INFORMETRICS 89/90
L. Egghe and R. Rousseau (Editors)

LOGLINEAR MODEL OF LIBRARY ACQUISITIONS

Abraham BOOKSTEIN*

University of Chicago
Center for Information Studies

Edward O'NEILL, Martin DILLON, David STEPHENS

OCLC - Office of Research

Abstract

Libraries vary greatly in size, but it is also widely believed that they vary in personality as well, the latter being manifested in their book buying patterns. But even if libraries had identical acquisitions profiles, the combination of size variation and random effects could give the impression of variation in intention. We explore here the possibility that, for certain classes of books, this is the case. We define a probabilistic model of library acquisition, indicate how the model's parameters can be estimated, and test the model on a number of libraries for their collections in calculus.

1. INTRODUCTION

Libraries and information systems generate masses of data. To understand and manage better the activities taking place in these systems, it is necessary that we learn how to exploit the data that are available to us. But data collection is expensive, and simply collecting and tabulating large quantities of data has not always proved to be an effective management aid. An alternative strategy is to focus on the most important activities of the system, conceptualizing and modeling those aspects of these activities that are most critical. Such a modeling process at once suggests what data to collect and points to how these data may be used. Because of the character of information use, often this will involve creating statistical models, which will then have to be analyzed and tested. In this paper we describe an effort to study the problem of book acquisition by libraries. We shall describe in some detail the models and testing procedures we used as we believe these to be more widely applicable to problems occurring in the information field. These statistical techniques are variously known as "loglinear data analysis", "logistic regression", or "quantal response modeling".

The need for sophisticated methods are a consequence of the nature of library data. Much of library data shares two characteristics that limit the applicability of the most heavily used data-analytical methods. These data are :
a) Multivariate, and
b) Categorical.

---

* The first author participated in this project while a visiting scholar at OCLC.

That is, much of the phenomena which occur in information systems intrinsically involve a number of variables that strongly interact with one another, and many of these variables are categorical in nature - that is, more like "sex", "subject of search", "satisfied or not satisfied", rather than like "age", "speed" or "distance", although, of course, both types occur.

Most commonly,. Chi-square techniques are used when treating categorical data. Pairs of variables are displayed in a contingency table, collapsing over other variables. For many kinds of experimentally derived data, these methods can be illuminating. Herbert Goldhor [1], for example, studied in this manner the effect of various methods of displaying books on whether or not they were checked out. The problem with approaching multivariate data in this way is that it often obscures relations among variables that might be interesting. In some cases these analyses can be seriously misleading. When a number of variables act cooperatively to produce a result, it is necessary to take their influences into account simultaneously in order to understand the subtle interactions among these variables, and so that we do not force a variable to act as a substitute for another variable which has been omitted.

When we have a quantitative dependent variable, regression models have very nicely satisfied this need [2]. We can, for example, study how *amount* of public library use is influenced by a range of contributory characteristics such as education level, income, age, sex, distance from library, level of social integration, etc. [3,4]. When the dependent variable is categorical, however, such an approach is no longer ideal (although attempts have been made to adapt regression methods to this context) [5].

The loglinear methods have been developed to bridge the gap between contingency table analysis and regression methods [6,7,8]. They are similar to regression analysis in that they are driven by models, often linear models, that, on substantive grounds, are expected to describe the data. However, since the dependent variable is categorical, it is the *probability* (or some function of the probability) that the dependent variable takes a particular value that is modeled rather than the value of the variable itself. Symbolically, if y, a two-valued categorical variable, is believed to be related to the independent variables $\{x_i\}$, we might test the model :

$$f(p) = \Sigma\ b_i x_i \ ,$$

where p is the probability that y takes one of its two values, and f is an appropriate function. The b's are parameters to be evaluated in the course of the analysis.

If the equation can be inverted, we can write :

$$p = g(\Sigma\ b_i x_i) \ .$$

Maximum likelihood techniques can be used to estimate the parameters $\{b_i\}$, to establish confidence intervals, and to assess the validity of the model itself. In such analysis, the "logistic" transformation $g(x) = \frac{1}{1 + \exp(-x)}$ has been found particularly attractive - hence the name of some of these techniques. These models are most frequently used for contingency table analyses, but can also be used to model response and choice data [9].

