MULTILEVEL MODELLING NEWSLETTER

Centre for Multilevel Modelling

Bedford Group for Lifecourse and Statistical Studies
Institute of Education, University of London
20 Bedford Way, London WC1H 0AL, ENGLAND
Web site: http://multilevel.ioe.ac.uk/
Enquiries about newsletter to Ian Plewis
E-mail: i.plewis@ioe.ac.uk
Tel: +44 (0) 20 7612 6688
Fax: +44 (0) 20 7612 6572

Vol. 15 No. 2 **December, 2003**

Forthcoming Workshops

29-31 March 2004. A three-day introductory workshop to multilevel modelling using *MLwiN* will take place at the University of Bristol. Enquiries to Teresa Nurser, Graduate School of Education, 35 Berkeley Square, Bristol BS8 1HJ, United Kingdom. Tel: +44 (0) 117 331 4289 Fax: +44 (0) 117 925 5412 Email: teresa.nurser@bristol.ac.uk.

If you plan to run any workshops using *MLwiN*, please notify Amy Burch and she will advertise these workshops on the multilevel web site.

Beta version of MLwiN 2.0

The beta version of *MLwiN* version 2.0, available from http://multilevel.ioe.ac.uk/beta/index.html is now in the final stages of testing. It will go on full release early in the New Year. There are several enhancements to *MLwiN* version 2.0 from version 1.10. MCMC estimation has been greatly enhanced; new features include:

1. The DIC diagnostic (Spiegelhalter et al., 2002. Bayesian measures of model complexity and fit (with discussion). Journal of the Royal Statistical Society, B, (64): 583-640).
2. Multilevel factor analysis models with multiple (correlated or uncorrelated) factors at each level.
3. Multicategory ordered and unordered response models.
4. Multivariate mixtures of continuous and binary responses.
5. Complex level one models.
6. Multivariate models, including those with missing responses.
7. Adjustments for measurement error in predictors.
8. Cross classified models.
9. Multiple membership models.
10. Autoregressive structures at level one.
11. Spatial data models.
12. An interface with the WINBUGS software package.

E·S·R·C ECONOMIC & SOCIAL RESEARCH COUNCIL

Changes to the user interface have been made to improve ease of use for the following features:

1. The multivariate window has been removed and multivariate models are now set up using the 'responses' button on the equations window.
2. There is now an 'add term' button on the equations window for adding continuous variables, categorical variables and interactions to a model.
3. A 'notation' button has been added to the equations window that allows switching between different notational representations.
4. A separate MCMC menu has been added.
5. Much improved interface for the specification of ordered and unordered categorical response models.
6. Single level models can be specified using standard notation.

MLwiN 2.0 documentation comes in two volumes. The first volume, "A User Guide to *MLwiN*" now has an extra introductory chapter to help users better understand the key differences between single and multilevel models; also chapters on modelling for ordered and unordered categorical responses have been added. The second volume, "MCMC estimation in *MLwiN*", demonstrates both the theory and practice of fitting MCMC models in *MLwiN*. Both volumes are approximately 300 pages long.

We hope you enjoy using *MLwiN* version 2.0.

Project Team, Centre for Multilevel Modelling, Institute of Education, University of London.

Bootstrapping the Effects of Measurement Errors

Dougal Hutchison, Jo Morrison and Rachel Felgate
National Foundation for Educational Research
d.hutchison@nfer.ac.uk

Introduction

This paper describes a method of allowing for measurement error in multilevel regression by using bootstrapping procedures. We illustrate this method on a simple two level model with two explanatory variables, one of which is measured with error.

A two level linear model for y_{ij} and true or 'latent' values x_{1ij}, x_{2ij}, where *i,j* refers to the i^{th} level-1 unit within the j^{th} level-2 unit, is given by

$$y_{ij} = \beta_0 + \beta_1 x_{1ij} + \beta_2 x_{2ij} + u_j + e_{ij} \qquad (1)$$

$$Cov(u_{j'}, u_j) = Cov(e_{ij}, e_{i'j'}) = Cov(u_j, e_{ij}) = 0$$
$, i' \neq i, j' \neq j$

$E(u_j) = E(e_{ij}) = 0$;

$\text{var}(u_j) = \sigma_u^2; \text{var}(e_{ij}) = \sigma_e^2$.

