MULTILEVEL MODELLING NEWSLETTER

Produced through the Multilevel Models Project:

c/o Math, Statistics & Computing Department
Institute of Education, University of London,
20 Bedford Way, London WC1H 0AL
e-mail: rprosser%uk.ac.lon.educ.isis@ac.uk
Telephone: 071-636-1500 ext 473

Vol. 2 No. 2 June 1990

NEW MULTILEVEL TEXT AVAILABLE SOON

Anthony Bryk and Steven Raudenbush have just submitted a manuscript to Sage for a book entitled *Hierarchical Linear Models for Social and Behavioral Research: Applications and Data Analysis Methods.* The monograph will be the first in Sage's new methodology series (edited by Jan de Leeuw of UCLA) and is expected to become available by early 1991. The contents of the nine chapters are outlined here.

The introductory chapter discusses the multilevel character of data in much social science research and illustrates three general objectives of multilevel modelling: improved estimation of effects within individual units, formulation and testing of hypotheses about cross-level effects, and the partitioning of variance and covariance components among levels. Basic notation for two-level modelling is introduced.

The logic of hierarchical linear models (HLMs) is the topic of chapter 2, and the focus is on formulation of simple two-level models and interpreting the parameter estimates and other statistics obtained in an analysis. Data from the *High School and Beyond* survey (National Center for Educational Statistics, 1981) form the basis for question-driven analyses involving models of increasing complexity, from a one way random effects ANOVA model through to a cross-level interaction model. The chapter concludes with a discussion of the task of fitting a regression for a single school, and compares OLS and two types of shrinkage estimators.

Chapter 3 deals with two topics: (a) the basics of estimation theory concerning fixed parameters, random effects, and dispersion components, and (b) hypothesis testing. The concept of a *precision-weighted average* is introduced in the context of estimating a grand mean, and the discussion leads up to examination of generalized least squares estimation of a set of fixed parameters. Shrinkage estimators for random effects are described, and a comparison is given of full maximum likelihood (MLF), restricted maximum likelihood (MLR), and Bayes estimation of covariance components. Individual and composite hypothesis tests are described, and likelihood ratio testing is compared with standard testing of the general linear hypothesis.

Chapters 4 and 5 are the core application sections of the book. The former deals with individual growth and the latter discusses organizational effects on individual-level processes. The first part of chapter 4 explains the formulation of within-person polynomial growth models and associated between-person models. Analyses based on Head Start program data and a study of effects of maternal speech on children's vocabulary are used to illustrate the interpretation of growth parameters, and the examination of individual variation around a mean growth trajectory, reliability of parameter estimates, correlation of change with initial status, and effects of person characteristics on growth. The final section of the chapter deals with several complexities including piecewise linear models, alternate age metrics, time-varying covariates, and prediction of future status.

Two classes of organizational situations are examined from a two-level modelling perspective in chapter 5: those in which an organization's features–structure or climate, say–"exert a common influence on each individual within it," and those in which these features "modify both the mean level of outcomes and how effects are distributed among individuals." In one illustration, outcomes of a multilevel analysis are compared with two OLS analyses–one at the class level and the other at the student level. Sorting out contextual and "frog-pond" effects is discussed as is the problem of estimating performance / effects of individual organizations.

Three-level analysis is the topic of chapter 6, and basic models are introduced in the context of organizational research: students are nested within classes which are in turn nested within schools. The

cont'd on p. 2

BOOKS

cont'd from p. 1

emphasis in the first part is on interpretation of models' components and the presentation of modelling possibilities. Hypothesis testing is discussed. The chapter's major illustration is an analysis of schools' impacts on individual growth in mathematics achievement, based on the Sustaining Effects Study (Carter, 1984).

In some studies, the variance of the level 1 random terms may be assumed to be known. Applications of such knowledge in multilevel modelling are discussed in chapter 7. The characteristics of *V known* situations are stated and three examples are provided–two concerned with meta-analysis and one with dispersion as an outcome.

The purpose of chapter 8 is to provide data analysts with information about the following: (a) the assumptions required in using HLMs (b) "the sensitivity of conclusions to possible violations of these assumptions, (c) techniques available to investigate the tenability of the assumptions, and (d) ameliorative strategies available when likely violations are discovered." These topics are discussed for the fixed and random parts of the within-unit model and the between-unit model in turn. The final section examines the question of validity of inferences when samples are small.