In the research being reported, data available at OCLC describing the book purchases of a group of libraries were analyzed by means of a special case of this model. Programs to carry out this analysis were written in the Gauss programming language, a very powerful PC-based language that is particularly strong in the matrix manipulation routines needed for complex statistical analysis. These are described in detail in reference [10].

In trying to discern patterns in how libraries select the books they purchase, we found particularly valuable the concept of a *peer group*, a group of libraries that are similar enough so that collectively the group as a whole can serve as a reference to guide each of its members in book selection. That is, a peer group is a group of libraries that may vary in size but share a book buying "personality". The validity of this concept is tested below. But first we must state more precisely what we mean by a peer group, and define a model that describes the consequences of peer groups existing.

The approach we shall take is based on a model of choice behavior, in which a peer group is defined as a group of libraries within which, except for *statistical fluctuation, the size of the library and the popularity of a book* among the peers are the only factors governing book purchasing decisions. In accordance with this model, each library is described by a single size (or purchasing strength) parameter, b, and each book by a restrictiveness parameter, d, measured on the same scale.

We will be interested in whether this model, described in detail below, does, at least approximately, describe book purchases of libraries that on subjective grounds seem to form a peer group.

We can now look more closely at the model itself.

2. CHOICE MODEL

2.1. Introduction

The specific model we studied here describes a "choice" situation in which only "propensity to choose" and object "desirability" influenced the model. That is, the subjects making the choice are assumed to have no individuating personality characteristics other than that of proclivity to select times : knowing how much they acquire tells us all that can be known about the subjects. In the example of the library collection development problem described above, the subjects were a group of libraries subjectively chosen as peers, though differing in size, and the entities chosen were books in specific subject categories : calculus and botany.

The model is defined as follows : Each subject has a "propensity to acquire" parameter : b_i for subject i, and each entity has a "difficulty" parameter : d_j for object j. The probability that subject-i will select item-j depends only on the degree to which the acquisition propensity exceeds the resistance of the item : $b_i - d_j$.

The specific function used was :

$$P_{ij} = \text{Prob}\ \{X_{ij}|b_i,d_j\} = \frac{\exp\ ((b_i - d_j)X_{ij})}{1 + \exp(b_i - d_j)}\ , \quad \text{where}$$

$$X_{ij} = \begin{cases} 1 & \text{if subject i acquires item j} \\ 0 & \text{otherwise} \end{cases}$$

If $\{b_i\}$ and $\{d_j\}$ were known, L, the probability of any given matrix of choices, X, would be given by ΠP_{ij}. Maximum likelihood estimation proceeds by finding thos values of b_i and d_j for which L is as large as possible. Although the *maximum likelihood equations have no closed-form solutions, simple iterative* procedures exist that permit numerical solution. Programs have been written in the Gauss-programming language to carry out this analysis.

2.2. Implications of model for collection analysis

Rather little is known about how libraries select material. It is widely assumed that each library has its own personality, and that this personality expresses itself as a library selects its books. An alternative hypothesis is that one can divide libraries into groups of peers, and that within such a group, only size influences what a library will buy (at least within specific classes of books). The notion of a "peer group" is a fundamental one in trying to discuss systematically how a library selects materials. This raises two questions :

1) Can the notion of a peer group be made precise?

2) How can the existence of a peer group structure assist libraries in selecting material?

The models discussed above offer an approach toward responding to both questions. We suggest that a group of libraries be considered a peer group, at least within a subject or format domain, if their book acquisitions can be described by the above baseline choice model. Maximum likelihood estimation allows us both to estimate the values of the model parameters and to assess the validity of the model. We suggest that if the model fits reasonably well, the group should be considered a peer group. Thus the model offers not only a means of analyzing selection data, but plays a central conceptual role as well.

But also, should the model be found to fit reasonably well, it would be interesting to examine discrepancies from the model's predictions, to see what these tell us about the models or about the libraries. For example, if the model fails for a particular library, but describes other libraries in the group, the breakdown can be interpreted as an indication that the library does not in fact belong in this group. Alternatively, especially if the breakdown can be traced to decisions on a small number of items, it could indicate an oversight on the part of the library. That is, should the model be effective as a description of library purchases, it could serve as a guide to libraries in that the model would allow us to note, for a library's consideration, that it has not purchased an item that it would have been expected to purchase on the basis of the model. (It would also indicate that a library might be overpurchasing certain categories of books.) Of course, book purchasing decisions are the responsibility of each library. The value of such tools is that it could bring to a library's attention books that the library might wish to have purchased but which may have been overlooked. In this sense, our analysis can be useful as a collection development tool.

