Module 13: Multiple Membership Multilevel Models

George Leckie
Centre for Multilevel Modelling

Pre-requisites

- Modules 1-5,11,12

Contents

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

Leckie, G. (2013). Multiple Membership Multilevel Models - Concepts. LEMMA VLE Module 13, 1-61.
http://www.bristol.ac.uk/cmm/learning/course.html

Address for correspondence:

George Leckie
Centre for Multilevel Modelling
University of Bristol
2 Priory Road
Bristol, BS8 1TX
UK

g.leckie@bristol.ac.uk

What are Multiple Membership Multilevel Models?

In Module 12, we illustrated multilevel models for analyzing cross-classified data structures. In this Module, we describe and illustrate another important form of non-hierarchical data, multiple membership data.

In multiple membership data, lower level units do not belong to one and only one higher level unit. Rather, lower level units are nested within multiple higher level units from the same classification. An example in educational research arises in studies of student attainment where students are often taught by more than one teacher and so students are described as being multiple members of those teachers. An example in health services research arises in studies of hospital patient outcomes. Hospital patients will often be seen and treated, not by one doctor, but by multiple doctors and so patients are described as being multiple members of doctors. There is nothing to stop data structures being even more complex and having two or more multiple membership structures and we shall consider examples of such data in this module. Many further examples of multiple membership data structures are described in C4.4.3 in Module 4.

An important feature of multiple membership data structures is that the degree to which each lower level unit belongs to each higher level unit will often vary across those higher level units. Multiple membership weights are used to quantify this and this information is used when fitting multiple membership models. In our educational research example, students may spend more time with some teachers than others. Here we might define multiple membership weights as the proportion of time spent with each teacher. Thus, if a student is taught for two lessons a week by teacher A and three lessons a week by teacher B, we would assign multiple membership weights to that student of 0.4 and 0.6 for teachers A and B respectively. These weights reflect the fact that we might expect teacher B to be more influential in determining the student's outcome than teacher A. Other weighting schemes may equally be applied, and it can be interesting to explore a range of weighting schemes as part of a sensitivity analysis for the model.

It is important to incorporate multiple membership structures in to our models when they arise in the data. In our educational research example, it seems intuitive that we should recognize the contribution or effect that every teacher has on student attainment. Similarly, in our health services research example, it seems intuitive to recognize the role that every doctor plays in contributing to patients' outcomes. Naively assigning each lower level unit to just one of their higher level units and then fitting the nearest equivalent hierarchical model to multiple membership data will lead us to misattribute response variation to the 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, assigning students to the first teacher that teaches them and then fitting a students-within-teachers two-level model of student attainment will likely lead us to underestimate the importance of teachers and overestimate the importance of students as sources of variation in student attainment. 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 teacher-level predictor variables in our students-within-teachers two-level model, but ignoring the fact that students are multiple members of teachers will typically lead us to underestimate the standard errors on these teacher-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

We will illustrate multiple membership models in the context of the same school effectiveness application which was analysed in Modules 11 and 12. Readers familiar with this application may wish to skip the next two paragraphs.

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 Modules 11 and 12 we used three-level and cross-classified models to explore the stability of school effects over time and the potential effects that the school attended in an earlier phase of education might have on students' academic progress. However, a potentially important problem that we did not consider was student mobility. When students move between schools during the period in which we measure student progress, we should model students as belonging to every school that they attend and not just their final school. When students move schools, students can no longer be viewed as being strictly nested within schools. Rather students (level 1) should be viewed as being multiple members of schools (level 2). In this module, we shall introduce multiple membership multilevel models to analyse such data and to therefore better measure the effects that schools have on student progress during secondary schooling in the English school system.

We shall then go on to consider the further nesting of students within residential neighbourhoods and additionally the further nesting of schools within administrative educational regions referred to as local authorities (LAs).[1] Accounting for neighbourhoods and LAs leads to an even more complex data structure which has both multiple membership and cross-classified sub structures within it.

Another potential influence on child learning is the residential neighbourhood within which a child lives. For example, we might ask: Do communities where adults have few educational qualifications entrench low academic aspirations in children growing up there? Alternatively, we might ask: Do gangs and children playing truant in deprived neighbourhoods disrupt not only their own education but also that of other children in the street? An important issue when considering neighbourhood effects is the spatial scale at which they are purported to operate. There are a multitude of spatial scales in UK geography. We focus on lower super output areas (LSOAs), which were designed using the 2001 UK Census and are defined to be fairly consistent in size, having a mean population of approximately 1,500, and to reflect as far as possible social homogeneity.

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, we 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.

A further source of clustering which may be important relates to the role that schools attended in an earlier phase of education may continue to exert on students after they have left these schools. Indeed, in Module 11 we examined whether primary school attended (ages 4 to 11) predicted student academic progress during secondary schooling (ages 11 to 16). While in that module we found some evidence that primary school effects do carry over into secondary schooling, for simplicity we will not consider the role of primary school further in this module.

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 the academic cohort of students who sat their General Certificate of Secondary Education (GCSE) examinations (age 16 years) in London schools in 2010 and their Key Stage 2 (KS2) examinations (age 11 years) five years earlier in 2005.[2] [3]

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

[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.

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.