The 'true' or latent values x_{1ij} in (1) are observed with measurement error m_{ij} giving observed values X_{1ij} where

$$X_{1ij} = x_{1ij} + m_{ij}.$$

$x_{2j} = X_{2j}$ is considered to be measured without error in this example.

$$Cov(m_{ij}, m_{i'j}) = Cov(m_{ij}, e_{ij}) = Cov(m_{ij}, u_j) = 0$$

$E(m_{ij}) = 0$; $\text{var}(m_{ij}) = \sigma_m^2$

m_{ij} is independent of x_{1ij}, x_{2ij}.

When y_{ij} is regressed on the observed X_{1ij} and X_{2ij} then

$$y_{ij} = \gamma_1 X_{1ij} + \gamma_2 X_{2ij} + u'_j + e'_{ij} \quad ,$$

and the resulting estimates γ_1, γ_2 are biased for β_1, β_2.

These are standard assumptions in this type of work, as defined by Goldstein (2003): for examples of other assumptions, see Fuller (1987). Theory has been developed in this area mainly for the situation where errors are normally distributed, but also for multinomial misclassification (Fuller, 1987; Goldstein, 2003). More general models have not been widely considered, though Woodhouse (1996) and Browne et al. (2001) have looked at the effect of errors in variables on slopes.

The use of the bootstrap to correct for biases

There are two main uses for bootstrapping techniques: estimation of sampling distributions and standard errors; and correction of biases. We discuss the use of the bootstrap for sampling distributions and standard errors in a later section. Bootstrapping techniques can be used to correct for biases in estimation techniques, using an iterative procedure. We illustrate on model 1 above.

Stage 1

Regress y on observed $\underline{X} = (X_1, X_2)$ to obtain $\underline{\hat{\gamma}_0} = (\hat{\gamma}_{10}, \hat{\gamma}_{20})$.

Stage 2

Simulate $\hat{y}$, using $\underline{\hat{\beta}_0} = (\hat{\beta}_{10}, \hat{\beta}_{20}) = \underline{\hat{\gamma}_0}$, and an estimated value of $\underline{\hat{x}} = (\hat{x}_1, \hat{x}_2)$ to be determined.

Add measurement error to $\underline{\hat{x}}$ to give $\underline{\hat{X}}$.

Regress simulated $\hat{y}$ on $\underline{\hat{X}}$ to obtain $\underline{\hat{\gamma}_{1b}}$.

Do this a large number B of times, to obtain $\underline{\hat{\gamma}_1}$ the mean of the $\underline{\hat{\gamma}_{1b}}$

Estimate bias $\underline{\hat{b}_1}$ by $\underline{\hat{b}_1} = (\underline{\hat{\gamma}_1} - \underline{\hat{\gamma}_0})$.

Estimate $\underline{\hat{\beta}_1}$ as $\underline{\hat{\beta}_0} - \underline{\hat{b}_1}$

Stage 3

Repeat stage 2, starting at $\underline{\hat{\beta}_1}$.

Keep iterating until process converges.

Example: Simulated data

Generating the simulated dataset

A dataset of 5000 cases of 200 level-2 units with 25 level-1 units in each, was constructed according to model (1) with $\beta_1 = \beta_2 = 1$ and x_1, x_2 correlated $\rho_{x_1 x_2}$. $\sigma^2_{b1}, \sigma^2_{b2}$ are the between-group variance and $\sigma^2_{w1}, \sigma^2_{w2}$ the within-group variance of x_1, x_2, respectively. $rho_1 = \sigma^2_{w1} / (\sigma^2_{w1} + \sigma^2_{b1})$ and $rho_2 = \sigma^2_{w2} / (\sigma^2_{w2} + \sigma^2_{b2})$ are the intracluster correlations for x_1, x_2 respectively.

Analyses were carried out for a range of values for rho_1, rho_2 for level 1 x-reliability ρ_1 (Woodhouse et al., 1996) and for $\rho_{x_1 x_2}$. The aim was, given the observed covariance matrix C_X, to produce two variables, $\hat{x}_1, \hat{x}_2$ which have the error-corrected covariance matrix c_x. Note that the variables individually do not need to be only linear transformations of the corresponding observed variables.

