analyses performed at a lower level is just as misleading. This fallacy is known as the *atomistic fallacy*. A related but different fallacy is known as 'Simpson's Paradox' (see Lindley & Novick, 1981). Simpson's paradox refers to the problem that completely erroneous conclusions may be drawn if grouped data, drawn from heterogeneous populations, are collapsed and analyzed as if they came from a single homogeneous population. An extensive typology of such fallacies is given by Alker (1969). When aggregated data are the only available data, King (1997) presents some procedures that make it possible to estimate the corresponding individual relationships without committing the ecological fallacy.

A more general way to look at multilevel data is to realize that there is not one 'proper' level at which the data should be analyzed. Rather, all levels present in the data are important in their own way. This becomes clear when we investigate cross-level hypotheses, or *multilevel* problems. A multilevel problem is a problem that concerns the relationships between variables that are measured at a number of different hierarchical levels. For example, a common question is how a number of individual and group variables influence one single individual outcome variable. Typically, some of the higher-level explanatory variables may be the aggregated group means of lower-level individual variables. The goal of the analysis is to determine the direct effect of individual and group level explanatory variables, and to determine if the explanatory variables at the group level serve as moderators of individual-level relationships. If group level variables moderate lower-level relationships, this shows up as a statistical interaction between explanatory variables from different levels. In the past, such data were usually analyzed using conventional multiple regression analysis with one dependent variable at the lowest (individual) level and a collection of explanatory variables from all available levels (cf. Boyd & Iversen, 1979; Roberts & Burstein, 1980; van den Eeden & Hüttner, 1982). Since this approach analyzes all available data at one single level, it suffers from all of the conceptual and statistical problems mentioned above. Much research has been directed at developing more appropriate analysis methods for this hierarchical regression model, and at clarifying the associated conceptual and statistical issues.

1.1. WHY DO WE NEED SPECIAL MULTILEVEL ANALYSIS TECHNIQUES?

A multilevel problem concerns a population with a hierarchical structure. A sample from such a population can be described as a multistage sample: first, we take a sample of units from the higher level (e.g., schools), and next we sample the sub-units from the available units (e.g., we sample pupils from the schools). In such samples, the individual observations are in general not completely independent. For instance, pupils in the same school tend to be similar to each other, because of selection processes (for instance, some schools may attract pupils from higher social economic status (SES) levels, while others attract more lower SES pupils) and because of the common history the pupils share by going to the same school. As a result, the average correlation (expressed in the so-called *intraclass correlation*) between variables measured on pupils from the same school will be higher than the average correlation between variables measured on pupils from different schools. Standard statistical tests lean heavily on the assumption of independence of the observations. If this assumption is violated (and in multilevel data this is almost always the case) the estimates of the standard errors of conventional statistical tests are much too small, and this results in many spuriously ‘significant’ results.

The problem of dependencies between individual observations also occurs in survey research, if the sample is not taken at random but cluster sampling from geographical areas is used instead. For similar reasons as in the school example given above, respondents from the same geographical area will be more similar to each other than respondents from different geographical areas are. This leads again to estimates for standard errors that are too small and produce spurious ‘significant’ results. In survey research, this effect of cluster sampling is well known (cf. Kish, 1965, 1987). It is called a ‘design effect’, and various methods are used to deal with it. A convenient correction procedure is to compute the standard errors by ordinary analysis methods, estimate the intraclass correlation between respondents within clusters, and finally employ a correction formula to the standard errors. A correction described by Kish (1965: p. 259) corrects the standard error using $s.e._{eff} = s.e. \cdot (1 + (n_{clus} - 1)\rho)$, where $s.e._{eff}$ is the effective standard error, n_{clus} is the cluster size, and ρ is the intraclass correlation. The formula assumes equal group sizes, which is not always realistic. A variation of this formula computes the effective sample size in two-stage cluster sampling as $n_{eff} = n / [1 + (n_{clus} - 1)\rho]$, where n is the total sample size and n_{eff} is the effective sample size. Using this formula, we can simply calculate the effective sample size for different situations.¹ For instance, suppose that we take a sample of 10 classes,

¹ The formulas given here apply to two-stage cluster sampling. Other sampling schemes, such as stratified sampling, require different formulas. See Kish (1965, 1987) for details. The symbol ρ (the Greek letter rho) was introduced by Kish (1965, p. 161) who called it *roh* for ‘rate of homogeneity’.

each with 20 pupils. This comes to a total sample size of 200, which is reasonable. Let us further suppose that we are interested in a variable, for which the intraclass correlation ρ is 0.10. This seems a rather low intraclass correlation. However, the effective sample size in this situation is $200/[1+(20-1)0.1]=69.0$, which is much less than the apparent total sample size of 200! Gulliford, Ukoumunne and Chin (1999) give an overview of estimates of the intraclass correlation to aid in the design of complex health surveys. Their data include variables on a range of lifestyle risk factors and health outcomes, for respondents clustered at the household, postal code, and health authority district levels. They report between-cluster variation at each of these levels, with intraclass correlations ranging from 0.0-0.3 at the household level, and being mostly smaller than 0.05 at the postal code level, and below 0.01 at the district level. Since the design effect depends on both the intraclass correlation and the cluster sample size, the large household intraclass correlations are partly compensated by the small household sizes. Conversely, the small intraclass correlations at the higher levels are offset by the usually large cluster sizes at these levels. Groves (1989) also discusses the effects of cluster sampling on the standard errors, and concludes that the intraclass correlation is usually small, but in combination with the usual cluster size still can lead to substantial design effects.

Some of the correction procedures developed for cluster and other complex samples are quite powerful (cf. Skinner, Holt & Smith, 1989). Actually, in principle these correction procedures could also be applied in analyzing multilevel data, by adjusting the standard errors of the statistical tests. However, in general the intraclass correlation and hence the effective N is different for different variables. In addition, in most multilevel problems we have not only clustering of individuals within groups, but we also have variables measured at all available levels. Combining variables from different levels in one statistical model is a different and more complicated problem than estimating and correcting for design effects. Multilevel models are designed to analyze variables from different levels simultaneously, using a statistical model that properly includes the various dependencies.

For example, an explicitly multilevel or contextual theory in education is the so-called 'frog pond' theory, which refers to the idea that a specific individual frog may either be a small frog in a pond otherwise filled with large frogs, or a large frog in a pond otherwise filled with small frogs. Applied to education, this metaphor points out that the effect of an explanatory variable such as 'intelligence' on school career may depend on the average intelligence of the other pupils in the school. A moderately intelligent pupil in a highly intelligent context may become demotivated and thus become an underachiever, while the same pupil in a considerably less intelligent context may gain confidence and become an overachiever. Thus, the effect of an individual pupil's intelligence depends on the average intelligence of the other pupils. A popular approach in educational research to investigate 'frog pond' effects has been

to aggregate variables like the pupils' IQ into group means, and then to disaggregate these group means again to the individual level. As a result, the data file contains both individual level (global) variables and higher-level (contextual) variables in the form of disaggregated group means. Cronbach (1976; cf. Cronbach & Webb, 1979) has suggested to express the individual scores as deviations from their respective group means, a procedure that has become known as *centering on the group mean*, or *group mean centering*. Centering on the group means makes very explicit that the individual scores should be interpreted relative to their group's mean. The example of the 'frog pond' theory and the corresponding practice of centering the predictor variables makes clear that combining and analyzing information from different levels within one statistical model is central to multilevel modeling.

