Published in final edited form as:

Biometrics. 2018 March ; 74(1): 354–361. doi:10.1111/biom.12707.

A Note on Marginalization of Regression Parameters from Mixed Models of Binary Outcomes

Donald Hedeker*,

Department of Public Health Sciences, University of Chicago, 5841 S. Maryland Ave., Room W254, MC2000, Chicago, IL 60637, USA

Stephen du Toit,

Scientific Software International, Inc., Skokie IL 60076, USA

Hakan Demirtas, and

School of Public Health, University of Illinois at Chicago, Chicago IL 60612, USA

Robert D. Gibbons

Department of Public Health Sciences, University of Chicago, Chicago IL 60637, USA

Summary

This paper discusses marginalization of the regression parameters in mixed models for correlated binary outcomes. As is well known, the regression parameters in such models have the “subject-specific” (SS) or conditional interpretation, in contrast to the “population-averaged” (PA) or marginal estimates that represent the unconditional covariate effects. We describe an approach using numerical quadrature to obtain PA estimates from their SS counterparts in models with multiple random effects. Standard errors for the PA estimates are derived using the delta method. We illustrate our proposed method using data from a smoking cessation study in which a binary outcome (smoking, Y/N) was measured longitudinally. We compare our estimates to those obtained using GEE and marginalized multilevel models, and present results from a simulation study.

Keywords

Clustered data; Longitudinal data; Multilevel models; Population-averaged estimates; Subject-specific estimates

1. Introduction

Mixed-effects models are a popular approach for analysis of longitudinal and/or clustered binary outcomes, with many texts either devoted to or with substantial treatment on the topic

* hedeker@uchicago.edu.

This paper has been submitted for consideration for publication in *Biometrics*

Supplementary Materials

The dataset and a SAS script using PROCs NLMIXED and IML that can be used to obtain the marginalized estimates are available at the *Biometrics* website on Wiley Online Library. These materials are also available from the first author upon request.

(Brown and Prescott, 2006; Demidenko, 2004; Fitzmaurice et al., 2011; Goldstein, 2011; Hedeker and Gibbons, 2006; McCulloch et al., 2008; Molenberghs and Verbeke, 2005; Raudenbush and Bryk, 2002; Skrondal and Rabe-Hesketh, 2004). Such models are often referred to as belonging to the class of models known as generalized linear mixed models (GLMMs). It is well known that the regression parameters in such models have the “subject-specific” (SS) or conditional interpretation (Neuhaus et al., 1991), representing effects of the covariates controlling for, or holding constant, the random effects. This is in contrast to the marginal or “population-averaged” (PA) estimates of, most notably, the GEE approach (Zeger and Liang, 1986) and the marginalized multilevel model (Heagerty and Zeger, 2000). The SS and PA regression parameters generally differ in meaning and value, and the SS regression parameters depend on how many random effects are included in the model (e.g., conditional on random intercepts only or conditional on random intercepts and trends). Thus, Zeger et al. (1988) recommend that SS models are useful when the focus is on making inference at the individual level, whereas PA models are most useful for population-based inference. Here, we describe a numerical approach to obtain PA estimates and standard errors from their SS counterparts for models with one or multiple random effects. We focus on longitudinal binary outcomes and present results primarily for the logit link function, but our approach applies to other GLMMs as well.

Analysts often decide a-priori on conducting a PA or SS analysis depending on their target of inference. Thus, if one is interested in obtaining PA estimates, then typically a GEE or marginalized multilevel model is used. As is well known for longitudinal data, GEE estimates are valid under the restrictive missing data mechanism of missing completely at random (MCAR). Of course, weighted GEE (Robins et al., 1995) can be used to allow for missing at random (MAR), but this requires a joint modeling of the longitudinal outcome and the missing data, and further requires that the missing data model is correctly specified (Preisser et al., 2002; Demirtas, 2004). Alternatively, our approach uses the SS estimates from a full-likelihood mixed model analysis to obtain their PA counterparts, and therefore provides valid inference under MAR, though this does assume that the variance-covariance structure is correctly specified. Also, mixed models are more flexible in handling measurements made at irregular time intervals, since the covariance structure can depend on a continuous time variable. Marginalized multilevel models are a viable alternative to obtain PA estimates under MAR, and can allow for irregular time intervals. Available software for such marginalized multilevel models is somewhat limited, though Griswold et al. (2013) includes SAS and R scripts for estimating such models.

As an example, Gruder et al. (1993) report on the effects of a televised smoking cessation intervention in which nearly 500 smokers were randomized to one of three conditions and then followed longitudinally in terms of their binary smoking status (Y/N). The three conditions were no-contact control, discussion, and social support, and subjects were assessed by self report if they were smoking or not at post-intervention, 6-, 12-, and 24-months. The analysis in the paper consisted of a random intercept and trend probit model, which found that there were strong group effects at post-intervention that were not maintained over the two-year follow-up. The estimates presented in the paper were SS estimates, and therefore conditional on the random subject intercepts and trends. Alternatively, PA estimates would help to describe the magnitude of, for example, the post-

intervention group differences in the population of subjects, rather than conditional on subject effects.

In terms of the organization of the paper, Section 2 briefly describes mixed and marginal models for longitudinal binary data, Section 3 describes the proposed marginalization approach, Section 4 presents results of simulation studies, Section 5 presents a re-analysis of the Gruder et al. (1993) smoking cessation data, and Section 6 discusses some aspects of the proposed approach. Supplementary materials are mentioned in Section 7, including the smoking cessation dataset and a SAS syntax file that can be used to perform the proposed marginalization.

