

Lecture Notes in Statistics

179

Edited by P. Bickel, P. Diggle, S. Fienberg, U. Gather,
I. Olkin, and S. Zeger

D.Y. Lin P.J. Heagerty
Editors

Proceedings of the Second Seattle Symposium in Biostatistics

Analysis of Correlated Data

 Springer

D.Y. Lin
School of Public Health
Department of Biostatistics, CB #7420
University of North Carolina
Chapel Hill, NC 27599-7420
USA
lin@bios.unc.edu

P.J. Heagerty
School of Public Health and Community
Medicine
Department of Biostatistics, Box 357232
University of Washington
Seattle, WA 98195
heagerty@biostat.washington.edu

Library of Congress Cataloging-in-Publication Data
Seattle Symposium in Biostatistics: Analysis of Correlated Data (2nd :
2000 : Seattle, Wash.)

Proceeding of the second Seattle Symposium in Biostatistics : analysis
of correlated data / edited by D.Y. Lin, P.J. Heagerty.

p. cm. — (Lecture notes in statistics, (Springer-Verlag) ; v.
179)

Includes bibliographical references (p.).

I. Biometry—Congresses. I. Lin, D. Y. (Danyu Y.) II. Heagerty, P. J.

III. Title. IV. Series.

QH323.5.S425 2000
570'.1'5195—dc22

2004052209

ISBN 978-0-387-20862-6 ISBN 978-1-4419-9076-1 (eBook)

DOI 10.1007/978-1-4419-9076-1

Printed on acid-free paper.

© 2004 Springer Science+Business Media New York

Originally published by Springer Science+Business Media Inc. in 2004

All rights reserved. This work may not be translated or copied in whole or in part without the
written permission of the publisher Springer-Science+Business Media, LCC except for brief
excerpts in connection with reviews or scholarly analysis.

Use in connection with any form of information storage and retrieval, electronic adaptation, com-
puter software, or by similar or dissimilar methodology now known or hereafter developed is for-
bidden.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if
they are not identified as such, is not to be taken as an expression of opinion as to whether or not
they are subject to proprietary rights.

9 8 7 6 5 4 3 2 1

SPIN 10963296

springeronline.com

Contents

Preface

N. Breslow. Whither PQL?	1
O. B. Linton, E. Mammen, X. Lin and R. J. Carroll. Correlation and marginal longitudinal kernel nonparametric regression.	23
G. Tang, R. J. A. Little and T. E. Raghunathan. Analysis of multivariate monotone missing data by a pseudolikelihood method.	35
L. Chen, L. J. Wei and M. I. Parzen. Quantile regression for correlated observations.	51
Z. Feng, T. Braun and C. McCulloch. Small sample inference for clustered data.	71
W. Jiang and B. W. Turnbull. Some applications of indirect inference to longitudinal and repeated events data.	89
D. Oakes. On characterizing joint survivor functions by minima.	107
R. L. Prentice, Z. Moodie and J. Wu. Nonparametric estimation of the bivariate survivor function.	113
J. A. Wellner, Y. Zhang and H. Liu. A semiparametric regression model for panel count data: When do pseudo-likelihood estimators become badly inefficient?	143
M. Gail and N. Chatterjee. Some biases that may affect kin-cohort studies for estimating the risks from indentified disease genes.	175
J. Robins. Optimal structural nested models for optimal sequential decisions.	189
Addresses for Contact Authors	327
List of Referees	329

