A Model for Repeated Measurements of a Multivariate Binary Response
Author(s): Alan Agresti
Source: *Journal of the American Statistical Association*, Vol. 92, No. 437 (Mar., 1997), pp. 315-321
Published by: Taylor & Francis, Ltd. on behalf of the American Statistical Association
Stable URL: https://www.jstor.org/stable/2291476

A Model for Repeated Measurements of a Multivariate Binary Response

Alan AGRESTI

This article presents a logit model for a vector of binary variables observed for each subject under multiple conditions. The model contains a vector of random effects to account for correlations among the repeated measurements. A nonparametric treatment of the random effects implies a multivariate log-linear model having quasi-symmetric structure for the cross-classification of responses at the various conditions. The fit yields estimates of within-subject effects comparing the conditions for each variable. The estimates are identical to conditional maximum likelihood estimates for a fixed-effects version of the logit model. Extensions incorporate independent groups or allow variables to have multiple response categories.

KEY WORDS: Item response model; Marginal homogeneity; Matched pair; McNemar test; Quasi-symmetry; Rasch model.

1. INTRODUCTION

When a study measures a binary variable for each subject under two conditions, inferential methods such as McNemar's test are well established for comparing the matched-pair responses. This article discusses a model that extends inference to a *vector* of categorical variables, with each observed under at least two conditions. The conditions refer to the separate situations under which measurements occur for a variable, such as distinct time points, different variations of a question in a survey, or different treatments. For instance, a crossover experiment in a biomedical study might measure several binary outcome measures under each of two or more treatments.

Table 1, first presented by Coleman (1964), is a simple example of repeated categorical measurement data of this type. A study interviewed a sample of schoolboys twice, several months apart, and asked about their self-perceived membership in the "leading crowd" (yes, no) and about whether one must sometimes go against his principles to be part of that leading crowd (agree, disagree). The table summarizes responses on two variables (membership in the leading crowd, attitude toward the leading crowd) under two conditions (the two interview times). Each subject has measurements at two times for each of the two binary variables.

More generally, different variables may have different numbers of response categories. The different variables may even refer to different numbers of conditions. For simplicity of notation, I express models for the case of a common number of conditions for each variable.

The model presented in this article provides comparisons of responses under the various conditions, simultaneously for each variable. The primary focus is on binary responses. I propose a logit model with a vector of random effects for each subject, to account for the correlation among the repeated measurements. A nonparametric approach with the random effects implies a multivariate marginal model that for each variable has quasi-symmetric log-linear structure for the cross-classification of responses among conditions. That model is simple to fit with standard software for log-linear models. In Table 1, for instance, the logit model describes subject-specific changes in membership and changes in attitude between the two interview times. I provide fitted values for a related marginal model that contains the same parameter describing changes in membership, the same parameter describing changes in attitude, and a parameter describing the association between membership and attitude at each interview time.

Section 2 introduces the multivariate logit model for the repeated responses. Section 3 discusses the log-linear model implied by a nonparametric random-effects treatment and illustrates the model by providing analyses for Table 1. Section 4 discusses special cases of the model and extensions to handle comparisons of response patterns for separate groups or to handle multiple-category responses. Finally, Section 5 explores connections with conditional maximum likelihood estimation for the logit model and raises questions for further research.

2. A MULTIVARIATE LOGIT MODEL FOR REPEATED MEASUREMENT

Suppose that subjects respond to I separate binary variables, each measured for T conditions. For a given subject, denote the response under condition t for variable i by Y_{it}, with observed value $y_{it} = 1$ or 0, $i = 1, \ldots, I, t = 1, \ldots, T$. I refer to the outcome categories 1 and 0 as *success* and *failure.* For subject $s, s = 1, \ldots, n$, let ϕ_{ist} denote the probability of success on variable i under condition t. Different conditions for a variable use the same scale, and the questions or instrument used to elicit the response are set up so that the successive responses on a variable are positively correlated.

Consider the model

$$\text{logit}(\phi_{ist}) = \alpha_{is} + \beta_{it}, \tag{1}$$

$i = 1, \ldots, I, t = 1, \ldots, T$, which assumes a lack of subject-by-condition interaction for each variable. For each vari-

Alan Agresti is Professor, Department of Statistics, University of Florida, Gainesville, FL 32611. This work was partially supported by a grant from the National Institutes of Health. The author thanks Martin Knott and Atalanta Ghosh for helpful discussions, and acknowledges the hospitality of the Statistics Department at London School of Economics and Political Science, where part of this work was conducted during a sabbatical.