1.2. MULTILEVEL THEORIES

Multilevel problems must be explained by multilevel theories, an area that seems underdeveloped compared to the advances made in the recently developed modeling and computing machinery (cf. Hüttner & van den Eeden, 1993). If there are effects of the social context on individuals, these effects must be mediated by intervening processes that depend on characteristics of the social context. Multilevel models so far require that the grouping criterion is clear, and that variables can be assigned unequivocally to their appropriate level. In reality, group boundaries are sometimes fuzzy and somewhat arbitrary, and the assignment of variables is not always obvious and simple. In multilevel problems, decisions about group membership and operationalizations involve a wide range of theoretical assumptions, and an equally wide range of specification problems for the auxiliary theory (Blalock, 1990). When the number of variables at the different levels is large, there is an enormous number of possible cross-level interactions. Ideally, a multilevel theory should specify which variables belong to which level, and which direct effects and cross-level interaction effects can be expected. Cross-level interaction effects between the individual and the context level require the specification of processes within individuals that cause those individuals to be differentially influenced by certain aspects of the context. Attempts to identify such processes have been made by, among others, Stinchcombe (1968), Erbring and Young (1979), and Chan (1998). The common core in these theories is that they all postulate one or more psychological processes that mediate between individual variables and group variables. Since a global explanation by 'group telepathy' is generally not acceptable, communication processes and the internal structure of groups become important concepts. These are often measured as a 'structural variable'. In spite of their theoretical relevance, structural variables are infrequently used in multilevel research. Another theoretical area that has been largely neglected by

multilevel researchers is the influence of individuals on the group. This is already visible in Durkheim's concept of sociology as a science that focuses primarily on the constraints that a society can put on its members, and disregards the influence of individuals on their society. In multilevel modeling, the focus is on models where the outcome variable is at the lowest level. Models that investigate the influence of individual variables on group outcomes are scarce. For a review of this issue see DiPrete and Forristal (1994), an example is discussed by Alba and Logan (1992).

1.3. MODELS DESCRIBED IN THIS BOOK

This book treats two classes of multilevel models: multilevel regression models, and multilevel models for covariance structures.

Multilevel regression models are essentially a multilevel version of the familiar multiple regression model. As Cohen and Cohen (1983), Pedhazur (1997) and others have shown, the multiple regression model is very versatile. Using dummy coding for categorical variables, it can be used to analyze analysis of variance (ANOVA)-type of models as well as the more usual multiple regression models. Since the multilevel regression model is an extension of the classical multiple regression model, it too can be used in a wide variety of research problems. It has been used extensively in educational research (cf. the special issues of the *International Journal of Educational Research*, 1990 and the *Journal of Educational and Behavioral Statistics* in 1995). Other applications have been in the analysis of longitudinal and growth data (cf. Bryk & Raudenbush, 1987; Goldstein, 1989; DiPrete & Grusky, 1990; Goldstein, Healy & Rasbash, 1994), the analysis of interview survey data (Hox, de Leeuw & Kreft, 1991; Hox, 1994a; O'Muirchartaigh & Campanelli, 1999; Pickery & Loosveldt, 1998), data from surveys with complex sampling schemes with respondents nested within sampling units (Goldstein & Silver, 1989; Snijders, 2001), and data from factorial surveys and facet designs (Hox, Kreft & Hermkens, 1991; Hox & Lagerweij, 1993). Raudenbush and Bryk have introduced multilevel regression models in meta-analysis (cf. Raudenbush & Bryk, 1985, 1987; Hox & de Leeuw, 1994; Raudenbush, 1994). Multilevel regression models for binary and other non-normal data have been described by Wong and Mason (1985), Longford (1988), Mislavy and Bock (1989) and Goldstein (1991).

Chapter Two of this book contains a basic introduction to the multilevel regression model, also known as the hierarchical linear model, or the random coefficient model. Chapters Three and Four discuss estimation procedures, and a number of important methodological and statistical issues. It also discusses some

technical issues that are not specific to multilevel regression analysis, such as coding categorical explanatory variables and interpreting interactions.

Chapter Five introduces the multilevel regression model for longitudinal data. The model is a straightforward extension of the standard multilevel regression model, but there are some specific complications, such as autocorrelated errors, which will be discussed.

Chapter Six treats the logistic model for dichotomous data and proportions. When the response (dependent) variable is dichotomous or a proportion, standard regression models should not be used. This chapter discusses the multilevel version of the logistic regression model.

Chapter Seven discusses cross-classified models. Some data are multilevel in nature, but do not have a neat hierarchical structure. Examples are longitudinal school research data, where pupils are nested within schools, but may switch to a different school in later measurements, and sociometric choice data. Multilevel models for such cross-classified data can be formulated, and estimated with standard software provided that it can handle restrictions on estimated parameters.

Chapter Eight describes a variant of the multilevel regression model that can be used in meta-analysis. It resembles the weighted regression model often recommended for meta-analysis. Using standard regression procedures, it is a flexible analysis tool.

Chapter Nine discusses multilevel regression models for multivariate outcomes. These can also be used to estimate models that resemble confirmative factor analysis, and to assess the reliability of multilevel measurements. A different approach to multilevel confirmative factor analysis is treated in chapter Eleven.

Chapter Ten deals with the sample size needed for multilevel modeling, and the problem of estimating the power of an analysis given a specific sample size. An obvious complication in multilevel power analysis is that there are different sample sizes at the distinct levels, which should be taken into account.

Chapter Eleven treats some advanced methods of estimation and assessing significance. It discusses the profile likelihood method, robust standard errors for establishing confidence intervals, and multilevel bootstrap methods for estimating bias-corrected point-estimates and confidence intervals. This chapter also contains an introduction into Bayesian (MCMC) methods for estimation and inference.

Multilevel models for covariance structures, or multilevel structural equation models (SEM), are a powerful tool for the analysis of multilevel data. Much fundamental work has been done on multilevel factor and path analysis (cf. Goldstein & McDonald, 1988; Muthén, 1989, 1990; McDonald & Goldstein, 1989). There is also a growing number of applications, for instance Härnqvist, Gustafsson, Muthén, and Nelson (1994), and Hox (1993). These applications require only conventional software for structural equation modeling (e.g., Amos, Eqs, Lisrel) with unusual setups.

Specialized software to analyze multilevel structural equation models is available as well (*Mplus*, Muthén & Muthén, 1998). The general statistical model for multilevel covariance structure analysis is quite complicated. Chapter Twelve in this book describes a simplified statistical model proposed by Muthén (1990, 1994), and explains how multilevel confirmatory factor models can be estimated with either conventional SEM software or using specialized programs like *Mplus*. It also describes a direct estimation approach, and deals with issues of calculating standardized coefficients and goodness-of-fit indices in multilevel structural models. Chapter Thirteen extends this to path models. Chapter Fourteen describes structural models for latent curve analysis. This is a SEM approach to analyzing longitudinal data, which is very similar to the multilevel regression models treated in Chapter Five.

This book is intended as an introduction to the world of multilevel analysis. Most of the chapters on multilevel regression analysis should be readable for social scientists who have a good general knowledge of analysis of variance and classical multiple regression analysis. Some of these chapters contain material that is more difficult, but this is generally a discussion of specialized problems, which can be skipped at first reading. An example is the chapter on longitudinal models, which contains a prolonged discussion of techniques to model specific structures for the covariances between adjacent time points. This discussion is not needed to understand the essentials of multilevel analysis of longitudinal data, but it may become important when one is actually analyzing such data. The chapters on multilevel structure equation modeling obviously require a strong background in multivariate statistics and some background in structural equation modeling, equivalent to, for example, the material covered in Tabachnick and Fidell's (1996) book. Conversely, in addition to an adequate background in structural equation modeling, the chapters on multilevel structural equation modeling do not require knowledge of advanced mathematical statistics. In all these cases, I have tried to keep the discussion of the more advanced statistical techniques theoretically sound, but non-technical.

Many of the techniques and their specific software implementations discussed in this book are the subject of active statistical and methodological research. In other words: both the statistical techniques and the software tools are evolving rapidly. As a result, increasing numbers of researchers will apply increasingly advanced models to their data. Of course, researchers still need to understand the models and techniques that they use. Therefore, in addition to being an introduction to multilevel analysis, this book aims to let the reader become acquainted with some advanced modeling techniques that might be used, such as bootstrapping and Bayesian estimation methods. At the time of writing, these are specialist tools, and certainly not part of the standard analysis toolkit. But they are developing rapidly, and are likely to become more popular in applied research as well.

