

Module 4: Multilevel structures and classifications

Jon Rasbash
Centre for Multilevel Modelling

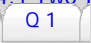
Contents

Aims	2
Introduction	2
C4.1 Two-level hierarchical structures	4
C4.1.1 Students within schools	4
C4.1.2 Issues of sample size	6
C4.1.3 Variables and levels, fixed and random classifications	7
C4.1.4 Other examples of a two-level structure	9
C4.1.5 Repeated measurements within individuals, panel data	10
C4.1.6 Multivariate responses within individuals	12
C4.1.7 Two-stage sample survey design	14
C4.1.8 An experimental design in which the intervention is at the higher level	17
C4.2 Three-level structures	20
C4.2.1 Students within classes within schools	20
C4.2.2 A repeated cross-sectional design: students within cohorts within schools	22
C4.3 Four-level structures	25
C4.3.1 Doubly nested repeated measures	25
C4.4 Non-hierarchical structures	28
C4.4.1 Cross-classifications: students cross-classified by school and neighbourhood	28
C4.4.2 Repeated measures within a cross classification of patients by clinician.	30
C4.4.3 Multiple Membership Structures	32
C4.5 Combining structures: hierarchies, cross-classifications and multiple membership relationships	34
C4.6 Spatial structures	37
C4.7 Summary	39

Some of the sections within this module have online quizzes for you to test your understanding. To find the quizzes:

EXAMPLE

From within the LEMMA learning environment

- Go down to the section for **Module 4: Multilevel Structures & Classifications**
- Click "[4.1 Two-level hierarchical structures](#)" to open Lesson 4.1
- Click  to open the first question

Aims

After completing this chapter you will be able to:

- Recognise a range of multilevel structures and classifications and how they correspond to real-world situations, research designs, and/or social-science research problems;
- Appreciate the different types of data frames associated with each structure and how subscripts are used to represent structure;
- Begin to appreciate 'targets of inference';
- Distinguish between levels and variables, and fixed and random classifications;
- Appreciate that multilevel structures are likely to generate dependent, correlated data that requires special modelling;
- Recognise the difference between long and wide forms of data structures;
- Begin to appreciate the advantages, both technical and substantive, of using a multilevel model, and the disadvantages of not doing so.

Introduction

Multilevel modelling is designed to explore and analyse data that come from populations which have a complex structure. In any complex structure we can identify *atomic* units. These are the units at the lowest level of the system. Often, but not always, these atomic units are individuals. Individuals are then grouped into *higher-level* units, for example schools. By convention we then say that students are at level 1 and schools are at level 2 in our structure.

This module aims to give a 'pictionary' of structures that underlie multilevel models. We give 'pictures' of common structures as *unit diagrams*, as *classification diagrams*, as data frames and in words. Note that the terms *classification* and *level* can be used somewhat interchangeably but level implies a nested hierarchical relationship of units (in which lower units nest in one, and only one, higher-level unit) whereas classification does not. The data frames, in addition to showing the structure, will also provide some example *explanatory* (predictor) variables and a *response* (y variable) as discussed in Module 2. We have

chosen the following examples to show a range of population structures where multilevel modelling is useful, and often necessary. We have also tried to introduce what are often seen as demanding and difficult concepts in a straightforward manner (e.g. fixed and random classifications, missing at random). While we have given the basic structures in a schematic and rather abstract form, we always point to published examples where the structure has been used in research.

C4.1 Two-level hierarchical structures

Hierarchical structures arise when the lower-level unit nests in one and only one higher-level unit. Such a relatively simple structure can, as we shall see, accommodate a wide range of study designs and research questions.

C4.1.1 Students within schools

Figure 4-1 is a *unit diagram* which aims to show the underlying structure of a research problem in terms of individual units; the nodes on the diagram are specific population units. In this case the units are students and schools which form two levels (or classifications). The lower units form the student classification (St1, St2 etc.) and the higher units form the school classification (Sc1, ... , Sc4). This unit diagram is just a schema to convey the essential structure of students nested within schools. In a real data set we would have many more than four schools and 12 students. The hierarchical structure means that a student only attends one school and has not moved about. Such a structure may arise when we are interested in school performance and we make repeated measurements of this by assessing student performance for multiple students from each school. This structure is likely to give rise to correlated or non-independent data, in the sense that students in the same school will often have a tendency to be similar on such variables as exam performance. Even if the initial allocation to a group was at random, social processes usually act to create this dependence. Traditionally, statistical modelling has faced difficulties with such dependence, indeed it has largely assumed it does not exist, but with multilevel modelling such correlation is expected and explicitly modelled.




Figure 4-1. Unit diagram of a two-level nested structure; students in schools

This two-level nested structure can also be represented by a *classification diagram* (Figure 4-2). Classification diagrams have one node per classification (or level). Nodes joined by a single arrow indicate a nested (strict hierarchical) relationship between the classifications.

C4.1.1 Students within schools




Figure 4-2. Classification diagram of a two-level nested structure; students in schools

Classification diagrams are more abstract than unit diagrams and are particularly useful, as we shall see, when the population being studied has a complex structure with many classifications.

Table 4.1 shows a data frame for the structure shown in Figure 4-1. We have also included a response (exam score in the current year), one school-level explanatory variable (school type), and two student-level explanatory variables (gender and previous exam score, say two years earlier). You will notice that the response is measured on the atomic unit, that is, level 1 (students); and that school 1 has three students, while school 4 has four students. That is, the data are not *balanced*; multilevel models do not require that there are the same number of lower level units in each and every higher level unit. In this example (and by common convention) the subscript *i* is used to index (represent) the lower level unit of the Student, while the subscript *j* indexes Schools.

With such a data frame we could ask a very rich set of questions by using a two-level multilevel model in which a student's current attainment is related to prior attainment (a previous test score) and there are data available on the gender of the student and the public/private nature of the school; these include

- i) Do males make greater progress than females?
- ii) Does the gender gap vary across schools?
- iii) Are males more or less variable in their progress than females?
- iv) What is the between-school variation in students' progress?
- v) Is School X (that is, a specific school) different from other schools in the sample in its effect?
- vi) Is there more variability in progress between schools for students with low prior attainment?
- vii) Do students make more progress in private than public schools?
- viii) Are students in public schools less variable in their progress?
- ix) Do girls make greater progress in state schools¹

¹ A classic study of school effects with an extended discussion of the issues involved is given by Aitkin, M. and Longford, N.T. (1986) Statistical modelling issues in school effectiveness studies (with Discussion). *J. Roy. Statist. Soc. A* 149, 1-43. Other examples include Goldstein, H.,

C4.1.1 Students within schools

Questions ii, iii, iv, vi, and viii can be addressed by modelling variability as functions of explanatory variables, whereas questions i, v, vii, and ix are about modelling the mean as a function of explanatory variables. The defining strength of multilevel modelling is that it can do both, that is, model the mean and the variance simultaneously (traditional techniques can only model the mean). This idea may seem a little confusing at the moment but it is a theme we will be returning to throughout these training materials.