2. Regression Models for Longitudinal Binary Outcomes

Consider the mixed logistic regression model for a binary outcome Y_{ij} from subject i ($i = 1, \dots, N$) at timepoint j ($j = 1, \dots, n_i$):

$$\log \left[\frac{P(Y_{ij}=1|\mathbf{v}_i)}{1 - P(Y_{ij}=1|\mathbf{v}_i)} \right] = \mathbf{x}'_{ij}\boldsymbol{\beta} + \mathbf{z}'_{ij}\mathbf{v}_i, \quad (1)$$

where the random effects $\mathbf{v}_i$ are normally distributed in the population with mean 0 and variance-covariance matrix $\boldsymbol{\Sigma}_v$. The total number of observations is denoted as $n = \sum n_i$. Using the Cholesky factorization of the random-effects variance-covariance matrix $\boldsymbol{\Sigma}_v = \mathbf{T}\mathbf{T}'$, the model can also be written as

$$\log \left[\frac{P(Y_{ij}=1|\boldsymbol{\theta}_i)}{1 - P(Y_{ij}=1|\boldsymbol{\theta}_i)} \right] = \mathbf{x}'_{ij}\boldsymbol{\beta} + \mathbf{z}'_{ij}\mathbf{T}\boldsymbol{\theta}_i, \quad (2)$$

where $\boldsymbol{\theta}_i$ are standardized random effects with mean 0 and variance-covariance matrix $\mathbf{I}$. The regression parameters $\boldsymbol{\beta}$ in this model have the “subject-specific” interpretation (Neuhaus et al., 1991), and so can be more precisely denoted as $\boldsymbol{\beta}^{ss}$ to reflect the fact that they represent effects of the covariates controlling for, or holding constant, the random subject effects. Thus, the meaning of $\boldsymbol{\beta}^{ss}$ depends on how many and which random effects are included in a given analysis.

In contrast to the mixed model, consider the marginal model

$$\log \left[\frac{P(Y_{ij}=1)}{1 - P(Y_{ij}=1)} \right] = \mathbf{x}'_{ij}\boldsymbol{\beta}^{pa}, \quad (3)$$

in which there are no random effects. Instead, marginal models separate the mean and association structures. The model above is for the mean structure and would be combined with some specification of the association of the longitudinal data, for example the specification of the working correlation in a GEE model. The regression parameters $\boldsymbol{\beta}^{pa}$ in

this model have the marginal or “population-averaged” interpretation (Neuhaus et al., 1991), as they are not conditional on subject effects. As mentioned, these are preferable in settings in which one wants to make inferences about the population, for example in epidemiological studies. From this model, population-averaged probabilities, defined as π_{ij}^{pa} , can be expressed via the logistic response function as $\pi_{ij}^{pa} = \Psi(\lambda^{pa}) = 1/(1 + \exp(-\lambda^{pa}))$, which translates the population-averaged logit λ^{pa} into a probability π^{pa} . Ψ is the CDF of the standard logistic distribution and is sometimes termed the “expit” function.

3. Marginalization of the Subject-Specific Regression Estimates

Suppose that one has estimated a mixed logistic regression model, with estimates $\hat{\beta}^{ss}$ and $\hat{\tau}$ (or $\hat{\sigma}_v$ for a random intercept model), and that one would also like to obtain the population-averaged regression coefficient estimates, denoted as $\hat{\beta}^{pa}$. For a random-intercept probit model, there is a closed form solution, namely $\hat{\beta}^{pa} = \hat{\beta}^{ss} / \sqrt{1 + \hat{\sigma}_v^2}$ (Heagerty and Zeger, 2000), while for a random-intercept logistic model, the following approximation can be

used: $\hat{\beta}^{pa} \approx \hat{\beta}^{ss} / \sqrt{\frac{\hat{\sigma}_v^2 + \pi^2/3}{\pi^2/3}}$, where $\pi^2/3$ is the variance of the standard logistic distribution (Agresti, 2002). Zeger et al. (1988) recommend multiplying this logistic variance by the factor 15/16, which results from using a cumulative Gaussian approximation to the logistic function. The square-root term on the right-hand side represents the ratio of the conditional to the marginal standard deviations, and can be viewed as the “marginalization” factor which transforms the subject-specific parameters into their population-averaged counterparts. In a random-intercepts model the marginalization factor is a scalar and equal across time. However, for models with multiple random effects this is generally not the case and so there is no simple relationship like the one above.

In the more general case of multiple random effects, following the results of Griswold and Zeger (2004) and others, population-averaged estimated probabilities for each observation can be expressed as

$$\hat{\pi}_{ij}^{pa} = \int_{\theta} \Psi(\mathbf{x}_{ij}' \hat{\beta}^{ss} + \mathbf{z}_{ij}' \hat{\tau} \boldsymbol{\theta}_i) dF(\boldsymbol{\theta}). \quad (4)$$

Gauss-Hermite quadrature can be used to approximate these probabilities as

$$\hat{\pi}_{ij}^{pa} \approx \sum_{q=1}^{Q^r} \Psi(\mathbf{x}_{ij}' \hat{\beta}^{ss} + \mathbf{z}_{ij}' \hat{\tau} \mathbf{B}_q) A(\mathbf{B}_q). \quad (5)$$

The integration is approximated by a summation on a specified number of quadrature points Q for each dimension of the integration. For the standard normal univariate density, optimal points and weights (which are denoted $\mathbf{B}_q$ and $A(\mathbf{B}_q)$, respectively) are given in Stroud and Sechrest (1966). For the multivariate density, the r -dimensional vector of quadrature points

is denoted by $B'_q = (B_{q1}, B_{q2}, \dots, B_{qr})$, with its associated (scalar) weight given by the product of the corresponding univariate weights, $A(B_q) = \prod_{h=1}^r A(B_{qh})$.