2

The Basic Two-Level Regression Model: Introduction

The multilevel regression model has become known in the research literature under a variety of names, such as ‘random coefficient model’ (de Leeuw & Kreft, 1986; Longford, 1993), ‘variance component model’ (Longford, 1987), and ‘hierarchical linear model’ (Raudenbush & Bryk, 1986, 1988). Statically oriented publications tend to refer to this model as a mixed-effects or mixed model (Littell, Milliken, Stroup & Wolfinger, 1996). The models described in these publications are not *exactly* the same, but they are highly similar, and I will refer to them collectively as ‘multilevel regression models’. They all assume that there is a hierarchical data set, with one single outcome or response variable that is measured at the lowest level, and explanatory variables at all existing levels. Conceptually, it is useful to view the multilevel regression model as a hierarchical system of regression equations. In this chapter, I will explain the multilevel regression model for two-level data. Regression models with more than two levels are used in later chapters.

2.1 EXAMPLE

Assume that we have data from J classes, with a different number of pupils n_j in each class. On the pupil level, we have the outcome variable ‘popularity’ (Y), measured by a self-rating scale that ranges from 0 (very unpopular) to 10 (very popular). We have one explanatory variable *gender* (X : 0=boy, 1=girl) on the pupil level, and one class level explanatory variable *teacher experience* (Z : in years). We have data from 2000 pupils from 100 classes, so the average class size is 20 pupils. The data are described in the Appendix.

To analyze these data, we can set up separate regression equations in each class to predict the outcome variable Y by the explanatory variable X as follows:

$$Y_{ij} = \beta_{0j} + \beta_{1j}X_{ij} + e_{ij} \quad (2.1)$$

Using variable labels instead of algebraic symbols, the equation reads:

$$\text{popularity}_{ij} = \beta_{0j} + \beta_{1j}\text{gender}_{ij} + e_{ij} \quad (2.2)$$

In this regression equation, β_{0j} is the usual intercept, β_{1j} is the usual regression coefficient (regression slope) for the explanatory variable gender, and e_{ij} is the usual residual error term. The subscript j is for the classes ($j=1 \dots J$) and the subscript i is for individual pupils ($i=1 \dots n_j$). The difference with the usual regression model is that we assume that each class has a different intercept coefficient β_{0j} , and a different slope coefficient β_{1j} . This is indicated in equations (2.1) and (2.2) by attaching a subscript j to the regression coefficients. The residual errors e_{ij} are assumed to have a mean of zero, and a variance to be estimated. Most multilevel software assumes that the variance of the residual errors is the same in all classes. Different authors (cf. Bryk & Raudenbush, 1992; Goldstein, 1995) use different systems of notation. This book uses σ_e^2 to denote the variance of the lowest level residual errors.¹

Since the intercept and slope coefficients are assumed to vary across the classes, they are often referred to as *random* coefficients.² In our example, the specific value for the intercept and the slope coefficient for the pupil variable ‘gender’ are a class characteristic. In general, a class with a high intercept is predicted to have more popular pupils than a class with a low value for the intercept. Similarly, differences in the slope coefficient for gender indicate that the relationship between the pupils’ gender and their predicted popularity is not the same in all classes. Some classes have a high value for the slope coefficient of gender; in these classes, the difference between boys and girls is relatively large. Other classes have a low value for the slope coefficient of gender; in these classes, gender has a small effect on the popularity, which means that the difference between boys and girls is small.

Across all classes, the regression coefficients β_j have a distribution with some mean and variance. The next step in the hierarchical regression model is to explain the variation of the regression coefficients β_{0j} and β_{1j} by introducing explanatory variables at the class level, as follows:

$$\beta_{0j} = \gamma_{00} + \gamma_{01}Z_j + u_{0j} \quad (2.3)$$

¹ At the end of this chapter, a section explains the difference between some commonly used notation systems. Models that are more complicated sometimes need a more complicated notation system, which is introduced in the relevant chapters.

² Of course, we hope to be able to explain at least some of the variation by introducing higher-level variables. Generally, we will not be able to explain all the variation, and there will be some unexplained residual variation. Hence the name ‘random coefficient model’: the regression coefficients (intercept and slopes) have some amount of (residual) random variation between groups. The name ‘variance component model’ refers to the statistical problem of estimating the amount of random variation.

and

$$\beta_{1j} = \gamma_{10} + \gamma_{11}Z_j + u_{1j} \quad (2.4)$$

Equation (2.3) predicts the average popularity in a class (the intercept β_{0j}) by the teacher's experience (Z). Thus, if γ_{01} is positive, the average popularity is higher in classes with a more experienced teacher. Conversely, if γ_{01} is negative, the average popularity is lower in classes with a more experienced teacher. The interpretation of equation (2.4) is a bit more complicated. Equation (2.4) states that the *relationship*, as expressed by the slope coefficient β_{1j} , between the popularity (Y) and the gender (X) of the pupil, depends upon the amount of experience of the teacher (Z). If γ_{11} is positive, the gender effect on popularity is larger with experienced teachers. Conversely, if γ_{11} is negative, the gender effect on popularity is smaller with experienced teachers. Thus, the amount of experience of the teacher acts as a *moderator variable* for the relationship between popularity and gender; this relationship varies according to the value of the moderator variable.

The u -terms u_{0j} and u_{1j} in equations (2.3) and (2.4) are (random) residual error terms at the class level. These residual errors u_j are assumed to have a mean of zero, and to be independent from the residual errors e_{ij} at the individual (pupil) level. The variance of the residual errors u_{0j} is specified as $\sigma_{u_0}^2$, and the variance of the residual errors u_{1j} is specified as $\sigma_{u_1}^2$. The *covariance* between the residual error terms u_{0j} and u_{1j} is σ_{u_0} , which is generally *not* assumed to be zero.

Note that in equations (2.3) and (2.4) the regression coefficients γ are not assumed to vary across classes. They therefore have no subscript j to indicate to which class they belong. Because they apply to *all* classes, they are referred to as *fixed* coefficients. All between-class variation left in the β coefficients, after predicting these with the class variable Z_j , is assumed to be residual error variation. This is captured by the residual error terms u_j , which do have subscripts j to indicate to which class they belong.

Our model with one pupil level and one class level explanatory variable can be written as a single complex regression equation by substituting equations (2.3) and (2.4) into equation (2.1). Rearranging terms gives:

$$Y_{ij} = \gamma_{00} + \gamma_{10}X_{ij} + \gamma_{01}Z_j + \gamma_{11}X_{ij}Z_j + u_{1j}X_{ij} + u_{0j} + e_{ij} \quad (2.5)$$

Using variable labels instead of algebraic symbols, we have

$$\begin{aligned} \text{popularity}_{ij} = & \gamma_{00} + \gamma_{10} \text{gender}_{ij} + \gamma_{01} \text{experience}_j + \gamma_{11} \text{experience}_j \times \text{gender}_{ij} \\ & + u_{1j} \text{gender}_{ij} + u_{0j} + e_{ij} \end{aligned}$$

The segment $[\gamma_{00} + \gamma_{10} X_{ij} + \gamma_{01} Z_j + \gamma_{11} Z_j X_{ij}]$ in equation (2.5) contains the fixed coefficients. It is often called the fixed (or deterministic) part of the model. The segment $[u_{0j} + u_{1j} X_{ij} + e_{ij}]$ in equation (2.5) contains the random error terms, and it is often called the random (or stochastic) part of the model. The term $Z_j X_{ij}$ is an interaction term that appears in the model as a consequence of modeling the varying regression slope β_{1j} of pupil level variable X_{ij} with the class level variable Z_j . Thus, the moderator effect of Z on the relationship between the dependent variable Y and the predictor X , is expressed in the single equation version of the model as a *cross-level interaction*. The interpretation of interaction terms in multiple regression analysis is complex, and this is treated in more detail in Chapter Three. In general, the point made in Chapter Three is that the substantive interpretation of the coefficients in models with interactions is much simpler if the variables making up the interaction are expressed as deviations from their respective means.

Note that the random error term u_{1j} is connected to X_{ij} . Since the explanatory variable X_{ij} and the error term u_{1j} are multiplied, the resulting total error will be different for different values of X_{ij} , a situation that in ordinary multiple regression analysis is called ‘heteroscedasticity’. The usual multiple regression model assumes ‘homoscedasticity’, which means that the variance of the residual errors is independent of the values of the explanatory variables. If this assumption is not true, ordinary multiple regression does not work very well. This is another reason why analyzing multilevel data with ordinary multiple regression techniques does not work well.