Journal of the American Statistical Association
March 1997, Vol. 92, No. 437, Theory and Methods

Table 1. Membership (M) and Attitude (A) Toward the "Leading Crowd" for Boys (Fitted Values for Multivariate Quasi-Symmetry Models in Parentheses)

(M,A) for first interview		(M,A) for second interview (Yes, Agree)	(Yes, Disagree)	(No, Agree)	(No, Disagree)	Total
Yes	Agree	458	140	110	49	757
		(458)	(141.8)	(119.5)	(49.1)	
		(453.4)	(143.8)	(121.7)	(49.1)	
Yes	Disagree	171	182	56	87	496
		(169.2)	(182)	(58.6)	(74.8)	
		(171.5)	(182.8)	(58.5)	(72.8)	
No	Agree	184	75	531	281	1,071
		(174.5)	(71.7)	(531)	(282.3)	
		(177.7)	(71.6)	(530.6)	(280.4)	
No	Disagree	85	97	338	554	1,074
		(85.6)	(109.2)	(336.7)	(554)	
		(85.5)	(106.3)	(334.5)	(557.9)	
Total		898	494	1,035	971	3,398

NOTE: Second set in parentheses constrain equal odds ratios between M and A for each interview.

able, this model has the form of the Rasch model (Rasch 1961) for responses under the various conditions. Given the model parameters, I treat the observations as independent Bernoulli variates. Identifiability requires a constraint such as $\beta_{i1} = 0$ for each variable. The $\{\beta_{i1}, \ldots, \beta_{iT}\}$ for each i are the parameters of interest for describing the condition effects for each variable. The $\{\alpha_{is}\}$ parameters reflect the heterogeneity among subjects that induces the correlations among repeated responses on a variable.

For subject s, the probability of a particular sequence of responses $\mathbf{y} = (y_{11}, \ldots, y_{IT})$ for the IT variable–condition combinations equals

$$\prod_i \prod_t \left(\frac{e^{\alpha_{is}+\beta_{it}}}{1+e^{\alpha_{is}+\beta_{it}}}\right)^{y_{it}} \left(\frac{1}{1+e^{\alpha_{is}+\beta_{it}}}\right)^{1-y_{it}}$$

$$= \frac{\exp[\sum_i \alpha_{is}(\sum_t y_{it}) + \sum_i \sum_t \beta_{it} y_{it}]}{\prod_i \prod_t [1+\exp(\alpha_{is}+\beta_{it})]}.$$

Denote the subject terms for subject s by $\boldsymbol{\alpha}_s = (\alpha_{1s}, \ldots, \alpha_{Is})$. I treat this vector as a random effect, permitting the components to be correlated. In Table 1, for instance, subjects having a relatively high random effect for the membership variable (thus having a propensity to be members regardless of the interview time) probably tend to have a relatively high random effect for the attitude variable.

Thus, the analysis using model (1) treats $\boldsymbol{\alpha}_s$ as a multivariate random effect with correlated components. Suppose that $\boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_n$ are independent with a cumulative distribution function denoted by F. Denote the marginal probability, averaged over the subjects, of a particular sequence of responses $\mathbf{y}$ by $\pi(\mathbf{y})$. The next section utilizes a log-linear model for these marginal probabilities implied by the logit model to estimate the condition effects $\{\beta_{it}\}$.

3. A MARGINAL MULTIVARIATE LOG-LINEAR MODEL

For model (1), the marginal probability equals

$$\pi(\mathbf{y}) = \exp\left(\sum_i \sum_t \beta_{it} y_{it}\right) \times \int_{\boldsymbol{\alpha}_s} \frac{\exp[\sum_i \alpha_{is}(\sum_t y_{it})]}{\prod_i \prod_t [1+\exp(\alpha_{is}+\beta_{it})]} \, dF(\alpha_{1s}, \ldots, \alpha_{Is}).$$

Regardless of the form of F, the integral determining this marginal probability yields a complex function of the $\{\beta_{it}\}$. Note, however, that this function depends on the data only through the values of $(\sum_t y_{1t}, \ldots, \sum_t y_{It})$. Thus model (1) implies that this marginal probability has a structure that is a special case of a model providing a separate parameter for each possible value of that vector of sums. This more general marginal model has the form

$$\pi(\mathbf{y}) = \exp\left(\sum_i \sum_t \beta_{it} y_{it}\right) \gamma \left(\sum_t y_{1t}, \ldots, \sum_t y_{It}\right),$$

where γ is an unspecified positive parameter that can assume a different value for each combination of the arguments.