Table 4.1. Data frame representation of Figure 4.1 and 4.2: a two-level study for examining school effects on student progress

Classifications or levels		Response	Explanatory variables		
Student i	School j	Student exam score $_{ij}$	Student previous examination score $_{ij}$	Student gender $_{ij}$	School type $_j$
1	1	75	56	M	State
2	1	71	45	M	State
3	1	91	72	F	State
1	2	68	49	F	Private
2	2	37	36	M	Private
3	2	67	56	M	Private
1	3	82	76	F	State
2	3	85	50	F	State
1	4	54	39	M	Private
2	4	91	71	M	Private
3	4	43	41	M	Private
4	4	66	55	F	Private

C4.1.2 Issues of sample size

A question that often comes up at this point is how many units are needed at each level. It is difficult to give specific advice but there are some general principles that are worth stating now. The key one is the *target of inference*: in other words, are the units in your dataset special ones that you are interested in in their own right, or are you regarding them as representatives of a larger population which you wish to use them to draw conclusions about? If the target

Rasbash, J., Yang, M., Woodhouse, G., et al. (1993). A multilevel analysis of school examination results. *Oxford Review of Education* 19: 425-433, Thomas, S (2001) Dimensions of Secondary School Effectiveness: Comparative Analyses Across Regions. *School Effectiveness and School Improvement* 12(3), 285 - 322

C4.1.2 Issues of sample size

of inference in an educational study is a particular school then you would need a lot of students in that school to get a precise effect. If the target of inference is between-school differences in general, then you would need a lot of schools to get a reliable estimate. That is, you could not sensibly use a multilevel model with only two schools even if you had a sample of 1000 students in each of them. In the educational literature it has been suggested that, given the size of effects that are commonly found for between-school differences, a minimum of 25 schools is needed to provide a precise estimate of between-school variance, with a preference for 100 or more schools.² You would not normally omit any school from the analysis merely because it has few students, but at the same time you will not be able to distinguish between-school and between-student variation if there is only one student in each and every school. Note that schools with only one pupil still add information to the estimates of the effects of the explanatory variables on the mean. There are, of course, some contexts where some or all of the higher-level units will have only a few lower-level units. An extreme and common case is when individuals are at level 1 and households are at level 2, because then the sample size within a level 2 unit is typically less than five people. This need not be a problem if the target of inference is households in general because the quality of estimates in this case is based on the total number of households in the sample and it should be possible to sample a large number of these. If the target of inference is a specific household, however, parameters will be poorly estimated because a single household has very few members. See Snijders and Bosker (1993)³ for more details on sample size issues for multilevel models.

C4.1.3 Variables and levels, fixed and random classifications

We now come straight up against an issue which causes a lot of confusion: When is a variable to be treated as a classification or level as opposed to an explanatory variable? For example, school type is a classification of schools so why not redraw Figure 4-1, Figure 4-2, and re-specify Table 4.1 as a three-level multilevel model (with the subscript ijk representing students in schools in type of school), as shown in Table 4.2 and Figure 4-3.

School type is certainly a way of classifying schools and as such it is a classification. However, we can divide classifications into two types which are treated in different ways when modelling:

- i) random classifications and
- ii) fixed classifications.

² L Paterson, H Goldstein (1991) New Statistical Methods for Analysing Social Structures: An Introduction to Multilevel Models, *British Educational Research Journal*, 17(4), 387-393; <http://www.jstor.org/view/01411926/ap050037/05a00080/0>

³ Snijders, T.A.B., and Bosker, R.J. (1993). Standard errors and sample sizes for two-level research. *J. Educational Statist.*, 18, 237-259

A classification is a random classification if its units can be regarded as a *random* sample from a wider population of units. For example the students and schools in our example are a random sample from a wider population of students and schools. However, school type or indeed student gender has a small *fixed* number of categories. There is no wider population of school types or genders to sample from. State and private are not two types sampled from a large number of school types, and male and female are not just two of a possibly large number of genders. Students and schools, however, can be treated as a sample of students and schools to which we want to generalise.

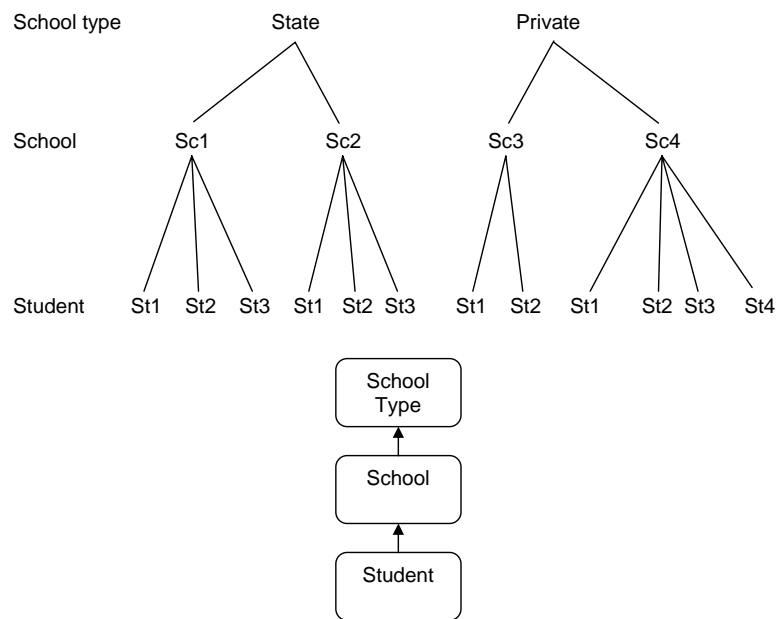


Figure 4-3. Unit and classification diagrams for a three-level nested structure; students in schools in school types

Table 4.2. Data frame representation of Figure 3. 3: a three-level study of students nested in schools in school type

Classifications or levels			Response	Explanatory Variables	
<i>Student i</i>	<i>School j</i>	<i>School type k</i>	<i>Student exam score_{ijk}</i>	<i>Previous exam score_{ijk}</i>	<i>Student gender_{ijk}</i>
1	1	State	75	56	M
2	1	State	71	45	M
3	1	State	91	72	F
1	2	Private	68	49	F
2	2	Private	37	36	M
Etc					

The distinction between fixed and random classifications has important implications for how we handle the classifying variable in a statistical analysis. Strictly speaking, for a classification to correspond to a level in a multilevel model, it must be a random classification. It turns out that school type is not a random classification; it is better to conceive of this as a fixed classification, and that is why we would treat school type as a variable not as a level.

The distinction between fixed and random classifications is also linked to the concept of target of inference. If you wanted to infer to each and every specific school as an island by itself, then this would be a fixed classification. Using a fixed classification you would be able to estimate whether School X is different in terms of progress from School Z and School W, etc. If, however, the target of inference is between-school differences then the random classification should be used. To take another example, if the target of inference is a specific country in comparison to others (see Module 3) then a fixed classification is needed, but if between-country variation is the target, then a random classification would be needed.

C4.1.4 Other examples of a two-level structure

Other common examples of two-level nested structures are people within households, patients within hospitals, and people within neighbourhoods. All these examples and the students-in-schools example arise when the real world has a multilayered structure, that is, the levels exist in the *population*. However, the multilevel structure might also be imposed through the study *design* and data collection. We now cover four interesting examples of two-level nested structures that arise from common types of research design that can be found in social science research:

- Repeated measures, panel data;
- Multivariate designs;

C4.1.4 Other examples of a two-level structure

- Multistage survey designs;
- Intervention studies where the intervention is made at the group level.