Consider the marginal model in Equation (3), and denote $\boldsymbol{\pi}^{pa}$ as the $n \times 1$ vector of probabilities that is obtained by stacking all of the n observations, and $\mathbf{X}$ as the $n \times p$ covariate matrix that is obtained by stacking the covariate vectors of all n observations (including a column of ones for the intercept). The marginal model can then be written as $\boldsymbol{\lambda}^{pa} = \mathbf{X}\boldsymbol{\beta}^{pa}$, where the $n \times 1$ vector of logits $\boldsymbol{\lambda}^{pa}$ is formed using elementwise division of $\boldsymbol{\pi}^{pa}$ by $1 - \boldsymbol{\pi}^{pa}$. Multiplying both sides of this equation by $(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$, yields the relationship between the population-averaged regression coefficients and the population-averaged logits:

$$\boldsymbol{\beta}^{pa} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\boldsymbol{\lambda}^{pa}. \quad (6)$$

Marginalized regression coefficients can be obtained using Equation (6), in conjunction with the quadrature method in Equation (5) which provides estimates of the marginal probabilities for each observation. Specifically, using estimates $\hat{\beta}^{ss}$ and $\hat{T}$, along with the known covariate and random effect matrices $\mathbf{X}$ and $\mathbf{Z}$, in Equation (5) yields $\hat{\pi}_{ij}^{pa}$ from which the $n \times 1$ vector $\hat{\lambda}^{pa}$ is constructed. Substituting $\hat{\lambda}^{pa}$ into Equation (6) yields $\hat{\beta}^{pa}$, which are estimated population-averaged regression coefficients based on a model with multiple random effects.

3.1 Marginalized Standard Errors

The delta method can be used to obtain standard errors associated with the estimates $\hat{\beta}^{pa}$. For this, let $\hat{\eta}$ represent the $(p + r(r+1)/2) \times 1$ stacked vector of estimates $\hat{\beta}^{ss}$ and $\hat{T}$. To simplify the notation, let $p^* = (p + r(r+1)/2)$, and denote the estimated variance-covariance matrix of the p^* parameter estimates as $V(\hat{\eta})$. Then, using the delta method,

$$V(\hat{\beta}^{pa}) = \frac{\partial \hat{\beta}^{pa}}{\partial \hat{\eta}'} V(\hat{\eta}) \frac{\partial (\hat{\beta}^{pa})'}{\partial \hat{\eta}}. \quad (7)$$

Here, $\frac{\partial \hat{\beta}^{pa}}{\partial \hat{\eta}'}$ is a matrix of size $p \times p^*$, $V(\hat{\eta})$ is of size $p^* \times p^*$, and $\frac{\partial (\hat{\beta}^{pa})'}{\partial \hat{\eta}}$ is of size $p^* \times p$. Using the chain rule,

$$\frac{\partial \hat{\beta}^{pa}}{\partial \hat{\eta}'} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}' \frac{\partial \hat{\lambda}^{pa}}{\partial \hat{\eta}'}, \quad (8)$$

where

$$\frac{\partial \hat{\lambda}^{pa}}{\partial \hat{\eta}'} = \frac{\sum_{q=1}^{Q^r} \Psi(\hat{\lambda}_{ijq})(1 - \Psi(\hat{\lambda}_{ijq}))A(\mathbf{B}_q) \left(\partial \hat{\lambda}_{ijq} / \partial \hat{\eta}' \right)}{\hat{\pi}^{pa} (1 - \hat{\pi}^{pa})} \quad (9)$$

with $\hat{\lambda}_{ijq} = \mathbf{x}_{ij}' \hat{\beta}^{ss} + \mathbf{z}_{ij}' \hat{T} \mathbf{B}_q$. In the above equation, the numerator is a $1 \times p^*$ row vector that is obtained for each of the n observations, which upon stacking yields a $n \times p^*$ matrix. The denominator is the elementwise product of the two $n \times 1$ vectors. Thus, to form this derivative matrix, each element in a row of the $n \times p^*$ matrix in the numerator is divided by the corresponding element of the $n \times 1$ denominator vector. Finally, the $1 \times p^*$ row vector is given as

$$\frac{\partial \hat{\lambda}_{ijq}}{\partial \hat{\eta}'} = \left[\mathbf{x}_{ij}' : (\mathbf{B}_q' \otimes \mathbf{z}_{ij}') \mathbf{J}_r' \right], \quad (10)$$

where $\otimes$ represents the Kronecker product, and $\mathbf{J}_r$ is the $(r \times (r+1)/2) \times r^2$ elimination matrix, which removes the elements above the main diagonal of a lower triangular matrix (Magnus, 1988).