Under the assumptions just made, the sample of n observations on the binary responses $\mathbf{y}$ for the IT variable–condition combinations forms a multinomial sample with probabilities $\{\pi(\mathbf{y})\}$. The form just derived that these probabilities satisfy can be expressed as a log-linear model for expected frequencies $\{\mu(\mathbf{y})\}$ in a 2^{IT} contingency table that cross-classifies the responses for the IT variable-condition combinations. This model has the form

$$\log[\mu(\mathbf{y})] = \sum_i \sum_t \beta_{it} y_{it} + \lambda \left(\sum_t y_{1t}, \ldots, \sum_t y_{It}\right). \tag{2}$$

For this model, the interaction term is invariant under permutations of the response outcomes for the conditions for a particular variable.

No matter what form the random effects distribution F takes, the implied marginal model has the same main effect structure, and it has an interaction term that is a special case of the one in (2). Thus one can consistently estimate the within-subject condition effects for each variable in a nonparametric manner using the ordinary maximum likelihood (ML) estimates obtained by fitting the log-linear model. The estimates pertain to comparisons of the form $\beta_{it_1} - \beta_{it_2}$ for each pair of conditions t_1 and t_2 for a variable. The interaction parameters in models such as (2) reflect the dependence in responses within and between variables. However, summary interpretations for the model refer to the $\{\beta_{it}\}$ marginal parameters rather than to these interaction parameters, the interpretations referring to odds ratios based on the original logit model (1).

An obvious question is whether estimating the condition effects using log-linear model (2) results in some efficiency loss, due to replacing the actual form of those parameters in the interaction term by a more general interaction term. I believe that any efficiency loss is likely to be minor, because of the argument (presented in Sec. 5) showing that the ML estimates for log-linear model (2) are also the conditional ML estimates for logit model (1) based on eliminating the subject parameters by conditioning on their sufficient statistics. In univariate problems with item response models, I am unaware of any studies showing significant efficiency gains over the conditional ML estimate. In fact, in the univariate case it follows from de Leeuw and Verhelst (1986) and Tjur (1982) that the actual nonparametric estimate is asymptotically identical to this extended nonparametric estimate based on the more general interaction term. Some efficiency loss may occur relative to the estimates for a particular parametric choice of random effects distribution, but those estimates have the disadvantage of potential inconsistency due to misspecification of that distribution.

One can fit model (2) using the standard Newton–Raphson algorithm for log-linear models; for instance, with software for generalized linear models. (An example is available from the author for Table 1 using SAS or GLIM.) For identifiability, one sets a constraint on the main effect parameters for each variable, such as $\beta_{iT} = 0$ for all i. The usual goodness-of-fit statistics have large-sample chi-squared distributions with df $= 2^{IT} - [I(T-1) + (T+1)^I]$. The data are often sparse, in which case such indices are mainly useful for comparing fits of models.

The likelihood equations induced by the first term in the log-linear model equate the fitted values to the observed data in the marginal distribution for each variable–condition combination. These are the likelihood equations for the model of mutual independence of the IT responses. For any combination of integer values $(u_1, \ldots, u_I), 0 \le u_i \le T$ for all i, let $n^*(u_1, \ldots, u_I)$ denote the sum of all cell counts in the contingency table having $\sum_t y_{1t} = u_1, \ldots, \sum_t y_{It} = u_I$, and let $\hat{\mu}^*(u_1, \ldots, u_I)$ denote the fitted total in these cells for a model. Then the interaction term in the model induces the likelihood equations $\hat{\mu}^*(u_1, \ldots, u_I) = n^*(u_1, \ldots, u_I)$ for all such combinations; in the univariate case, this second set of likelihood equations are the ones for the complete symmetry model, for which each cell with the same number of successes has the same probability. Only one cell has a particular combination $(u_1, \ldots, u_I)$ when $u_i = 0$ or T for each i; that is, when for each variable all responses are successes or all responses are failures. In those cells the model provides a perfect fit.