C4.1.5 Repeated measurements within individuals, panel data

Table 4.1 shows the case when there are measures on an individual on two time occasions, e.g. a prior test score and a current test score for students. We can analyse change (that is, progress) in this situation by specifying the current attainment as the response and including prior score as an explanatory variable. However, when there are measurements on more than two occasions there are advantages in treating occasion as a level nested within individuals. Such a two-level strict hierarchical structure is known as a repeated measurement or panel design (Figure 4-4). This occurs when we have repeated measurements (level 1) over time on a number of people (level 2). We think of measurement occasion as being nested within individuals and this is an example of when the atomic units are not individuals. These structures apply when we are analysing the extent and nature of variation between individuals in their patterns of growth. With repeated measurements we can often expect quite strong correlations across time within individuals. For example, a tall person is likely to continue to be tall across time and their value this year is likely to be related to the value last year. Again multilevel models anticipate this dependency across time and explicitly model it.⁴




Figure 4-4. Classification and unit diagram for a two-level repeated measures design

The data frame for the repeated measures structure shown in Table 4.3 a) is in *long form*- i.e. it has one row per measurement occasion. However, repeated measures data often come in *wide form*- i.e. there is one row per individual with a column for each measurement occasion. Many packages for multilevel analyses, including *MLwiN*, require repeated measures data in the long form. The wide form tends to correspond to the traditional multivariate framework for analysing repeated measures data (where all individuals must be measured at the same time points) and the long form tends to correspond to the multilevel framework for analysing repeated measures data (where individuals can be measured on a different set of occasions from each other). Importantly, the multilevel approach

⁴ The most extended treatment of this model is given by Singer, JD and Willett, JB (2003). *Applied Longitudinal Data Analysis: Modeling Change and Event Occurrence*, Oxford University Press, New York; which is accompanied by a very useful website at <http://gseacademic.harvard.edu/alda/>

C4.1.5 Repeated measurements within individuals, panel data

does not require that each individual is measured on every occasion, that is the multilevel framework does not require balance. The spacing between occasions can also differ between individuals, so that, for example, some are measured annually but others are measured less frequently. The explanatory variables can be fixed across time, for example gender. Or they can be time varying as in the case of age. Clearly the data in Table 4.3 are from an annual survey.

Table 4.3. Data frame for a two-level repeated measures study

a) in long form

Classifications or levels		Response	Explanatory variables	
Occasion <i>i</i>	Person <i>j</i>	Height _{<i>ij</i>}	Age _{<i>ij</i>}	Gender _{<i>j</i>}
1	1	75	5	F
2	1	85	6	F
3	1	95	7	F
1	2	82	7	M
2	2	91	8	M
1	3	88	5	F
2	3	93	6	F
3	3	96	7	F

b) in wide form

Person	H-Occ1	H-Occ2	H-Occ3	Age-Occ1	Age-Occ2	Age-Occ3	Gender
1	75	85	95	5	6	7	F
2	82	91	*	7	8	*	M
3	88	93	96	5	6	7	F

In this example, we have height as the response, but in other examples the response might be cognitive, emotional or educational scores; indeed any construct where we might wish to investigate patterns of change over time, such as income or voting choice of individuals at different times.

An important advantage of the multilevel approach is that incompleteness of the data on the dependent variable does not complicate the analysis, provided that the data can be assumed *missing at random* (MAR). That is, the modelling can proceed without explicitly and additionally modelling the dropout mechanism. The missing observations can simply be omitted from the data. In going from the wide to the long form we can omit the rows with a missing response as we have

done in moving from *Table 4.3 b)* to *Table 4.3 a)*. The dropout mechanism is MAR when the distribution of the response for cases where it is missing is identical to the distribution of the response for cases where it is observed, after taking account of predictor variables. For example, if older people are more likely to refuse to respond to an income question, then age needs to be included as a predictor in the model.⁵

C4.1.6 Multivariate responses within individuals

Sometimes we may wish to model more than one response. For example, we may wish to consider jointly English and mathematics exam scores for students because the two responses are likely to be related. We can regard this as a multilevel structure with subjects (English and maths) nested within students as shown in *Figure 4-5*.




Figure 4-5. Classification and unit diagrams for a multivariate response model

The associated data frame in *Table 4.4 a)* shows the multivariate structure in wide form. There is a single row for each individual and the response and predictor variables form columns. There are two columns for the response, an English exam score and a maths exam score, and a single column for the predictor at the student level, their gender.

Table 4.4 b) gives the multilevel equivalent which requires one row per level 1 unit, so that the responses are interleaved in a long column. Here the level 1 units are now examinations taken (English and maths) where the *i* subscript indexes the relevant response variable. The subscript *j* indexes the individual student. We would also anticipate that the responses are to some extent correlated; if they were not, we would fit a separate model for each response. As is clear from *Table 4.4 b)*, the multivariate multilevel structure requires that the explanatory variables take on a particular form. First there must be a set of 0/1 variables that

⁵ Further discussion of missing data mechanisms, including MAR, can be found in Little, R.J.A. & Rubin, D.B. (1987). *Statistical analysis with missing data*. New York: Wiley. The benefits of the multilevel approach are explicitly considered in Maas, CJM and Snijders, TAB (2003) The multilevel approach to repeated measures for complete and incomplete data *Quality & Quantity* 37: 71-89, 2003. The latter, if available to you, is at: <http://www.springerlink.com/content/gn565733566tn062/fulltext.pdf>

distinguish each outcome variable, these are called *dummy* or *indicator variables* (e.g. Eng-Indic and Math-Indic); then, for each student-level predictor (e.g. Gender), there must be a variable associated with each response that takes on the value of that predictor for rows which have that response and 0 for rows which have a different response (e.g. Gender-Eng and Gender-Math).⁶ It is perfectly acceptable for some students to have missing data, for example student 2 has no English score and student 4 has no maths score, that is, the data can again be unbalanced.⁷

Indeed, such an approach is essential where the data are missing by design as in a matrix-sampling design. In such a design, all respondents are asked a set of core questions but additional questions are asked of random subsets of the total sample. For example, all students could be asked a set of core mathematics questions but, because of time constraints, subsets of students are asked detailed questions on trigonometry, matrix algebra, or calculus. In the multilevel design all the responses could be modelled simultaneously. The required MAR assumption is reasonable because of the random way in which questions are allocated to students.⁸

Table 4.4. Data frame for a multivariate response example

a) in wide form

Student	English score	Maths score	Gender
1	95	75	M
2	55	*	F
3	65	40	F
4	*	75	M

⁶ The predictor variables are formed from an interaction of the dummy indicator variable and the long-form explanatory variable; on interactions see Module 3.

⁷ The assumption is again that the data are missing at random.

⁸ For an application of this model see Yang, M., Goldstein, H., Browne, W. J. and Woodhouse, G. (2002). Multivariate multilevel analysis of examination results. *Journal of the Royal Statistical Society, A*, 165: 137-153.


b) in long form

Classifications or levels		or	Response	Explanatory variables			
Exam subject i	Student j		Exam score _{ij}	Eng-Indic _{ij}	Math-Indic _{ij}	Gender-Eng _j	Gender-Math _j
Eng	1	1	95	1	0	M	0
Math	2	1	75	0	1	0	M
Eng	1	2	55	1	0	F	0
Eng	1	3	65	1	0	F	0
Math	2	3	40	0	1	0	M
Math	2	4	75	0	1	0	M

C4.1.7 Two-stage sample survey design

In many large scale face-to-face surveys, the data collectors may adopt a two-stage design to minimise interviewer costs. In a study of voting, for example, the researchers may first select constituencies, the so-called *primary sampling unit* (PSU), and then select individuals within these areas. This selection procedure leads to a sample that is geographically clustered (*Figure 4-6*). There would of course normally be many more than the four clusters shown here.⁹ Table 4.5 shows the data frame for a two-level study of voter participation. Respondents are at level 1 and constituencies are at level 2 (*Figure 4-7*). The response is a binary indicator of whether or not an individual voted in the last general election. The predictors at the individual level are age and gender, while at the constituency level there is a single categorical predictor indicating whether the seat is marginal as opposed to safe.

a) Simple random sample



b) Two stage sample




Figure 4-6. A schematic map of respondents obtained from a simple random sample contrasted with that from a two-stage sample design