3.2 Other link functions

The procedure can be extended to other link functions as well. For example, for a probit mixed model, population-averaged probabilities are obtained as:

$$\hat{\pi}_{ij}^{pa} \approx \sum_{q=1}^{Q^r} \Phi(\mathbf{x}_{ij}' \hat{\beta}^{ss} + \mathbf{z}_{ij}' \hat{T} \mathbf{B}_q) A(\mathbf{B}_q), \quad (11)$$

where $\Phi(\cdot)$ is the CDF of the standard normal distribution, and regression coefficients are obtained as:

$$\hat{\beta}^{pa} = (\mathbf{X}' \mathbf{X})^{-1} \mathbf{X}' \Phi^{-1}(\hat{\pi}^{pa}), \quad (12)$$

with $\Phi^{-1}(\cdot)$ as the probit function (inverse of the normal CDF). Standard errors are obtained using Equation (7) with

$$\frac{\partial \hat{\beta}^{pa}}{\partial \hat{\eta}'} (\mathbf{X}' \mathbf{X})^{-1} \mathbf{X}' \frac{\sum_{q=1}^{Q^r} \phi(\hat{\lambda}_{ijq}) A(\mathbf{B}_q) \left(\partial \hat{\lambda}_{ijq} / \partial \hat{\eta}' \right)}{\phi(\Phi^{-1}(\hat{\pi}^{pa}))}, \quad (13)$$

with $\phi(\cdot)$ is the PDF of the standard normal distribution, and again using Equation (10) for the derivatives $(\partial \hat{\lambda}_{ijq} / \partial \hat{\eta}')$.

More generally, denoting $F(\cdot)$ and $F^{-1}(\cdot)$ as the CDF and the link function (inverse CDF), the following general forms are:

$$\hat{\pi}_{ij}^{pa} \approx \sum_{q=1}^{Q^r} F(\mathbf{x}'_{ij} \hat{\beta}^{ss} + \mathbf{z}'_{ij} \hat{T} \mathbf{B}_q) A(\mathbf{B}_q), \quad (14)$$

$$\hat{\beta}^{pa} = (\mathbf{X}' \mathbf{X})^{-1} \mathbf{X}' F^{-1}(\hat{\pi}^{pa}), \quad (15)$$

$$\frac{\partial \hat{\beta}^{pa}}{\partial \hat{\eta}'} = (\mathbf{X}' \mathbf{X})^{-1} \mathbf{X}' \partial F^{-1}(\hat{\pi}^{pa}) \sum_{q=1}^{Q^r} A(\mathbf{B}_q) \frac{\partial F(\hat{\lambda}_{ijq})}{\partial \hat{\eta}'} \left(\frac{\partial \hat{\lambda}_{ijq}}{\partial \hat{\eta}'} \right). \quad (16)$$

The specific elements for several link functions are listed in Table 1. It should be noted that the loglog link is sometimes presented as the negative loglog link (*i.e.*, $-\log(-\log(\pi))$), in which case the elements would be of opposite sign.

4. Simulation Studies

To evaluate and compare the performance of the proposed marginalization method, two small-scale simulation studies were conducted. The first considered a random-intercept probit model in which marginalized estimates can be expressed in closed form. This allows us to compare our approach to the marginalized (closed form) estimates of the random-intercept probit model. The second simulation study considered a logistic mixed model with random intercepts and time trends. Here, we compared results of the proposed method with estimates from GEE and marginalized multilevel models. We also investigated the role of MAR missing data on our proposed approach, and contrasted it with GEE estimates.

4.1 Random-intercept probit model

The specifications for this simulation study were based on the example and data analysis presented in Section 5. To examine the effect of sample size, 1000 datasets were simulated each with 100 or 500 subjects and 4 repeated observations according to a marginal probit regression model

$$\Phi^{-1}(\pi_{ij}^{pa}) = \beta_0 + \beta_1 G_i + \beta_2 T_j + \beta_3 (G_i \times T_j), \quad (17)$$

where G_i is a subject-level grouping variable coded 0 or 1 (half of the 100 or 500 subjects had $G_i = 0$, while the other half were coded as 1), T_j represents time and is coded 0, 1, 2, 4 for the four timepoints, respectively. The values of the regression coefficients on the marginal scale (PA) were set as: $\beta_0 = -1.1259$, $\beta_1 = 0.4505$, $\beta_2 = 0.027$, and $\beta_3 = -0.1147$. These were selected to yield response proportions of .130, .136, .142, and .155 ($G_i = 0$) and .

250, .223, .198, and .155 ($G_i = 1$), for the two groups and four timepoints, respectively. Additionally, the random intercept variance was set as $\sigma_v^2 = 1$ to reflect an intraclass correlation, $r = \sigma_v^2 / (\sigma_v^2 + 1)$, of 0.5 for the repeated observations. Thus, the SS regression coefficients would be equal to the PA coefficients multiplied by a factor of $\sqrt{2}$.

To generate the binary data, four random normal variates were obtained for each of the 100 or 500 subjects with mean 0 and variance-covariance matrix $\sigma_v^2 \mathbf{1}\mathbf{1}' + \mathbf{I}$. These random variates were then dichotomized in accordance with the marginal response probabilities for the two groups across the four timepoints given in the preceding paragraph.