Estimating the coefficients

We carried out the procedure separately for levels one and two.

We estimated

$$\tilde{C}_X = \begin{bmatrix} \sigma^2_{\tilde{X}_1} & \sigma_{\tilde{X}_1 \tilde{X}_2} \\ \sigma_{\tilde{X}_1 \tilde{X}_2} & \sigma^2_{\tilde{X}_2} \end{bmatrix}$$

at level one. The program *MLwiN* (Rasbash et al., 2000) was used for this. Since we were not looking at level 2 error, only the level 1 covariance matrix $\tilde{C}_X$ had to be corrected. The matrix $\tilde{c}_x = \begin{bmatrix} \rho_1 \sigma^2_{\tilde{X}_1} & \sigma_{\tilde{X}_1 \tilde{X}_2} \\ \sigma_{\tilde{X}_1 \tilde{X}_2} & \sigma^2_{\tilde{X}_2} \end{bmatrix}$ was taken as the target for the simulation.

The command MRAN in *MLwiN* is now used on matrix $\tilde{c}_x$ to create the level 1 part $(\hat{x}_1, \hat{x}_2)$ of the simulation data set. The resulting estimated matrix $\hat{\hat{c}}_x$ is only equal to the required quantity in expectation and is subject to sampling fluctuation. The variables are transformed to make the estimated $\hat{\tilde{c}}_x$ precisely equal to $\tilde{c}_x$. This is done by multiplying the data $(\hat{x}_1, \hat{x}_2)$ by $\tilde{M}\tilde{L}^{-1}$, where $\tilde{L}$ is the Choleski decomposition of $\hat{\hat{c}}_x$ and $\tilde{M}$ is the Choleski decomposition of $\tilde{c}_x$. Variables will be $\hat{\tilde{x}}_1, \hat{\tilde{x}}_2$.

Similarly we created the level 2 data, $\hat{\bar{x}}_1, \hat{\bar{x}}_2$.

Add $\hat{\tilde{x}}_1, \hat{\bar{x}}_1$ and $\hat{\tilde{x}}_2, \hat{\bar{x}}_2$ to form the total $\hat{x}_1, \hat{x}_2$. Measurement error was added to $\hat{x}_1$ to give $\hat{X}_1, \hat{X}_2$.

The bootstrap procedure described in the previous section was used to estimate β_1, β_2. Sets of analyses using 2000 replications were carried out. Ten iterations were used to investigate the convergence of the procedure.

Figure 1 shows one example of the convergence of the process for a correlation $\rho(x_1,x_2)=0.4$, $rho_1 = rho_2 = 0.2$, and $\rho_1 = 0.8$. The true value for both β_1 and β_2 coefficients is 1.0. The level-2 variation is set at $\sigma_u^2 = 4$ and the level-1 variation at $\sigma_e^2 = 25$. It can be seen that at the first iteration, the coefficient of x_1 is below the 'true' value, and that of x_2 is above. From about iteration 4, values stabilise to values slightly below 1.0. However, one would not expect the process to converge precisely to this value because of the random quantities introduced in generating the original data. We consider whether this is a reasonable expectation below.

