10 Attainment in Reading and Maths: Variance Components Analysis

This chapter and the next one present the results of a series of multivariate analyses of the data for the first two years of the study. This analysis has two general objectives. First, multivariate models can help to show how certain outcomes of schooling are determined. We already know, for example, that the second-year scores are related to social class, to ethnic group and to the first-year scores, but we need to build a model that describes this whole pattern of inter-relationships, so as to show, for example, how far ethnic group is related to second-year scores after taking account of the effect of other variables. The second objective is to assess the nature and size of differences between schools in the outcomes they achieve, after taking account of differences in the inputs.

To achieve these objectives, we adopted the method of 'variance components analysis' developed by the Department of Applied Statistics at the University of Lancaster. Models that aim to explain second-year attainment in reading and maths are described in the present chapter. The next chapter considers models in which factors other than attainment (for example, participation in school activities) are taken as the outcomes of the educational process. A technical description of the form of analysis appears at Appendix A.

Introduction to variance components analysis

In any particular analysis, we start by choosing one variable as the criterion of success to be considered. This is called the 'dependent variable'. Many of the analyses take the second year reading or maths score as the dependent variable, but the next chapter describes further analyses that take enthusiasm for school or participation in school activities as the criterion of success. Next, a number of other variables are chosen (called 'independent variables') and a 'model' is constructed which exhibits the relationships between the various independent variables and the one dependent variable or criterion of success. This model takes account of the separate and joint effects of the independent variables.

The variables, such as reading score or social class, refer to individual children and not to schools. Where the 'outcome' or dependent variable in the analysis is a second year attainment score, the corresponding first year attainment score is always included as one of the independent variables. That means that the standard the children had reached at the time when they entered the schools is always taken into account. Strictly speaking, the analysis does not consider 'rate of progress'; instead, it tries to explain attainment at a given point after taking into account the level of attainment of the same child at an earlier point.

The ‘fixed part’ of the model describes the relationships without taking account of the way that children are grouped into schools. The ‘random part’ of the model describes the way in which this pattern of relationships is modified depending on which school a child belongs to. The ‘random part’ of the model always allows for the possibility that the level of the outcome variable (say the second-year reading score) may be higher or lower depending on the school, and there is a method of testing whether there are significant differences of this kind between schools. In addition, it may be possible to improve the model (that is, make a better prediction of what the outcome variable – say the second-year reading score – will be) by allowing the *nature* of the relationships to vary from one school to another. For example, West Indians might tend overall to achieve poorer second-year reading scores than white children (after taking account of the effect of other variables); but there might be some schools where they do better than whites. In that case, the nature of the relationship between ethnic group and reading score would vary between one school and another.

Numbers

There were about 3,100 children who attended the study schools for all or part of the two-year period under study. The number who completed the various tests varies from 2,331 (second year reading test) to 2,731 (first year maths test). Information on attendance is available for 2,644 children in the first year and 2,559 in the second. The second-year pupil questionnaire was completed by 2,526 children, and the parental survey covered 2,074 families. The Rutter B2 score is available for 1,763 children.

When each instrument is considered on its own, the number of children covered is satisfactory, and except in the case of the Rutter B2 score there is a reasonable coverage of every school. However, difficulties do arise when information from several different instruments is to be brought together in the same analysis. The set of children for whom information is available changes, depending on the instrument being considered, so when information from different sources is combined the number of complete records is considerably reduced.

To minimise this problem, we have included in each particular analysis the children (and families) for which all the data required in that analysis are present. This means that the set of children included varies somewhat from one analysis to another. There are strong indications that those excluded in each case are close to being a randomly selected set. For example, average scores do not vary significantly between the subsets included in different analyses. The number of children that can be included in any particular model is usually between 1,200 and 1,400, which is, of course, low in relation to the 2,500 or so who completed each attainment test. This acts as an important limitation on what can be done with the data.

The variables

All of the variables to be used in the variance components analysis have already been introduced, but it is worth briefly reviewing them.

Ethnic group

The classification by country of origin is essentially based on the country the family ‘came from originally’. Across all schools, 58 per cent of children come from families originating from the UK or Eire, 27 per cent from south Asian families (originating from the Indian sub-continent), seven per cent from West Indian families, four per cent from families of mixed origins, and five per cent from other countries. Even on the basis of a crude classification like the one above (which throws some highly disparate groups together) this is a very fragmented pattern. The only sizeable group, apart from UK/Eire children, is south Asians, and they are in fact many disparate groups. West Indians are an important group, but there are too few in our sample for them to be shown separately in a variance components analysis. We are obliged to use a crude three-way classification into 'UK/Eire', 'South Asians' and 'Others'. The 'others' group includes West Indians, along with children originating from many other countries, and those of mixed parentage. An alternative principle of ethnic classification is by religion, but country of origin has been preferred in these analyses.