For Table 1, log-linear model (2) fits fairly well. The goodness-of-fit statistics are $G^2 = 4.92$ for the likelihood-ratio statistic and $X^2 = 4.95$ for the Pearson statistic, based on df $= 5$. Table 1 also displays the fitted values for this model. The ML estimates of the condition effects are $\hat{\beta}_{A1} - \hat{\beta}_{A2} = .176$ (asymptotic standard error (ASE) $=$ .058) for attitude and $\hat{\beta}_{M1} - \hat{\beta}_{M2} = .379$ (ASE $= .075$) for membership. These are interpreted using odds ratios. For instance, for each subject, the estimated odds of membership in the leading crowd at the first interview equal $\exp(.379) = 1.46$ times the estimated odds of membership at the second interview.

To test the significance of these condition effects simultaneously, we can test within-variable marginal homogeneity for the two variables by comparing this model to the simpler model that forces the effects to be zero. The likelihood ratio statistic of 35.4 with df $= 2$ provides extremely strong evidence against the hypothesis that $\beta_{A1} - \beta_{A2} = \beta_{M1} - \beta_{M2} = 0$. Individual tests provide strong evidence of an effect for each variable, particularly of a decrease in the odds of membership in the leading crowd.

Goodman (1974a,b) and Haber (1985) presented alternative models for these data. Goodman used a latent class model with four latent classes that cross-classify two associated binary latent variables, one affecting the membership responses and one affecting the attitude responses. Common elements are shared by this latent class approach and models containing a pair of correlated random effects, though parameter interpretations for our model apply directly to the observed variables rather than to relationships between the latent variables or between the latent variables and the observed variables.

Haber (1985) fitted a model that assumes solely that the odds ratio between attitude and membership is identical for each interview. That is, the model applies to the 2×2 marginal table of membership and attitude at the first interview and the 2×2 marginal table of membership and attitude at the second interview. The sample odds ratios in these tables are 1.53 and 1.71, and Haber's model yielded fitted odds ratios in each table of 1.62.

The fit of model (2) also suggests that these marginal odds ratios are similar, as their estimates based on the fitted values for that model equal 1.63 and 1.61. Using the methodology and the algorithm described by Lang and Agresti (1994) for simultaneous fitting of generalized log-linear models to joint and marginal distributions of contingency tables, we fitted the slightly simpler version of model (2) that constrains these marginal odds ratios to be identical. The fit, also shown in Table 1, has $G^2 = 5.31$ and $X^2 = 5.41$ with df $= 6$. The fitted common odds ratio equals 1.62, and the estimated condition effects are $\hat{\beta}_{A1} - \hat{\beta}_{A2} = .176$ (ASE $= .058$) and $\hat{\beta}_{M1} - \hat{\beta}_{M2} = .378$ (ASE $= .075$). In summary, this analysis describes Table 1 using three parameters. One parameter compares the attitude responses at the two interviews, estimated by an odds ratio of $\exp(.176) = 1.19$; a second parameter compares the membership responses at the two interviews, estimated by an odds ratio of $\exp(.378) = 1.46$;

and a third parameter describes the association between the attitude and membership responses at each interview, estimated by an odds ratio of 1.62.

4. SPECIAL CASES AND EXTENSIONS

For the single-variable case $I = 1$, model (2) is the quasi-symmetry model (Bishop, Fienberg, and Holland 1975; Caussinus 1966). Conaway (1989), Darroch (1981), Darroch, Fienberg, Glonek, and Junker (1993), Fienberg (1981), Kelderman (1984), Tjur (1982), and others have discussed connections between the logit and log-linear models in this case. For arbitrary I, the likelihood equations for the general log-linear model (2) imply that the fit in the 2^T marginal table for variable i is identical to the fit of the ordinary quasi-symmetry model to that marginal table alone. I refer to the full model (2) as a *multivariate quasi-symmetry model.*

In the multivariate matched-pairs case (arbitrary I, but $T = 2$), model (2) has fitted values in the 2×2 marginal table for each variable that are identical to the observed counts. The estimate of $\exp(\beta_{i2} - \beta_{i1})$ then equals the number of cases with $(y_{i1}, y_{i2}) = (0, 1)$ divided by the number of cases with $(y_{i1}, y_{i2}) = (1, 0)$. This is precisely the information used in the univariate case with methods such as McNemar's test (Cox 1958).

For four special cases of logit model (1), nonparametric marginal ML solutions relate to log-linear models that are special cases of model (2). First, suppose that the logit model has a degenerate random effects distribution; that is, the variance equals zero for each component. Then the marginal model is precisely the special case of (2) without the interaction term. This is the log-linear model of mutual independence among the responses for all the variable–condition combinations.