2.3. Results

The choice model was tested on data collected earlier by Sanders et al. [11]. The data consisted of choices made by a group of 11 large, Mid-West research libraries of books on calculus. The data set was relatively small, but carefully collected and verified, so that we can have confidence in the results of our analysis.

The results of our analysis are presented in the Appendix. We first note that for the model we are using, "score" (for libraries : the total number of books selected; for books : the total number of libraries selecting a book) is a sufficient statistic. This means that once we know the score for an item (library) or subject (book) we have all the information that the model can use. A consequence of this observation is that all objects with the same score will be assigned the same value for the parameters describing them. For this reason, the program collects objects into "score groups", that is, objects having identical scores, before beginning the analysis. For our data, each library constitutes a separate score group; the parameter describing a library is denoted by Dxxx, where xxx is the score of that library (i.e., the number of items in the set of books being studied that it acquired). Similarly, books are divided into groups, with parameters denoted by Bx. Since there are only 11 libraries in our database, x could take only the values 0 (acquired by no library) to 11 (acquired by each library).

The appendix displays the results of the analysis for the complete data set. The heading includes information indicating the fit of the model : since the statistical test used is a chi-square test, we need the degrees of freedom (here 110 df) and the value of the chi-square statistic. The latter is computed using two methods : these are given as the (preferred) G-square value and the alternative Pearson Chi-square value. In this data set, as is usually the case, the two values are close to each other. The heading also shows the significance level (p-level) of the chi-square statistic. For the complete data set, the significance level, p, is .015 for the G-square statistic. Thus the model does not seem to fit the data.

Below the heading, the values of the parameters are given, along with statistics that can be used to compute confidence intervals. In particular we see that the data does not permit a clean computation of the parameter values : as indicated by the t-statistic, the standard errors are large compared to the actual parameter values, and the results are consistent with all the parameter values being the same.

We then include a breakdown of the data by individual entity. For example, for the item (library) groups, each cell indicates 1) The actual number of acquisitions made by that library within the specified subject group, 2) the number expected by the model, and 3) the standardized residual, which is a measure of cell fit (these can be thought of as normal deviates - values much larger than two indicate lack of fit). Finally, the output produces translation tables that associate each object with its class, though, to save space, this is not included.

A couple of comments are in order. We have analyzed the full data set in part to provide a baseline to compare with subsequent analyses, in part to illustrate the output of our programs. The large standard errors of the parameters are explainable by our including the extreme cases of books acquired by all or none of the libraries. It is very difficult for the model to fit such cases : a book acquired by all libraries, for example, is easy to acquire, but *how* easy? Its parameter would be well below those of the libraries, but that still leaves many possibilities. The problem is that the model has no basis for bracketing the value. Here the lowest value of D is -3. B0 might be -10, since this would result in every library acquiring the book, but it could also be -20 or -100. Without a library whose parameter value is so low that it does not buy the book, the model has no acceptable way of establishing a value for these extreme cases. Instead it forces a value, but indicates its discomfort by computing large standard errors. It is standard in analyses of this kind to remove these extreme groups before proceding, and in our subsequent analyses we do this. Thus, below, it will be understood that these extremes are absent from the analyses.

It is also interesting to study the lack of fit of the model. If we examine the individual cells, we find that for the most part the model has done quite well, with relatively few residuals exceeding two in absolute value. An example of misfit appears in subject group 1 : the model predicts that the library in item group 105 would acquire about .48 books of the 62 books in subject group 1. Books in group 1 (by definition of "group 1") have been acquired by only one library each, so they are not popular books; and the library in group 105 has acquired the fewest calculus books overall, so it is not a strong collecter in this area. It is unlikely that the weakest library would get a book from the class of books that most resist acquisition. Yet this library in fact has 2 books from this class. An error of two books is not large, but the model finds this too unlikely not to raise a warning flag in the form of a t-value of 2.19.