Social class

Families are classified on the basis of the father's job, if there is a resident father who is working or has worked in the past five years, and if not, on the basis of the mother's job. The jobs are coded into five groups, which are aggregates of the Registrar General's 17 socio-economic groups. In addition, there is a sixth group of families where neither parent (or the one parent in single-parent families) had had a job in the previous five years. Earlier analyses have shown that the families where the parents had not worked are an 'underclass' group, fitting at the bottom of the scale below the families of unskilled manual workers. In the variance components analysis, social class is treated as a six-point scale (not as six separate categories).

Qualifications of parents

This reflects the highest educational or job qualification of either parent. There are five groups: none; job qualification short of an apprenticeship; school leaving qualification; apprenticeship; professional or post-school academic qualification. The variable is treated as a five-point scale.

Sex

Overall the study children are divided almost evenly between the sexes, but there are two girls' schools (15 and 41) and two boys' schools (12 and 42).

Attainment tests

For the most part, the tests used in the variance components analysis are those of maths and reading completed in September-November 1981 and in July 1983. Full details are given in Chapter 9. The maths and reading tests used in the first and second years were not the same nor strictly equivalent, but none of the analytic techniques used assumes that they were.

Attendance

There is a separate variable for number of half days off the register in each school year.

Rutter B2 score

The purpose of including this measure was to control for the possibility that some schools have a larger intake than others of children with disturbed behaviour. Although the data are incomplete, it is possible to examine the hypothesis that apparent school differences are explained by differences in the proportion of children showing behaviour disturbances.

Enthusiasm for school
As the first question in the second-year pupil questionnaire, children were asked: 'Imagine that you are lying in bed and you start to wake up, and you think to yourself "It's the first day of term, I'm going to school today". How do you feel?' They answered by choosing one of four points on a scale. This scale has been retained for the variance components analysis.

Participation in school activities
This is a score from 0 to 4 derived from four items in the pupil questionnaire covering playing in teams, taking part in plays, concerts and special evenings, going on school trips or visits, and doing something special in assembly or a house or year meeting.

Index of praise and index of blame
Each of these indices is based on four items from the pupil questionnaire, producing a range of scores from 0 to 4. The items in the index of praise are about whether the child has been told that he or she has done good work or done well in any way. Each of the items in the index of blame is the mirror image of an item in the index of praise; the items in the index of blame are about whether the child has been told that he or she has done poor work or been 'told off'.

Parents' contact with school
The number of visits that parents have made to school in the past year is the main variable used here.

Parents' views about the school and the child's progress
The four ratings discussed in Chapter 6 have been used in variance components analysis. These are: overall satisfaction with the school; satisfaction with standards of behaviour; how well the child is getting on with school subjects; how happy the child is at school.

Basic model
The models considered in this section take the second-year reading or maths score as the outcome, and include four other basic variables: sex, social class, country of origin, and the first-year test score in reading or maths. In interpreting the results, it should be remembered that detailed information about small or localised groups has to be sacrificed to allow the multivariate analysis to be carried out. For example, the ethnic classification is necessarily crude. Diverse groups originating from the Indian sub-continent have had to be lumped together as 'south Asians', and West Indians have been combined with members of other minority groups. Also, the ethnic groups are very unevenly distributed across the schools, so that it is difficult for any method of analysis to inter-relate school effects and ethnic differences in a wholly satisfactory way. The analysis is not a self-validating mechanism, but just a way of looking at the results that may be helpful if the problems and difficulties are always borne in mind.

Second-year reading score as the outcome
The model gives a method of predicting the second-year reading score of an individual child from the child's characteristics on four basic variables (sex, social class, country of origin, and the first-year reading score). The final model, which is found after examining

a series of possible models, is the one that is most successful in predicting the actual scores of the individual children (often called the model that provides the ‘best fit’). The analytic method provides a criterion of the fit of any model (this criterion is the deviance) and a test to show whether the fit has been significantly improved by some change to the model, for example by introducing another variable. Using this criterion, the ‘best fit’ model is found by a process of trial and error, and by making use of a prior knowledge of the relationships.

In every variance components model, there is a ‘fixed part’ and a ‘random part’. The fixed part describes the relationships without taking account of how the children are grouped into schools. The random part shows how the results predicted by the fixed part vary, depending on which school the child belongs to.

Information about the best fit model is summarised in Table 10.1. It is important to remember that the range of scores in the Edinburgh reading test that was used in the second year is 0 - 155. The table gives the ‘grand mean’ as 38.92. This is the score a child is predicted to achieve if he falls into all of the categories for which the estimated effect is zero. This would apply to a boy whose first-year reading score was zero, who belonged to the underclass group and originated from UK/Eire.[1] The column of ‘estimates’ shows what is the effect on this prediction of the second-year reading score if the child belongs to each particular category. The only complication here is that the variables may be ‘continuous’ (a numerical quantity like the reading scores) or they may be ‘categorical’ (the child is or is not a member of each of a number of categories like ethnic groups). In the case of continuous variables, the estimate is a factor to be multiplied by the child’s actual score, and the product is then to be added to the grand mean. In other words, the table shows that for each additional point scored on the first-year reading test, the model estimates that 5.316 additional points will be scored on the second-year reading test. (The first-year reading test has a range of scores from 0 to 15, while the second-year test has a range of 0 to 155; that is why one point on the first test counts for as many as five on the second.) In the case of ‘categorical’ variables, the estimates are just amounts to be added to the grand mean.