As explained in the introduction in Chapter One, multilevel models are needed because with grouped data observations from the same group are generally more similar than the observations from different groups, which violates the assumption of independence of all observations. The amount of dependence can be expressed as a correlation coefficient: the intraclass correlation. The methodological literature contains a number of different formulas to estimate the intraclass correlation ρ . For example, if we use one-way analysis of variance with the grouping variable as independent variable to test the group effect on our outcome variable, the intraclass correlation is given by $\rho = [\text{MS}(A) - \text{MS}(\text{error})] / [\text{MS}(A) + (n-1) \times \text{MS}(\text{error})]$, where n is the common group size. Shrout and Fleiss (1979) give an overview of formulas for the intraclass correlation for a variety of research designs.

If we have simple hierarchical data, the multilevel regression model can also be used to produce an estimate of the intraclass correlation. The model used for this purpose is a model that contains no explanatory variables at all, the so-called *intercept-only* model. The intercept-only model is derived from equations (2.1) and (2.3) as follows. If there are no explanatory variables X at the lowest level, equation (2.1) reduces to

$$Y_{ij} = \beta_{0j} + e_{ij} \quad (2.6)$$

Likewise, if there are no explanatory variables Z at the highest level, equation (2.2) reduces to

$$\beta_{0j} = \gamma_{00} + u_{0j} \quad (2.7)$$

We find the single equation model by substituting (2.7) into (2.6):

$$Y_{ij} = \gamma_{00} + u_{0j} + e_{ij} \quad (2.8)$$

We could also have found equation (2.8) by removing all terms that contain an X or Z variable equation (2.5). The intercept-only model of equation (2.8) does not explain any variance in Y . It only decomposes the variance into two independent components: σ_e^2 , which is the variance of the lowest-level errors e_{ij} , and $\sigma_{u_0}^2$, which is the variance of the highest-level errors u_{0j} . Using this model, we can define the intraclass correlation ρ by the equation

$$\rho = \frac{\sigma_{u_0}^2}{\sigma_{u_0}^2 + \sigma_e^2} \quad (2.9)$$

The intraclass correlation ρ indicates the proportion of the variance explained by the grouping structure in the population. Equation (2.9) simply states that the intraclass correlation is the proportion of group level variance compared to the total variance.¹ The intraclass correlation ρ can also be interpreted as the expected correlation between two randomly chosen units that are in the same group.

2.2 AN EXTENDED EXAMPLE

Ordinary multiple regression analysis uses an estimation technique called Ordinary Least Squares, abbreviated as OLS. The statistical theory behind the multilevel regression model is more complex, however. Based on observed data, we want to estimate the parameters of the multilevel regression model: the regression coefficients and the variance components. The usual estimators in multilevel regression analysis are

¹ Note that the intraclass correlation is an estimate of the proportion of explained variance *in the population*. The proportion of explained variance in the *sample* is given by the correlation ratio η^2 (eta-squared, cf. Tabachnick & Fidell, 1996, p. 335): $\eta^2 = \text{SS(A)}/\text{SS(Total)}$.

Maximum Likelihood (ML) estimators. Maximum Likelihood estimators estimate the parameters of a model by providing estimated values for the population parameters that maximize the so-called Likelihood Function: the function that describes the probability of observing the sample data, given the specific values of the parameter estimates. Simply put, ML estimates are those parameter estimates that maximize the probability of finding the sample data that we have actually found. For an accessible introduction to maximum likelihood methods see Eliason (1993).

Maximum Likelihood estimation includes procedures to generate standard errors for most of the parameter estimates. These can be used in significance testing, by computing the test statistic Z : $Z = \text{parameter} / (\text{st. error param.})$. This statistic is referred to the standard normal distribution, to establish a p -value for the null-hypothesis that the population value of that parameter is zero. The Maximum Likelihood procedure also produces a statistic called the *deviance*, which indicates how well the model fits the data. In general, models with a lower deviance fit better than models with a higher deviance. If two models are *nested*, meaning that a specific model can be derived from a more general model by removing parameters from that general model, the deviances of the two models can be used to compare their fit statistically. For nested models, the difference in deviance has a chi-square distribution with degrees of freedom equal to the difference in the number of parameters that are estimated in the two models. The deviance test can be used to perform a formal chi-square test, in order to test whether the more general model fits significantly better than the simpler model. The chi-square test of the deviances can also be used to good effect to explore the importance of a set of random effects, by comparing a model that contains these effects against a model that excludes them. If the models to be compared are not nested models, the principle that models should be as simple as possible (theories and models should be parsimonious) indicates that we should generally keep the simpler model.

The intercept-only model is useful as a null-model that serves as a benchmark with which other models are compared. For our example data, the intercept-only model is written as

$$Y_{ij} = \gamma_{00} + u_{0j} + e_{ij}$$

The model that includes pupil gender and teacher experience, but not the cross-level interaction between those two, is written as

$$Y_{ij} = \gamma_{00} + \gamma_{10} X_{ij} + \gamma_{01} Z_j + u_{1j} X_{ij} + u_{0j} + e_{ij}$$

or, using variable names instead of algebraic symbols,

$$\text{popularity}_{ij} = \gamma_{00} + \gamma_{10} \text{gender}_{ij} + \gamma_{01} \text{experience}_j + u_{1j} \text{gender}_{ij} + u_{0j} + e_{ij}$$

Model:	M0: intercept-only		M1: + pup. gender and t. exp.	
Fixed part				
Predictor	coefficient	standard error	coefficient	standard error
intercept	5.31	0.10	3.34	0.16
pupil gender			0.84	0.06
teacher exp.			0.11	0.01
Random part				
σ_e^2	0.64	0.02	0.39	0.01
σ_{u0}^2	0.87	0.13	0.40	0.06
σ_{u1}^2			0.27	0.05
σ_{u01}			0.02	0.04
Deviance	5112.7		4261.2	

Table 2.1 presents the parameter estimates and standard errors for both models.¹ In this table, the intercept-only model estimates the intercept as 5.31, which is simply the average popularity across all schools and pupils. The variance of the pupil level residual errors, symbolized by σ_e^2 , is estimated as 0.64. The variance of the class level residual errors, symbolized by σ_{u0}^2 , is estimated as 0.87. All parameter estimates are much larger than the corresponding standard errors, and calculation of the Z-test shows that they are all significant at $p < 0.005$. The intraclass correlation, calculated by equation (2.9) $\rho = \sigma_{u0}^2 / (\sigma_{u0}^2 + \sigma_e^2)$, is $0.87/1.52$, which equals 0.58. Thus, 58% of the variance of the popularity scores is at the group level, which is very high. Since the intercept-only model contains no explanatory variables, the residual variances represent unexplained error variance. The deviance reported in Table 2.1 is a measure of model misfit; when we add explanatory variables to the model, the deviance is expected to go down.

The second model includes pupil gender and teacher experience as explanatory variables. The regression coefficients for both variables are significant. The regression coefficient for pupil gender is 0.84. Since pupil gender is coded 0=boy, 1=girl, this means that on average the girls score 0.84 points higher on the popularity measure. The regression coefficient for teacher experience is 0.11, which means that for each year of experience of the teacher, the average popularity score of the class goes up with 0.11 points. This does not seem very much, but the teacher experience in our example data ranges from 2 to 25 years, so the predicted difference between the least experienced and

¹ For reasons to be explained later, different options that can be chosen for the details of the Maximum Likelihood procedure may result in slightly different estimates. So, if you re-analyze the example data from this book, your results may differ slightly from the results given here. However, these differences should never be so large that you would draw entirely different conclusions.

the most experienced teacher is $(25-2) \times 0.11 = 2.53$ points on the popularity measure. We can use the standard errors of the regression coefficients reported in Table 2.1 to construct a 95% confidence interval. For the regression coefficient of pupil gender, the 95% confidence interval runs from 0.72 to 0.96, and the 95% confidence interval for the regression coefficient of teacher experience runs from 0.09 to 0.13.

The model with the explanatory variables includes a variance component for the regression coefficient of pupil gender, symbolized by σ_{u1}^2 in Table 2.1. The variance of the regression coefficients for pupil gender across classes is estimated as 0.27, with a standard error of 0.05. The covariance between the regression coefficient for pupil gender and the intercept is very small and obviously not significant.