Figure 1. Convergence of x_1, x_2 coefficients, N = 5000, corr = 0.8

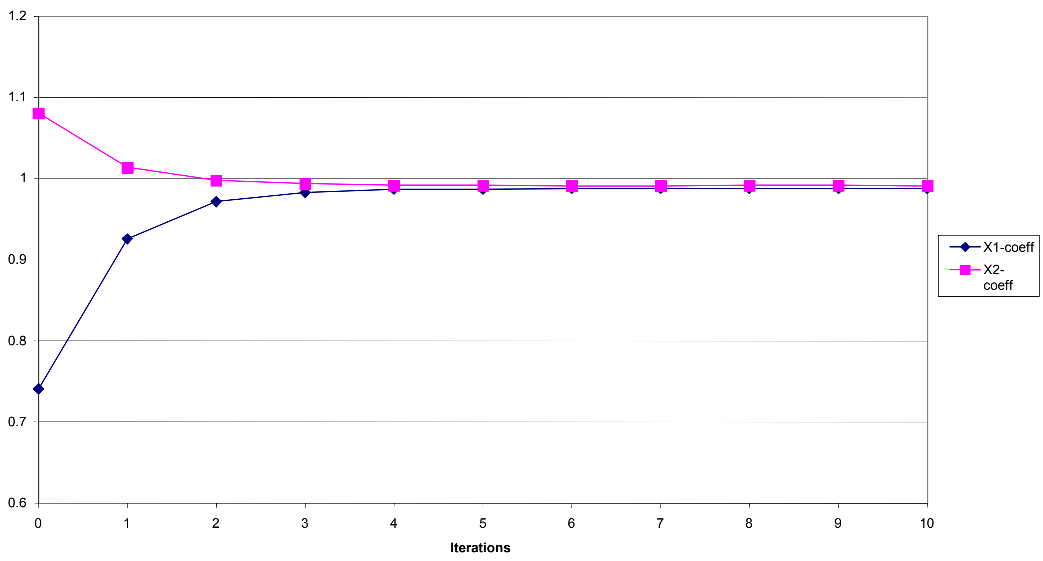

Estimating standard errors by the bootstrap

We estimated the standard errors as follows. We produce a bootstrap replication of the original sample, and carry out the procedure on the resample. This was replicated many times: i.e. we bootstrapped the bootstrap. There are two types of possibility for the 'outer' bootstrap:

a) Whole case resampling (resampling of level-2 units is recommended).

b) Residuals resampling. This can be either parametric or non-parametric (Carpenter et al., 1999; Hutchison, 1999).

Here we present results using whole case resampling. This means that we have three levels of looping. Some kind of convention on nomenclature is obviously necessary.

a) Resamples from the original (actual or generated) data set. The literature suggests that it would be necessary to take of the order of 2000 resamples to get reliable estimates of the percentiles, confidence intervals, etc.

b) Iterations to convergence reduce the (expected) bias in the estimate. Preliminary investigation has shown that the bias is effectively removed in four iterations for the simple model considered here and that the variation remaining is in the nature of oscillation rather than bias-correction. (The required number of iterations may be larger on more complex problems.)

c) We have found that 2000 replications within each iteration does a good job in reducing oscillation, though obviously this depends on the degree of precision required.

This suggests a total of 16,000,000 bootstrap analyses (per problem).

However, if we have a large number of equivalent resamples, then this should provide a large set of comparable estimates. If we know when the iterations have converged, then we can focus attention on the replications within each converged analysis. This is a three level structure (replications within iterations within resamples). However, in the middle level, the iterations are not exchangeable if the procedure has not converged. Consequently we focus on a single iteration at or beyond the convergence stage. This gives a two level model (replications within converged-iteration-within-resample). We can feed these results into a multilevel model.

In fact we shouldn't need a very large number of replications within each iteration, since the multilevel structure means that we can handle a degree of variation. With a smaller number of replications, this would have the drawback that we wouldn't necessarily know that the iterations have converged. We would need to have some kind of prior idea of the total number of iterations under a wide range of resamples. Alternatively one could run a larger number of iterations than strictly necessary, and examine the convergence behaviour.

We then have a two level model for β_{bcd}, the d^{th} replication within the c^{th} iteration of the b^{th} resample.

$$\beta_{bcd} = \beta_0 + \beta_{bc} + e_{bcd}$$

$$V[\beta_{bc}] = \sigma_\beta^2, V[e_{bcd}] = \sigma_e^2$$

$sqrt\{V[\beta_{bc}]\}$ can be taken as an estimate of the standard error of the estimate of β_0. A normal approximation to confidence intervals could be taken from the highest level variation. For a more general result, shrunken top level residuals could be partially re-inflated to give the appropriate variance. Then the percentiles of these partially reinflated residuals could be used to give percentiles of the distributions.

Example of implementation

We created a basic data set of 5000 cases, as above: 2000 resamples were drawn in each analysis, 10 iterations and 10 replications.

Results of a set of simulations are shown in Table 1.