An example will illustrate how these estimates are to be used. Taking the case of a girl who scored 9 on the first-year reading test, and who comes from a white-collar family that originates from the Indian sub-continent, the estimate of her second-year reading score will be as follows.

	39.987	grand mean
	-2.087	female
9 x 5.316 =	47.844	first-year reading score was 9
4 x 1.530 =	6.12	from skilled manual family
	-4.595	south Asian
	87.269	estimated second-year reading score

The standard errors provide a basis for applying tests of statistical significance. Where the estimate is twice as great as the standard error, the effect is significant at the 95 per cent level of confidence. Where it is more than twice as great, the level of confidence is higher.

Table 10.1 therefore shows that social class, the first-year reading score and country of origin are significantly related to the second-year reading score. The (negative) estimate for females just fails to reach significance at the 95 per cent level; if this estimate were significant, it would indicate a slight tendency for girls to progress more slowly than boys in reading over the first two years. It is clear that the first-year reading score has a much greater effect than the other two variables (social class and country of origin) either severally or jointly. In the context of the model, there is a difference of 8 points in the second-year reading score between the top and bottom social class groups, and a difference of 4.5 points between south Asian children and those originating from the UK or Eire (the estimate for the 'other' ethnic group is not significant). By contrast the first-year reading score can make a difference of up to 80 points in the estimate of the second-year reading score (the maximum first-year reading score, which is 15, multiplied by 5.316).

Of course, social class and ethnic group are more closely related to reading scores than these findings seem to imply, but the point is that they are related both to the first-year *and* to the second-year reading scores. So after allowing for the (enormous) effect of the first-year reading score in determining the second-year score, there is not much further effect to be contributed by these other variables. There is another way of putting the same point. From a child's reading score on entering secondary school, a good prediction can be made of his or her score at the end of the second year. Further information about the child's social class and ethnic group will allow the prediction to be improved a bit (but not much). This is like saying that a child's social class and ethnic group will be just a bit of help in making a prediction about how far the child will *make progress* in reading (since 'progress' is like 'attainment now, after taking account of attainment at an earlier period'). If the first-year reading score were not included in the model, then social class and ethnic group would appear to be related to the second-year reading score much more strongly. Thus, social class and ethnic group are much more strongly related to *attainment at a given time* than they are to *progress in attainment*.

At the point of entry to secondary school, certain categories of south Asian children (those of Bangladeshi and Pakistani origin) scored substantially below average in reading. Children of West Indian origin also scored below average at the point of entry, but higher than the low-scoring south Asian groups. The findings from the basic multivariate model suggest that south Asians (who have had to be treated as a single group) progressed rather more slowly in reading over the first two years, from this lower starting point. If anything, therefore, the findings suggest that the gap between the south Asians and other children was getting wider. However, this finding must be treated with caution, since the low- and high-scoring south Asian groups have been lumped together in the multivariate analysis. The evidence suggests that West Indian children progressed at a similar rate to white children over the first two years, though again this result must be treated with caution, since West Indians have been lumped with various other minority groups.

So far, the discussion has been confined to the fixed part of the model, which provides a method for predicting the second-year reading score of an individual child, but without taking account of which school the child belongs to. The random part deals with the possibility that the outcome may vary between schools. The general approach is to show whether the predictions of the fixed part can be significantly improved by assuming that there are school differences. The simplest variance components model allows only for the possibility that the 'grand mean' may vary between schools: that is, it allows for the possibility that the second-year reading scores tend overall to be higher in some schools than in others, after taking account of the four independent variables (sex, first-year reading score, social class and ethnic group). More complex models allow for the possibility that the 'slopes' of the relationships between the four variables and the outcome may vary between schools: that is, they allow for the possibility that there is variation between schools in the performance relative to each other of different social class or ethnic groups or of children with high versus low first-year reading scores. If the grand mean varies between schools, this implies that some schools tend to achieve better results than others overall. If the slope, for example between the first-year and second-year reading scores, varies between schools, this implies that some schools achieve their best results with initially low-scoring children, and others with initially high-scoring children. If all of the slopes vary significantly between schools, then there is effectively a different model of the relationships for children in every school.