Estimation is presented conceptually in the initial chapters. Chapter 9–a technical appendix–details the estimation theory for hierarchial linear models and is developed from a Bayes linear model perspective. Estimation of covariance components via the EM algorithm is also described.

The book is intended to serve as a text and reference, and it should be of interest to researchers and graduate students in education and the social sciences. A good grounding in applied multiple regression is assumed.

References

Carter, L. F. (1984). The sustaining effects of compensatory and elementary education. *Educational Researcher, 13,* 4-13.

National Center for Educational Statistics. (1981). *High School and Beyond: A National Longitudinal Study for the 1980's.* Washington, DC: Author.

MULTILEVEL CONFERENCE PAPERS TO BE PUBLISHED

Until recently, most published applications of multilevel statistical methods have appeared as illustrative examples in methodological articles. However, during the past three or four years, an increasing number of researchers in various countries have begun to use multilevel methods for substantive research in education. A new book entitled *Schools, Classrooms, and Pupils: International Studies of Schooling from a Multilevel Perspective* edited by Steven Raudenbush and Douglas Willms presents some of the best of the "first wave" of these substantive applications in educational research.

The volume, to be published this autumn by Academic Press, is based on papers read at last summer's *International Conference on Application of Multilevel Methods in Educational Research* at the Centre for Educational Sociology in Edinburgh. The chapters shed light on a variety of problems: curriculum coverage and curricular reform (chapters by Adam Gamoran; by Ian Plewis; and by Ruth Zuzovsky & Murray Aitkin); the stability and validity of educational indicators (Roel Bosker & H. Guldemond; Carol Fitz-Gibbon; Nicholas Longford); the organizational context of teaching (Anthony Bryk & Ken Frank; Valerie Lee & Julia Smith; Brian Rowan, Stephen Raudenbush & Sang Jin Kang); the evaluation of innovative programs (Suzanne Jacobson; David Raffe); the changing social distribution of achievement in Scotland (Lindsay Paterson); and school effects in a developing nation (Marlaine Lockheed & Nicholas Longford).

The chapters illustrate application of the three most commonly used multilevel computer programs: HLM, *ML2*, and VARCL. They provide interesting insights into the educational systems of countries as diverse as Canada, England, Israel, the Netherlands, Scotland, the United States, and Thailand. The chapters are accessible to many researchers, requiring a minimum of technical background. Lindsay Paterson's introductory chapter presents a primer on multilevel statistical methods for the newcomer.

A future issue of the *Newsletter* will carry a review of this book.

PROJECT NEWS

FIVE MULTILEVEL PROJECTS AT UCLA

Jan de Leeuw, Ita Kreft, & Kyung-Sung Kim

As version 3.0 of the report comparing GENMOD, HLM, *ML2*, and VARCL nears completion, the Multilevel Group at UCLA is starting five new projects. (There are some loose ends to tie up on the software testing such as comparison of actual likelihood values at convergence, but most of the work is done.)

For the Netherlands Institute for Educational Research, the team is developing *MULTIPATH*, a program which performs multilevel path analysis. The software uses a general formulation of path analysis in which the joint distribution of the observed and latent variables is decomposed using conditional independence and the notion of various levels. This makes it possible to construct nested path models in which the path coefficients of the first level model are variables on the second level. These and other level 2 variables are connected in a second order path model on the second level.

In cooperation with Peter Bentler of UCLA's Psychology Department, the team is working on a preprocessor for the *EQS* program that takes correlations among clustered observations into account and handles multiple dependent measures. Assuming a simple dependence structure between individuals makes it possible to apply existing theory on the matrix normal distribution. For this distribution, the covariance between individual i on variable j and individual k on variable l is given by $Cov\ (Y_{ij}, Y_{kl}) = \omega_{ik}\sigma_{jl}$. The matrix Ω of order n has an intraclass correlation structure, and the matrix Σ of order m satisfies a factor analysis or path analysis model. This work provides a second and much simpler approach for the development of multilevel path analysis. GAUSS versions of the factor analysis implementation are available. A FORTRAN module to be used with *EQS* is in preparation.

Nick Longford and Bengt Muthen, are working on a multilevel version of factor analysis that is more general than the one based on the matrix normal. A FORTRAN program to perform this analysis and a report describing the method are almost complete.