⁹ For an online guide to survey design and survey analysis see the Practical Exemplars and Survey Analysis site at <http://www.napier.ac.uk/depts/fhls/peas/>




Figure 4-7. A classification diagram for a two-stage sample design to assess the effects of constituency characteristics on individual voter turnout

Table 4.5. Data frame representation for a two-stage design for studying voter participation with categorical response and explanatory variables included at each level

Classifications or levels		Response	Explanatory variables		
Respondent i	Constituency j	Respondent voted _{ij}	Voter age _{ij}	Voter gender _{ij}	Constituency type _j
1	1	Yes	56	M	Safe seat
2	1	No	45	M	Safe seat
3	1	Yes	72	F	Safe seat
1	2	No	49	F	Marginal
2	2	No	36	M	Marginal
3	2	Yes	56	M	Marginal
1	3	Yes	76	F	Safe seat
2	3	No	50	F	Safe seat
1	4	No	39	M	Marginal
2	4	Yes	71	M	Marginal
3	4	Yes	41	M	Marginal
4	4	Yes	55	F	Marginal

Traditionally multistage designs are used to keep costs down as interviewers could be based in a particular PSU and not require extensive travel time. However, such a design was often seen as nuisance in the analysis.¹⁰ The problem is that

¹⁰ Jones, K (1997) Multilevel approaches to modelling contextuality: from nuisance to substance in the analysis of voting behaviour. In: G.P. Westert & R.N. Verhoeff (eds). *Places and people: multilevel modelling in geographical research*. The Royal Dutch Geographical Society Utrecht.

multistage designs (usually) generate dependent data so that respondents living within the same PSU can be expected to be more alike than respondents selected at random from different PSUs. This is not just the case for variables such as voting intention where it might seem obvious that there will be similarity between people living in the same area- it is an issue for anything we wish to measure. This so-called *design effect* would result in incorrect estimates of precision (standard errors being too low) and an increased risk of Type 1 errors, that is, finding a 'significant' relationship where none actually exists. These problems result when any dependent data are modelled as if they are independent. Multilevel models explicitly model this dependency and automatically correct for the design effect. There are in fact two broad ways in which the design effect can be taken into account. In one approach the dependency between outcomes within higher-level units is taken into account in the estimation of regression coefficients and their standard errors, but the source and nature of this dependency is not directly investigated and is regarded as a nuisance to be corrected for.¹¹ In contrast, in the multilevel approach, the nature of the dependency is of direct interest and we can investigate how this dependency changes in relation to explanatory variables at each level. When the response is a continuous one, there is no difference between the estimated coefficients and their standard errors, but the multilevel model provides greater flexibility. When there is a categorical response the estimates can be quite different and this depends on the size of the between-group variance.

The multistage design is the preferred design for a multilevel study of place or neighbourhood effects to ensure that more than one respondent is sampled from each local place, thereby allowing the estimation of the between-people and between-place effects. For example, the Millennium Cohort study in the UK adopted a multistage design in order to study the effects of local communities on children and their development.¹² Earlier studies of the 1946, 1958 and 1970 birth cohorts were based on children born in the UK in a particular week giving in effect a random sample across space. It is of course very important that the PSUs should be representative of particular types of places and, given the relatively small number of PSUs that are likely to be sampled, it is important to stratify on the process under study. For example, in studying voting behaviour in the UK, you may wish to stratify on whether the seats of the various parties were marginal or safe.¹³

¹¹ One well known software product that uses a marginal approach to multi-stage survey analysis is Sudaan to be found at <http://www.rti.org/sudaan>. Procedures for correcting standard errors for design effects are also available in Stata and SAS.

¹² For details on the design of the Millennium Cohort, see <http://www.cls.ioe.ac.uk/studies.asp?section=000100020001>

¹³ An excellent discussion of the issues involved in designing such studies is Stoker, L, and Bowers, J (2002) Designing multilevel studies: Sampling voters and electoral contexts. *Electoral Studies* 21:235-267.

C4.1.8 An experimental design in which the intervention is at the higher level

So far, we have considered hierarchical structures that arise in observational studies. In our final example, we will consider an experimental design in which the intervention is randomly allocated. The purpose of an intervention study is usually to evaluate the effectiveness of some policy instrument or treatment. The most common type of intervention is a randomised control trial in which the intervention is randomly allocated at the individual level, for example to compare a new drug treatment against the control (a placebo or the current drug in use). However, in randomised cluster trials (also known as randomised community trials), the intervention is at the higher level of a group or a place. In both designs randomisation is used to ensure that any possible confounders that may affect the result have been balanced in both the control and the intervention so that any effects of the treatment that are found are due to the intervention and not to any other (even unknown) factor. The community trial is growing in importance as evidence-based policies are developed and this approach is used in areas other than biomedical intervention¹⁴. Some examples include:

- A study of the effects of fluoridisation of the water supply on dental caries, where the domestic water supply is usually to an area;
- An evaluation of the introduction of a neighbourhood watch scheme to reduce both crime and the fear of crime;
- An assessment of a mass media campaign targeted at different areas.

In the above examples, the higher-level unit is an area. An example of an intervention implemented at the classroom level is:

- An assessment of the effect of introducing interactive whiteboard technology to a whole class on individual progress in learning.

Figure 4-8 gives the classification diagram and *Table 4.6* gives the data frame for such a study. The current exam score for students is the response, and the predictors are previous score, gender and, at the group level, a binary variable indicating whether or not the class has received the whiteboard intervention. The multilevel approach is needed to analyse the data from such a design, otherwise the anticipated dependency within the class will lead to an overestimate of the statistical significance of the intervention. In substantive terms, we can estimate whether the classes receiving the intervention are more consistent in their progress than the control, whether the intervention has a differential effect for those with low prior ability as compared to those with high, and whether the gender gap in progress is increased or decreased by the intervention.

¹⁴ More details on such cluster trials are given on Martin Bland's website on Clustered study designs at <http://www-users.york.ac.uk/~mb55/clust/clustud.htm>; in the book length treatment: Donner A, Klar N (2000). *Design and Analysis of Cluster Randomization Trials in Health Research*. London: Arnold; and in the overview article: Donner, A and Klar, N (2004) Pitfalls of and Controversies in Cluster Randomization Trials *American Journal of Public Health*, 94(3), 416-422




Figure 4-8. A classification diagram of a two-level nested structure for a group based intervention

Don't forget to take the online quiz!

From within the LEMMA learning environment


- Go down to the section for **Module 4: Multilevel Structures & Classifications**
- Click "[4.1 Two-level hierarchical structures](#)" to open Lesson 4.1
- Click  to open the first question

Table 4.6. Data frame representation for an intervention study at the group level to assess the effect of the introduction of interactive whiteboard technology on student progress

Classifications or levels		Response	Explanatory variables		
<i>Student i</i>	<i>Class j</i>	<i>Student exam score_{ij}</i>	<i>Student previous examination score_{ij}</i>	<i>Student gender_{ij}</i>	<i>Classroom intervention_j</i>
1	1	75	56	M	Whiteboard
2	1	71	45	M	Whiteboard
3	1	91	72	F	Whiteboard
1	2	68	49	F	Control
2	2	37	36	M	Control
3	2	67	56	M	Control
1	3	82	76	F	Whiteboard
2	3	85	50	F	Whiteboard
1	4	54	39	M	Control
2	4	91	71	M	Control
3	4	43	41	M	Control
4	4	66	55	F	Control

C4.2 Three-level structures

In this section we consider a range of three-level strict hierarchical structures in terms of unit and classification diagrams and associated data frames. Again the emphasis is on typical research problems that give rise to such structures.