In the case under discussion, there is an improvement in the predictive power of the model (significant at a very high level of confidence) by placing the grand mean in the random part. This shows that there are clear differences between schools in the reading scores recorded at the end of the second year by children with the same scores and other characteristics at the beginning of the first year. There is a further improvement, significant at the 99 per cent level of confidence, when ethnic group is placed in the random part. This suggests that some schools achieve their best results with white British children, while others achieve their best results with south Asians or others. However, this latter result should be treated with caution. The (inevitable) crudity of the ethnic classification has been emphasised: it is quite possible that the results are influenced by the more particular ethnic groups that are present in individual schools, and there is not enough room in the sample to compare like with like (the south Asians in one school may be radically different in origin from those in another with which it is being compared). There is also an improvement in the predictive power of the model, significant at the 95 per cent level of confidence, when socio-economic group is placed in the random part. This result should, again, be treated with some caution, but it seems to imply that the tendency for middle-class children to progress faster than working-class children is stronger in some schools than in others. The best fit model, then, is one with both country of origin and socio-economic group in the random part. All of the estimates shown are derived from this model.

Table 10.2 illustrates the scores predicted by the fixed part of the model. It shows that a child's second-year reading score varies much more with respect to his first-year reading score than with respect to socio-economic group or country of origin.

In the best fit model, the proportion of the variance attributable to the school level depends on the socio-economic group and country of origin of the child under consideration (see Appendix 2). In the case of a child of UK origin, for example, the school level accounts for 14 per cent of the variance within the 'underclass' group, for 17 per cent within the professional and managerial group and for about 9 per cent within the skilled manual group. Among those belonging to the 'other' origin group, the proportion of the variance attributable to the school level is about the same as among the UK origin group; it is rather smaller among south Asians (see Table A1 in Appendix 2).

These findings show that shool differences are of great importance. This is illustrated by Table 10.3, which shows the second-year reading scores predicted by the model for a boy belonging to a skilled manual family originating from the UK or Eire, and who had a first-year reading score of 6 (which is close to the average). The model predicts that if he went to school 43, this boy would obtain a second-year reading score of 62, but if he went to school 33 he would obtain a score of 91.[2] This is a substantial difference in the light of the stability of individual performance over time (see the discussion in Chapter 3). A difference of 29 points on the Edinburgh Reading Test (which has a range of 155) may be small compared with the huge variations associated with the accidents of birth and upbringing, but it could well be enough to have a critical influence on a person's life chances. Also, differences of this kind would build up over successive periods (the present model is confined to a period of less than two years).

Table 10.3 illustrates the differences between the results produced by the schools for a child with a given set of characteristics. The findings imply, however, that the performance of the schools relative to each other varies depending on the kind of child that is being considered. There is, however, a set of schools that is found to perform better and one that is found to perform worse than average with most kinds of child. It is fairly clear, for example, that schools 33 and 25 achieve better than average results, while schools 35 and 43 achieve worse than average results.

Previous research has found that the balance of the intake is related to children's rate of progress; children in schools where the average scores of entrants are relatively high tend to progress better than those in schools the initial scores are lower.[3] However, this factor does not seem to explain differences between schools in the present study in terms of progress in reading over the first two years. There is no significant correlation between the balance of intake and the second-year reading scores predicted by the model.[4] Nor is there a significant correlation between the second-year reading scores and the proportion of pupils belonging to ethnic minority groups.[5]

Second-year maths score as the outcome

An analysis was carried out on the pattern of the one just described, but taking the second-year maths score as the outcome instead of the second-year reading score. Of course, when trying to predict the second-year maths score, the first-year maths rather than reading score is used as an independent variable. Thus, the four independent variables included in the model are sex, first-year maths score, social class and country of origin. Information about the best fit model is summarised in Table 10.4.

The range of the second-year maths score is 0 to 60, and the mean score across the whole sample is about 31. The range of the first-year maths score is smaller (0 to 30). The estimates shown in Table 10.4 imply that there is a very strong relationship between the two maths scores. For each additional point scored on the first-year test, it is estimated that 1.297 additional points will be scored on the second-year test. From the standard error it is clear that this relationship is significant at a very high level of confidence. Social class is also clearly related to the second-year maths score, but the effect is fairly small in the context of the effect of the first-year score. Children belonging to the 'other' ethnic group (which includes West Indians) score significantly lower than the rest in the context of the model; this effect is small, but significant at a high level of confidence. The difference between boys and girls is not significant.

The scores predicted by the fixed part of the model are illustrated in Table 10.5.

When the grand mean is placed in the random part, there is an increase in the predictive power of the model which is significant at a very high level of confidence. The best fit model is one that also has the first-year maths score in the random part. The improvement in the predictive power of the model achieved by placing the first-year maths score in the random part is significant at the 99 per cent level of confidence. These findings show that there are significant differences between the results achieved in the second-year maths test by children in different schools, after taking account of their earlier scores and other factors; and that there are significant differences between schools in the slope of the relationship between the first and second-year maths scores. This second finding implies that some schools do relatively well with initially low-scoring pupils, and others with initially high-scoring pupils.

The schools where initially low-scoring children improved most were schools 12, 14 and 45. Those where initially high-scoring children improved most were 31, 33 and 42.