The closed form marginal estimates of the random intercept probit model were obtained using SAS PROC NLMIXED, while the proposed marginalization parameters were estimated with Supermix (Hedeker et al., 2008) based on a random intercept probit model. For both, adaptive 7-point Gauss-Hermite quadrature was used. Table 2 presents the simulation results for both sample sizes of 100 and 500 subjects. In Table 2, bias is the difference between the average parameter estimate and the true value, coverage is the proportion of 1000 datasets in which the 95% confidence interval included the true value, width is the average width of the 95% confidence interval, and root mean squared error (RMSE) for a given parameter θ equals $\sqrt{E[(\hat{\theta} - \theta)^2]}$. As can be seen, for both sample sizes of 100 and 500, the proposed marginalization method provides near-identical results to the closed form marginal probit estimates. Both approaches recover the parameter values well with low bias and decent coverage.

4.2 Random-intercept and trend logistic model

The specifications for this simulation study were also based on the example and data analysis presented in Section 5. Here, 1000 datasets were simulated, each with 500 subjects and 4 repeated observations according to a marginal logistic regression model

$$\log \left[\frac{\pi_{ij}^{pa}}{1 - \pi_{ij}^{pa}} \right] = \beta_0 + \beta_1 G_i + \beta_2 T_j + \beta_3 (G_i \times T_j), \quad (18)$$

where G_i and T_j were as previously defined. Again, the marginal regression coefficients were determined based on the same response proportions for the two groups across the four timepoints. Here, on the logistic scale, they were: $\beta_0 = -1.9$, $\beta_1 = 0.8$, $\beta_2 = 0.05$, and $\beta_3 = -0.2$. Additionally, the variance-covariance structure for the mixed model allowed for

random intercepts ($\sigma_I^2 = 3$), time-trends ($\sigma_T^2 = 1$), and covariance ($\sigma_{IT} = 0.08$); these values were based on and similar to the variance covariance estimates for the example in Section 5. Similar to the probit simulation, four random normal variates were obtained for each of the 500 subjects with mean 0 and variance-covariance matrix $\mathbf{Z}\mathbf{\Sigma}_v\mathbf{Z}' + \pi^2/3\mathbf{I}$, which were then dichotomized according to the same marginal response probabilities for the two groups across the four timepoints.

Table 3 presents results from the following: the GEE approach using an unstructured working correlation matrix, the proposed marginalization method, and a marginalized logistic-probit-normal model as described by Griswold and Zeger (2004) including random subject intercepts and trends. In this latter model, the regression coefficients are on the marginal logistic scale, while the random effects variance-covariance parameters are on the probit scale. As Griswold and Zeger (2004) note, the probit is a convenient choice for the variance-covariance parameters for computational simplicity (*i.e.*, to avoid an additional numerical integration). The parameters of the GEE and marginalized multilevel models were estimated in SAS (PROC GENMOD and NLMIXED, respectively), while the proposed marginalization parameters were estimated with Supermix based on a random intercept and trend logistic model. For both the marginalized multilevel model and the mixed logistic model, adaptive 7-point Gauss-Hermite quadrature was used. For the GEE, robust standard errors were used to obtain the coverage probabilities.

As can be seen from Table 3, the results from all three models are near-identical, with low levels of bias and good coverage levels (*i.e.*, near the nominal 95% level). As one would expect, the widths are appreciably smaller for the time-related parameters than for the intercept-related parameters. Finally, RMSE levels are relatively small, especially for the time-related parameters.

To examine the effect of missing data on the proposed approach, we generated data under a MAR mechanism. Using the same data generation procedure described above, we additionally imposed the following specifications. For a $G_i = 0$ (or $G_i = 1$) observation, if their first outcome was 0 (or 1), they had a 60% (or 80%) chance of dropping out (*i.e.*, being missing at the next three timepoints). Similarly, for a $G_i = 0$ (or $G_i = 1$) observation, if their second outcome was 0 (or 1), they had a 80% (or 85%) chance of dropping out (*i.e.*, being missing at the next two timepoints). Finally, for a $G_i = 0$ (or $G_i = 1$) observation, if their third outcome was 0 (or 1), they had a 90% (or 90%) chance of being missing at the last timepoint. Thus, the two groups were treated in an opposite manner based on their observed outcomes, and the amount of missing data was rather large. The observed proportions were .129, .170, .216, and .298 ($G_i = 0$) and .251, .176, .123, and .067 ($G_i = 1$), for the two groups and four timepoints, respectively, indicating very different trends across time for the two groups, relative to the complete data scenario. The retention rates for the two groups were 100%, 48%, 25%, and 14% ($G_i = 0$) and 100%, 80%, 69%, and 63% ($G_i = 1$) for the two groups and four timepoints, respectively. Thus, both the reasons and amount of missing data varied considerably for the two groups.

Table 4 lists the simulation results for the GEE and proposed marginalization approach. As can be seen, the time-related parameters are rather poorly estimated by GEE, and the coverage rates are off as well, especially for the group by time interaction. Conversely, the proposed marginalization approach does reasonably well with low bias and close to nominal coverage rates. For both, the widths are appreciably larger than their counterparts under complete data, reflecting the effect of missing data. RMSE values are also larger than under the complete data scenario, and tend to favor the proposed marginalization approach.

