

Module 11: Three-Level Multilevel Models

George Leckie
Centre for Multilevel Modelling

Pre-requisites

- Modules 1-5

Contents

What are Three-Level Multilevel Models?	2
Introduction to the Example Dataset	3
C11.1 Understanding Three-Level Data Structures	6
C11.1.1 Three-level data structures	6
C11.1.2 Four-level data structures	8
C11.2 A Three-Level Variance Components Model	11
C11.2.1 Specifying the three-level model	11
C11.2.2 Interpretation of the intercept and the random effects	12
C11.2.3 Testing for cluster effects	13
C11.2.4 Calculating coverage intervals, variance partition coefficients (VPCs) and intraclass correlation coefficients (ICCs)	16
C11.2.5 Predicting and examining cluster effects	21
C11.2.6 Example: Students nested within school-cohorts nested within schools	21
C11.3 Adding Predictor Variables	24
C11.3.1 Adding level 1, level 2 and level 3 predictor variables	24
C11.3.2 Example: Students nested within school-cohorts nested within schools	25
C11.4 Adding Random Coefficients	29
C11.4.1 Adding higher-level random coefficients	29
C11.4.2 Example: Students nested within school-cohorts nested within schools	30
C11.5 Adding Further Levels	37
C11.5.1 A four-level model	37
C11.5.2 Example: Students nested within school-cohorts nested within schools nested within LAs	37

Further reading	39
References	40

If you find this module helpful and wish to cite it in your research, please use the following citation:

Leckie, G. (2013). Three-Level Multilevel Models - Concepts. LEMMA VLE Module 11, 1-41.
<http://www.bristol.ac.uk/cmm/learning/course.html>

Address for correspondence:

George Leckie
Centre for Multilevel Modelling
University of Bristol
35 Berkeley Square
Bristol, BS8 1JA
UK

g.leckie@bristol.ac.uk

What are Three-Level Multilevel Models?

In the previous modules we illustrated two-level multilevel models for analysing two-level data structures where units (level 1) are nested within groups (or clusters) (level 2). When the groups are themselves nested within supergroups (or superclusters) (level 3), the data form a three-level hierarchy and three-level models can be fitted to account for the additional level. Examples of three-level data structures include: students (level 1) nested within classes (level 2) nested within schools (level 3); voters nested within counties nested within states; and patients nested within doctors nested within clinics. In this module, we describe three-level data structures and multilevel models which can be used to analyse them. Of course, there is nothing to stop data structures being even more complex and having four or more levels and we shall also consider examples of such data structures in this module. Many further examples of three- and four-level data structures are described in C4.2 and C4.3 of Module 4, respectively.

It is important to incorporate three-level structures in to our models when they arise in the data and lead the higher level clusters to differ substantially from one another on the response variable under study. Naively fitting two-level models to three-level data will lead us to misattribute response variation to the two included levels (van Landeghem et al., 2005; Moerbeek, 2004; van den Noortgate et al., 2005; Tranmer and Steele, 2001). This in turn may lead us to draw misleading conclusions about the relative importance of different sources of influence on the response. For example, fitting a students-within-classes two-level model of student attainment and ignoring the fact that classrooms are further nested within schools will likely lead us to overstate the importance of classrooms as a source of variation in student attainment. That is, much of the variation that we attribute to classrooms may be driven by school-to-school differences in attainment. Our naïve analysis would therefore overstate the importance of classrooms on student attainment and would ignore the role of schools (i.e. school policies, practices, context and compositional effects). Furthermore, by incorrectly modelling the dependency in the data we will likely obtain biased standard errors for the predictor variables, particularly those measured at higher levels. We therefore run the risk of making incorrect inferences and drawing misleading conclusions about the relationships being studied. For example, including school-level predictor variables in our students-within-classes two-level model, but ignoring school as a level in the model will typically lead us to severely underestimate the standard errors on these school-level variables. When we then go on to test the significance of these variables, we will run the risk of making type 1 errors of inference.

Introduction to the Example Dataset

In educational research, there is considerable interest in measuring the effects that schools have on students' educational achievements. Measuring the effects that schools have on their students is after all a necessary first step to learning how schools' policies and practices combine to generate differences between schools. Governments are also often interested in measuring school effects, typically for school accountability purposes, but often to also provide parents with information to help guide school choice. However, in nearly all education systems, there are substantial differences between schools in their students' attainments at intake (i.e. when students first arrive at their schools). For the purposes of researching the effects of schools' policies and practices, holding schools accountable, or informing school choice, schools should not be compared simply in terms of their average exam results as these differences will, at least in part, be driven by these initial differences.

Traditional studies of school effects attempt to measure the 'true' effects that schools have on their students by fitting two-level students-within-schools multilevel models to students' exam scores where covariate adjustments are made for students' initial scores, and typically for a range of other student background characteristics. The school-level residuals from these models are then argued to measure the effects that schools have on their students having adjusted for the non random selection of students into schools. These effects are interpreted as measuring the influences schools have on their students' academic progress (improvement or change in attainment) while they attend their schools. In school effectiveness research these influences are referred to as 'value-added' effects.

In terms of studying students' academic progress, there are many other potential sources of clustering or influence which may also be important determinants of student progress. For example, where data contain multiple academic cohorts of students, we can think of schools as potentially having different effects in different academic cohorts. This leads students from the same school-cohort to appear more alike than students from different school-cohorts. The data are then three-level with students (level 1) nested within school-cohorts (level 2) nested within schools (level 3). In this module, we shall introduce three-level multilevel models to explore such data. In particular, we shall focus on the stability of school effects over time by examining the extent to which school effects change from cohort to cohort.