Together with Rien van der Leeden of the University of Leiden, the team is writing a report comparing random coefficient software for growth curve and repeated measurement models. This is similar to the report comparing cross-sectional multilevel programs but concentrates on special purpose software such as BMDP5V and GGCAMOV.

In the multilevel literature one often finds the claim that random coefficient models are better than fixed coefficient models, that restricted maximum likelihood estimation is superior to unrestricted, and that weighted least squares outperforms unweighted. These claims are usually substantiated by references to general theoretical or statistical considerations rather than empirical evidence–an unsatisfactory situation. One project in the planning stage is a comparison of random coefficient and fixed coefficient models by means of leave-one-out cross-validation. Diagnostics similar to influence measures and residual plots would be developed in the process. Collaborators, programming assistance, support and computer time remain to be found, but such a study is long overdue.

LONDON WORKSHOP IN OCTOBER

The Multilevel Models Project team will conduct a workshop on three-level analysis in London on October 10–12. This training session will be of interest to researchers in education and the social sciences, and will provide participants with an opportunity to gain some hands-on experience conducting multilevel analyses with their own data. Basic models for organizational and longitudinal data will be covered as will treatment of multivariate and categorical multilevel data.

Participants will receive a complimentary copy of *ML3*, and there will be no charge for tuition. (This workshop is being conducted as part of the Project's mandate to disseminate information about multilevel modelling.) To obtain further information and to register, please contact Bob Prosser at the address on the cover page.

The Project team has begun to conduct **custom workshops** for organizations such as local education authorities interested in doing specialized multilevel analyses. For further details about this service, please contact Professor Goldstein at the address on the cover page.

THEORY &

MULTILEVEL MODELS WITH KNOWN LEVEL 1 VARIANCE STRUCTURES

Stephen W. Raudenbush & Anthony S. Bryk

A fundamental goal in two-level analyses is to estimate variation and covariation at each level. Usually the first stage variance structure is assumed to be quite simple, and data are pooled across level 2 units to estimate the variances at each level. For example, in a school effects study, one pools the data from each school to estimate variability among students within schools; and one estimates the variation and covariation of the coefficients which vary randomly across schools.

However, there are a number of interesting hierarchical data problems in which the structure of variation-covariation at the first level is arbitrarily complex but can be assumed known. These *V-known* cases include meta-analysis and other studies in which a goal is to compare proportions, variances, correlations, log-linear effects or other interesting parameters estimated for each of many independent samples. This note describes a set of procedures for V-known applications using the HLM program (Bryk, Raudenbush, Seltzer, & Congdon, 1989). The defining characteristics of a V-known problem are the following:

(1) For each level 1 unit, there must be available a statistical estimator (typically a maximum likelihood estimator) of an interesting parameter (or parameter vector) and an estimate of the sampling dispersion of that estimator.

(2) The parameters estimated for each group are assumed to vary randomly across groups. We shall refer to these as *micro-parameters.* It must be reasonable to assume that these micro-parameters are normally distributed.

(3) The level 1 units must all be large enough to justify the assumption that the estimators of the micro-parameters are approximately normally distributed with a variance which is approximately equal to its sample estimate.

These assumptions are not as restrictive as they may seem. In many instances where samples seem too small to justify the normality or known variance assumptions, the use of a suitable transformation of the sample estimator will render these assumptions quite reasonable.

When the assumptions above are reasonable, the class of problems accessible to the investigator by means of hierarchical models will be broadened in the following ways:

(1) Two level HLMs can be used in cases where the within-unit data are unavailable. An example is meta-analysis in which one typically has access only to summary statistics from each study but not to the original data.

(2) The V-known approach is applicable in some cases where the outcome variable is clearly non-normally distributed. An example is the two-level log-linear model described by Goldstein in *Fitting Loglinear ML Models* (Multilevel Modelling Newsletter, January 1990, page 3, Equation 2). Note that in that example it would be possible to estimate the log-linear effects within each school; and the varince-covariance matrix of those effects could be assumed known if enough data per school were available.

(3) A wider variety of micro-parameters (e.g., correlations, variances, standardized mean differences) become accessible to investigation. Note that in most multilevel applications the micro-parameters have been regression coefficients.

(4) A wide variety of level 1 variance structures can be studied. For example, in meta-analysis, the outcome variable is typically measured on a different scale for each study. In a multinomial application, the variance-covariance structure of the outcomes is non-diagonal with a special form.