The significant and quite large variance of the regression slopes for pupil gender implies that we should not interpret the estimated value of 0.84 without considering this variation. In an ordinary regression model, without multilevel structure, the value of 0.84 means that girls are expected to differ from boys by 0.84 points, for all pupils in all classes. In our multilevel model, the regression coefficient for pupil gender varies across the classes, and the value of 0.84 is just the expected value across all classes. In multilevel regression analysis, the varying regression coefficients are assumed to follow a normal distribution. The variance of this distribution is in our example estimated as 0.27. Interpretation of this variation is easier when we consider the standard deviation, which is the square root of the variance or 0.52 in our example data. A useful characteristic of the standard deviation is that with normally distributed observations about 67% of the observations lie between one standard deviation below and above the mean, and about 95% of the observations lie between two standard deviations below and above the mean. If we apply this to the regression coefficients for pupil gender, we conclude that about 67% of the regression coefficients are expected to lie between $(0.84-0.52) = 0.32$ and $(0.84+0.52) = 1.36$, and about 95% are expected to lie between $(0.84-1.04) = -0.20$ and $(0.84+1.04) = 1.88$. Using the more precise value of $Z_{.975} = 1.96$ we calculate the limits of the 95% interval as -0.18 and 1.86 . We can also use the standard normal distribution to estimate the percentage of regression coefficients that are negative. As it turns out, even if the mean regression coefficient for pupil gender is 0.84, about 5% of the classes are expected to have a regression coefficient that is actually negative. Note that the 95% interval computed here is totally different from the 95% confidence interval for the regression coefficient of pupil gender, which runs from 0.72 to 0.96. The 95% confidence interval applies to γ_{10} , the mean value of the regression coefficients across the classes. The 95% interval calculated here is the 95% *predictive interval*, which expresses that 95% of the regression coefficients of the variable ‘pupil gender’ in the classes are predicted to lie between -0.20 and 1.88 .

Given the large and significant variance of the regression coefficient of pupil gender across the classes it is attractive to attempt to predict its variation using class level variables. We have one class level variable: teacher experience. The individual level regression equation for this example, using variable labels instead of symbols, is given by equation (2.2), which is repeated below:

$$popularity_{ij} = \beta_{0j} + \beta_{1j} \text{gender}_{ij} + e_{ij} \quad (2.2, \text{repeated})$$

The regression equations predicting β_{0j} , the intercept in class j , and β_{1j} , the regression slope of pupil gender in class j , are given by equation (2.3) and (2.4), which are rewritten below using variable labels

$$\beta_{0j} = \gamma_{00} + \gamma_{01} \text{t.exp}_j + u_{0j} \quad (2.10)$$

$$\beta_{1j} = \gamma_{10} + \gamma_{11} \text{t.exp}_j + u_{1j} \quad (2.11)$$

By substituting (2.10) and (2.11) into (2.2) we get

$$\begin{aligned} popularity_{ij} = & \gamma_{00} + \gamma_{10} \text{gender}_{ij} + \gamma_{01} \text{t.exp}_j + \gamma_{11} \text{gender}_{ij} \times \text{t.exp}_j \\ & + u_{1j} \text{gender}_{ij} + u_{0j} + e_{ij} \end{aligned} \quad (2.12)$$

The algebraic manipulations of the equations above make clear that to explain the variance of the regression coefficients β_{1j} , we need to introduce an interaction term in the model. This interaction, between the variables pupil gender and teacher experience, is a cross-level interaction, because it involves explanatory variables from different levels. Table 2.2 presents the estimates from a model with this cross-level interaction. For comparison, the estimates for the model without this interaction are also included in Table 2.2.

The estimates for the fixed coefficients in Table 2.2 are similar for both models, except the regression slope for pupil gender, which is considerably larger in the cross-level model. The interpretation remains the same: girls are more popular than boys are. The regression coefficient for the cross-level interaction is -0.03 , which is small but significant. This interaction is formed by multiplying the scores for the variables ‘pupil gender’ and ‘teacher experience,’ and the negative value means that with experienced teachers, the advantage of being a girl is smaller than expected from the direct effects only. Thus, the difference between boys and girls is smaller with more experienced teachers.

Model:	M1: + pup. gender and t. exp.		M2: + cross-level interaction	
Fixed part				
Predictor	coefficient	standard error	coefficient	standard error
intercept	3.34	0.16	3.31	0.16
pupil gender	0.84	0.06	1.33	0.13
teacher exp.	0.11	0.01	0.11	0.01
pup. gender × teacher exp.			-.03	0.01
Random part				
σ_e^2	0.39	0.01	0.39	0.01
σ_{u0}^2	0.40	0.06	0.40	0.06
σ_{u1}^2	0.27	0.05	0.22	0.04
σ_{u01}	0.02	0.04	0.02	0.04
Deviance	4261.2		4245.9	

Comparison of the other results between the two models shows that the variance component for pupil gender goes down from 0.27 in the direct effects model to 0.22 in the cross-level model. Apparently, the cross-level model explains some of the variation of the slopes for pupil gender. The deviance also goes down, which indicates that the model fits better than the previous model.

The coefficients in Tables 2.1 and 2.2 are all unstandardized regression coefficients. To interpret them properly, we must take the scale of the explanatory variables into account. In multiple regression analysis, and structural equation models, for that matter, the regression coefficients are often standardized because that facilitates the interpretation when one wants to compare the effects of different variables within one sample. Only if the goal of the analysis is to compare parameter estimates from different samples to each other, should one always use unstandardized coefficients. To standardize the regression coefficients, as presented in Table 2.1 or Table 2.2, one could standardize all variables before putting them into the multilevel analysis. However, this would in general also change the estimates of the variance components. This may not be a bad thing in itself, because standardized variables are also centered on their overall mean. Centering explanatory variables has some distinct advantages, which are discussed in Chapter Four. Even so, it is also possible to derive the standardized regression coefficients from the unstandardized coefficients:

$$\text{Standardized coefficient} = \frac{(\text{unstandardized coeff.}) \times (\text{stand. dev. explanatory var.})}{\text{stand. dev. outcome variable}} \quad (2.13)$$

In our example data, the standard deviations are: 1.23 for popularity, 0.50 for gender, and 6.55 for teacher experience. Table 2.3 presents the unstandardized and standardized coefficients for the second model in Table 2.1. It also presents the estimates that we obtain if we first standardize all variables, and then carry out the analysis

Model:	Standardization after estimation		Using standardized variables		
Fixed part	unstandardized	standardized			
Predictor	coefficient	s.e.	coefficient	coefficient	s.e.
intercept	3.34	0.16	-	-	-
pupil gender	0.84	0.06	0.34	0.34	0.02
teacher exp.	0.11	0.01	0.59	0.58	0.05
Random part					
σ_e^2	0.39	0.01		0.26	0.01
σ_{u0}^2	0.40	0.06		0.32	0.05
σ_{u1}^2	0.27	0.05		0.05	0.01
σ_{u01}	0.02	0.04		0.05	70.02
Deviance	4261.2		3446.5		

Table 2.3 shows that the standardized regression coefficients are almost the same as the coefficients estimated for standardized variables. The small differences in Table 2.3 are simply rounding errors. However, if we use standardized variables in our analysis, we find very different variance components. This is not only the effect of scaling the variables differently, which becomes clear if we realize that the covariance between the slope for pupil gender and the intercept is significant for the standardized variables. This kind of difference in results is general. The fixed part of the multilevel regression model is invariant for linear transformations, just as the regression coefficients in the ordinary single-level regression model. This means that if we change the scale of our explanatory variables, the regression coefficients and the corresponding standard errors change by the same multiplication factor, and all associated p -values remain exactly the same. However, the random part of the multilevel regression model is not invariant for

linear transformations. The estimates of the variance components in the random part can and do change, sometimes dramatically. This is discussed in more detail in section 4.2 in Chapter Four. The conclusion to be drawn here is that, if we have a complicated random part, including random components for regression slopes, we should think carefully about the scale of our explanatory variables. If our only goal is to present standardized coefficients in addition to the unstandardized coefficients, applying equation (2.13) is safer than transforming our variables. On the other hand, we may estimate the unstandardized results, including the random part and the deviance, and then re-analyze the data using standardized variables, merely using this analysis as a computational trick to obtain the standardized regression coefficients without having to do hand calculations.