Preface

The First Seattle Symposium in Biostatistics: Survival Analysis was held on November 20 and 21, 1995 in honor of the twenty-fifth anniversary of the University of Washington (UW) School of Public Health and Community Medicine. The event was a big success. Exactly 5 years later, the *Second Seattle Symposium in Biostatistics: Analysis of Correlated Data* was held on November 20 and 21, 2000, and it was also very successful. The event was sponsored by Pfizer and co-sponsored by the UW School of Public Health and Community Medicine and the Division of Public Health Sciences, the Fred Hutchinson Cancer Research Center (FHCRC). The symposium featured keynote lectures by Norman Breslow, David Cox and Ross Prentice, as well as invited talks by Raymond Carroll, Peter Diggle, Susan Ellenberg, Ziding Fan, Mitchell Gail, Stephen Lagakos, Nan Laird, Kung-Yee Liang, Roderick Little, Thoms Louis, David Oakes, Robert O'Neill, James Robins, Bruce Turnbull, Mei-Cheng Wang and Jon Wellner. There were 336 attendees. In addition, 100 people attended the short course *Analysis of Longitudinal Data* taught by Patrick Heagerty and Scott Zeger on November 18, and 96 attended the short course *Analysis of Multivariate Failure Time Data* taught by Danyu Lin, Lee-Jen Wei and Zhiliang Ying on November 19.

When the UW School of Public Health and Community Medicine was formed in 1970, biostatistics as a discipline was only a few years old. In the subsequent thirty years, both the field and the UW Department of Biostatistics have evolved in many exciting ways. The Department had only seven faculty when it moved from the School of Medicine to the new School of Public Health and Community Medicine in 1970. The faculty roster currently lists 43 regular and research faculty and 34 adjunct and affiliate faculty. Ed Perrin was the Department Chair in 1970, succeeded by Donovan Thompson, Norman Breslow and presently Thomas Fleming. The faculty have been actively involved in methodological and collaborative research in addition to graduate teaching. The choice of *Analysis of Correlated Data* as the theme for the *Second Seattle Symposium in Biostatistics* was a tribute to the significant contributions made by the UW and FHCRC faculty to this important area of statistical science.

The Symposium Organizing Committee consisted of Thomas Fleming, Patrick Heagerty, Gordon Lan, Danyu Lin (Chair), Art Peterson and Lianng Yuh. The staff of the Biostatistics Department, including Diane Ames, David Fetrow, Michael Heroux, Barbara Jensen, Cynthia Marks, Alexandra MacKenzie, Jamie Miller, Rachel Rodke, Elaine Riot, Lian Schmidt

and Eleanor Schweihs, provided great administrative support to the symposium. The UW President Richard McCormick, the School Dean Patricia Wahl, and the Department Chair Thomas Fleming delivered the opening remarks. The scientific sessions were chaired by Norman Breslow, Patrick Heagerty, Gordon Lan, Danyu Lin, Yasuo Ohashi, Masahiro Takeuchi and Elizabeth Thompson. We are grateful to the aforementioned people as well as all the speakers and participants for making the symposium a great success.

This volume contains most of the papers presented at the symposium. These papers encompass recent methodological advances on several important topics, such as longitudinal data, multivariate failure time data and genetic data, as well as innovative applications of the existing theory and methods. This collection serves as a reference for those working in the area of correlated data analysis.

Each of the 11 papers in this volume was refereed by two peer reviewers, and their comments were incorporated by the authors into the final versions of the papers. The referees are listed at the end of this book. We are indebted to them for their time and efforts. We also appreciate the guidance and assistance by John Kimmel of Springer-Verlag as well as the secretarial support by April Smyth and Pete Mesling during the preparation of this volume.

Finally, we would like to acknowledge the scientific and financial contributions by Pfizer. Without their generous support, it would have been impossible to hold this symposium.

D. Y. Lin
P. J. Heagerty

Whither PQL?

Norman Breslow

Department of Biostatistics, University of Washington

ABSTRACT Generalized linear mixed models (GLMM) are generalized linear models with normally distributed random effects in the linear predictor. Penalized quasi-likelihood (PQL), an approximate method of inference in GLMMs, involves repeated fitting of linear mixed models with “working” dependent variables and iterative weights that depend on parameter estimates from the previous cycle of iteration. The generality of PQL, and its implementation in commercially available software, has encouraged the application of GLMMs in many scientific fields. Caution is needed, however, since PQL may sometimes yield badly biased estimates of variance components, especially with binary outcomes.