(5) The V-known computer program can efficiently summarize very large amounts of data. Again, meta-analysis is an example where a set of studies may be based on many thousands of data points. There are, however, other examples, including international studies of educational achievement where large samples are collected in scores of countries and the goal is to compare results across countries.

APPLICATIONS

Below we consider a univariate example based on a meta-analysis of the effect of teacher expectancy on pupil IQ. We then consider generalizations to other univariate and multivariate examples.

Example: The Effect of Teacher Expectancy on Pupil IQ

Raudenbush and Bryk (1985) summarize the results of 19 experiments that attempted to assess the effects of teacher expectations on pupil IQ. From each study it was possible to estimate a standardized mean difference $\delta_j = (\mu_{Ej} - \mu_{Cj})/\sigma_j$ with the statistic $d_j = (\overline{Y}_{Ej} - \overline{Y}_{Cj})/S_j$ where μ_{Ej} and μ_{Cj} are the population mean IQ scores for the experimental and control groups, respectively, in study j, $j = 1, \ldots, 19$; $\overline{Y}_{Ej}$ and $\overline{Y}_{Cj}$ are the corresponding sample means; and σ_j and S_j are, respectively, the population and sample standard deviations, pooled within groups. Hedges (1981) proved that, for a fixed value of δ_j, the statistic d_j is approximately unbiased and normally distributed with $V_j = (n_{Ej} + n_{Cj})/(n_{Ej}n_{Cj}) + \delta_j^2/[2(n_{Ej} + n_{Cj})]$.

Within-study model. For each study j, the micro-parameter is the effect size, δ_j, which has been estimated by the statistic d_j, and we assume that d_j, given δ_j, is normally distributed with a known variance V_j. Then the simple linear model

$$d_j = \delta_j + e_j, \quad e_j \sim \mathcal{N}(0, V_j)$$

applies, where d_j, δ_j, and V_j are defined above.

Between-study model. A simple model for the micro-parameters assumes that they vary around the grand mean with a variance to be estimated:

$$\delta_j = \gamma_0 + u_j, \quad u_j \sim \mathcal{N}(0, \tau)$$

where γ_0 is the grand mean effect size and τ is the variance of the true effect size δ_j about that mean.

Results. In the present example, the grand mean effect size estimate is small, eg, $\hat{\gamma}_0 = 0.08$, which represents just eight percent of a standard deviation advantage for the experimental group over the control group in mean IQ. However, the grand mean alone could be misleading if the size of the treatment effect varied substantially from study to study. The maximum likelihood estimate of τ is 0.019, corresponding to a standard deviation of 0.14, meaning that a study with an effect one standard deviation above the average would be expected to have an effect size of about 0.22, a non-trivial effect.

In this case it is important to test the significance of the variability of the effect sizes. If the true variance were null, all studies would have the grand mean effect size and it would be reasonable then to infer that in every study, the effect of teacher expectancy is small. Here the null hypothesis is $H_0 : \tau = 0$, and the test statistic suggested by Hedges (1982) is $H = \sum V_j^{-1}(d_j - \hat{\gamma}_0)^2$, where $\hat{\gamma}_0 = \sum(V_j^{-1}d_j)/\sum V_j^{-1}$, which has a χ^2 distribution with 18 degrees of freedom when the null hypothesis is true. In this case H takes on a value of 35.85, $p < 0.01$ implying that effect sizes do vary significantly from study to study. Hence, the mean effect size reported above, by itself, is misleading.

A subsequent analysis, reported in Raudenbush and Bryk (1985), used a single predictor (the extent of teacher-pupil contact prior to the experiment) in the between study model

$$\delta_j = \gamma_0 + \gamma_1 W_j + u_j, \quad u_j \sim \mathcal{N}(0, \tau)$$

where W_j is the weeks of prior contact in study j. This model accounted for nearly all of the variability among the true effect sizes.