2.3 INSPECTING RESIDUALS

Inspection of residuals is a standard tool in multiple regression analysis to examine whether assumptions of normality and linearity are met (cf. Stevens, 1996; Tabachnick & Fidell, 1996). Multilevel regression analysis also assumes normality and linearity, and inspection of the residuals can be used for the same goal. There is one important difference from ordinary regression analysis; we have more than one residual, in fact, we have residuals for each random effect in the model. Consequently, many different residuals plots can be made.

2.3.1 Examples of Residuals Plots

The equation below represents the one-equation version of the direct effects model for our example data. This is the multilevel model without the cross-level interaction.

$$popularity_{ij} = \gamma_{00} + \gamma_{10} gender_{ij} + \gamma_{01} experience_j + u_{1j} gender_{ij} + u_{0j} + e_{ij}$$

In this model, we have three residual error terms: e_{ij} , u_{0j} , and u_{1j} . The e_{ij} are the residual prediction errors at the individual level, similar to the prediction errors in ordinary single-level multiple regression. A simple boxplot of these residuals will enable us to identify extreme outliers. An assumption that is usually made in multilevel regression analysis is that the variance of the residual errors is the same in all groups. This can be assessed by computing a one-way analysis of variance of the groups on the absolute values of the residuals, which is the equivalent of Levene's test for equality of variances in Analysis of Variance (Stevens, 1996). Bryk and Raudenbush (1992) describe a chi-square test that can be used for the same purpose, which is an option in the program HLM (Raudenbush, Bryk, Cheong, & Congdon, 2000).

The u_{0j} are the residual prediction errors at the group level, which can be used in ways analogous to the analysis of the individual level residuals e_{ij} . The u_{1j} are the residuals of the regression slopes across the groups. By plotting the regression slopes for the various groups, we get a visual impression of how much the regression slopes actually differ, and we may also be able to identify groups which have a regression slope that is wildly different from the others.

To test the normality assumption, we can plot standardized residuals against their normal scores. If the residuals have a normal distribution, the plot should show a straight diagonal line. Figure 2.1 is a scatterplot of the standardized level-1 residuals (denoted by 'const' in the graph) against their normal scores. The graph indicates close conformity to normality, and no extreme outliers. Similar plots can be made for the level-2 residuals.




Figure 2.1. Plot of level 1 standardized residuals against normal scores

We obtain a different plot, if we plot the residuals against the predicted values of the outcome variable popularity, using the fixed part of the multilevel regression model for the prediction. Such a scatter plot of the residuals against the predicted values provides information about possible failure of normality, nonlinearity, and heteroscedasticity. If these assumptions are met, the plotted points should be evenly divided above and below their mean value of zero, with no strong structure (cf. Tabachnick & Fidell, 1996, p. 137). Figure 2.2 shows this scatter plot. For our example data, the scatter plot in Figure 2.2 does not indicate strong violations of the assumptions.




Figure 2.2. Level 1 standardized residuals plotted against predicted popularity




Figure 2.3. Level 2 residuals plotted against predicted popularity

Similar scatter plots can be made for the second level residuals for the intercept and the slope of the explanatory variable pupil gender. As an illustration, Figure 2.3 shows the scatterplots of the level-2 residuals around the average intercept (denoted ‘const’ in the graph) and around the average slope of pupil gender against the predicted values of the outcome variable popularity.

The spread of the plotted points for pupil gender (denoted ‘sex’ in the plot) around their mean value of 0.0 suggests some degree of heterogeneity for the residuals around the slope of pupil gender. In our case, this heterogeneity is caused by a misspecification of the model, which is the result of omitting the cross-level interaction to explain the variance of the regression slopes of pupil gender.

An interesting plot that can be made using the level-2 residuals, is a plot of the residuals against their rank order, with an added error bar. In Figure 2.4, an error bar surrounds each point estimate, and the classes are sorted in rank order of the residuals. The error bars represent the confidence interval around the individual estimate, constructed by multiplying its standard error by 1.39. This results in confidence intervals that have the property that two classes have significantly different residuals (at the 5% level), if their error bars do not overlap (Goldstein, 1995).




Figure 2.4. Error bar plot of level 2 residuals

In our example, we see large differences between the classes. A logical next step would be to identify the classes at the extremes of the rank order, and to seek for a post hoc interpretation of what makes these classes different. For a discussion of the construction and use of these error bars see Goldstein and Healy (1995) and Goldstein and Spiegelhalter (1996).

Examining residuals in multivariate models presents us with a problem. For instance, the residuals should show a nice normal distribution, which implies absence of extreme outliers. However, this applies to the residuals after including all important explanatory variables and relevant parameters in the model. If we analyze a sequence of models, we have a series of different residuals for each model, and scrutinizing them all at each step is not always practical. On the other hand, our decision to include a specific variable or parameter in our model might well be influenced by a violation of some assumption. Although there is no perfect solution to this dilemma, a reasonable approach is to examine the two residual terms in the intercept-only model, to find out if there are gross violations of the assumptions of the model. If there are, we should accommodate them, for instance by applying a normalizing transformation, by deleting certain individuals or groups from our data set, or by including a dummy variable that indicates a specific outlying individual or group. When we have determined our final model, we should make a more thorough examination of the various residuals. If we detect gross violations of assumptions, these should again be accommodated, and the model should be estimated again. Of course, after accommodating an extreme outlier, we might find that we should now change our model again. Procedures for model exploration and detection of violations in ordinary multiple regression are discussed, for instance, in Tabachnick and Fidell (1996) or Stevens (1996). In multilevel regression, the same procedures apply, but the analyses are more complicated because we have to examine more than one set of residuals, and must distinguish between multiple levels.

As mentioned in the beginning of this section, graphs can be useful in detecting outliers and nonlinear relations. However, an observation may have an undue effect on the outcome of a regression analysis without being an obvious outlier. Figure 2.5, a scatter plot of the so-called Anscombe data (Anscombe, 1973), illustrates this point. There is one data point in Figure 2.5, which by itself almost totally determines the regression line. Without this one observation, the regression line would be very different. Yet, it does not show up as an obvious outlier.




Figure 2.5. Regression line determined by one single observation

In ordinary regression analysis, various measures have been proposed to indicate the influence of individual observations on the outcome (cf. Tabachnick & Fidell, 1996). In general, such *influence* or *leverage* measures are based on a comparison of the estimates when a specific observation is included in the data or not. Langford and Lewis (1998) discuss extensions of these influence measures for the multilevel regression model. Since most of these measures are based on comparison of estimates with and without a specific observation, it is difficult to calculate them by hand. However, if the software offers the option to calculate influence measures, it is advisable to do so. If a unit (individual or group) has a large value for the influence measure, that specific unit has a large influence on the values of the regression coefficients. It is useful to inspect cases with extreme influence values for possible violations of assumptions, or even data errors.

2.3.2 Examining Slope Variation: OLS and Shrinkage Estimators

The residuals can be added to the average values of the intercept and slope, to produce predictions of the intercepts and slopes in different groups. These can also be plotted.




Figure 2.6. Plot of the 100 class regression slopes for pupil gender

For example, Figure 2.6 plots the 100 regression slopes for the explanatory variable pupil gender in the 100 classes. It is clear that for most classes the effect is positive: girls are more popular than boys are. It is also clear that in some classes the relationship is the opposite: boys are more popular than girls are. Most of the regression slopes are not very different from the others, although there is one slope that appears to be much steeper than the others are. It could be useful to examine the data for that one class in more detail, to find out if there is a reason for this steeper slope.

The predicted intercepts and slopes for the 100 classes are not identical to the values we would obtain, if we carry out 100 separate ordinary regression analyses in each of the 100 classes, using standard Ordinary Least Squares (OLS) techniques. If we would compare the results from 100 separate OLS regression analyses to the values obtained from a multilevel regression analysis, we would find that the results from the separate analyses are more variable. This is because the multilevel estimates of the regression coefficients of the 100 classes are weighted. They are so-called Empirical Bayes (EB) or *shrinkage* estimates; a weighted average of the specific OLS estimate in each class and the overall regression coefficient, estimated for all similar classes.