This should be an indication that schools in the first group emphasise remedial maths, while those in the second group emphasise maths for high achievers.

The school differences account for between 6 and 10 per cent of the variances in second-year maths scores (see Table A2 in Appendix 2). Taking as an example a boy belonging to a skilled manual family originating from the UK and who achieved a middling score of 16 on the first-year maths score, the model predicts that if he went to school 12 he would score 27, while if he went to school 33 he would score 36. These two schools are at the extremes of the range. Table 10.6 shows further examples of scores predicted within different schools. The analysis suggests that the relative performance of the different schools will vary depending on the prior attainment of the child in maths. For any given child, the schools are spread out fairly evenly across the range of performance.

Because of the complexity of the relationships, it is difficult to say whether the same schools achieve good and poor results with both reading and maths. It is, however, quite clear that school 33 does particularly well at both. This school has a lower proportion of children from ethnic minority groups (12 per cent) than any other school in the study. It has no West Indian children. Its social class composition is close to the average for schools in the study. The reading scores for children at this school on entry were close to the average for all children in the study, but the maths scores of entrants were well above the average, as were their non-verbal reasoning scores.

Further models

Taking the second-year reading and maths scores as the outcomes of schooling, the last section described statistical models that show how far and in what way these outcomes can be predicted from the sex, social class and ethnic group of the child and from the first-year score in either reading or maths. It went on to consider how much these predictions can be improved by allowing for the possibility that children do better in some schools than in others, and, further, for the possibility that the relative progress of different groups varies between schools. The provisional conclusion was that there are important differences between schools, both in the general level of progress in reading and maths, and in the ability groups, and possibly the ethnic groups, that achieve the best progress.

This section examines the results of adding further variables to the model. One reason for doing this is to check that the provisional conclusions are sound. There might be other characteristics of children not so far taken into account, such as disturbed behaviour, which would reduce or increase the estimated school differences. A second reason for including further variables is to explain and understand more of the differences in achievement between children and, possibly, schools.

Each further analysis starts from the basic model with the four independent variables (sex, first-year test score, social class and ethnic group) then adds the one further variable that is under consideration. If the further variables were added cumulatively, the model would soon become too complex. Instead, the procedure starts from the assumption that the four basic variables are basic and must always be taken into account. The further variables are tested one by one, to see if they have a significant effect in combination with and on top of the basic variables.

Rutter B2 score

Children having emotional or behavioural problems may be expected to progress more slowly as a result, and schools may differ in the proportion of children entering them who have problems of this kind. We have used the Rutter B2 scale (filled in by primary school teachers about four months after the children had left primary school) as an indicator of behavioural problems. Table 5.6 (on page 83) shows that there are some notable differences between schools in the scores on the behaviour scale. It is therefore important to establish whether the conclusions would be affected by taking account of these scores.

Unfortunately, there was a substantial shortfall of response for certain schools (see Chapter 5). The B2 scores are available for 1,854 children altogether, but when this variable is combined with others in the model, the number of children included in the analysis is only 714. Consequently, our estimate of the effect of additionally controlling for the B2 score is not very robust. In the model with the second-year reading score as the outcome, we do find a significant relationship with the Rutter B2 score; the estimate is -0.367, and the standard error of this estimate is 0.116. In other words, it is estimated from the model that for each additional point on the B2 score there is a decrease of about one-third of a point on the second-year reading score. This is not much, as the reading score has an upper limit of 155. The effect is very small indeed compared with the effect of the first-year reading score, for which the estimate is 5.078. In principle, this effect seems much too small to have a significant influence on the random part of the model (that is, on the estimates of school differences). However, this cannot be tested directly, because of the large amount of missing data in certain schools. These findings suggest that the provisional conclusions do not have to be modified after taking account of differences between schools in the number of children entering them with behavioural problems.

In the model with the second-year maths score as the outcome, there is no significant relationship with the Rutter B2 score, and no indication, therefore, that the provisional conclusions need to be modified.

A different analysis shows that there is a straightforward correlation between the B2 score and the test score achieved by a child at any particular stage. This can be illustrated by the correlation coefficients between the B2 scores and the test scores, which are shown below.

First year	
Writing	-0.26
Reading	-0.23
NVR	-0.24
VR	-0.24
Maths	-0.25
Second year	
VR	-0.26
Maths	-0.25
Reading	-0.21

A high B2 score indicates behaviour disturbance, so the negative correlation coefficients mean that disturbed behaviour is associated with low scoring on the tests, as might be expected. The coefficients shown are not very large: they indicate that the B2 score is associated with about six per cent of the variance in test sores. But the result obtained from adding the B2 score to the model suggests that this fairly low level of correlation is mediated by the basic variables: in other words, the B2 score adds very little when the basic variables are already taken into account. We cannot draw this conclusion with great confidence, however, because of the response shortfall on the B2 scores.