The greatest discrepancy is the t-value of -5.05, for library group 301 and book group 6. Book group 6 has 24 members, and examining the counts for this group, we see that it is quite easy for libraries to select books from this group. Yet the library in group 301, which on the basis of the model would be expected to have acquired at least 23 of these books, actually has acquired only 19. The model finds the discrepancy of 4 books unacceptably large.

We thus find that the model does fairly well at explaining how the libraries are making acquisitions. As such, it might serve quite well as a means for signalling libraries about purchases they might have overlooked. Nontheless, the accumulation of numerically small discrepancies is enough to indicate that the data is not consistent with the model.

We can interpret this outcome in several ways within the framework of the model. For one, we can simply conclude that these libraries do not form a peer group, and hunt for other groupings that are more consistent with the model. Alternatively, we can ask whether these libraries would form a peer group if we modified the book collection being studied. Both approaches will be taken below. A further possibility would be to accept the libraries as forming a peer group for this collection, and check with each library whether the discrepancies were an oversight - that is, to test the possibility that the model describes the items the libraries would have wished to acquire, given full information, rather than what they have in practice acquired. Finally, we can accept the model as an approximation of library collection development and use it for purposes of rough description and guidance rather than as a strictly accurate description of reality and as a basis for statistical tests.

In the second analysis, we removed the two-worst fitting libraries (Ohio State University and the University of Illinois), as well as the extreme cases discussed above. In effect, we are testing the first of the interpretations mentioned above. We shall not present the full tables - these are available in ref [10]. Instead we summarize the results. The G-squared value is now 84, a value that, while large, is not significant. Thus, the data do not give us any basis for rejecting the model for the smaller group of libraries - they do seem to form a peer group. Further investigation is required to see whether the two deviant libraries do indeed have individuating personalities that remove them from the group, or whether a different explanation is called for (data transcription errors, acquisition errors on the part of the libraries, etc.).

We also note the effect of removing the extreme cases of items acquired by all or none of the libraries : except for B5, the B-value closest to zero, all B-values are now statistically distinct from zero. Similarly, the model is able to distinguish from zero all but the three classes of items closest to zero. Thus, removing the extreme cases permits the model to differentiate between various classes of books and libraries.

There is similarly little that is striking about the actual cell values. The library constituting group 105, the weakest of the libraries, has gotten substantially more (two items) of item group 1, the least likely to be acquired class, than the model expects (.32 items). But actually the discrepancy is less than two books. Such discrepancies should be tested individually by identifying the books and perhaps communicating with the library to see if this is more than a chance effect. The translation tables described above make such an identification possible.

The only other noteworthy datum involves the library with class number 170, which acquired five less than expected of item class 7 (i.e. ten rather than

16 items : an error of 2.54 standard deviations). Again, only detailed investigation can explain the error. Overall, however, the model fits quite well, with only a hint that the model works least well at the extremes, perhaps underpredicting the extent to which small libraries acquire rarely acquired books and overpredicting the extent to which the big libraries acquire all the books.

We pursued the second explanation of model breakdown by studying the acquisition pattern of all the libraries for the restricted class of only English language materials (65 % of the collection). The G-square and Pearson chi-square values indicate an even better fit than before. The individual cells are similarly uninteresting, except perhaps again for the weakest library over-acquiring rare items.

The improvement in fit can be explained by

- the model indeed more precisely describing the English language acquisitions; or
- the effect of having a smaller amount of data.

We were partially able to test these alternatives by isolating the foreign language acquisitions and testing the model on these alone. It turned out that in this class the model is thoroughly routed ($G^2 = 91$, $p = 0.00$), even though much less data is available than for the English-language material. An examination of the cells reveals that the greatest discrepancies result from the purchases (or lack of purchases) by the largest library. For example, this library was expected to buy the one item in group 8, but did not. The error is but a single item, but this discrepancy is enough to trigger the very large standard error of -36. The reason is that the book in subject group 8 is very easy to acquire - 8 of the 11 libraries acquired it - yet the strongest of these libraries did not : the model is telling us that this is a very surprising result. Similarly, the defecit of a single item for class 6 triggers a standard error of -10. Perhaps the analysis should be redone with the largest library removed, but we were concerned about overly manipulating the data at this point, especially before the discrepancies were investigated individually.