Table 1. Results of bootstrap estimation of standard error (5000 resamples)			
Quantity	Estimate	Standard Error	Generating value
β_1	0.99	0.015	1
β_2	0.99	0.011	1
σ_u^2	4.54	.76	4
σ_e^2	25.93	1.08	25

The standard error of the x_1 coefficient is estimated as 0.015, and that for the x_2 coefficient is rather smaller at 0.011. This would be expected, since there is no measurement error in x_2.

Comparing the estimated and generating values of the coefficients, we see that in every case the difference is less than the estimated standard error. Further investigations (below) considered whether the procedure is unbiased, by importing yet another level of bootstrapping, so that the entire procedure is replicated on a number of data sets generated according to the model.

Consistency of estimates

Obviously we cannot show that a method is unbiased or consistent using a single example. The values of the estimates are dependent not simply on the true model and on the method, but also on the precise values of the measurement error added to give the observed values. Preliminary investigations, however, seemed to show that larger samples gave results that were closer to the true generating values, but still slightly biased. Accordingly, a series of simulations were carried out for 50, 100 and 200 level–2 units, and 25, 50, 100 and 200 level-1 units within each, to observe the behaviour of the estimated parameters for increasing sample size. 10,000 simulations were carried out for each such combination, each with 10 iterations and 5 replications. The same values of coefficients β_{x1}, β_{x2}, σ_u^2, σ_e^2 were used as before.

For all of these combinations of level-2 and level-1 numbers, even the smallest, the resultant coefficient estimates were extremely close to their generating values. The relative bias $= \frac{mean(bias)}{generating\ value}$ was of the order of .002 for β_{x1}, less than .001 for β_{x2}, .005 for σ_u^2, and .004 for σ_e^2. It was also found that the mean value of the four estimated coefficients β_{x1}, β_{x2} σ_u^2, σ_e^2 tended towards their true values with increasing numbers of level-2 and level-1 units. They did this, moreover, in a manner suggesting a consistent asymptotic convergence. (The results of the analyses described here, in graphical form, may be obtained from the first author on request.)

Computational efficiency

Even with the short cuts outlined above, the process is highly computationally intensive. For a single analysis, the bootstrap error correction takes of the order of 3 hours on a Dell Optiplex 300, and the standard error estimate nearer 20. It is obviously necessary to investigate how to cut these times, especially the standard error procedure, as much as possible if this approach is going to be widely used. The procedure may be compared with the MCMC procedure described by Browne et al (2001). While both methods appear to work well within the range of applications considered, the MCMC procedures appear to be considerably faster.

Accordingly, we investigated the possibility of reducing:

a) the number of iterations

b) the number of replications within each iteration.

Figure 2 shows the convergence of the standard error estimate with the number of iterations (2000 resamples, 10 replications). It can be seen that there is still a degree of oscillation after even 10 iterations, but for practical purposes 3 or 4 iterations would probably be sufficient.