Recent developments in numerical integration, including adaptive Gaussian quadrature, higher order Laplace expansions, stochastic integration and Markov chain Monte Carlo (MCMC) algorithms, provide attractive alternatives to PQL for approximate likelihood inference in GLMMs. Analyses of some well known datasets, and simulations based on these analyses, suggest that PQL still performs remarkably well in comparison with more elaborate procedures in many practical situations. Adaptive Gaussian quadrature is a viable alternative for nested designs where the numerical integration is limited to a small number of dimensions. Higher order Laplace approximations hold the promise of accurate inference more generally. MCMC is likely the method of choice for the most complex problems that involve high dimensional integrals

1 Introduction

Penalized Quasi-Likelihood is a technique for approximate inference in GLMMs and is not a rigorous statistical method in its own right.[33, p. 390, emphasis added]

The generalized linear model or GLM [35] is a prime tool of the applied statistician. It brings the power and flexibility of linear regression modeling to the analysis of data with outcomes, particularly discrete outcomes, that do not satisfy the conventional assumptions of least squares. The linear mixed model or LMM, with its multiple levels of random variation and best linear unbiased prediction of random effects [19], dominates statistical theory and applications in diverse fields including animal breeding and education. During the past decade these two models have been fused into a hybrid body of statistical theory and methodology known as the generalized linear mixed model or GLMM.[46, 40, 5, 49, 12, 31, 16, 29, 30] An even more general formulation, known as the hierarchical generalized linear model or HGLM, encompasses both normal and non-normal probability distributions for the random effects.[22, 23]

GLM and LMM parameter estimates are obtained from estimating equations that are unbiased under simple moment conditions and that may be solved by iterative solution of systems of linear equations. For the GLMM, by contrast, the specification of normally distributed random effects intrinsically defines the marginal likelihood and its logarithmic derivatives. The fact that the integrals in the GLMM estimating equations cannot be evaluated in closed form has seriously limited GLMM applications. Until recently the only available commercial software was the EGRET program [9] that implemented the logistic-normal model for clustered binary outcomes, unit level covariates and a cluster level random intercept. Thus substantial interest was generated by the work of Schall [40], Breslow and Clayton [5], Wolfinger [48] and others who developed a general approach to approximate inference. Their “penalized quasi-likelihood” or PQL procedure involved repeated fitting of the LMM using a working outcome variable and iterative weights that mimicked the standard iterative least squares algorithm used to fit the GLM.[28, §2.5] It was disseminated in macros written for several commercially distributed LMM programs: the GLIMMIX macro for PROC MIXED in SAS [26]; the PQL option for MLwiN [37]; and the HLM series distributed by SSI [38]. The IR-REML macro in GENSTAT [32] facilitated fitting of both GLMMs and HGLMs. This stimulated increasing use of these procedures in old disciplines such as sociology, where hierarchical models were already familiar, and in new ones like epidemiology [17], where they were just being discovered.

As usual when software for complicated statistical inference procedures is broadly disseminated, there is potential for abuse and misinterpretation. In spite of the fact that PQL was initially advertised as a procedure for *approximate* inference in GLMMs, and its tendency to give seriously biased estimates of variance components and *a fortiori* regression parameters with binary outcome data was emphasized in multiple publications [5, 6, 24], some statisticians seemed to ignore these warnings and to think of PQL as synonymous with GLMM.[7] In an apparent reaction to these developments, and to the algorithm’s acknowledged shortcomings for binary

outcome data, the authors of one recent textbook have recommended that PQL “not be used in practice”. [30, p.234]

The purpose of this review is to take stock of PQL as a tool of the applied statistician now that some years have passed since it was first implemented in commercial software. In the interim, substantial advances have taken place in statistical computing. “True” maximum likelihood (ML) estimation is now available for a much wider range of problems by using numerical integration to calculate marginal likelihoods and solve score equations. In particular, the adaptive Gaussian quadrature methods [27, 36] implemented in SAS PROC NL MIXED [45] apply to clustered data problems where the dimensionality of the required integrations is in the low single digits. Higher order Laplace approximations [39], implemented for the logistic-normal model in the latest HLM program [38], may prove to be just as accurate as quadrature and more widely applicable.