As a result, the regression coefficients are *shrunk* back towards the mean coefficient for the whole data set. The shrinkage weight depends on the reliability of the estimated coefficient. Coefficients that are estimated with small accuracy shrink more than very accurately estimated coefficients. Accuracy of estimation depends on two factors: the group sample size, and the distance between the group-based estimate and the overall estimate. Estimates in small groups are less reliable, and shrink more than estimates from large groups. Other things being equal, estimates that are very far from the overall estimate are assumed less reliable, and they shrink more than estimates that are close to the overall average. The statistical method used is called *empirical Bayes estimation*. Due to this shrinkage effect, empirical Bayes estimators are biased. However, they are often more precise, a property that is often more useful than being unbiased (cf. Kendall, 1959).

For instance, in an intercept-only model the equation to form the empirical Bayes estimate of the intercept is given in the equation

$$\hat{\beta}_{0j}^{\text{EB}} = \lambda_j \hat{\beta}_{0j}^{\text{OLS}} + (1 - \lambda_j) \gamma_{00} \quad (2.14)$$

where λ_j is the reliability of the OLS estimate $\hat{\beta}_{0j}^{\text{OLS}}$ as an estimate of β_{0j} , which is given by the equation $\lambda_j = \sigma_{u_0}^2 / (\sigma_{u_0}^2 + \sigma_e^2 / n_j)$ (Bryk & Raudenbush, 1992, p. 39), and γ_{00} is the overall intercept. The reliability λ_j is close to 1.0 when the group sizes are large and/or the variability of the intercepts across groups is large. In these cases, the overall estimate γ_{00} is not a good indicator of each group's intercept. If the group sizes are small and have only small variation across groups, the reliability λ_j is close to 0.0, and more weight is put on the overall estimate γ_{00} . Equation (2.14) makes clear that, since the OLS estimates are unbiased, the empirical Bayes estimates $\hat{\beta}_{0j}^{\text{EB}}$ must be biased towards the overall estimate β_{00} . They are *shrunk* towards the average value γ_{00} . For that reason, the empirical Bayes estimators are also referred to as shrinkage estimators. Although the empirical Bayes or shrinkage estimators are biased, they are also in general closer to the (unknown) values of β_{0j} (Bryk & Raudenbush, 1992, p. 40). If the regression model includes a group level model, the shrinkage estimators are conditional on the group level model. The advantages of shrinkage estimators remain, *provided the group-level model is well specified* (Bryk & Raudenbush, 1992, p. 80). This is especially important if the estimated coefficients are used to describe specific groups. For instance, we can use estimates for the intercepts of the schools to rank order them on their average outcome. If this is used as an indicator of the quality of schools, the shrinkage estimators introduce a bias, because high scoring schools will be presented too negatively, and low scoring schools will be presented too positively. This is offset by the advantage of having a smaller standard error (Carlin & Louis, 1996; Lindley

& Smith, 1972). Bryk and Raudenbush discuss this problem in an example involving the effectiveness of organizations (Bryk & Raudenbush, 1992, chapter 5); see also the cautionary points made by Raudenbush and Willms (1991) and Snijders and Bosker (1999, pp. 58-63). All stress that the higher precision of the Empirical Bayes residuals is bought at the expense of a certain bias. The bias is largest when we inspect groups that are both small and far removed from the overall mean. In such cases, inspecting residuals should be supplemented with other procedures, such as comparing error bars for all schools (Goldstein & Healy, 1995). Error bars are illustrated in this chapter in Figure 2.4.

2.4 THREE- AND MORE-LEVEL REGRESSION MODELS

2.4.1 Multiple-level Models

In principle, the extension of the two-level regression model to three and more levels is straightforward. There is an outcome variable at the first, the lowest level. In addition, there may be explanatory variables at all higher levels. The problem is that three- and more-level models can become complicated very fast. In addition to the usual fixed regression coefficients, we must entertain the possibility that regression coefficients for first-level explanatory variables may vary across units of both the second and the third level. Regression coefficients for second-level explanatory variables may vary across units of the third level. To explain such variation, we must include cross-level interactions in the model. Regression slopes for the cross-level interaction between first-level and second-level variables may themselves vary across third-level units. To explain such variation, we need a second-order interaction involving variables at all three levels.

The equations for such models are complicated, especially when we do not use the more compact summation notation but write out the complete single equation-version of the model in an algebraic format (for a note on notation see section 2.5).

The resulting models are not only difficult to follow from a conceptual point of view, they may also be difficult to estimate in practice. The number of estimated parameters is considerable, and at the same time the highest level sample size tends to become relatively smaller. As DiPrete and Forristal (1994, p. 349) put it, the imagination of the researchers "...can easily outrun the capacity of the data, the computer, and current optimization techniques to provide robust estimates."

Having said that, three- and more-level models have their place in multilevel analysis. Intuitively, three-level structures such as pupils in classes in schools, or respondents nested within households, nested within regions, appear to be both conceptually and empirically manageable. If the lowest level is repeated measures over

time, having repeated measures on pupils nested within schools again does not appear to be overly difficult. In such cases, the solution for the conceptual and statistical problems mentioned is to keep models reasonably small. Especially specification of the higher-level variances and covariances should be driven by theoretical considerations. A higher-level variance for a specific regression coefficient implies that this regression coefficient is assumed to vary across units at that level. A higher-level covariance between two specific regression coefficients implies that these regression coefficients are assumed to covary across units at that level. Especially when models become large and complicated, it is advisable to avoid higher-order interactions, and to include in the random part only those elements for which there is strong theoretical or empirical justification. This implies that an exhaustive search for second-order and higher-order interactions is not a good idea. In general, we should seek for higher-order interactions only if there is strong theoretical justification for their importance, or if an unusually large variance component for a regression slope calls for explanation. For the random part of the model, there are usually more convincing theoretical reasons for the higher-level variance components than for the covariance components. Especially if the covariances are small and insignificant, analysts sometimes do not include all possible covariances in the model. This is defensible, with some exceptions. First, it is recommended that the covariances between the intercept and the random slopes are always included. Second, it is recommended to include covariances corresponding to slopes of dummy-variables belonging to the same categorical variable, and for variables that are involved in an interaction or belong to the same polynomial expression (Longford, 1990, p. 79-80).

2.4.2 Intraclass-correlations in three-level models

In a two-level model, the intraclass correlation is calculated in the intercept-only model using equation (2.9), which is repeated below:

$$\rho = \frac{\sigma_{u_0}^2}{\sigma_{u_0}^2 + \sigma_e^2} \quad (2.9, \text{repeated})$$

The intraclass correlation is an indication of the proportion of variance at the second level, and it can also be interpreted as the expected correlation between two randomly chosen individuals within the same group.

If we have a three-level model, for instance pupils nested within classes, nested within schools, there are several ways to calculate the intraclass correlation. First, we estimate an intercept-only model for the three-level data, for which the single-equation model can be written as follows:

$$Y_{ijk} = \gamma_{000} + v_{0k} + u_{0jk} + e_{ijk} \quad (2.15)$$

The variances at the first, second, and third level are respectively σ_e^2 , $\sigma_{u_0}^2$, and $\sigma_{v_0}^2$. The first method (cf. Davis & Scott, 1995) defines the intraclass correlations at the class and school level as

$$\rho_{class} = \frac{\sigma_{u_0}^2}{\sigma_{v_0}^2 + \sigma_{u_0}^2 + \sigma_e^2} \quad (2.16)$$

and

$$\rho_{school} = \frac{\sigma_{v_0}^2}{\sigma_{v_0}^2 + \sigma_{u_0}^2 + \sigma_e^2} \quad (2.17)$$

The second method (cf. Siddiqui, Hedeker, Flay & Hu, 1996) defines the intraclass correlations at the class and school level as

$$\rho_{class} = \frac{\sigma_{v_0}^2 + \sigma_{u_0}^2}{\sigma_{v_0}^2 + \sigma_{u_0}^2 + \sigma_e^2} \quad (2.18)$$

and

$$\rho_{school} = \frac{\sigma_{v_0}^2}{\sigma_{v_0}^2 + \sigma_{u_0}^2 + \sigma_e^2} \quad (2.19)$$

Actually, both methods are correct (Algina, 2000). The first method identifies the proportion of variance at the class and school level. This should be used if we are interested in a decomposition of the variance across the available levels, or if we are interested in how much variance is explained at each level (a topic discussed in section 4.4). The second method represents an estimate of the expected correlation between two randomly chosen elements in the same group. So ρ_{class} as calculated in equation (2.18) is the expected correlation between two pupils within the same class, and it correctly takes into account that two pupils who are in the same class must also be in the same school. For this reason, the variance components for classes and schools must both be in the numerator of equation (2.18). If the two sets of estimates are different, which may happen if the amount of variance at the school level is large, there is no contradiction involved. Both sets of equations express two different aspects of the data, which happen to coincide when there are only two levels.