Figure 2. Estimated standard error of β_x versus number of iterations

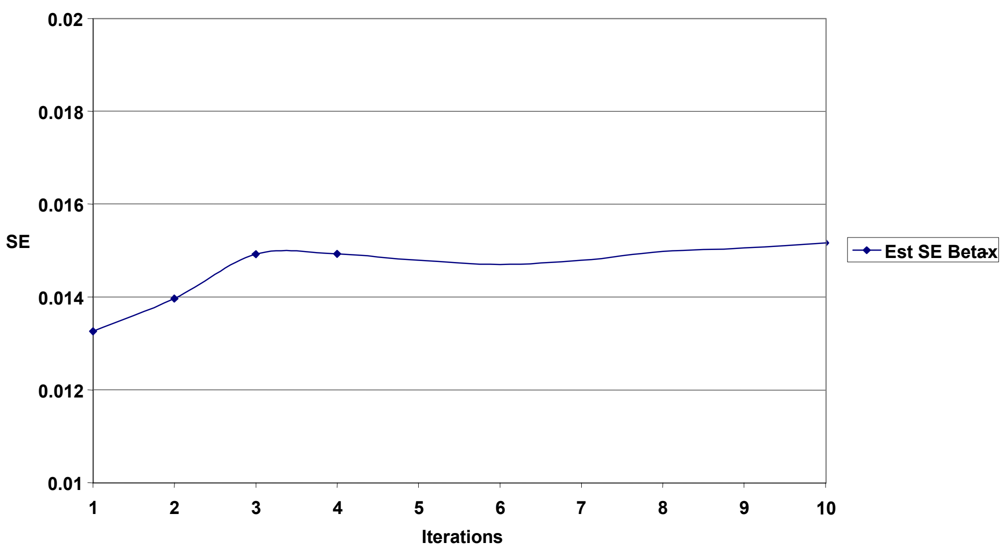

Similarly Figure 3 shows the extent to which the results stabilise compared with the number of replications. The results have not completely settled down by 10 iterations, but, as in Figure 2, the bias, as opposed to the oscillation, has probably been removed after four. The line for 10 replications is consistently the highest, suggesting that more replications would be required to get the estimate of standard error close to the asymptotic value to the third decimal place. The method as applied lets each individual analysis iterate to convergence: further speeding up of the entire process might be possible if we do a fixed small number of iterations, say one or two.

Figure 3. Convergence of SE estimates

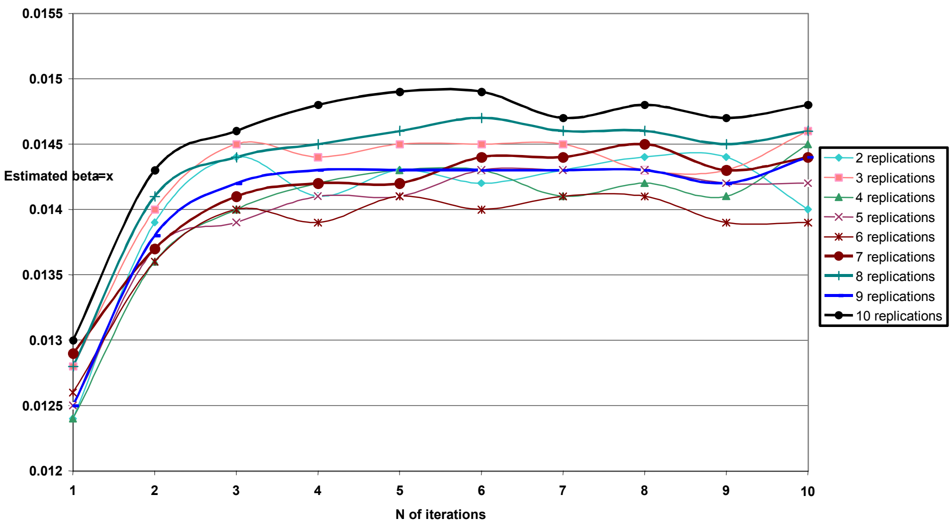

Conclusions

The results presented in this paper have provided an example of a potentially generalisable procedure for estimating the sampling behaviour of multilevel regression models under measurement error. Further investigations will be required to see whether the short cuts described would work on other models and whether further speeding up is possible. It would also be valuable to compare the results with those obtained by the MCMC procedure described by Browne et al. (2001) to see whether both methods work comparably well on a wide range of models.

References

Browne, W. J., Goldstein, H., Woodhouse, G., and Yang, M. (2001). An MCMC algorithm for adjusting for errors in variables in random slopes models. *Multilevel Modelling Newsletter*, **13**: 1.

Carpenter, J., Goldstein, H., and Rasbash, J. (1999). A nonparametric bootstrap for multilevel models. *Multilevel Modelling Newsletter*, **12**: 2.

Fuller, W.A. (1987). *Measurement Error Models*. Wiley, New York.

Goldstein, H. (2003). *Multilevel Statistical Models, 3rd Edition*. Kendall's Library of Statistics, 3. Edward Arnold, London.

Hutchison, D. (1999). *The effect of group-level influences on pupils' progress in reading*: a doctoral thesis submitted in partial fulfilment of the requirements for the award of Doctor of Philosophy of the University of London.

Kuk, A. (1995). Asymptotically unbiased estimation in generalized linear models with random effects. *Journal of the Royal Statistical Society, B*, **57**: 395-407.