Recent Monte Carlo approaches to numerical integration include Monte Carlo relative likelihood [13], Monte Carlo EM [29, 3] and Monte Carlo Newton-Raphson [20]. Kuk and Cheng [21] provide an excellent, comprehensive review of these stochastic procedures. Their use in practice to date has been limited by their longer computing times and the fact that none have yet been implemented in standard software packages. Booth and Hobert [3] argue that their “automated” Monte Carlo EM algorithm is an improvement on the Markov chain Monte Carlo (MCMC) version. It facilitates assessment of convergence and thus removes one of the main impediments to commercial implementation. Hierarchical Bayes procedures, which also depend on MCMC to evaluate posterior distributions, have been implemented in available, supported software and are increasingly used in applications. [8, 44, 34] These Monte Carlo methods will undoubtedly see much greater use with continuing improvements in computing technology. In view of their greater complexity, however, and the desire to keep this review focussed on the most immediate competitors to PQL, further discussion of Monte Carlo methods is left to investigators who are more familiar with their properties. Comparisons with the “h-likelihood” methodology of Lee and Nelder [22, 23] for inference in HGLMs also have been left for others.

2 GLMMs and PQL

The GLMM is a model for the hierarchical regression analysis of a series of n univariate response measurements y_i on p -dimensional covariates x_i associated with fixed effects and q -dimensional covariates z_i associated with random effects of interest ($i = 1, \dots, n$). Conditional on the unobserved values of a q -vector b of random effects, and on all the covariates, the y_i are assumed to be independent observations with means and variances specified

by a GLM.[28] Specifically we suppose

$$E(y_i|b) = \mu_i^b = h(\eta_i^b) = h(x_i^T \alpha + z_i^T b)$$

$$\text{Var}(y_i|b) = \frac{\phi}{a_i} v(\mu_i^b)$$

where $g = h^{-1}$ is the link function that relates the conditional means μ_i^b to the linear predictors η_i^b ; $v(\cdot)$ is the variance function that relates the conditional means and variances to each another; ϕ is a scale factor assumed equal to one for the standard binomial and Poisson models; and a_i is a prior weight such as a binomial denominator. Specification of the model is completed by the assumption that b follows a q -dimensional normal distribution with mean 0 and variance matrix $D(\theta)$ depending on a vector of dispersion parameters θ . Examples of typical GLMM applications are considered in Sections 4 and 5.

The objective function for estimation of the GLMM parameters is the integrated quasi-likelihood $L(\alpha, \theta)$ given by

$$L = \frac{1}{\sqrt{(2\pi)^q |D(\theta)|}} \int_{\mathbb{R}^q} \exp \left[-\frac{1}{2\phi} \sum_{i=1}^n d_i(y_i, \mu_i^b) - \frac{1}{2} b^T D^{-1}(\theta) b \right] db \quad (1)$$

where

$$d_i(y, \mu) = -2a_i \int_y^\mu \frac{y-u}{v(u)} du$$

denotes the weighted deviance.[28] If Y is Gaussian and $g(\cdot)$ the identity, the integral in (1) is normal and may be evaluated in closed form. Otherwise, maximization of this expression is intrinsically complicated by the integrations that must be performed numerically at each cycle of iteration. One approach to the integration, which eventually leads to the PQL algorithm, is to make a Laplace approximation. The term in square brackets in (1), the logarithm of the ‘‘penalized quasi-likelihood’’, is replaced by its quadratic expansion in b about the value \bar{b} at which it is maximized. Components of \bar{b} serve as predictors of the random effects. After some adjustments to the resulting normal integral, application of Fisher scoring to determine $(\hat{\alpha}, \bar{b})$ as a function of θ leads to the familiar mixed model equations for joint estimation of fixed and random effects, as originally derived by Henderson [19], but now involving a working vector Y^* and iterative weights w_i . Further approximations lead to the standard REML equations for θ . Specifically, with $\hat{\mu}_i^b = h(x_i^T \hat{\alpha} + z_i^T \bar{b})$,