Second, suppose that the components of $\boldsymbol{\alpha}_s = (\alpha_{1s}, \ldots, \alpha_{Is})$ are mutually independent. Then the marginal probability of a particular sequence of responses satisfies the log-linear model

$$\log[\mu(\mathbf{y})] = \sum_i \sum_t \beta_{it} y_{it} + \sum_i \lambda_i \left(\sum_t y_{it} \right). \quad (3)$$

This model satisfies the restrictive, and typically unrealistic, structure whereby responses on variable a for any condition t_a and on a different variable b for any condition t_b are independent, both marginally and also conditionally on other responses.

Third, suppose that the components of $\boldsymbol{\alpha}_s = (\alpha_{1s}, \ldots, \alpha_{Is})$ are perfectly positively correlated. Then the marginal probability of a particular sequence of responses satisfies the log-linear model

$$\log[\mu(\mathbf{y})] = \sum_i \sum_t \beta_{it} y_{it} + \lambda \left(\sum_i \sum_t y_{it} \right). \quad (4)$$

In this case the logit model (1) treats all the variable–condition combinations symmetrically. This model is identical to the Rasch model applied to the IT separate responses. Similarly, the derived log-linear model is identical to the quasi-symmetry model for the 2^{IT} contingency table that cross-classifies those responses.

Finally, suppose that $\{\beta_j\}$ in the logit model (1) are identical. Then they are also identical in the log-linear model. Model (2) then exhibits within-variable symmetry. Specifically, each cell having the same value of $(\sum_t y_{1t}, \ldots, \sum_t y_{It})$ has the same probability. The fitted value for each such cell is

$$\hat{\mu}(\mathbf{y}) = n^*(y_{1+}, \ldots, y_{I+}) \Big/ \prod_i \binom{T}{y_{i+}}.$$

Each of these four simpler models is typically too simplistic to fit well. For instance, Table 2 also summarizes the fit of these models to Table 1. All of them fit poorly.

It is straightforward to extend log-linear model (2) to incorporate a group factor or to handle multiple-category responses. I first extend the model to provide comparisons of G groups on their within-subject condition effects, using independent samples of subjects from the groups. For subject s in group $g, g = 1, \ldots, G$, let $\phi_{is(g)t}$ denote the probability of success. In general form, the logit model (1) extends to

$$\text{logit}(\phi_{is(g)t}) = \alpha_{is(g)} + \beta_{itg}. \quad (5)$$

The model maintains additivity of subject and condition effects for each variable, but it permits the condition effects to vary among groups. Assuming that $(\alpha_{1s(g)}, \ldots, \alpha_{Is(g)})$ are iid among subjects in group g with unknown distribution, this structure implies a marginal model that satisfies

$$\log[\mu_g(\mathbf{y})] = \sum_i \sum_t \beta_{itg} y_{it} + \lambda_g \left(\sum_t y_{1t}, \ldots, \sum_t y_{It} \right), \quad g = 1, \ldots, G. \quad (6)$$

Table 2. Summary of Log-Linear Model Fits to Table 1

Model	G^2	X^2	*df*	*A effect*	*ASE*	*M effect*	*ASE*
a. Mutual independence	1,421.7	1,572.6	11	.125	.048	.172	.050
b. 4-item quasi-symmetry	616.6	680.3	8	.159	.055	.224	.057
c. Independent random effects	97.5	96.8	9	.176	.058	.379	.075
d. Multivariate symmetry	40.3	40.0	7	0		0	
e. Multivariate quasi-symmetry	4.9	5.0	5	.176	.058	.379	.075
f. Multivariate quasi-symmetry and common odds ratio	5.3	5.4	6	.176	.058	.378	.075

NOTE: Models result from logit model (1) with (a) degenerate random effects, (b) perfectly correlated random effects, (c) independent random effects, (d) identical condition effects for each variable, (e) unspecified distribution of random effects, and (f) case (e) with identical odds ratio between variables at each time.

Fitting this general form of the model simultaneously for all g is equivalent to fitting model (2) separately for each group. Special cases in which condition effects are homogeneous across groups for certain variables yield special cases of log-linear model (6) in which $\beta_{it1} = \cdots = \beta_{itG}$ for all t for those variables. Such a model provides a decent fit to an expanded version of Table 1 that also contains data for schoolgirls from Coleman's (1964) study; the homogeneous attitude effect equals .126 (ASE = .043), and the homogeneous membership effect equals .384 (ASE = .057).