Other Univariate Applications

Other examples of δ_j might be a mean for each group, a correlation, a regression coefficient, a proportion, or a variance. In some of these cases, the normality and V-known assumptions will be questionable unless sample sizes are very large in every group. For example, the distribution of the sample variance converges to normality by virtue of the Central Limit Theorem; but convergence is very slow, and the distribution may be very non-normal unless sample sizes are very large. Also the sampling variance of the statistic $\hat{\sigma}_j^2$ is proportional to the square of the true variance, σ_j^2, which is the quantity being estimated! However, a simple log transformation (see Table 1) significantly improves the normality approximation and also stabilizes the variance. Thus, the V-known methodology could be profitably applied to study variances even when samples are as small as 10 per group (see Bartlett & Kendall, 1946; Raudenbush & Bryk, 1987).

Table 1 (p. 7) presents some examples of parameters which are likely to be of interest along with their estimators, sampling variances, and useful variance stabilizing or normalizing transformations.

Multivariate V-known Model

The multivariate V-known model is a straightforward extension of the univariate case. A vector of parameters is estimated separately for each of many groups; and the sampling dispersion of the vector of estimates is assumed unbiased and multivariate normally distributed with known variance. Some useful examples are discussed in Bryk and Raudenbush (forthcoming). Again, transformations are sometimes useful to increase the tenability of the assumption of multivariate normality and of known variances and covariances.

Example. Becker (1988) utilized the multivariate V-known method to compare experimental and control group pre-post change scores across five experiments. Because the five experiments utilized outcome variables measured on different scales, it was necessary to standardize the change scores. Within each study, j, she computed the standardized mean change scores $d_{Ej} = (\overline{Y}_{Ej} - \overline{X}_{Ej})/S_j$ and $d_{Cj} = (\overline{Y}_{Cj} - \overline{X}_{Cj})/S_j$ where $\overline{Y}_{Ej}$ and $\overline{Y}_{Cj}$ are the posttest means, $\overline{X}_{Ej}$ and $\overline{X}_{Cj}$ are the pretest means and S_j is a measure of the scale of the outcome.

Becker derived the sampling variance of each mean change d_{ij}, for $i = E$ or $i = C$:

$$V_{ij} = 2(1 - \rho_{ij})/n_{ij} + \delta_{ij}^2/[2(n_{ij} - 1)],$$

where ρ_{ij} is the pre-post correlation in sample i of study j. In each study the samples are independent, so the sampling covariance between d_{Ej} and d_{Cj} is zero. Thus, the within-study model is

$$d_{ij} = \delta_{ij} + e_{ij}, \quad e_{ij} \sim \mathcal{N}(0, V_{ij})$$

and $\mathrm{Cov}(e_{Ej}, e_{Cj}) = 0$. However, looking across studies, it might be that the true effect sizes are correlated because of study-to-study differences in treatment implementation, outcome measures, and subject populations. Hence, Becker formulated a between-study model

$$\delta_{ij} = \gamma_i + u_{ij}, \quad u_{ij} \sim \mathcal{N}(0, \tau_{ii})$$

and $\mathrm{Cov}(u_{Ej}, u_{Cj}) = \tau_{EC}$. Becker found that, indeed, the correlation between E and C group changes was high (0.91). A major advantage of the multilevel formulation concerned the test of the key null hypothesis of no difference in average change between E and C, $H_0 : \gamma_E = \gamma_C$. Using the multilevel formulation via the V-known model, this test takes into account interstudy differences in mean change not attributable to the treatment. Becker found no significant E-C difference. This contradicts the result which would have occurred had the study-to-study variation been ignored.

References

Bartlett, M. S., & Kendall, D. G. (1946). The statistical analysis of variances-heterogeneity and the logarithmic transformation. *Journal of the Royal Statistical Society, Suppl. 8,* 128-138.

Becker, B. (1988). Synthesizing standardized mean-change measures. *British Journal of Mathematical and Statistical Psychology, 41,* 257-278.

Bryk, A. S. & Raudenbush, S. W. (forthcoming). *Hierarchical linear models for social and behavioral research: Applications and data analysis methods.* Newbury Park, CA: Sage.

Bryk, A. S., Raudenbush, S. W., Seltzer, M., & Congdon, R. (1989). *An introduction to HLM: Computer program and user's guide.* (2nd ed.). Chicago: University of Chicago, Department of Education.

Hedges, L. V. (1982). Fitting continuous models to effect size data. *Journal of Educational Statistics, 7,* 245-270.

Hedges, L. V. (1981). Distribution theory for Glass's estimator of effect size and related estimators. *Journal of Educational Statistics, 6,* 107-128.