On the basis of the above analysis, it seems reasonable to conclude that the model shows considerable promise, at minimum, as a tool for suggesting a second look by libraries of items they might haved wished to acquire but didn't. But more interesting theoretically, the model seems to describe reasonably well, though by no means perfectly, how this group of libraries develop their collections. The fit is especially good if we restrict ourselves to English language purchases. Given the character of the misfits, it is reasonable to suggest that libraries may display greater individuating qualities for non-English purchases than they do for English languages materials, and that the model is most appropriate for English-language material. But an alternative explanation is that the model well describes data in the center, but breaks down at the extremes by exaggerating the implications of small discrepancies.

3. ALGORITHMS

3.1. Estimation

In this section we give an overview of the mathematical considerations underlying the model. We first discuss the general loglinear model, and then restrict ourselves to the choice model. In the general model, we have a binary, dependent variable y, taking values 0 and 1, and a relation between the probability that $y = 1$ and a number of independent variables x : $p = \Pr\{y=1\} = f(x;b)$ and $q = \Pr\{y=0\} = 1 - f(x;b)$. Thus $\Pr\{y\} = f^y(1-f)^{1-y}$. Here b is a (vectorvalued) parameter whose value isn't known. To evaluate $b = (b_1, b_2, \ldots, b_m)$, we construct the logarithm of the likelihood function :

$\ell \equiv \sum_{i=1}^{n} [y_i \log f(x_i, b) + (1 - y_i) \log (1 - f(x_i, b))]$, where the i refers to individual cases and the logarithm is taken to base e. The maximum likelihood estimate of the parameters, b, is the value at which ℓ takes its maximum, i.e., the value of b at which :

$$\frac{\partial \ell}{\partial b_j} = \sum_i \left[\frac{y_i}{f(x_i,b)} - \frac{1 - y_i}{1 - f(x_i,b)}\right] \frac{\partial f(x_i,b)}{\partial b_j} \equiv g_j(X,b) = 0 \qquad (1)$$

(X in g_j is the matrix whose rows are the vectors x_i).
Let b_M denote the value of b at which (1) is satisfied.

In general, equation (1) cannot be solved in closed form, and an iterative process is used. Suppose $b^{(i)}$ is the current estimate of b_M. Considering g_j as a function of b, we can expand $g_j^{(b)}$ around $b^{(i)}$; then if b_M is the value at which equation (1) is satisfied : $g_j(b_M) = g_j(b^{(i)}) + \sum_{j'} \frac{\partial g_j(b^{(i)})}{\partial b_{j'}} (b_M - b^{(i)})_{j'}$;
thus, since $g(b_M) = 0$, we wish to solve

$$\sum_{j'} \frac{\partial g_j}{\partial b_{j'}} (b^{(i+1)} - b^{(i)})_{j'} = -g_j(b^{(i)}) \qquad (2a)$$

for $b^{(i+1)}$, the next approximation to b_M. This process is continued until the change in the estimated value of b between iterations is small. Replacing g by $\frac{\partial \ell}{\partial b}$, we conclude :

$$-\frac{\partial \ell}{\partial b_j} = \sum_{j'} \frac{\partial^2 \ell}{\partial b_j \, \partial b_{j'}} (b^{(i+1)} - b^{(i)})_{j'} \qquad (2b)$$

Equation (2b) can be rewritten in matrix notation as

$$-g = H(b^{(i+1)} - b^{(i)}) \; ; \qquad (2c)$$

or

$$b^{(i+1)} = b^{(i)} - H^{-1} g \qquad (2d)$$

Thus the following sequence of steps is involved :

1. Estimate $b^{(0)}$ at stage 0

At each stage, i :

2. Evaluate g, H at $b^{(i)}$
3. Compute $H^{-1}g$
4. Compute $b^{(i+1)} = b^{(i)} - H^{-1}g$

If $|b^{(i+1)} - b^{(i)}|$ is not adequately small, continue at step 2.