The results in this article also extend directly to more general models that permit some or all variables to have nominal or ordinal scales. Let $\phi_{ij(i)st}$ denote the probability of response in category $j(i)$ on variable i for subject s under condition t, for $i = 1, \ldots, I, j(i) = 1, \ldots, J_i, s = 1, \ldots, n$ and $t = 1, \ldots, T$. For variable i and condition t for a given subject, let $y_{ij(i)t} = 1$ if the response falls in category $j(i)$ and 0 otherwise. A general extension of model (1) is given by

$$\phi_{ij(i)st} = \frac{\exp(\alpha_{ij(i)s} + \beta_{ij(i)t})}{\sum_{j'(i)} \exp(\alpha_{ij'(i)s} + \beta_{ij'(i)t})}, \tag{7}$$

where the parameters equal zero for a baseline response category (e.g., $j(i) = 1$) and a baseline condition for each variable.

For a nonparametric treatment of the random-effects distribution, the marginal probability of a particular sequence of responses on the variables and conditions is again a special case of a multivariate quasi-symmetric type of log-linear model. For a particular sequence $\mathbf{y}$ and expected frequencies $\{\mu(\mathbf{y})\}$, the log-linear model has form

$$\log[\mu(\mathbf{y})] = \sum_i \sum_t \sum_{j(i)} \beta_{ij(i)t} y_{ij(i)t} + \lambda\left(\sum_t y_{11t}, \ldots, \sum_t y_{1,J_1-1,t}, \ldots, \sum_t y_{I,J_I-1,t}\right) \tag{8}$$

for the $\prod_i J_i^T$ contingency table. The interaction term does not require elements for the final category of each variable, which would be redundant. For the single-variable case, this is the ordinary quasi-symmetry model (see, e.g., Conaway 1989).

When a particular variable i is ordinal, a simpler model is often adequate, replacing the parameters $\{\beta_{ij(i)t}, j(i) = 1, \ldots, J_i\}$ in the logit model (7) by $\{v_{j(i)}\beta_{it}, j(i) = 1, \ldots, J_i\}$ for a set of fixed monotone scores $\{v_{j(i)}\}$. A nonparametric random effects approach with this ordinal structure relates to a log-linear model in which the main effect terms $\sum_t \sum_{j(i)} \beta_{ij(i)t} y_{ij(i)t}$ for variable i in model (8) are replaced by $\sum_t \sum_{j(i)} \beta_{it} y^*_{ij(i)t}$, where $y^*_{ij(i)t} = v_{j(i)}$ if the response falls in category $j(i)$ and equals 0 otherwise. For a given variable i, the parameters $\{\beta_{i1}, \ldots, \beta_{iT}\}$ for the different conditions provide a stochastic ordering of the response distributions. The sufficient statistics for those parameters are sample mean scores for the various conditions, and the ML estimates of $\{\beta_{it}\}$ have the same ordering as those means. For the single-variable case, I have discussed this ordinal type of quasi-symmetry model in earlier work (Agresti 1993a, 1993b).

5. COMMENTS AND CONCLUSIONS

This article applied a nonparametric random-effects approach to the subject term in logit model (1). A fixed-effects approach to handling logit models with subject-specific terms uses conditional ML to eliminate the nuisance parameters. For model (1), under the independent Bernoulli assumption, the sufficient statistic for α_{is} for a subject with data $\mathbf{y}$ equals $\sum_t y_{it}$. The contribution to the conditional likelihood of that subject equals

$$\frac{\exp[\sum_i \sum_t \beta_{it} y_{it}]}{\sum_D \exp[\sum_i \sum_t \beta_{it} y^*_{it}]},$$

where the index set D for the denominator refers to the set of all $\mathbf{y}^*$ such that $\sum_t y^*_{it} = \sum_t y_{it}$ for all i. The conditional likelihood is the product of such terms for all subjects in the sample. It factors into a product of I terms, one for each variable. It follows that the conditional ML estimates of $\{\beta_{it}\}$ are identical to those obtained using conditional ML separately with the data for each variable. From work of Tjur (1982), those estimates are identical to the regular ML estimates obtained from fitting the quasi-symmetry model to the cross-classification among conditions for each variable. Thus the conditional ML estimates are identical to the ordinary ML estimates of $\{\beta_{it}\}$ obtained by fitting the multivariate quasi-symmetry model (2). See Fischer (1989) for a discussion of conditional ML estimation for multivariate models of type (1).