2.5 A NOTE ABOUT NOTATION AND SOFTWARE

2.5.1 Notation

In general, there will be more than one explanatory variable at the lowest level and more than one explanatory variable at the highest level. Assume that we have P explanatory variables X at the lowest level, indicated by the subscript p ($p=1\dots P$). Likewise, we have Q explanatory variables Z at the highest level, indicated by the subscript q ($q=1\dots Q$). Then, equation (2.5) becomes the more general equation:

$$Y_{ij} = \gamma_{00} + \gamma_{p0} X_{pij} + \gamma_{0q} Z_{qj} + \gamma_{pq} Z_{qj} X_{pij} + u_{pj} X_{pij} + u_{0j} + e_{ij} \quad (2.20)$$

Using summation notation, we can express the same equation as

$$Y_{ij} = \gamma_{00} + \sum_p \gamma_{p0} X_{pij} + \sum_q \gamma_{0q} Z_{qj} + \sum_p \sum_q \gamma_{pq} X_{pij} Z_{qj} + \sum_p u_{pj} X_{pij} + u_{0j} + e_{ij} \quad (2.21)$$

The errors at the lowest level e_{ij} are assumed to have a normal distribution with a mean of zero and a common variance σ_e^2 in all groups. The u -terms u_{0j} and u_{pj} are the residual error terms at the highest level. They are assumed to be independent from the errors e_{ij} at the individual level, and to have a multivariate normal distribution with means of zero. The variance of the residual errors u_{0j} is the variance of the intercepts between the groups; it is symbolized by $\sigma_{u_0}^2$. The variances of the residual errors u_{pj} are the variances of the slopes between the groups; they are symbolized by $\sigma_{u_p}^2$. The *covariances* between the residual error terms $\sigma_{u_{pp}}$ are generally not assumed to be zero; they are collected in the higher level variance/covariance matrix Ω .¹

Note that in equation (2.15), γ_{00} , the regression coefficient for the intercept, is not associated with an explanatory variable. We can expand the equation by providing an explanatory variable that is a constant equal to one for all observed units. This yields the equation

$$Y_{ij} = \gamma_{p0} X_{pij} + \gamma_{pq} Z_{qj} X_{pij} + u_{pj} X_{pij} + e_{ij} \quad (2.22)$$

where $X_{0ij}=1$, and $p=0\dots P$. Equation (2.22) makes clear that the intercept is a regression coefficient, just like the other regression coefficients in the equation. Some multilevel software, for instance HLM (Raudenbush, Bryk, Cheongh & Congdon, 2000) puts the

¹ We may attach a subscript to Ω to indicate to which level it belongs. As long as there is no risk of confusion, the simpler notation without the subscript is used.

intercept variable $X_0=1$ in the regression equation by default. Other multilevel software, for instance MLwiN (Goldstein et al., 1998), requires that the analyst includes a variable in the data set that equals one in all cases, which must be added explicitly to the regression equation. In some cases, being able to eliminate the intercept term from the regression equation is a convenient feature.

Equation (2.22) can be made very general if we let \mathbf{X} be the matrix of all explanatory variables in the fixed part, symbolize the residual errors at all levels by $u^{(l)}$ with l denoting the level, and associate all error components with predictor variables \mathbf{Z} , which may or may not be equal to the \mathbf{X} . This produces the very general matrix formula $\mathbf{Y}=\mathbf{X}\boldsymbol{\beta}+\mathbf{Z}^{(l)}\mathbf{u}^{(l)}$ (cf. Goldstein, 1995, appendix 2.1). Since this book is more about applications than about mathematical statistics, it generally uses the algebraic notation, except when multivariate procedures such as structural equation modeling are discussed.

The notation used in this book is close to the notation used by Goldstein (1987, 1995), Hox (1995), and Kreft and de Leeuw (1998). The most important difference is that these authors indicate the higher-level variance by σ_{00} instead of our $\sigma_{u_0}^2$. The logic is that, if σ_{01} indicates the covariance between variables 0 and 1 , then σ_{00} is the covariance of variable 0 with itself, which is its variance. Bryk and Raudenbush (1992), and Snijders and Bosker (1999) use a different notation; they denote the lowest level error terms by r_{ij} , and the higher-level error terms by u_j . The lowest level variance is σ^2 in their notation. The higher-level variances and covariances are indicated by the Greek letter τ ; for instance, the intercept variance is given by τ_{00} . The τ_{pp} are collected in the matrix TAU, symbolized as T. The HLM program and manual in part use a different notation, for instance when discussing longitudinal and three-level models.

2.5.2 Software

Multilevel models can be formulated in two ways: (1) by presenting separate equations for each of the levels, and (2) by combining all equations by substitution into a single model-equation. The software HLM (Raudenbush et al., 2000) requires specification of the separate equations at each available level. Most other software (e.g., MLwiN; Rasbash et al., 2000), SAS Proc Mixed (Littell et al., 1996)) uses the single equation representation. Both representations have their advantages and disadvantages. The separate-equation representation has the advantage that it is always clear how the model is built up. The disadvantage is that it hides from view that modeling regression slopes by other variables results in adding an interaction to the model. As will be explained in Chapter Four, estimating and interpreting interactions correctly requires careful thinking. On the other hand, while the single-equation representation makes the existence of interactions obvious, it conceals the role of the complicated error components that are created by modeling varying slopes. In practice, to keep track of

the model, it is recommended to start by writing the separate equations for the separate levels, and to use substitution to arrive at the single-equation representation.

To take a quote from Singer's excellent introduction to using SAS Proc Mixed for multilevel modeling (Singer, 1998, p. 350): "Statistical software does not a statistician make. That said, without software, few statisticians and even fewer empirical researchers would fit the kinds of sophisticated models being promulgated today." Indeed, software does not make a statistician, but the advent of powerful and user-friendly software for multilevel modeling has had a large impact in research fields as diverse as education, organizational research, demography, epidemiology, and medicine. This book focuses on the conceptual and statistical issues that arise in multilevel modeling of complex data structures. It assumes that researchers who apply these techniques have access to and familiarity with *some* software that can estimate these models. Software is mentioned in various places, especially when a technique is discussed that is only available in a specific program. In addition to the relevant program manuals, several software programs have been discussed in introductory articles. Using SAS Proc Mixed for multilevel and longitudinal data is discussed by Singer (1998). Both Arnold (1992), and Heck and Thomas (2000) discuss multilevel modeling using HLM as the software tool. Sullivan, Dukes and Losina (1999) discuss HLM and SAS Proc Mixed. Hox (1995) applies the programs HLM, MLn and Varcl to the same data set, to highlight their similarities and differences. Kreft, de Leeuw and van der Leeden (1994) compare the programs BMDP-5V, Genmod, HLM, ML3 (a precursor to MLn/MLwiN) and Varcl on a variety of criteria, ranging from user-interface to statistical methods implemented. The multilevel procedure in SPSS is relatively new, and has not appeared in any published comparisons.

Since statistical software evolves rapidly, with new versions of the software coming out much faster than new editions of general handbooks such as this, I do not discuss software setups or output in detail. As a result, this book is more about the possibilities offered by the various techniques than about the specifics of how these things can be done in a specific software package. The various techniques are explained using analyses on small but realistic data sets, with examples of how the results could be presented and discussed. At the same time, if the analysis requires that the software used have some specific capacities, these are pointed out. This should enable interested readers to determine whether their software meets these requirements, and assist them in working out the software setups for their favorite package.

The data used in the various examples are described in the appendix, and are available through the Internet.