C4.2.1 Students within classes within schools

As a first example we will consider an analysis of students nested within classes in schools. The structure is a strict hierarchy if each student belongs to one and only one classroom and each classroom group is found in one and only one school, as is often the case. The unit and classification diagrams are given in *Figure 4-9*. The data frame of

Table 4.7 shows the structure that could be used to assess the effect of the teacher’s classroom style (formal or informal) on progress and to assess whether this progress is different in state and private schools. In this particular study each class is taught by a single teacher (as is often the case in a primary school) so it is impossible to separate out between-teacher differences and between-class differences; they are confounded.




Figure 4-9. Unit diagram and classification diagram for the three-level structure of students within classes within schools

C4.2.1 Student within Classes within Schools

Table 4.7. Data frame representation of Figure 4.9, with response and explanatory variables added, three-level model of students within classes within schools

Classifications or levels			Response	Explanatory variables			
Student i	Class j	School k	Current exam score $_{ijk}$	Student previous examination score $_{ijk}$	Student gender $_{ijk}$	Class teaching style $_{jk}$	School type $_k$
1	1	1	75	56	M	Formal	State
2	1	1	71	45	M	Formal	State
3	1	1	91	72	F	Formal	State
1	2	1	68	49	F	Informal	State
2	2	1	37	36	M	Informal	State
1	1	2	67	56	M	Formal	Private
2	1	2	82	76	F	Formal	Private
3	1	2	85	50	F	Formal	Private
1	1	3	54	39	M	Informal	State
2	1	3	91	71	M	Informal	State
3	1	3	43	41	M	Informal	State
4	1	3	66	55	F	Informal	State

In a three-level study there can be imbalances at each of the higher levels, so there may be a different number of students in each class and a different number of classes in each school. There is also likely to be dependency within each of the higher levels, so that students in the same class are likely to be more alike than students selected at random from different classes, and classes within a school are also likely to be more alike than classes in different schools. In a major study, Bennett (1976) uses a single-level model to assess whether “teaching styles” affected test scores for English, reading and mathematics at age 11. He found progress was significantly influenced by teaching style, resulting in a call for a return to ‘traditional’ or formal methods. However, this study did not take account of dependency in the scores of students from the same classes. In a multilevel analysis, it was subsequently found that the effects were not significant¹⁵.

¹⁵ The original research is Bennett, N (1976) *Teaching styles and student progress* Open Books, London. The multilevel analysis is given by Aitkin, M., Anderson, D.A. and Hinde, J.P. (1981) Statistical modelling of data on teaching styles (with Discussion). *J. Roy. Statist. Soc. A* 144, 419-461. <http://www.jstor.org/view/00359238/di993039/99p0188p/0>. A more recent re-analysis is given by Spencer, N.H. (2002) Combining modelling strategies to analyse teaching styles data. *Quality and Quantity*, 36, 2, pp 113-127.

C4.2.1 Student within Classes within Schools

Other three-level nested structures include

- Repeated measures within students within schools. This allows us to look at how learning trajectories vary across students and schools.
- Multivariate responses on four health behaviours (drinking, smoking, exercise and diet) on individuals within communities. Such a design allows the assessment of the correlation between these behaviours at the individual level and at the community level, and to do so taking account of other characteristics at both the individual and community level. We can also assess the extent to which there are unhealthy communities as well as unhealthy individuals.
- A study to investigate individual-level determinants of self-rated health and happiness, as well as the extent of community level co-variation in health and happiness. The data have a three-level multivariate structure with the measures of self-rated poor health and unhappiness at level 1, nested within people at level 2, nested within communities at level 3¹⁶.

Another three-level design is worth examining in a little more detail as it shows the usefulness of the multilevel approach, namely a repeated cross-sectional design with students nested within cohorts within schools.

C4.2.2 A repeated cross-sectional design: students within cohorts within schools

Figure 4-10 shows the unit and classification diagrams where we have exam scores for groups of students who entered school in 1990 and a further group who entered in 1991. The model can be extended to handle an arbitrary number of cohorts. Table 4.8 shows the structure of the data frame, with the addition of a variable indicating whether in a given year the school adopted a selective policy.




Figure 4-10. Classification and unit diagrams for students within cohorts within schools

¹⁶ Subramanian,SV, Kim, D and Kawachi, I (2005) Covariation in the socioeconomic determinants of self rated health and happiness: a multivariate multilevel analysis of individuals and communities in the USA. *Journal of Epidemiology and Community Health*; 59:664-669; doi:10.1136/jech.2004.025742

C4.2.2 A repeated cross-sectional design: students within cohorts within schools

Table 4.8. Data frame for the structure shown in Figure 4.10 in which the selection policy of a school changes over time


Classifications or levels			Response		Explanatory variables		
Student <i>l</i>	Cohort <i>j</i>	School <i>k</i>	Exam score aged 15 _{ijk}	Exam score aged 11 _{ijk}	Student gender _{ijk}	Selection _{jk}	School type _k
1	1990	1	75	46	M	No	State
2	1990	1	71	32	M	No	State
1	1991	1	91	57	F	Yes	State
2	1991	1	68	61	F	Yes	State
1	1990	2	80	39	F	Yes	Private
2	1990	2	55	44	F	Yes	Private
1	1991	2	65	56	M	Yes	Private
2	1991	2	89	31	M	Yes	Private
1	1990	3	76	46	F	No	State
2	1990	3	54	50	F	No	State
1	1991	2	62	43	M	No	State
2	1991	2	78	60	M	No	State

Again the data need not be balanced: a school does not have to have children from both cohorts, within-school cohort groups can be different sizes, and the total number of students per school can differ from school to school. An example of exactly this structure is Nuttall *et al's* (1989) analysis of changing school performance in London.¹⁷ They were able to assess changing school performance as different cohorts of students passed through the schools. This type of design is often called a *repeated cross-sectional design* and could be used, for example, to examine community change in smoking behaviour. Level 3 would be community, level 2 the year in which the survey is carried out (the cross-section), and level 1 would be individuals. This is unlike a panel study, which we looked at in C 4.1.5, in which occasion is at the lowest level nested within individuals. Consequently the repeated cross-sectional design does not permit the study of individual change, but on the other hand it is useful because we don't need to have sampled any of the same individuals for each occasion (which may make it cheaper).

¹⁷ Nuttall, D.L., Goldstein, H., Prosser, R. and Rasbash, J. (1989). Differential school effectiveness. *International Journal of Educational Research* 13: 769-776. See also Gray, J, Goldstein, H & Thomas, S (2003) Of Trends and Trajectories: searching for patterns in school improvement *British Educational Research Journal*. Vol 29(1):83-88.

Don't forget to take the online quiz!

From within the LEMMA learning environment

- Go down to the section for **Module 4: Multilevel Structures & Classifications**
- Click "[4.2 Three-level structures](#)" to open Lesson 4.2
- Click  to open the question

C4.3 Four-level structures

By now you should be getting a feel for how basic random classifications such as people, time, multivariate responses and institutions can be combined within a multilevel framework to model a wide variety of nested population structures. Here are some examples of four-level nested structures:

- Student within classes within schools within local education authorities;
- Multivariate responses within repeated measures within students within schools¹⁸;
- Repeated measures within patients within doctors within hospitals;
- People within households within postcode sectors within regions.

As a final example of a strict hierarchy we will consider a doubly nested repeated measures structure.