5. Example

As mentioned in Section 1, Gruder et al. (1993) present data from a smoking-cessation study in which 489 subjects were assessed for their smoking status (y/n) at four timepoints: post-intervention, and 6, 12, and 24 months later. Subjects were randomized to either a control, discussion, or social support condition. Here, for simplicity, subjects in the two treatment conditions of discussion and social support are combined into one group. The missing data pattern across time was monotonic, in that subjects who were missed at a given follow-up were also missing thereafter. The number of subjects measured at the four timepoints was 489, 454, 429, and 372, respectively. Smoking abstinence rates for the two groups across time are presented in Table 5. These show that the control rates are somewhat constant across time, while the treatment group rates decline after the end-of-program timepoint. Also included in Table 5 are the observed correlations of the smoking abstinence outcomes across time. These pairwise correlations are clearly not equal, which would violate the assumption of a random-intercept model.

The outcome Y_{ij} was coded 0=smoking and 1=quit, and several logistic regression models were applied to these data with covariates for time, group, and group by time interaction. The time variable was coded so that a unit represents 6 months ($T_j = 0, 1, 2, 4$), and the group variable was an indicator of treatment ($G_j = 0$ for control, and $G_j = 1$ for treatment). Table 6 presents estimates from the following: the GEE approach using an independent working correlation matrix (GEE IND) and an unstructured working correlation matrix (GEE UN), a mixed effects model with random subject intercepts and time trends (MRM), a marginalized logistic-probit-normal model including random subject intercepts and trends (MM), and the proposed marginalization approach (PM).

As can be seen in Table 6, the subject-specific MRM estimates are not on the same numerical scale as the population-averaged estimates of the GEE. However, the proposed marginalization of these MRM estimates yields results that are in close agreement with those obtained from both the GEE and the marginalized logistic-probit-normal models. An advantage for the user of the proposed marginalization is that both subject-specific and population-averaged estimates and inferences can be obtained from the same analysis. In the present example, both approaches indicate a significant group effect such that there is a benefit for treatment at post-intervention (e.g., from GEE UN, $z = 0.905/0.275 = 3.29$). However, the marginal models indicate a significant group by time interaction such that the treatment benefit diminishes across time. Interestingly, the MRM does not yield a significant group by time interaction. Thus, the conclusions are different depending on whether one controls for the subject effects or not. Such situations in which the conclusions vary between marginal and conditional approaches have been previously reported in the literature (Agresti, 1989; Lindsey and Lambert, 1998).

In Table 6, we have only presented the regression coefficient estimates, but the random intercept and trend variance estimates of the MRM were large ($\hat{\sigma}_I^2 = 3.774$, standard error = 1.214, and $\hat{\sigma}_T^2 = 1.302$, standard error = 0.451) reflecting a good deal of subject heterogeneity in the data. The covariance between the two random effects was small ($\hat{\sigma}_{IT} = 0.170$, standard

error = 0.359) indicating that a subject's post-intervention smoking status was not related to their trend across time.

6. Discussion

In this paper, we have presented a simple approach using numerical quadrature to produce marginal estimates and standard errors for the regression coefficients from mixed model estimates. Specifically, we have presented results for several link functions applied to binary outcome data. We have also derived results (not shown) for proportional and non-proportional odds models of ordinal data under all of the presented link functions, as well as for nominal outcomes under the logistic link. An advantage of the proposed approach is that both subject-specific and population-averaged estimates (and standard errors) can be obtained from one analysis with minimal additional computation. In addition, although we have described our approach in terms of two-level (longitudinal) data, extending to three or more levels is relatively straightforward.

The limited simulation studies confirmed that for a random intercept model the marginalized estimates were near-identical to closed form marginal probit estimates. Extending to a random intercept and trend logistic model, the marginalized estimates also agreed quite closely to marginalized multilevel and GEE estimates. A further simulation allowing for missing data, under MAR missingness, showed that the GEE estimates of the time-related parameters were poorly estimated, whereas the proposed approach produced reasonable results. This agrees with the theory that indicates that mixed models using full likelihood estimation provide valid results under MAR, whereas GEE only does so under MCAR. It should be emphasized that the amount of missing data in the simulation was rather large and strictly followed an MAR mechanism. In addition, extensions to allow MAR missingness under the GEE approach have been developed based on weighting (Robins et al., 1995) and imputation (Paik, 1997). Nonetheless, the proposed approach yields reasonable marginal estimates under MAR missingness with rather minimal effort. An important caveat is that the variance-covariance structure of the data must be properly specified in a mixed model or biased results can be obtained (Hedeker and Gibbons, 2006; Kwok et al., 2007).

In the example presented, we found close agreement of our marginalized estimates to those obtained from marginalized multilevel models. This is not surprising as our approach can be viewed as essentially an alternative way of producing marginalized multilevel estimates. However, given that software for marginalized multilevel models is somewhat limited, whereas software for mixed models is widely available, our approach can be used in conjunction with mixed model estimates from the many different programs that exist. The marginalized results presented in this article were obtained using Supermix (Hedeker et al., 2008), which produces both SS and PA estimates automatically.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was supported by National Heart, Lung, and Blood Institute grant R01 HL121330 (Hedeker). The authors thank Dr. Robin Mermelstein for use of the smoking cessation dataset. The authors also thank the Editor, Associate Editor, and Referee for several excellent points raised in the review of this paper which greatly improved it.