Rasbash, J., Browne, W. J., Healy, M., Cameron, B, and Charlton, C. (2000). *MLwiN* v1.10.006.

Woodhouse, G., (1996). *Errors in variables in multilevel models*. PhD dissertation, Institute of Education: University of London.

Woodhouse, G., Yang, M., Goldstein, H., Rasbash, J., and Pan, H. Q. (1996). Adjusting for measurement error in multilevel analysis. *Journal of the Royal Statistical Society, A*, **159**:, 201-212.

Acknowledgements

The authors would like to express their appreciation for helpful comments from members of the Centre for Multilevel Modelling at the Institute of Education, University of London, in particular William Browne, Harvey Goldstein and Ian Plewis. This paper is a development of one presented at the 2000 American Statistical Association Conference.

Review of 'Applied Longitudinal Data Analysis'. Singer, J. D. and Willett, J. B. (2003). New York: OUP ISBN: 0-19-515296-4, pp. 644.

Ian Plewis

Institute of Education, University of London

Anyone teaching courses on the analysis of repeated measures data or on the analysis of survival data in the social sciences will find this book extremely helpful. It is thorough, well written and the associated web site (http://www.oup-usa.org/alda) provides useful back-up material in the form of datasets used in the book, and example programs using a number of multilevel and statistical packages.

The book is divided into two substantial parts, each of book length. The first part deals with the multilevel approach to the analysis of growth and change, the second with techniques and models for 'time to event' data, firstly in discrete time and then in continuous time. The two parts are essentially self-contained and this is rather a pity because an opportunity is lost to show how repeated episodes in event history data can be handled within a multilevel framework. Instead, the authors restrict themselves to single events and thus to single level logistic regression (proportional odds) models and Cox proportional hazards models.

The great strength of the first part of the book is its painstaking description of the specification and interpretation of models that relate a single response to time or age. The authors' approach will appeal particularly to quantitative social scientists who have had just a little exposure to statistical modelling; those students looking for a more rigorous statistical approach might be put off by the high ratio of words to symbols. The reader is taken slowly through ideas about age and time-related change to (in Chapter 3) a model with a random intercept and a random slope, ways of accounting for these variances surface in Chapter 4, different ways of dealing with time are covered in Chapter 5, extensions to discontinuous and non-linear change in Chapter 6, implications of different assumptions about change in terms of the underlying covariance structure in Chapter 7, and finally there is a chapter discussing latent growth modelling from a structural equations modelling perspective.

The authors' approach – admirable as it is – does mean that a number of topics are not covered. So there is nothing on models with more than two levels, on multivariate models with more than one outcome, on spline functions in Chapter 5, and on models with binary or categorical outcomes. There is also perhaps less than might have been expected on the analysis of residuals, outliers and influence. Finally – and this is an important omission – there is essentially no discussion of the contrast between the unconditional or time-related approach to the analysis of change favoured by the authors, and the conditional or regression approach in which earlier values of the outcome are used as explanatory variables. Also, although time-varying predictors are covered in Chapter 5, there is no discussion of models in which y_{ti} (for example, political values) are related to x_{ti} (income, say) in a regression framework favoured by econometricians when analysing panel data.

And so, to sum up, this is a fine book but one which is just a little limited in its coverage, especially for more experienced longitudinal researchers.

Some Recent Publications Using Multilevel Models

Adams, G, Gulliford, M. C., Ukoumunne, O. C., Chinn, S., and Campbell, M. J. (2003). Geographical and organisational variation in the structure of primary care services: implications for study design. *Journal of Health Services Research and Policy*, **8**: 87-93.

Bolaños M. V., Ocaña-Riola, R., and Prados-Torres, A (2002). Variations in health services utilization by primary care patients. *Health Services Management Research*, **15**: 116-125.

Carpenter, J., Goldstein, H., and Rasbash, J. (2003). A novel bootstrap procedure for assessing the relationship between class size and achievement. *Journal of the Royal Statistical Society, Series C*, **52**: 431-443.

Fielding, A., Yang, M. and Goldstein, H. (2003). Multilevel ordinal models for examination grades. *Statistical Modelling*, **3**: 127-153.