C4.3.1 Doubly nested repeated measures

Suppose we have repeated measures within students within cohorts within schools (*Figure 4-11*). Cohorts are now repeated measures on schools and tell us about stability of school effects over time. Measurement occasions are repeated measures on students and can tell us about students' learning trajectories. The measurement occasions (in this example) refer to the age of the student rather than the time, so that measurement occasion 1 for all students is when they are in the first year of school, for example, rather than measurement occasion 1 being 1991, say, for all students. Note that there is no requirement for this to be the case in general: in a different example, measurement occasions could be calendar years, for example, with different individuals having different ages on the same measurement occasion. The data frame for this structure is shown in *Table 4.9*.

¹⁸ Repeated measures within multivariate responses within students within schools would also be a valid structure, although the technical details of setting up models for this and the multivariate responses within repeated measures within students within schools structure would differ in ways that we will not go into here

C4.3.1 Doubly nested repeated measures

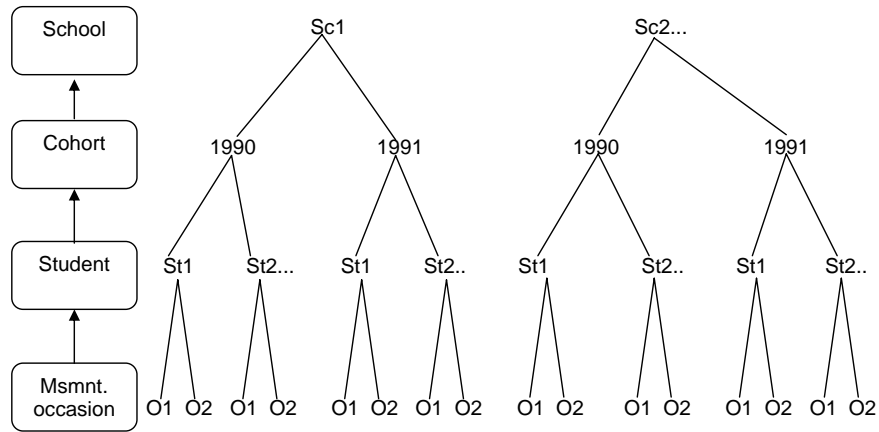


Figure 4-11. Unit and classification diagrams for a doubly nested repeated measures structure

Table 4.9. Data frame for doubly nested repeated measures structure

Classifications or levels				Response	Explanatory variables	
Occasion <i>i</i>	Student <i>j</i>	Cohort <i>k</i>	School <i>l</i>	Test score _{ijkl}	Gender _{ijkl}	School Type _l
1	1	1990	1	66	M	State
2	1	1990	1	76	M	State
1	2	1990	1	55	F	State
2	2	1990	1	86	F	State
1	1	1991	1	91	M	State
2	1	1991	1	93	M	State
1	2	1991	1	50	M	State
2	2	1991	1	70	M	State
1	1	1990	2	72	F	Private
2	1	1990	2	68	F	Private
1	2	1990	2	81	M	Private
2	2	1990	2	84	M	Private
1	1	1991	2	55	F	Private
2	1	1991	2	66	F	Private
1	2	1991	2	82	F	Private
2	2	1991	2	84	F	Private

C4.3.1 Doubly nested repeated measures

To clarify this doubly nested repeated measures structure, consider extending the above to four cohorts starting secondary school in 1990, 1991, 1992, and 1993 where children within cohorts are measured every year between the ages of 7 and 11. The non-empty cells in the body of Table 4.10 show the ages of the student from a particular cohort at a particular time.

Table 4.10. Student ages by cohort and time

Cohort	Time							
	1990	1991	1992	1993	1994	1995	1996	1997
1990	7	8	9	10	11	-	-	-
1991	-	7	8	9	10	11	-	-
1992	-	-	7	8	9	10	11	-
1993	-	-	-	7	8	9	10	11

C4.4 Non-hierarchical structures

So far all our examples have been of an exact nesting of a lower level unit in one and only one higher-level unit. That is, we have been dealing with strict hierarchies. But social reality can be more complicated than that! In fact it turns out that we need two non-hierarchical structures which, in combination with strict hierarchies, will be able to deal with all the different types of designs, realities and research questions that we will meet:

- Cross-classified structures;
- Multiple membership structures.

The existence of these non-hierarchical structures accounts for our preference for the term ‘multilevel models’ in preference to the more limited ‘hierarchical models’.

C4.4.1 Cross-classifications: students cross-classified by school and neighbourhood

Consider the unit diagram in *Figure 4-12* which links students to their area of residence and to the school they attend. You can see that there is a non-nested structure, for example:

- Students 1 and 2 attend School 1 but come from different areas;
- Students 6 and 10 come from the same area but attend different Schools.

Area is not nested within school and school is not nested within area. Students lie within a cross-classification of school by area. Students are nested within schools and students are nested within areas. However, schools and areas are cross-classified, so the nodes for school and area are not connected on the classification diagram of

Figure 4-13. The data frame is shown in *Table 4.11* with predictors of achievement (Test score $i_{(jk)}$) at each of the three levels or classifications: student gender, school type, and area deprivation score (the index of multiple deprivation, IMD). Notice that in terms of subscripts we use brackets to signify that areas and schools are cross-classified. We also now index the students by a unique number in sequence (not Student 1, 2 3 in School 1) to emphasize that they are their own classification.

Another way of looking at this structure is shown in *Table 4.12* which is a cross-tabulation of students by areas and schools. If schools were nested within areas, all the students in a row of *Table 4.12* would lie in a single column. If areas were nested within schools, all the students in a column of *Table 4.12* would lie in a single row. They do not follow this pattern; therefore the data are cross-classified.

C4.4.1 Cross-classifications: students cross-classified by school and neighborhood




Figure 4-12. Unit diagram for a two-way cross-classified structure




Figure 4-13. Classification diagram for students cross-classified by school and area

Table 4.11. Data frame for a cross-classified structure

Classifications or levels			Response	Explanatory variables		
Student i	School j	Area k	Exam score $e_{i(jk)}$	Student gender $i_{(jk)}$	Area IMD $_k$	School type j
1	1	1	75	M	24	State
2	1	2	71	F	46	State
3	1	1	91	F	24	State
4	2	2	68	M	46	Private
5	2	1	37	M	24	Private
6	3	2	67	F	46	Private
7	3	2	82	F	46	State
8	3	3	85	M	11	State
9	4	3	54	M	11	Private
10	4	2	91	M	46	Private
11	4	3	43	F	11	Private
12	4	3	66	M	11	Private

C4.4.1 Cross-classifications: students cross-classified by school and neighborhood

Table 4.12. Tabulation of students by school and area to reveal a cross-classified structure

	Area 1	Area 2	Area 3
School 1	St1, St3	St2	
School 2	St5	St4	
School 3		St6, St7	St8
School 4		St10	St9, St11, St12

C4.4.2 Repeated measures within a cross classification of patients by clinician.

As a final example of a cross-classification, we consider a health setting with repeated measures on patients but patients being assessed by different clinicians at different times. This is depicted in Figure 4-14 where, for example, patient 1 is seen by clinician 1 at occasions 1 and 2 then by clinician 2 at occasion 3. The data frame for this structure is shown in Table 4.13.