References

- Agresti A. A survey of models for repeated ordered categorical response data. *Statistics in Medicine*. 1989; 8:1209–1224. [PubMed: 2814070]
- Agresti, A. *Categorical Data Analysis*. 2nd. Wiley; New York: 2002.
- Brown, H., Prescott, R. *Applied Mixed Models in Medicine*. 2nd. Wiley; Chichester, UK: 2006.
- Demidenko, E. *Mixed Models Theory and Applications*. Wiley; Hoboken, NJ: 2004.
- Demirtas H. Assessment of relative improvement due to weights within generalized estimating equations framework for incomplete clinical trials data. *Journal of Biopharmaceutical Statistics*. 2004; 14:1085–1098. [PubMed: 15587981]
- Fitzmaurice, GM., Laird, NM., Ware, JH. *Applied Longitudinal Analysis*. 2nd. Wiley; Hoboken, NJ: 2011.
- Goldstein, H. *Multilevel Statistical Models*. 4th. Wiley; Chichester UK: 2011.
- Griswold ME, Swihart BJ, Caffo BS, Zeger SL. Practical marginalized multilevel models. *Stat*. 2013; 2:129–142.
- Griswold, ME., Zeger, SL. On marginalized multilevel models and their computation. John Hopkins University; 2004. Department of Biostatistics Working Papers, Working Paper 99
- Gruder CL, Mermelstein RJ, Kirkendol S, Hedeker D, Wong SC, Schreckengost J, Warnecke RB, Burzette R, Miller TQ. Effects of social support and relapse prevention training as adjuncts to a televised smoking cessation intervention. *Journal of Consulting and Clinical Psychology*. 1993; 61:113–120. [PubMed: 8450096]
- Heagerty PJ, Zeger SL. Marginalized multilevel models and likelihood inference. *Statistical Science*. 2000; 15:1–26.
- Hedeker, D., Gibbons, RD. *Longitudinal Data Analysis*. Wiley; Hoboken, NJ: 2006.
- Hedeker, D., Gibbons, RD., du Toit, M., Cheng, Y. *SuperMix: Mixed Effects Models*. Scientific Software International; Lincolnwood, IL: 2008.
- Kwok O, West SG, Green SB. The impact of misspecifying the within-subject covariance structure in multiwave longitudinal multilevel models: A Monte Carlo study. *Multivariate Behavioral Research*. 2007; 42:557–592.
- Lindsey JK, Lambert P. On the appropriateness of marginal models for repeated measurements in clinical trials. *Statistics in Medicine*. 1998; 17:447–469. [PubMed: 9496722]
- Magnus, JR. *Linear Structures*. Oxford University Press; New York: 1988.
- McCulloch, CE., Searle, SR., Neuhaus, JM. *Generalized, Linear, and Mixed Models*. 2nd. Wiley; Hoboken, NJ: 2008.
- Molenberghs, G., Verbeke, G. *Models for Discrete Longitudinal Data*. Springer-Verlag; New York, NY: 2005.
- Neuhaus JM, Kalbfleisch JD, Hauck WW. A comparison of cluster-specific and population-averaged approaches for analyzing correlated binary data. *International Statistical Review*. 1991; 59:25–35.
- Paik MC. The generalized estimating equation approach when data are not missing completely at random. *Journal of the American Statistical Association*. 1997; 92:1320–1329.
- Preisser JS, Lohman KK, Rathouz PJ. Performance of weighted estimating equations for longitudinal binary data with drop-outs missing at random. *Statistics in Medicine*. 2002; 21:3035–3054. [PubMed: 12369080]
- Raudenbush, SW., Bryk, AS. *Hierarchical Linear Models: Applications and Data Analysis*. 2nd. Sage; Thousand Oaks, CA: 2002.

- Robins JM, Rotnitzky A, Zhao LP. Analysis of semi-parametric regression models for repeated outcomes in the presence of missing data. *Journal of the American Statistical Association*. 1995; 90:106–121.
- Skrondal, A., Rabe-Hesketh, S. *Generalized Latent Variable Modeling*. Chapman and Hall/CRC; Boca Raton, FL: 2004.
- Stroud, AH., Sechrest, D. *Gaussian Quadrature Formulas*. Prentice Hall; Englewood Cliffs, NJ: 1966.
- Zeger SL, Liang KY. Longitudinal data analysis for discrete and continuous outcomes. *Biometrics*. 1986; 42:121–130. [PubMed: 3719049]
- Zeger SL, Liang KY, Albert PS. Models for longitudinal data: A generalized estimating equation approach. *Biometrics*. 1988; 44:1049–1060. [PubMed: 3233245]

Table 1

Corresponding elements for several link functions, where $F = CDF$, $\pi = Pr(Y = 1)$, and $\lambda = x' \beta^{ss} + z' \mathbf{TB}$.

link	F^{-1}	F	F^{-1}	F
logit	$\log(\pi(1 - \pi))$	$1 / (1 + \exp(-\lambda))$	$1 / (\pi(1 - \pi))$	$F(\lambda)(1 - F(\lambda))$
probit	$\Phi^{-1}(\pi)$	$\Phi(\lambda)$	$1/(\phi(\Phi^{-1}(\pi)))$	$\phi(\lambda)$
cloglog	$\log(-\log(1 - \pi))$	$1 - \exp(-\exp(\lambda))$	$1/((\pi - 1) \log(1 - \pi))$	$(F(\lambda) - 1)(-\exp(\lambda))$
loglog	$\log(-\log(\pi))$	$\exp(-\exp(\lambda))$	$1/(\pi \log(\pi))$	$F(\lambda)(-\exp(\lambda))$

Table 2

Random-intercept probit model simulation results (1000 datasets with either $N=100$ or $N=500$). CF: closed form marginalization; PM: proposed marginalization.