$$Y_i^* = x_i^T \hat{\alpha} + z_i^T \bar{b} + (y_i - \hat{\mu}_i^b) g'(\hat{\mu}_i^b)$$

and

$$w_i = \phi a_i [g'(\hat{\mu}_i^b)]^2 v(\hat{\mu}_i^b)^{-1},$$

the algorithm repeatedly applies mixed model REML estimation to the normal theory problem

$$Y^* = X\alpha + Zb + \varepsilon, \quad b \sim \mathcal{N}(0, D(\theta)), \quad \varepsilon \sim \mathcal{N}(0, W^{-1})$$

where $W = \text{diag}(w_i)$. See Breslow and Clayton [5] for details.

Although PQL yields REML estimates of variance components and regression coefficients in the Gaussian linear case, in general it only provides an approximation to these quantities. For the simplest GLMM involving clustered data with a single dispersion component θ , Breslow and Lin [6] expanded both the efficient score based on the true profile log-likelihood function, and the PQL variance estimating equation, in Taylor series about $\theta = 0$. They thereby showed that the asymptotic bias in the PQL estimator $\hat{\theta}_p$ was a nearly linear function of θ in a neighborhood of the origin. By determining the slope of this linear relationship, which is estimable from the standard GLM fit assuming $\theta = 0$, they derived a correction factor for $\hat{\theta}_p$ that removed the asymptotic bias for small θ at the cost of some increase in variability. Lin and Breslow [24] extended this work for models with multiple variance components, deriving a matrix correction factor, and termed the resulting procedure corrected PQL or CPQL.

An alternative derivation of the PQL algorithm developed by Schall [40] and others uses a linearization of the conditional mean as a function of fixed and random effects. Consider, for example, the two-level model with I clusters having n_i observations per cluster, $i = 1, \dots, I$, and random effects b_i assumed independent between clusters. The j^{th} observation in cluster i may be written

$$y_{ij} = \mu_{ij}^b + \varepsilon_{ij} = h(x_{ij}^T \alpha + z_{ij}^T b_i) + \varepsilon_{ij}$$

with $\text{var}(\varepsilon_{ij}) = \phi v(\mu_{ij}^b)/a_i$, $j = 1, \dots, n_i$. Expanding h about the current estimates $(\hat{\alpha}, \hat{b})$ based on the current $\hat{\theta}$ gives

$$y_{ij} \approx \hat{\mu}_{ij}^b + h'(\hat{\eta}_{ij}^b)[x_{ij}^T(\alpha - \hat{\alpha}) + z_{ij}^T(b_i - \hat{b}_i)] + \varepsilon_{ij} \quad (2)$$

which implies that the “working” observation $Y_{ij}^* = \hat{\eta}_{ij}^b + g'(\hat{\mu}_{ij}^b)(y_{ij} - \hat{\mu}_{ij}^b)$ satisfies

$$Y_{ij}^* = x_{ij}^T \alpha + z_{ij}^T b_i + \varepsilon_{ij}^* \quad (3)$$

where, at least to an approximation for the ε_{ij}^* ,

$$b_i \sim \mathcal{N}(0, D(\theta)) \quad \text{and} \quad \varepsilon_{ij}^* \sim \mathcal{N}(0, \phi[g'(\hat{\mu}_{ij}^b)]^2 v(\hat{\mu}_{ij}^b)/a_i). \quad (4)$$

Updated estimates of (α, b, θ) are obtained by solving for them in the LMM defined by (3) and (4), *i.e.*, by using the PQL algorithm.

A further expansion of the conditional mean in terms involving b_i alone adds $\frac{1}{2} h''(\hat{\eta}_{ij}^b) z_{ij}^T (b_i - \hat{b}_i)(b_i - \hat{b}_i)^T z_{ij}$ to the right hand side of (2). Goldstein and Rasbash [14, 16, 15] suggested that one ignore the cross-products

involving different components of b_i , add the mean values of the resulting quadratic terms as offsets to the regression model and treat their residuals as additional random error terms with known variance. This modified procedure, implemented as PQL2 in MLwiN [37], is also intended to improve the estimates of variance components.