Figure 4-14. Classification and unit diagrams for cross-classification of measurement occasion by patient and clinician

C4.4.2 Repeated measures within a cross-classification of patients by clinician

Table 4.13. Data frame for cross-classification of measurement occasion by clinician and patient

Classifications or levels			Response	Explanatory variables	
Occasion <i>i</i>	Patient <i>j</i>	Clinician <i>k</i>	Cholesterol level (<i>ijk</i>)	Clinician type <i>k</i>	Patient gender <i>j</i>
1	1	1	3.5	Doctor	M
2	1	1	3.8	Doctor	M
3	1	2	3.4	Nurse	M
1	2	1	4.1	Doctor	F
2	2	3	4.8	Nurse	F
3	2	2	4.8	Nurse	F
1	3	2	6.9	Nurse	F
2	3	3	6.2	Nurse	F
3	3	4	6.5	Doctor	F
1	4	3	5.1	Nurse	M
2	4	4	5.3	Doctor	M
3	4	4	5.5	Doctor	M

Other examples of two-way cross-classifications include:

- Exam marks within a cross-classification of student and examiner, where a student's paper is marked by more than one examiner to get an indication of examiner reliability.
- Students within a cross-classification of primary school by secondary school. We may have students' exam scores at age 16 and wish to assess the relative effects of primary and secondary schools on attainment at age 16
- Patients within a cross-classification of general practitioner (GP) practice and hospital.

In many of these examples, individuals are seen as occupying more than one set of contexts, for example students may be influenced by residential setting and school setting.¹⁹

¹⁹ A large scale study is Simonite, V. and Browne, W.J. (2003) Estimation of a large cross-classified multilevel model to study academic achievement in a modular degree course. *Journal of the Royal Statistical Society, Series A*, 166(1) 119-134. A comprehensive review of educational research using these models is given by Fielding, A. and Goldstein, H. (2006) *Cross-classified and Multiple Membership Structures in Multilevel Models: An Introduction and Review*, DFES report no. 791, http://www.economics.bham.ac.uk/people/fielding/Cross_classified_review_RR791.pdf

C4.4.3 Multiple Membership Structures

We now come to our final structure, a multiple membership structure, in which the atomic units are seen as nested within more than one unit from a higher-level classification. A simple example (Figure 4-15) is of students nested within primary school teachers. Some students are taught by more than one teacher and it is possible to define a 'weight' based on the proportion of time the student spent with the teacher; these will of course sum to 1. Notice that in the classification diagram, double arrows are used to signify multiple membership. As the data frame in Table 4.14 shows, while student 2 is taught only by Ms Mayer, student 1 is taught 20 percent of the time by Mr Edgar and 80 percent of the time by Ms Mayer. We have to create two teacher columns (and two associated weight columns) to reflect that, in this study, the maximum number of teachers that can teach any single student is two. If a student is taught by six teachers, we would need six identifying columns and associated weights. Explanatory variables relating to teachers should then also be entered in multiple columns with the values of the explanatory variable multiplied by the teacher weights. This has been done for Teacher Style in this example, resulting in the two columns Teacher 1 style and Teacher 2 style. For example, student 3 is taught by Ms Mayer (as Teacher 1) and Mr Forbes (as Teacher 2), with weights 0.6 and 0.4 respectively. To obtain the value of Teacher 1 style for student 3, we note that Teacher 1 for student 3 is Ms Mayer. Ms Mayer has a value of Informal for Teacher Style. The weight for Teacher 1 for student 3 is 0.6, so Teacher 1 style for student 3 is 0.6×Informal. Now to obtain the value of Teacher 2 style for student 3, we note that Teacher 2 for student 3 is Mr Forbes. Mr Forbes has a value of Formal for Teacher Style. The weight for Teacher 2 for student 3 is 0.4, so Teacher 2 style for student 3 is 0.4×Formal. Note that this only applies to explanatory variables relating to teachers: we do not do it for explanatory variables relating to students such as Student gender and Student previous exam score in this example. Hill and Goldstein (1997) and Browne et al (2001) provide more details on the formulation and development of this type of structure.²⁰ It is also worth noting that multiple membership models can handle the situation where there are missing identifiers, provided it is possible to specify the weights as the anticipated probability of membership of a lower-level unit to a higher-level unit. This could be used when you do not know exactly which school a student attends but you could specify it as being a possible member of three schools with a 60 percent chance of going to School A (the nearest one), and 20 percent for Schools B and C (two equally distant ones).

²⁰ Hill, P. W. and Goldstein, H. (1998). Multilevel modelling of educational data with cross classification and missing identification of units. *Journal of Educational and Behavioural statistics* 23: 117-128.
 Browne, W., Goldstein, H. and Rasbash, J. (2001). Multiple membership multiple classification (MMMC) models. *Statistical Modelling* 1: 103-124.

Other examples of multiple membership relationships are:

- Health outcomes where patients are treated by a number of nurses, and patients are multiple members of nurses;
- People are multiple members of households (for example in the case of children with separated parents).²¹




Figure 4-15. Unit and classification diagrams for a two-level multiple membership structure with weights; students nested within teacher

Table 4.14. Data table for two-level multiple membership structure with weights

Classifications or levels		Weights		Response	Explanatory variables				
Student <i>i</i>	Teacher 1	Teacher 2	Wt1	Wt2	Student exam score _{ij}	Student previous exam score _{ij}	Student gender _{ij}	Teacher 1 style	Teacher 2 style
1	Mr Edgar	Ms Mayer	0.2	0.8	75	56	M	0.2*Formal	0.8*Informal
2	Ms Mayer		1	0	71	45	M	1*Informal	
3	Ms Mayer	Mr Forbes	0.6	0.4	91	72	F	0.6*Informal	0.4*Formal
4	Mr Forbes		1	0	68	49	F	1*Formal	
5	Mr Forbes		1	0	37	36	M	1*Formal	

²¹ Goldstein, H., Rasbash, J., Browne, W. J., Woodhouse, G. and Poulain, M. (2000). Multilevel Models in the Study of Dynamic Household Structures. *European Journal of Population*, 16.

C4.5 Combining structures: hierarchies, cross-classifications and multiple membership relationships

It is possible to combine the three types of structure (strict hierarchy, cross-classification and multiple membership) to reflect social reality and the research designs employed to study it. Consider the structure depicted in Figure 4.16. Students are nested within a cross-classification of school by area. However

- Student 1 moves in the course of the study from residential area 1 to 2 and from school 1 to 2;
- Student 8 has moved schools but still lives in the same area;
- Student 7 has moved areas but still attends the same school.

Now in addition to schools being crossed with residential areas, students are *multiple members* of both areas and schools. In the classification diagram, students are connected to schools and areas by double arrows to represent two sets of multiple membership relationship. School and Area are not connected by arrows indicating that the lower level units, students, lie within a cross-classification of school by area.

Table 4.15 gives the associated data frame.

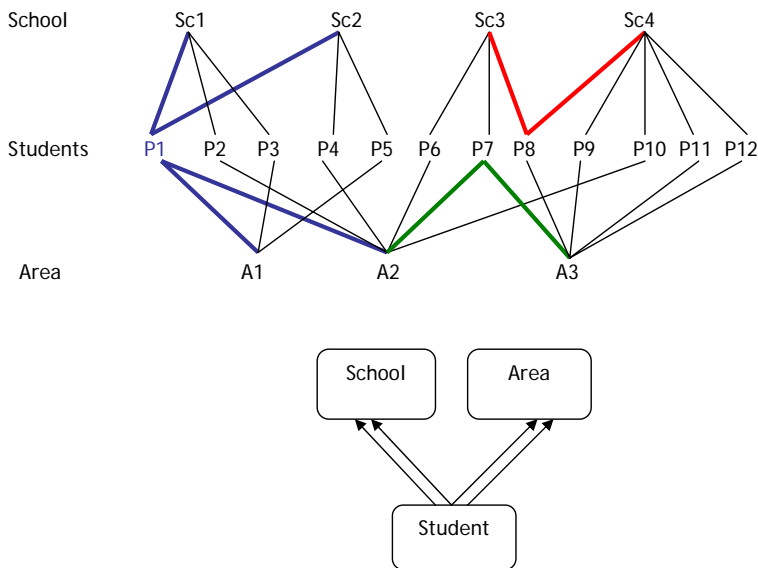


Figure 4-16. Unit and classification diagrams for cross classifications and multiple memberships: students, areas and schools