Non-verbal reasoning

The reason for including a test of non-verbal reasoning among those taken at the beginning of the first year was that this is a reliable test of an aspect of intellectual performance that is related less than any other to knowledge of a particular language. We would expect children whose first language is not English to be at less of a disadvantage on the non-verbal reasoning test than on the other tests. Results already shown in the Chapter 9 confirm that this is so, though the difference is not very large.

For this reason, including the first-year score on non-verbal reasoning in the model could make an important difference to the result, because it might improve our capacity to explain or predict the later performance of Asians who initially had poor English. In fact, we do find that non-verbal reasoning has a considerable and statistically significant effect on second-year reading and maths scores.

We find that each point on the non-verbal reasoning (NVR) score is worth 1.16 points on the second-year reading score, and this estimate is significant at a very high level of confidence (the standard error is 0.11). With NVR included in the model, the estimated effect of the first-year reading score is a bit lower than without it (4.58 instead of 5.37). The effect of the first-year reading score is nearly two and a half times as great as that of the NVR score, after allowing for the difference in the range of the two scores (0-15 and 0-25 respectively).

Each point on the NVR score is worth 0.32 points on the second-year maths score, and this estimate is also significant at a very high level of confidence (the standard error is 0.046). With NVR included in the model, the estimated effect of the first-year maths score is a bit lower than without it (1.15 instead of 1.31). The effect of the first-year maths score is, therefore, about four times as great as that of the NVR score, allowing for a difference in the range of the two scores (0-30 and 0-25 respectively).

From these findings, NVR has more importance in predicting reading than maths (after taking account of the first year reading or maths score). This confirms that the test is doing what we expected of it. Some children are put at a disadvantage in the reading test specifically because they have not yet adapted to using English (instead of another language). Adding the NVR score allows us to make a better assessment of how they will later do in reading. The effect is less pronounced in the case of maths, because performance on the initial maths test is less closely related to knowledge of English than performance on the reading test; also NVR and maths are more closely related than NVR and reading, so when seeking to explain maths scores, NVR adds less.

These findings show that NVR has some importance when explaining second-year reading and maths scores. But when we go through the rest of the analysis with a model including NVR, the general pattern of the results is very little affected. School differences still account for the same proportion of the variation and the same schools are identified as performing better or worse than average. Thus, although NVR does have some importance in itself, including it in the analysis does not affect the conclusions reached so far.

Use of English score among south Asians

We have shown in Chapter 9 that the test scores, and the reading scores in particular, vary considerably among south Asians depending on how far they use languages other than English in various contexts. It is important to establish whether *progress* in reading among south Asians is similarly related to use of other languages. Table 10.7 sets out the main results of a variance components analysis carried out among south Asians only, taking the second-year reading score as the outcome, and including the use of English score in the model in addition to the basic variables. The use of English score is shown to be significantly related to the second-year reading score in the context of the model (95 per cent level of confidence). The estimate is negative, which means that south Asian children tend to progress more slowly in reading over the first two years to the extent that they use languages other than English in various contexts. This is a fairly important factor: there is a predicted difference of between 8 and 9 points on the second-year reading score between south Asian children scoring 0 on the use of English score and those scoring 7. This finding shows that the south Asian children who speak mainly other languages are tending to fall further behind in reading over the first two years.

The main model with the second-year reading score as the outcome apparently shows that the relative progress of different ethnic groups varies between schools. This result must be interpreted with caution, since disparate ethnic groups have had to be lumped together; for example, south Asians in one school may have different ethnic origins in detail from those in another. In particular, their linguistic background may be widely different. The need for caution is underlined by the finding that progress in reading is related to the use of languages other than English.

Other background variables

We found in Chapter 9 that the test scores are associated with parents' qualifications as well as their socio-economic group. Within particular socio-economic groups, there was a small difference in the children's test scores according to the parents' qualifications. In line with this, when the highest qualification of the parents is added to the variance components model, it is found to have a small but statistically significant effect on top of the effect of the basic variables. However, the addition of this variable does not affect the results or conclusions in an important way.

Further models have been set up with each of the following variables added, in turn, to the basic four: absences from school in 1981/2 and in 1982/3; whether one parent at least was in work in the summer of 1982; whether there is a single parent or two parents. None of these background variables is significantly related either to the second-year reading or to the second-year maths score in the context of the model.

The analysis in Chapter 7 suggested that school attendance is a reflection of several contrasting influences, such as rebelliousness, dislike of school, presence of an adult at home, demands to help out at home, degree of emphasis placed by the parents on education. It appears to be a 'hard' measure, but in fact it does not relate to a single dimension. Also, levels of attendance are rarely so low as to represent an important educational handicap. These two points probably explain why attendance is not a significant influence on progress in reading and maths over the first two years.

Whether or not the parents are unemployed would be related to test scores on its own, but something like this variable is already incorporated in the analysis, since one of the social class groups is families where the parents or parent had not had a job in the previous five years. That is probably why no further effect is shown.