We shall then go on to consider the further nesting of schools within administrative educational regions referred to as local authorities (LAs) (level 4).¹ In England, secondary schools are organised into 150 LAs. Traditionally, LAs controlled the distribution of government funds across schools, co-ordinated school admissions, and were the direct employers of all teachers and staff in many schools. While over the last few years there has been a reduction of LAs' powers, one might still expect to identify LA effects in the data. If nothing else, we would expect LA effects to pick up geographic variation in student attainment that exists across England.

¹ LAs correspond to school districts in the U.S.

We shall use data from England's National Pupil Database (NPD), a census of all students in state (i.e. government funded) schools in England. The data are provided by the Department for Education (<http://www.education.gov.uk>). The NPD records students' academic attainments and a limited number of background characteristics. We focus on three consecutive academic cohorts of students who sat their General Certificate of Secondary Education (GCSE) examinations (age 16 years) in London schools in 2008, 2009 and 2010, respectively. These students sat their Key Stage 2 (KS2) examinations (age 11 years) five years earlier in 2003, 2004 and 2005, respectively.^{2 3}

Table 11.1 presents the number of units at each level of this data hierarchy.

Table 11.1 Number of units at each level of the data hierarchy

Level number	Level	Number of units
4	LAs	32
3	Schools	427
2	School-cohorts	1,232
1	Students	189,940

Thus, there are 32 LAs at level 4, 427 schools at level 3, 1,232 school-cohorts at level 2 and 189,940 students at level 1 of the data hierarchy. At this point it is helpful to explicitly define 'cohort' and 'school-cohort'. When we say 'cohort' we are referring to the three academic cohorts in the data: 2008, 2009 and 2010. When we refer to 'school-cohorts' we are referring to the 1,232 groups, or school-by-cohort combinations of students, in the data which are formed by crossing the 427 schools by the three cohorts. The number of schools and students present in the data for each cohort are as follows. In 2008 there were 412 schools and 63,208 students. In 2009 there were 410 schools and 63,072 students. In 2010 there were 410 schools and 63,660 students.⁴ Three hundred and ninety five schools had all three cohorts represented in the data, 15 schools had only two of the three cohorts, while a further 17 schools had only one of the cohorts present. The 32 schools which were not present for one or more cohorts reflect the opening of new schools and the closing of old schools.

The response variable for all our analyses is a continuous point score summarising students' overall attainment in their GCSE examinations.⁵ To ease the interpretation of this variable, and so that the residuals at each level better approximate the normality assumptions of the models, we transform it to a standard normal score

² GCSE examinations are taken in the last year of secondary schooling. Successful GCSE results are often a requirement for taking A-level examinations (age 18 years) which in turn are a common type of university entrance determinant. For those who leave school at 16 years of age, GCSE results are their main job market qualification.

³ KS2 examinations are taken in the last year of primary schooling.

⁴ The 2010 cohort of 410 schools and 63,660 students will provide the example dataset in Modules 11 and 12.

⁵ Specifically, the response variable is the student's capped 'best 8' total point score at GCSE with an additional bonus for attainment in each of English and Mathematics, and is the same measure as that published in Government school performance tables (see <http://www.education.gov.uk/performance-tables>).

which has the property of being more normally distributed with mean zero and variance one.⁶ This transformation allows the effects of the covariates in our multilevel models to be interpreted in terms of standard deviation units of the response. As our focus is on the stability of school effects across cohorts and not on any overall, London-wide, trend in student attainment over time, we carry out this transformation separately for each cohort. Put differently, in this analysis we are interested in the relative performance of schools to one another; we are not interested in the average absolute performance of schools.

We consider eight student-level predictor variables

- Attainment at age 11 (average point score across English, maths and science) (transformed to a standard normal score)
- Female (0 = male; 1 = female)
- Age (ranges from 0 to 1 where higher values correspond to older children; specifically, 0 corresponds to the youngest child in the data, born on the last day of the academic year, while 1 corresponds to the oldest child in the data, born on the first day of the academic year)
- Eligible for free school meals (FSM) (0 = no FSM; 1 = FSM)
- Special education needs (SEN) (0 = no SEN; 1 = SEN)
- English as an additional language (EAL) (0 = no EAL; 1 = EAL)
- Ethnicity (1 = White; 2 = Mixed; 3 = Asian; 4 = Black; 5 = Chinese; 6 = Other)
- Index of deprivation affecting children index (IDACI) a measure of residential neighbourhood social deprivation (transformed to a standard normal score)

and one school-cohort-level variable

- Cohort (1 = 2008; 2 = 2009; 3 = 2010)

⁶ The transformation is carried out by first ranking the N students by their original scores. The standard normal score for the i th ranked student in the data is then $\Phi^{-1}\{(i - 0.5)/N\}$, where Φ^{-1} denotes the inverse of the standard normal cumulative distribution function. This transformation is order preserving and students with the same original scores will also be tied in terms of their standard normal scores.

This document is only the first few pages of the full version.

To see the complete document please go to learning materials and register:

<http://www.cmm.bris.ac.uk/lemma>

The course is completely free. We ask for a few details about yourself for our research purposes only. We will not give any details to any other organisation unless it is with your express permission.