3.2. Model Evaluation

A number of approaches are available for evaluating the above model. We used the following.

i) To evaluate the model overall : given our estimate for b_M, we can compute p_i for each configuration x_i. If there are n_i cases satisfying the configuration x_i, we expect $E_i = n_i p_i$ of these cases to have $y = 1$, and $E_i = n_i(1-p_i)$ cases to have $y = 0$. If in fact O_i of these cases are in the i^{th} possible cell (i.e. $y = 0$ or $y = 1$ for any x value), $-2\,\Sigma\, O_i \ln \frac{O_i}{E_i}$ is a measure of the degree to which the predicted and actual counts disagree. It can be shown that this value, often called G^2, is approximately described by the Chi-square distribution when the model is valid; the degrees of freedom is given by (the number of cells - the number of independent parameters that are estimated).

ii) Should the model fit, we can assess how well each cell conforms to the model : If p_i is the probability that $y = 1$, given x_i, then if n_i items have value x_i we expect $n_i p_i$ of these to have $y = 1$, with a standard deviation of $\sqrt{n_i p_i(1-p_i)}$. Thus $d_i \equiv \frac{y_i - n_i p_i}{\sqrt{n_i p_i\,(1-p_i)}}$ is approximately normally distributed, with mean of zero and unit standard deviation. Since p_i can be estimated once b is, we can evaluate each cell in this manner. A similar measure, d_i, applies for $y = 0$, based on the probability $1 - p_i$. A workable procedure for finding badly fitting cells is to search for values of d_i much greater than two in absolute value.

iii) Finally, a general property of maximum likelihood estimation is that $-E\left(\frac{\partial^2 \ell}{\partial b\, \partial b^T}\right)$ is the inverse of the covariance matrix, Σ. Thus, Σ can be estimated by $-\left(\frac{\partial^2 \ell}{\partial b\, \partial b^T}\right)^{-1}$. The square roots of the diagonal values give us estimates for the standard errors of the components of b, permitting us to test hypotheses and compute confidence intervals.

3.3. Choice model

The above equations are general. We now summarize these results for the model we are examining.

1) $p_{rs} = \frac{\exp(b_r - d_s)}{1 + \exp(b_r - d_s)}$

2) $\ell = \sum_{rs} n_{rs}\,(b_r - d_s) - \Sigma\, N_{rs} \log\,(1 + e^{(b_r - d_s)})$

where, of N_{rs} opportunities for subjects in class r (having identical values for b) to select items in class s, n_{rs} actually do.

For technical reasons, we imposed the constraint $\Sigma\, b_i = 0$. This can be realized by using only the values $b_2,\ldots,b_n$ as parameters to be solved for, and substituting $-(b_2+\ldots+b_n)$ for b_1. If we do this, we can continue :

$$2a)\ \frac{\partial \ell}{\partial b_i} = (S_i - S_1) - (\sum_j N_{ij}p_{ij} - \sum_j N_{1j}p_{1j})$$

$$\frac{\partial \ell}{\partial d_j} = -I_j + \sum_i N_{ij}p_{ij}\ ,$$

where $S_i = \sum_j n_{ij}$ and $I_j = \sum_i n_{ij}$.

$$3a)\ \frac{\partial^2 \ell}{\partial b_i\, \partial b_{i'}} = -\ \delta_{ij} \sum_j N_{ij}p_{ij}q_{ij} - \sum_j N_{1j}p_{1j}q_{1j}\ ,$$

$$\frac{\partial^2 \ell}{\partial b_i\, \partial d_j} = -N_{1j}p_{1j}q_{1j} + N_{ij}p_{ij}q_{ij}\ ;\ \text{and}$$

$$\frac{\partial^2 \ell}{\partial d_i\, \partial d_j} = -\delta_{jj'} \sum_i N_{ij}p_{ij}q_{ij}\ .$$

All of the above can be converted into matrix form and evaluated using Gauss.

REFERENCES

[1] Goldhor, Herbert. "The Effect of Prime Display Location on Public Library Circulation of Selected Adult Titles", Library Quarterly, Vol 42,4 (Oct. 1972), p.371-89.

[2] Draper, N. and Smith, H. Applied Regression Analysis (2nd ed.), New York: Wiley, 1981.

[3] Zweizig, D. and Dervin, B. "Public Library Use, Users, Uses : Advances in Knowledge of the Characteristics and Needs of Adult Clientele of American Public Libraries", Advances in Librarianship, Vol.7, New York : Academic Press, 1977.