Given this result that the within-variable estimates from fitting the multivariate model are no different from those obtained by analyzing the variables separately, one might question the utility of the multivariate model. Though the full model is not needed to estimate $\{\beta_{it}\}$, treating the I variables simultaneously in this way has certain advantages including

a. providing fitted values for the complete cross-classification that satisfy marginally the model for each separate variable
b. enabling tests of fit comparing these fitted values to the observed counts
c. reflecting the dependence that exists between variables and permitting additional structure pertaining to their associations, such as the special case of model (2) for Table 1 that has equal membership-attitude odds ratios for each interview;
d. allowing comparisons with simpler models, such as log-linear models resulting from degenerate or independent or perfectly correlated random effects or identical fixed effects for different conditions or even different variables in logit model (1).

In particular, the multivariate quasi-symmetry model (2) provides joint fitted values that imply standard marginal analyses. Because the logit model (2) implies log-linear model (3), severe lack of fit of the log-linear model casts strong doubt on the applicability of the logit model.

An alternative approach worth pursuing with model (1) is fitting it using a parametric rather than a nonparametric structure for the random effects vector. It would be interesting to analyze whether results tend to agree for the nonparametric and parametric formulations, such as often happens exactly in the matched-pairs case for a single variable (Neuhaus, Kalbfleisch, and Hauck 1994). A fully nonparametric approach provides so much freedom for the joint relationship that the fits for individual variables are no different from those obtained by analyzing the data separately for each variable; this might not happen for a narrower restriction on the joint distribution of the random effects, particularly with further structure connecting the components in the random effect, such as equal correlations for all the pairs.

Assuming additional parametric structure for the random-effects distribution raises other questions, of course. If the specification is correct, do the nonparametric estimates suffer a substantive efficiency loss? Because the nonparametric estimates are also conditional ML estimates, our intuition is that the answer is negative. If the specification is incorrect, could this introduce much bias? For a parametric marginal ML approach, it is important to check the degree to which the estimates depend on the choice of subject distribution, and to develop diagnostics that could help with that choice. Neuhaus, Hauck, and Kalbfleisch (1992) suggested that the bias is small, but previous work (e.g., Heckman and Singer 1984) in a somewhat different context has shown that results may depend strongly on the choice, and this is an advantage of the nonparametric approach (see also Aitkin 1996). In particular, under the assumption that logit model (1) holds, log-linear model (2) is valid and provides consistent estimates of the condition effects regardless of the true distribution for the random effects. Thus one informal diagnostic for the parametric marginal ML approach would be to compare those estimates under various distributional assumptions to the nonparametric estimates; substantial deviations from the nonparametric estimates provide evidence of a possibly inappropriate choice.

Logit model (1) and the multicategory extensions presented in this article describe how responses for each variable depend on the condition, possibly within levels of a group factor, but otherwise they contain no explanatory variables and have no provision for missing data. Other important problems for future work relate to extensions of such models for more complex data structures. In the single-variable case, for instance, Fischer (1974) and Hatzinger (1989) modeled the main-effect item parameters directly in terms of explanatory variables. Perhaps multivariate models such as those presented by Glonek and McCullagh (1995) could be extended to this repeated-measurement setting with random effects.

One could also generalize to this multivariate form of data other analyses that differ from the traditional item-response form of analysis. For instance, Haberman and Gilula (1995) presented an information-theoretic approach that provides summaries of predictive power associated with various log-linear models. Finally, Ten Have and Becker (1995) have explored an alternative variety of log-linear models with quasi-symmetric structure.

[Received November 1995. Revised June 1996.]

REFERENCES

Agresti, A. (1993a), "Computing Conditional Maximum Likelihood Estimates for Generalized Rasch Models Using Simple Loglinear Models with Diagonals Parameters," *Scandinavian Journal of Statistics*, 20, 63–71.

——— (1993b), "Distribution-Free Fitting of Logit Models With Random Effects for Repeated Categorical Responses," *Statistics in Medicine*, 12, 1969–1987.

Aitkin, M. (1996), "A General Maximum Likelihood Analysis of Overdispersion in Generalized Linear Models," *Statistics and Computing*, 6, 251–262.