N	Model	Parameter	True value	Bias	Coverage	RMSE	Width
100	CF	Intercept	-1.1259	-0.0157	0.951	0.1940	0.7385
		T_j	0.0270	0.0019	0.950	0.0651	0.2530
		G_i	0.4505	0.0093	0.947	0.2556	0.9832
		$G_i \times T_j$	-0.1147	-0.0030	0.956	0.0891	0.3465
	PM	Intercept	-1.1259	-0.0150	0.950	0.1944	0.7413
		T_j	0.0270	0.0019	0.951	0.0651	0.2529
		G_i	0.4505	0.0088	0.949	0.2556	0.9856
		$G_i \times T_j$	-0.1147	-0.0028	0.957	0.0891	0.3462
500	CF	Intercept	-1.1259	-0.0052	0.952	0.0839	0.3276
		T_j	0.0270	-0.0002	0.951	0.0286	0.1113
		G_i	0.4505	-0.0022	0.960	0.1098	0.4368
		$G_i \times T_j$	-0.1147	0.0021	0.941	0.0393	0.1523
	PM	Intercept	-1.1259	-0.0048	0.952	0.0839	0.3280
		T_j	0.0270	-0.0002	0.952	0.0285	0.1111
		G_i	0.4505	-0.0027	0.960	0.1097	0.4371
		$G_i \times T_j$	-0.1147	0.0023	0.941	0.0393	0.1521

Table 3

Logistic regression simulation results (1000 datasets with N=500). GEE UN: GEE with unstructured working correlation structure; MM: marginalized multilevel model with random intercept and time effects; PM: proposed marginalization from a mixed model with random intercepts and time effects.

Model	Parameter	True value	Bias	Coverage	RMSE	Width
GEE UN	Intercept	-1.90	-0.0032	0.949	0.152	0.606
	G_i	0.80	0.0002	0.944	0.204	0.781
	T_j	0.05	-0.0033	0.950	0.056	0.214
	$G_i \times T_j$	-0.20	0.0007	0.947	0.077	0.292
MM	Intercept	-1.90	0.0014	0.947	0.156	0.605
	G_i	0.80	-0.0034	0.942	0.203	0.777
	T_j	0.05	-0.0078	0.933	0.057	0.216
	$G_i \times T_j$	-0.20	0.0030	0.953	0.077	0.291
PM	Intercept	-1.90	-0.0055	0.956	0.148	0.595
	G_i	0.80	-0.0112	0.946	0.200	0.770
	T_j	0.05	-0.0062	0.947	0.056	0.215
	$G_i \times T_j$	-0.20	0.0096	0.944	0.076	0.288

Table 4

Logistic regression simulation results (1000 datasets with N=500) with missing data. GEE UN: GEE with unstructured working correlation structure; PM: proposed marginalization from a mixed model with random intercepts and time effects

Model	Parameter	True value	Bias	Coverage	RMSE	Width
GEE UN	Intercept	-1.90	-0.0141	0.9481	0.1889	0.7335
	G_i	0.80	-0.0155	0.9481	0.2380	0.8865
	T_j	0.05	-0.0576	0.9888	0.1544	0.6245
PM	$G_i \times T_j$	-0.20	0.1303	0.8576	0.2518	0.6507
	Intercept	-1.90	-0.0121	0.9500	0.1767	0.6745
	G_i	0.80	-0.0089	0.9490	0.2270	0.8662
	T_j	0.05	-0.0030	0.9540	0.1133	0.4437
	$G_i \times T_j$	-0.20	0.0115	0.9340	0.1573	0.5952

Table 5

Point prevalence rates and correlations of smoking abstinence over time.

	Post-intervention	6 months	12 months	24 months
Control group ($N = 109, 97, 92, 77$)	0.174	0.072	0.185	0.182
Treatment group ($N = 380, 357, 337, 295$)	0.345	0.182	0.196	0.217
<i>Correlations</i>				
Post-intervention	1.000	0.332	0.287	0.263
6 months	0.332	1.000	0.481	0.337
12 months	0.287	0.481	1.000	0.493
24 months	0.263	0.337	0.493	1.000

Table 6

Logistic regression estimates (standard errors), and p-values. GEE IND: GEE with independent working correlation structure; GEE UN: GEE with unstructured working correlation structure; MRM: Mixed model with random intercept and time effects; MM: Marginalized multilevel model with random intercept and time effects; PM: Proposed marginalized estimates from MRM.

Model	Intercept	T	G	$G \times T$
GEE IND	-1.839	0.074	0.902	-0.221
	(0.256)	(0.094)	(0.275)	(0.103)
	0.001	0.431	0.001	0.032
GEE UN	-1.841	0.075	0.905	-0.201
	(0.256)	(0.092)	(0.275)	(0.101)
	0.001	0.418	0.001	0.047
MRM	-2.785	-0.479	1.478	-0.331
	(0.430)	(0.270)	(0.410)	(0.247)
	0.001	0.076	0.001	0.180
MM	-1.889	0.109	0.941	-0.211
	(0.240)	(0.086)	(0.260)	(0.095)
	0.001	0.206	0.001	0.027
PM	-1.885	0.092	0.945	-0.232
	(0.236)	(0.093)	(0.257)	(0.102)
	0.001	0.321	0.001	0.023