Raudenbush, S. W., & Bryk, A. S. (1987). Examining correlates of diversity. *Journal of Educational Statistics, 12,* 241-269.

Raudenbush, S. W., & Bryk, A. S. (1985). Empirical Bayes meta-analysis. *Journal of Educational Statistics, 10,* 75-98.

APPLICATIONS

Table 1: Some Univariate V-known Cases

Micro-parameter, ϵ_j	*Sample Estimator,* d_j	*Approximate Variance,* V_j
Regression Slope,[a] β	$b = rS_y/S_x$	$\sigma^2/\sum(X_j - \overline{X})^2$
Standardized Mean Difference,[b] $(\mu_E - \mu_C)/\sigma$	$(\overline{Y}_E - \overline{Y}_C)/S$	$(n_E + n_C)/(n_E n_C) + d^2/[2(n_E + n_C)]$
Correlation $0.5 \log[(1+\rho)/(1-\rho)]$	$0.5 \log[(1+r)/(1-r)]$	$1/(n-3)$
Proportion,[c] π	p	$p(1-p)/n$
Logit, $\log[\pi/(1-\pi)]$	$\log[p/(1-p)]$	$[p^{-1} + (1-p)^{-1}]/n$
Transformed S.D.,[d] $\log(\sigma)$	$\log(S) - 1/(2f)$	$1/(2f)$

[a] Here r is the sample correlation between x and y and S_x and S_y are the sample standard deviations.

[b] Typically σ is the pooled, within-treatment standard deviation, μ_E and μ_C are the experimental and control population means, respectively, and $\overline{Y}_E$ and $\overline{Y}_C$ are the corresponding sample estimates.

[c] We denote π as the proportion of subjects in the population with a given characteristic; p is the sample proportion.

[d] We denote σ as a s. d., S as the sample estimate, and f as the degrees of freedom associated with S. Thus $f = n - 1$ when σ is a s. d., but $f = n - q$ when σ is the residual s. d. estimated in a regression model with q parameters.

Figures for Tymms & Fitz-Gibbon, (p. 8)