Bishop, Y. V. V., Fienberg, S. E., and Holland, P. W. (1975), *Discrete Multivariate Analysis*, Cambridge, MA: MIT Press.

Caussinus, H. (1966), "Contribution a l'Analyse Statistique des Tableaux de Correlation," *Annales de la Faculte des Sciences de l'Universite de Toulouse*, 29, 77–182.

Coleman, J. S. (1964), *Introduction to Mathematical Sociology*, London: Free Press of Glencoe.

Conaway, M. (1989), "Analysis of Repeated Categorical Measurements With Conditional Likelihood Methods," *Journal of the American Statistical Association*, 84, 53–62.

Cox, D. R. (1958), "Two Further Applications of a Model for Binary Regression," *Biometrika*, 45, 562–565.

Darroch, J. N. (1981), "The Mantel–Haenszel Test and Tests of Marginal Symmetry; Fixed Effects and Mixed Models for a Categorical Response," *International Statistical Review*, 49, 285–307.

Darroch, J. N., Fienberg, S. E., Glonek, G. F. V., and Junker, B. W. (1993), "A Three-Sample Multiple-Recapture Approach to Census Population Estimation With Heterogeneous Catchability," *Journal of the American Statistical Association*, 88, 1137–1148.

Fienberg, S. E. (1981), "Recent Advances in Theory and Methods for the Analysis of Categorical Data: Making the Link to Statistical Practice," *Bulletin of the International Statistical Institute*, 49, Book 2, 763–791.

Fischer, G. H. (1974), *Einfuhrung in die Theorie Psychologischer Tests*, Bern: Huber.

Fischer, G. H. (1989), "An IRT-Based Model for Dichotomous Longitudinal Data," *Psychometrika*, 54, 599–624.

Glonek, G., and McCullagh, P. (1995), "Multivariate Logistic Models," *Journal of the Royal Statistical Society*, Ser. B, 57, 533–546.

Goodman, L. A. (1974a), "The Analysis of Systems of Qualitative Variables When Some of the Variables are Unobservable. I. A Modified Latent Structure Approach," *American Journal of Sociology*, 79, 1179–1259.

——— (1974b), "Exploratory Latent Structure Analysis Using Both Identifiable and Unidentifiable Models," *Biometrika*, 61, 215–231.

Haber, M. (1985), "Maximum Likelihood Methods for Linear and Log-Linear Models in Categorical Data," *Computational Statistics and Data Analysis*, 3, 1–10.

Haberman, S. J., and Gilula, Z. (1995), "Information-Theoretic Approach to Analysis and Prediction of Categorical Response Profiles," unpublished manuscript.

Hatzinger, R. (1989), "The Rasch Model, Some Extensions and Their Relation to the Class of Generalized Linear Models," in *Statistical Modelling: Proceedings of GLIM89 and the 4th International Workshop on Statistical Modelling*, Lecture Notes in Statistics, 57, Berlin: Springer.

Heckman, J. J., and Singer, B. (1984), "A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models of Duration," *Econometrica*, 52, 271–320.

Kelderman, H. (1984), "Loglinear Rasch Model Tests," *Psychometrika*, 49, 223–245.

Lang, J. B., and Agresti, A. (1994), "Simultaneously Modeling Joint and Marginal Distributions of Multivariate Categorical Responses," *Journal of the American Statistical Association*, 89, 625–632.

Neuhaus, J. M., Hauck, W. W., and Kalbfleisch, J. D. (1992), "The Effects of Mixture Distribution Misspecification When Fitting Mixed-Effects Logistic Models," *Biometrika*, 79, 755–762.

Neuhaus, J. M., Kalbfleisch, J. D., and Hauck, W. W. (1994), "Conditions for Consistent Estimation in Mixed-Effects Models for Binary Matched-Pairs Data," *Canadian Journal of Statistics*, 22, 139–148.

Rasch, G. (1961), "On General Laws and the Meaning of Measurement in Psychology," in *Proceedings of the 4th Berkeley Symposium in Mathematical Statistics and Probability*, Vol. 4, ed. J. Neyman, Berkeley, CA: University of California Press, pp. 321–333.

Ten Have, T. R., and Becker, M. P. (1995), "Multivariate Contingency Tables and the Analysis of Exchangeability," *Biometrics*, 51, 1001–1016.

Tjur, T. (1982), "A Connection Between Rasch's Item Analysis Model and a Multiplicative Poisson Model," *Scandinavian Journal of Statistics*, 9, 23–30.