[4] D'Elia, G. "The Development and Testing of a Conceptual Model of Public Library User Behavior", Library Quarterly, Vol.50 (October 1981), p. 410-430.

[5] Grizzle, J.E. et al. "Analysis of Categorical Data by Linear Models", Biometrics, Vol.25 (1969), p.489-504.

[6] Bishop, Yvonne M.M. et al. Discrete Multivariate Analysis : Theory and Practice, Cambridge, Mass., MIT Press, 1975.

[7] Fienberg, Stephen. The Analysis of Cross-Classified Categorical Data (2nd ed.), Cambridge, Mass., MIT Press, 1980.

[8] Freeman, Daniel H. Applied Categorical Data Analysis, New York : M. Dekker, 1987.

[9] Amemiya, Takeshi. "Qualitative Response Models : A Survey", J. Economic Literature, Vol.19,4 (Dec.1981), p.1483-1536.

[10] Bookstein, Abraham. "Loglinear Analysis of Library Data", Research Report, OCLC, Office of Research, 1988.

[11] Sanders, Nancy P., O'Neill, Edward T., and Weible, Stewart. "Automated Collection Analysis Using the OCLC and RLG Bibliographic Databases", College and Research Libraries, Vol.49,4 (July, 1988), p.305-14.

APPENDIX

OUTPUT FROM GROUP RASCH ANALYSIS
CALCULUS DATA : ALL CASES, ALL LANGUAGES

Number of Subject groups : 12 Degrees of freedom : 110
(454 subjects)

Number of Item groups : 11
(11 items)

-2*Log Likelihood : 2739.543982
G-Squared : 144.423787 (p = 0.015)
Pearson Chi-square: 150.052508 (p = 0.007)

Convergence after iteration : 15 (Max value exceeded.)
Maximum change in params : 1.000000

Var	Coef	Std. Error	T-Stat	P-Value
B0	-19.272903	151.567261	-0.127157	0.899047
B1	-2.996520	20.822976	-0.143905	0.885838
B2	-1.881695	20.822790	-0.090367	0.928159
B3	-1.122421	20.822767	-0.053904	0.957109
B4	-0.528326	20.822951	-0.025372	0.979804
B5	-0.018690	20.822949	-0.000898	0.999285
B6	0.453261	20.822801	0.021768	0.982673
B7	0.922373	20.822852	0.044296	0.964748
B8	1.425874	20.822958	0.068476	0.945531
B9	2.025791	20.822936	0.097286	0.922675
B10	2.898007	20.823042	0.139173	0.889567
B11	18.095251	172.982664	0.104607	0.916877
D105	1.849958	20.823069	0.088842	0.929369
D124	1.320778	20.823004	0.063429	0.949540
D141	0.879854	20.822967	0.042254	0.966373
D142	0.854615	20.822965	0.041042	0.967337
D147	0.729394	20.822957	0.035028	0.972120
D163	0.338040	20.822938	0.016234	0.987077
D170	0.170514	20.822932	0.008189	0.993481
D174	0.075620	20.822929	0.003632	0.997109
D198	-0.483845	20.822922	-0.023236	0.981504
D252	-1.726192	20.822973	-0.082898	0.934082
D301	-3.001024	20.823175	-0.144119	0.885669

ACTUAL AND FITTED CELL COUNTS (Successes)