C4.5 Combining Structures: hierarchies, cross-classifications & multiple membership relationships

Table 4.15. Data table for multiple membership model

Classifications or levels			Response	Explanatory variable
Student	School	Area	Exam score	Student gender
1	1,2	1,2	75	M
2	1	2	71	M
3	1	1	91	F
4	2	2	68	F
5	2	1	37	M
6	3	2	67	M
7	3	2,3	82	F
8	3,4	3	85	F
9	4	3	54	M
10	4	2	91	M
11	4	3	43	M
12	4	3	66	F

Another example of a non-hierarchical structure relating patients to nurses, hospitals and GPs is shown in Figure 4.17:




Figure 4-17. Unit diagram showing a complex mixture of multiple membership, hierarchical and crossed relationships

In summary:

- Patients are multiple members of nurses;
- Nurses are nested within hospitals;

C4.5 Combining Structures: hierarchies, cross-classifications & multiple membership relationships

- Patients are nested within GPs;
- Nurses are crossed with GPs;
- Hospitals are crossed with GPs.

This is all quite a mouthful and is perhaps more succinctly expressed in the classification diagram in *Figure 4-18*.

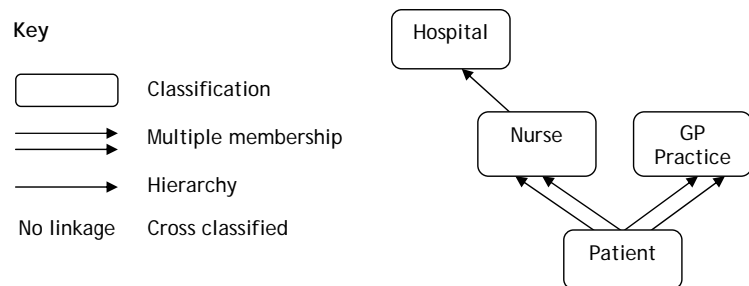


Figure 4-18. Classification diagram showing complex mixture of multiple membership, hierarchical and crossed relationships

C4.6 Spatial structures

Figure 4- shows a region carved up into thirteen districts labelled A to M. We are interested in whether an individual nested within a district is going to be affected by a contagious disease. We have predictors for the individual, say whether or not they are of school age, and we have a predictor for the district they live in, e.g. the number of cases last week in that area. But a contagious disease spreads as a spatial process and an outbreak may spill over from one community to another. Therefore we should include the structure of the neighbouring districts in the modelling framework. We can do so as a combination of a strict hierarchy and multiple membership (

Figure 4-19). Thus for example:

- An individual in region A is affected by region A as a strict hierarchical relationship, but also by a multiple membership relation to sub-regions B,C, and D;
- An individual in region H is affected by region H as a strict hierarchical relationship, but also by a multiple membership relation to sub-regions E, G, I and K.

Here, the multiple membership is defined by whether the districts have a common boundary. We could also include weights such as the inverse of a function of the distance between the centroids of the districts. Another idea is to examine the nature of the spatial process by estimating models with different weight structures. For example, we could define the weights as just suggested to specify distance-based diffusion so that nearby places are given a higher weight. But we could alternatively test for hierarchical diffusion in which the spread of the disease does not depend on distance but on the size of the population of the area, with disease visiting places in approximate order of size. Such spatial models can be more generally applicable with for example a set of weights being used to define which schools are in competition for attracting students.²²

²² For examples of spatial models as multilevel structures see Lawson, A.B., Browne W.J., and Vidal Rodeiro, C.L. (2003). *Disease Mapping using WinBUGS and MLwiN* Wiley. London

C4.6 Spatial Structures

A	B		
	C		
D	E	F	
G	H	I	J
	K	L	M

Figure 4-19. Schematic structure of a region for a study of contagious diffusion of disease





Figure 4-19. Spatial structure a combination of a strict hierarchy and a multiple membership

Don't forget to take the online quiz!

From within the LEMMA learning environment

- Go down to the section for **Module 4: Multilevel Structures & Classifications**
- Click "[4.6 Spatial structures](#)" to open Lesson 4.6
- Click  to open the first question

C4.7 Summary

The major goal of this module has been to present tools for thinking about complex social structures, and we hope you will find the unit and classification diagrams and the accompanying data frames useful in structuring your own research problems. We end with a final classification diagram (*Figure 4-20*) of a large-scale and complex research problem. The aim is to examine student progress using the Avon Longitudinal Study of Parents and their Children (ALSPAC) which has followed longitudinally all children born in Avon (the greater Bristol area) in 1990. It can be seen that:

- The atomistic unit is the measurement occasion, with students measured on multiple occasions. Thus, there is a strict hierarchy of occasion nested within student;
- Students span three school-year cohorts (1995, 1996, and 1997) as some go early to primary school, some go late, but the majority go when they are aged 6;
- Students can move between teachers, schools, and neighbourhoods so that there are sets of multiple memberships relation involved.

Students' progress is potentially affected by their own changing characteristics, the students around them, their current and past teachers, schools and neighbourhoods, and we wish to estimate the contributions of each from the data.

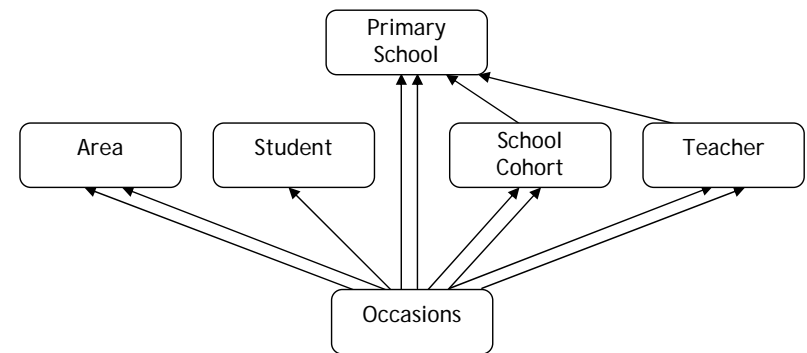


Figure 4-20. The ALSPAC study as a classification diagram

You may well be asking whether all this complexity is needed. The answer is that it might be and you will only be able to tell by estimating such complex models. Complex models are not always reducible to simpler models, and there may be

confounding of variation across levels which if not taken into account will give incorrect estimates. What may appear to be school differences may in fact be between-teacher effects. In Module 5 of the training materials we will be dealing only with two-level nested models. A thorough grounding in the theory and practice of these two-level multilevel models then provides a firm base from which to progress to some of the more complex multilevel structures described in this module. Some of these more complex structures will be dealt with in future training materials.

In the modules to come we will show how to construct and run multilevel models. We will demonstrate the benefits to be gained from multilevel modelling, in terms of a wider range of questions that can be addressed, and we will explain the problems that can be encountered if multilevel modelling is not used for data with structures of the kinds presented in this module.

Don't forget to take the online quizzes for this module!
(See page 2 for details of how to find the quizzes)

What should I have learnt?

- The real world has a complex structure and/or we impose one through our research design.
- Complex structure tends to create dependencies between observations.
- Most (all?) social science research problems and designs are a combination of strict hierarchies, cross-classifications and multiple memberships.
- Multilevel modelling deals with complex structure deriving from reality and the study design; it explicitly models dependencies and copes with imbalanced data structures.

References

Snijders, T.A.B., and Bosker, R.J. (1993). Standard errors and sample sizes for two-level research. *J. Educational Statist.*, 18, 237-259.