Goldstein, H. and Browne, W. (2003). Multilevel factor analysis models for continuous and discrete data. In Olivares, A. (ed). *Advanced Psychometrics. A Festschrift to Roderick P. McDonald*. Lawrence Erlbaum, Mahwah, NJ.

Goldstein, H. (2003). *Multilevel Statistical Models*. 3rd edition. Edward Arnold, London.

Goldstein, H. and Noden, P. (2003). Modelling social segregation. *Oxford Review of Education*, **29**: 225-237.

*Gorard, S. (2003). What is multilevel modelling for? *British Journal of Educational Studies*, **51** (1): 46-63.

*Gorard, S. (2003). In defence of a middle way: A reply to Plewis and Fielding. *British Journal of Educational Studies*, **51** (4): 420-426.

Grilli, L., and Rampichini, C. (2003). Alternative specifications of multivariate multilevel probit ordinal response models. *Journal of Educational and Behavioral Statistics*, **28**: 31-44.

Gulliford, M. C., Mahabir, D., and Rocke, B. (2003). Food insecurity, food choices and body mass index in adults: nutrition transition in Trinidad and Tobago. *International Journal of Epidemiology*, **32**: 508-515.

Gulliford, M. C., and Ukoumunne, O. C. (2001). Determinants of glycated haemoglobin in the general population: associations with diet, alcohol and cigarette smoking. *European Journal of Clinical Nutrition*, **55**: 615-623.

Heo, M., Faith, M. S., Mott, J. W., Gorman, B. S., et al. (2003). Hierarchical linear models for the development of growth curves: an example with body mass index in overweight/obese adults. *Statistics in Medicine*, **22**: 1911-1942.

Hernández-Lloreda, M. V., Colmenares, F., and Martínez-Arias, R. (2003). Application of hierarchical linear modelling to the study of trajectories of behavioural development. *Animal Behaviour*, **65**: 607-613.

Jack, R., Gulliford, M., Ferguson, J., and Moller, H. (2003). Geographical inequalities in lung cancer management and survival in South East England: Evidence of variation in access to oncology services? *British Journal of Cancer*, **88**:1025-1031.

Leyland, A. H., and Groenewegen, P. P. (2003). Multilevel modelling and public health policy. *Scandinavian Journal of Public Health*, **31**: 267-274.

Lindqvist, P. G., Marsal, K., and Pirhonen, J. (2003). Thermal response to submaximal exercise before, during, and after pregnancy: a longitudinal study. *Journal of Maternal-Fetal and Neonatal Medicine*, **13**: 152-156.

Pirhonen, J., Lindqvist, P. G., and Marsal, K. (2003). A longitudinal study of maternal oxygen saturation during short-term submaximal exercise. *Clinical Physiology*, **23**: 37-41.

*Plewis, I., and Fielding, A. (2003). What is multilevel modelling for? A critical response to Gorard. *British Journal of Educational Studies*, **51** (4): 408-419.

Singer, J. D., and Willett, J. B. (2003). *Applied Longitudinal Data Analysis*. OUP, New York.

Zaslavsky A. M., Landon B. E., Beaulieu, N. D., and Cleary P. D. (2000). How consumer assessments of Medicare managed care vary within and among markets. *Inquiry*, **37** (2): 146-161.

Zaslavsky, A. M., and Buntin M. J. B. (2002). Using survey measures to assess risk selection among Medicare managed care plans. *Inquiry*, **39** (2): 138-151.

Zaslavsky, A. M., and Cleary, P. D. (2002). Dimensions of plan performance for sick and healthy members on the Consumer Assessments of Health Plans Study 2.0 survey. *Medical Care*, **40** (10): 951-964.

*These could be read together.

Please send us your new publications in multilevel modelling for inclusion in this section in future issues.

MLwiN Clinics in London

Wednesday 4 February 2004
Wednesday 3 March 2004

at

Centre for Multilevel Modelling
11 Woburn Square, London WC1H 0NS

Contact *MLwiN* Technical Support for appointments
Tel: +44 (0) 20 7612 6688
mlwin.support@ioe.ac.uk

Future clinic dates will be announced at:
http://multilevel.ioe.ac.uk/support/clinics.html