Itgp. Number No. Members	105 1	124 1	141 1	142 1	147 1	163 1	170 1	174 1	198 1	252 1	301 1
Subject gp 0 (101 members)											
Actual Count	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Fitted Count	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Std. resid.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Subject gp 1 (62 members)											
Actual Count	2.00	1.00	0.00	0.00	0.00	3.00	2.00	3.00	4.00	10.00	37.00
Fitted Count	0.48	0.82	1.26	1.29	1.46	2.13	2.51	2.74	4.65	13.59	31.07
Std. resid.	2.19	0.21	-1.13	-1.15	-1.22	0.60	-0.33	0.16	-0.31	-1.10	1.51
Subject gp 2 (49 members)											
Actual Count	0.00	1.00	2.00	0.00	3.00	0.00	7.00	3.00	16.00	26.00	40.00
Fitted Count	1.15	1.91	2.91	2.98	3.35	4.80	5.58	6.06	9.71	22.60	36.94
Std. resid.	-1.08	-0.67	-0.55	-1.78	-0.20	-2.31	0.64	-1.33	2.25	0.97	1.02
Subject gp 3 (37 members)											
Actual Count	1.00	2.00	3.00	5.00	2.00	4.00	6.00	10.00	21.00	27.00	30.00
Fitted Count	1.80	2.96	4.40	4.50	5.02	6.97	7.97	8.58	12.79	23.92	32.10
Std. resid.	-0.61	-0.58	-0.71	0.25	-1.45	-1.25	-0.79	0.55	2.84	1.06	-1.02
Subject gp 4 (20 members)											
Actual Count	3.00	1.00	6.00	3.00	6.00	8.00	7.00	9.00	7.00	13.00	17.00
Fitted Count	1.70	2.72	3.93	4.01	4.43	5.92	6.64	7.07	9.78	15.36	18.44
Std. resid.	1.05	-1.12	1.16	-0.56	0.85	1.02	0.17	0.90	-1.24	-1.25	-1.21
Subject gp 5 (18 members)											
Actual Count	0.00	3.00	5.00	6.00	6.00	8.00	8.00	10.00	12.00	14.00	18.00
Fitted Count	2.41	3.74	5.21	5.30	5.78	7.41	8.15	8.58	11.06	15.24	17.13
Std. resid.	-1.67	-0.43	-0.11	0.36	0.11	0.28	-0.07	0.67	0.46	-0.81	0.96
Subject gp 6 (24 members)											
Actual Count	6.00	8.00	10.00	12.00	11.00	15.00	17.00	15.00	10.00	21.00	19.00
Fitted Count	4.76	7.10	9.48	9.62	10.35	12.69	13.69	14.24	17.24	21.56	23.26
Std. resid.	0.63	0.40	0.22	0.99	0.27	0.94	1.37	0.32	-3.29	-0.38	-5.05

Cont.

ACTUAL AND FITTED CELL COUNTS (Successes) (Cont.)

Itgp. Number No. Members	105 1	124 1	141 1	142 1	147 1	163 1	170 1	174 1	198 1	252 1	301 1
Subject gp 7 (22 members)											
Actual Count	6.00	10.00	12.00	13.00	15.00	17.00	10.00	15.00	15.00	20.00	21.00
Fitted Count	6.24	8.84	11.23	11.37	12.06	14.13	14.95	15.40	17.67	20.55	21.57
Std. resid.	-0.11	0.51	0.33	0.69	1.26	1.28	-2.26	-0.18	-1.43	-0.47	-0.89
Subject gp 8 (20 members)											
Actual Count	6.00	13.00	11.00	14.00	14.00	13.00	18.00	16.00	16.00	20.00	19.00
Fitted Count	7.91	10.52	12.66	12.78	13.35	14.96	15.56	15.88	17.42	19.18	19.76
Std. resid.	-0.87	1.11	-0.77	0.57	0.31	-1.01	1.31	0.06	-0.95	0.92	-1.58
Subject 9 (27 members)											
Actual Count	19.00	19.00	20.00	18.00	20.00	25.00	23.00	21.00	24.00	27.00	27.00
Fitted Count	14.68	18.07	20.49	20.61	21.20	22.79	23.35	23.64	24.97	26.38	26.82
Std. resid.	1.67	0.38	-0.22	-1.18	-0.56	1.17	-0.20	-1.54	-0.71	0.80	0.42
Subject gp 10 (39 members)											
Actual Count	27.00	31.00	37.00	36.00	35.00	35.00	37.00	37.00	38.00	39.00	38.00
Fitted Count	28.88	32.32	34.42	34.53	35.00	36.20	36.61	36.81	37.72	38.62	38.89
Std. resid.	-0.69	-0.56	1.28	0.74	0.00	-0.75	0.26	0.13	0.25	0.62	-2.74
Subject gp 11 (35 members)											
Actual Count	35.00	35.00	35.00	35.00	35.00	35.00	35.00	35.00	35.00	35.00	35.00
Fitted Count	35.00	35.00	35.00	35.00	35.00	35.00	35.00	35.00	35.00	35.00	35.00
Std. resid.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00