From the results shown in Chapter 9, it is not surprising to find that there is no clear influence according to whether there is a single parent or two parents. This is because the prospects of children of single parents seem to vary markedly depending on whether or not the parent has a job, yet the numbers are too small for us to take account of this point in the multivariate analysis.

Participation

There is a fairly strong relationship between the level of attainment in absolute terms and the level of participation in school activities. At the same time, as described in Chapter 7, there are fairly large differences between schools in the level of participation (this point will be further pursued in Chapter 11). This opens up the possibility that participation in school activities may be bound up in some way with differences between schools in attainment at the end of the second year.

Taking the second-year reading score as the outcome, Table 10.8 shows the fixed part of a model with the participation score added as one of the independent variables. This model shows a significant relationship between the participation score and the second-year reading score. The relationship is not particularly strong. The figures imply that a child with the maximum participation score of 4 would achieve a second-year reading score five points higher than a child with the minimum participation score of 0. Still, this is greater than the differences shown between ethnic groups and about half as great as the social class differences.

A similar model can be set up with the second-year maths score as the outcome and with participation added as one of the independent variables. The results are shown in Table 10.9. This model again shows a significant relationship between the participation score and the outcome – this time the second-year maths score. Again the relationship is not particularly strong, but the influence of participation is about as great as that of ethnic group.

There is little evidence that the differences between schools in rates of progress achieved are associated with differences in levels of participation, for with the participation score controlled, school differences in progress in reading and maths apparently remain the same.

Further evidence can be provided by plotting the average participation score at the school level against a measure of the progress achieved by the school in maths or reading. There is no significant correlation.[6]

These findings suggest that participation is somehow bound up with academic performance at the individual level. Children with high participation scores tend to attain well in absolute terms and show better progress in attainment. However, there is no evidence that participation helps to explain differences in academic performance between schools. This may be because the sample of schools is too small to demonstrate a fairly weak effect.

Index of praise and index of blame

When the index of praise is added to the model, we find no significant relationship with the second-year test scores in either maths or English. With reading as the outcome, the estimate for the effect of the index of praise is negative (-0.46 against a standard error of 0.45) which certainly makes the point that there is no tendency for praise to be associated with good progress.

But there is a relationship between the index of blame and progress in reading and maths, a relationship that is significant at a high level of confidence in both cases (see Tables 10.10 and 10.11). In the case of maths, the index of blame seems to be more strongly related to progress than participation. Children belonging to the 'other' ethnic group (which includes those of West Indian origin) had a significantly poorer rate of progress in maths than the rest, and also had higher scores on the index of blame. This suggests the possibility that poor progress in maths among children of West Indian origin may be related to discouragement by teachers.

One of the items in the index of blame is about being ‘told you have done poor work’, while the other three are about having been ‘told off’. It seems very unlikely that children whose attainment in the second year is poor *compared with their attainment at the point of entry* are being identified and ‘told off’. For this reason, the theory that the process runs from poor progress to negative messages from teachers seems unattractive. It seems much more likely that the process runs from negative messages to poor progress: that children who are often criticised tend to progress slowly as a result.

There is, however, no evidence that the level of negative messages to children helps to explain differences between schools in progress in reading or maths. When the average index of blame at the school level is plotted against a measure of progress, no significant relationship is found, though the non-significant relationship is in the expected direction.[7]

These results demonstrate that the level of negative messages is bound up with individual progress, but do not show that it helps to explain differences in rate of progress at the school level. This may be because the sample of schools is too small to demonstrate a fairly weak effect.

Enthusiasm for school

Enthusiasm for school is not significantly related to progress in reading or maths over the first two years.

Number of English and maths teachers

To obtain an indication of the stability and continuity of the teaching they had received, children were asked ‘How many English teachers have you had at this school?’ and the question was repeated for maths teachers. With the second-year reading score as the outcome, the number of English teachers was added to the model, and with the second-year maths score as the outcome, the number of maths teachers was added.

The number of maths teachers is not significantly related to the outcome in any way. The number of English teachers is not significantly related to the second-year reading score in the fixed part of the model, which means that there is no significant association when the grouping of children into schools is ignored. The weak relationship shown, which just fails to reach significance at the 95 per cent level of confidence, is in the expected direction (more teachers is associated with poorer performance). When the slope of the relationship between the number of English teachers and the reading score is allowed to vary between schools, the model is significantly improved (95 per cent level of confidence).

On the whole, these findings suggest that stability of English teachers is mildly associated with progress in reading, but the relationship is weak.

Parental attitudes and contact

There are some relationships between children’s attainment in absolute terms and their parents’ attitudes and their level of contact with the school. These will be further discussed in Chapter 11. However, when variables to do with parents’ attitudes and contact with the school are added to the basic model, the results show that these variables are not significantly related to children’s *progress* in reading or maths.

Conclusions

At the point of entry to secondary school, certain categories of south Asian children scored substantially lower in reading and maths than the average for the children tested. Children of Bangladeshi origin achieved the lowest scores (lower in maths than in reading) but children of Pakistani origin also scored substantially below average. Children of West Indian origin also scored below average at the point of entry, but considerably higher than the low-scoring south Asian groups.