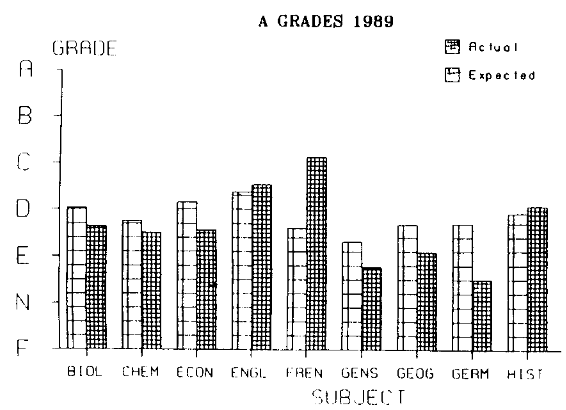

Figure 1

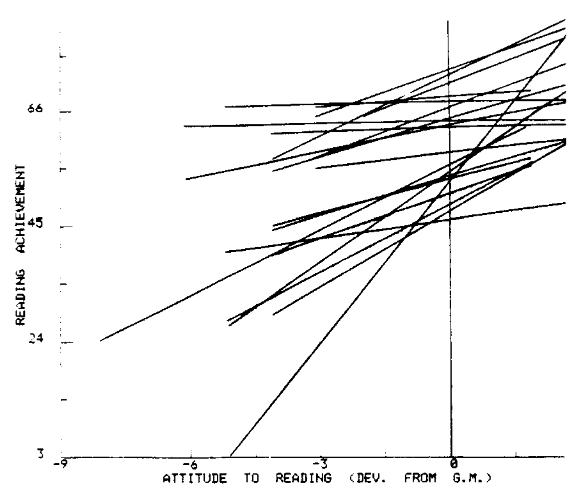

Figure 2

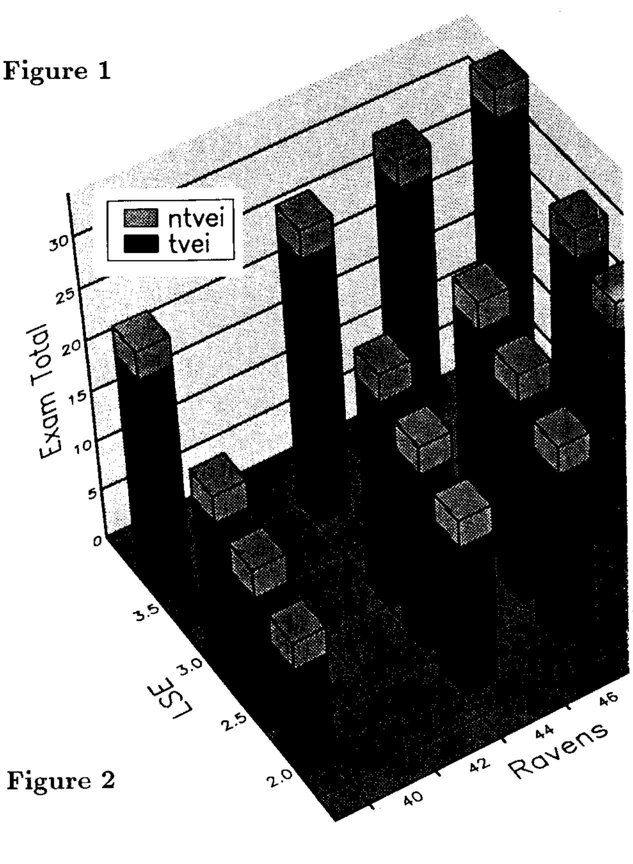

Figure 3

A GRAPH IS WORTH A THOUSAND NUMBERS

P. Tymms & C. Fitz-Gibbon

The raw output from a multilevel analysis cannot be understood by the average layperson, yet if the results are to improve education, they must be communicated in an understandable manner to advisors, inspectors, headteachers, and teachers. Findings must, therefore, be presented graphically since the eye can take in those complexities which are inaccessible through tables of figures. From our experience, two types of presentation are needed: one to show departmental and school data and the other to show general patterns found in the analysis.

Consider a comparison of schools, by department, in which separate models have been fitted for each subject. In the simplest situation, these models might have one explanatory variable with a fixed slope and a random intercept. In presenting school differences, we have found bar graphs to be effective, and they are well received by headteachers.

Figure 1 shows the kind of chart we have used in the A Level Information System (Fitz-Gibbon, 1989) to summarize data for a school or college. The expected (average) achievement corresponds to $\hat{Y}_j$, and the actual level is given by $\overline{Y}_j$. If the slopes are allowed to vary, then the charts could still be used, but it would then be important to base the expected and actual values on a hypothetical average pupil for a particular department and to explain that a different pattern might be seen if more/ less able students were considered.

Charts such as these could be used to convey a variety of findings by presenting breakdowns by gender, ability band, and so on. But there is an uncomfortable simplification if there is evidence that slopes vary across schools, and in fact it might be better to use a more complex graph.

To show the whole picture to school personnel, one might use a graph like the one in Figure 2, but with many schools the presentation could become quite confusing. An alternative would be to indicate the general trend with a single regression line and to place a particular school's results on the same graph. It would be desirable to indicate confidence limits on such a graph.

Various elaborations are possible. The performance of different subgroups within institutions could be compared with the overall average patterns for these subgroups. When a context effect has been identified, it would be necessary to show a general line which took that effect into account, and so a series of diagrams would be appropriate.

A third type of graphical output could be used to portray some overall patterns in the data. Figure 3, for example, shows relationships among three explanatory variables–two continuous and one categorical.

There are clearly some very exciting possibilities for the graphical presentation of findings from a multilevel analysis such as showing interactions in three dimensions and picturing the achievements of various subgroups, but it must be recognized that a single picture would take considerable time and effort to produce at the moment. In our view, an important addition to multilevel software would be the capacity to produce the kinds of output described above and/ or to produce output which could be read directly into graphing packages.

Reference

Fitz-Gibbon, C. (1989). *Multilevel modelling in an indicator system.* Paper presented to the ESRC Conference on Multilevel Methods in Educational Research. University of Edinburgh.

CONTRIBUTORS

Thanks very much to the people who provided articles for this issue.

Tony Bryk

Faculty of Education
University of Chicago

Carol Fitz-Gibbon & Peter Tymms

Faculty of Education
University of Newcastle-upon-Tyne

Jan de Leeuw, Kyung-Sung Kim, & Ita Kreft

UCLA, Los Angeles, California

Steve Raudenbush

Faculty of Education,
Michigan State University