By the end of the second year, the relative position of the different ethnic groups in maths and reading was much the same as at the point of entry: if anything, the gap had grown wider. When the second-year reading score is regarded as the outcome, a multivariate analysis shows that, after taking account of the first-year reading score, sex and social class, south Asians scored significantly lower than children of UK origin, while the 'other' group (which includes West Indians) scored about the same. This is roughly the same as saying that *progress* in reading was slower over the first two years of secondary school for south Asians than for white children. In the case of maths, the progress of south Asian children (from a substantially lower baseline) was not significantly different from that of children originating from the UK. However, the progress of other ethnic minorities (including West Indians), from a lower starting point than that of children originating from the UK, was rather slower.

South Asians who use languages other than English in various contexts tend to achieve lower test scores at the point of entry to secondary school than those who only use English. A multivariate analysis shows that, in addition, the south Asian children who speak mainly other languages are tending to fall further behind in reading over the first two years. This suggests that linguistic background may explain the general tendency for south Asians to progress more slowly than white British children. It also underlines the need for caution in making comparisons between the results achieved by different schools with south Asians, as there is not room to take account of the particular ethnic and linguistic background of south Asians in each particular school.

There is clear evidence of substantial differences between schools in the level of attainment in reading and maths at the end of the second year, after allowing for attainment at the point of entry, and for other factors. In other words, the rate of progress in reading and maths over the first two years is substantially better in some schools than in others. Also, the findings show that some schools achieve their better results in maths with initially low-scoring children, and others with initially high-scoring children. From this it seems that certain schools concentrate on remedial maths, while others concentrate on maths for high achievers. These school differences do not arise because of differences in ethnic composition, and they are larger than the differences in progress between ethnic groups. Therefore, which school a child goes to is more important than which ethnic group he or she belongs to as an influence on the rate of progress in reading and maths.

The main multivariate models on which these conclusions are based take account of sex, the first-year test score, social class and ethnic group. Further analyses have been carried out to check the effects of other background factors (the level of disturbed behaviour, the non-verbal reasoning score and parents' qualifications, among others). The findings do not alter any of the main conclusions.

Participation in school activities is somehow bound up with academic performance at the individual level. Children with high participation scores tend to attain well in absolute terms and show better progress in attainment. However, there is no evidence that participation helps to explain differences in academic performance between schools. This may be because the sample of schools is too small to demonstrate a fairly weak effect.

There is no tendency for praise from teachers to be associated with good progress in reading or maths, but there is a clear tendency for children to progress more slowly to the extent that they are criticised by teachers. It seems very unlikely that children whose attainment in the second year is poor *compared with their attainment at the point of entry* are being identified and 'told off'. It seems much more likely that children who are often criticised tend to progress slowly as a result. There is, however, no evidence that the level of negative messages to children helps to explain differences between schools in progress in reading or maths. Again, this may be because the sample of schools is too small to demonstrate a fairly weak effect.

Notes

1. In fact, the socio-economic scale is scored from 1 (underclass) to 5 (professional or managerial), so there is no zero value. Thus 1 x the estimate for social class has to be added to the grand mean even for the underclass group. This does not affect the general principle that the estimates are to be added to the grand mean.
2. These two schools are the extreme cases. The figures quoted are only illustrative, since there is no method for determining the confidence limits of these predicted values.
3. For a review of the evidence on this point, see Rutter (1983).
4. To test the hypothesis about balance of intake, the posterior means for the 18 schools were plotted against the mean first year reading scores. The correlation between these two statistics ($r = 0.247$) is not significant at the 95 per cent level of confidence. For the purpose of this test, the random effect of ethnic group at the school level was ignored, as this finding is not robust (see the text). The school means of the first-year reading scores were used as the measure of balance of intake.
5. The percentage of pupils not originating from the UK or Eire was plotted against the posterior means for the second-year reading scores (see note 4 above). The value of r (the correlation coefficient) was -0.299, which is not significant at the 95 per cent level of confidence. Of course, a significant result might be obtained with a larger sample of schools. A negative correlation would indicate that schools with a large proportion of ethnic minority pupils tend to achieve slower progress than those with a small proportion.
6. As before, the measure of progress is the posterior means for the 18 schools, not taking account of the random effect of ethnic group (in the case of reading) or of the first-year maths score (in the case of maths). When the posterior mean is plotted against the mean participation score for the school, the value of r (the correlation coefficient) is -0.066 in the case of reading and -0.148 in the case of maths. Neither result is significant at the 95 per cent level of confidence.
7. As before, the measure of progress was the posterior mean for each school (see note 6 above). This was plotted against the mean score on the index of blame for each school. The value of r (the correlation coefficient) was -0.256 in the case of reading and -0.102 in the case of maths. Neither value is significant at the 95 per cent level of confidence.