3 Adaptive Gauss-Hermite Quadrature

Consider the two-level GLMM with I independent clusters of observations $\{y_{ij}, j = 1, \dots, n_i\}, i = 1, \dots, I$ and a random intercept so that

$$\mu_{ij}^b = E(y_{ij}|b_i) = h(x_{ij}^T \alpha + b_i), \quad b_i \stackrel{i.i.d}{\sim} \mathcal{N}(0, \theta).$$

To simplify matters, suppose $g = h^{-1}$ is the canonical link function so that $v(\mu) = [g'(\mu)]^{-1}$ and furthermore that the scale factor and prior weights are all unity. This setup applies, for example, to two-level log-linear modeling of Poisson data and to logistic regression for clustered binary outcome data. The contribution to the marginal likelihood (integrated quasi-likelihood) for the i^{th} cluster is

$$\begin{aligned} L_i &= \frac{1}{\sqrt{2\pi\theta}} \int L_i^c(b) e^{-\frac{b^2}{2\theta}} db \\ &= E_{\mathcal{N}(0,\theta)} L_i^c(b) \end{aligned} \tag{5}$$

where $E_{\mathcal{N}(\mu,\theta)}$ denotes expectation with respect to the $\mathcal{N}(\mu, \theta)$ distribution and L_i^c is the conditional quasi-likelihood contribution

$$L_i^c(b) = \exp\left\{-\frac{1}{2} \sum_{j=1}^{n_i} d_{ij}(y_{ij}, \mu_{ij}^b)\right\}.$$

Ordinary Gauss-Hermite quadrature approximates the integral in (5) with the sum

$$L_i \simeq \frac{1}{\sqrt{\pi}} \sum_{r=1}^R \omega_r L_i^c(\sqrt{2\theta} t_r)$$

where the t_r are the R quadrature points, roots of the R -degree Hermite polynomial, and the ω_r denote the associated weights.[11, p. 924] The problem with this approach is that the same quadrature points are used for each cluster, irrespective of the cluster outcomes. Thus, for some i , the conditional quasi-likelihoods $L_i^c(b)$ may take large values for b well outside the range covered by the points $\{\sqrt{2\theta} t_r, r = 1, \dots, R\}$.

Let $\phi(b; \mu, \sigma^2)$ denote the density of the normal distribution with mean μ and variance σ^2 . The basic idea behind adaptive quadrature as introduced by Liu and Pierce [27] is the same one that underlies the Laplace integral

approximation, namely, to determine the normal density $\phi(b; \bar{b}_i, \bar{\sigma}_i^2)$ that best approximates the entire integrand $L_i^c(b)\phi(b; 0, \theta)$ in (5). The value that maximizes the integrand, \bar{b}_i , is obtained as the solution (in b) to $\sum_j (y_{ij} - \mu_{ij}^b) + b/\theta = 0$. The curvature in the log integrand at its maximum is the inverse of $\bar{\sigma}_i^2 = [\sum_j v(\mu_{ij}^{\bar{b}_i}) + \theta^{-1}]^{-1}$ [5, §2.1]. Once these are computed, the marginal likelihood contribution is approximated via

$$\begin{aligned} L_i &= \mathbb{E}_{\mathcal{N}(\bar{b}_i, \bar{\sigma}_i^2)} \left[\frac{L_i^c(b)\phi(b; 0, \theta)}{\phi(b; \bar{b}_i, \bar{\sigma}_i^2)} \right] \\ &\approx \frac{1}{\sqrt{\pi}} \sum_{r=1}^R \omega_r \frac{L_i^c(\bar{b}_i + \sqrt{2}\bar{\sigma}_i t_r)\phi(\bar{b}_i + \sqrt{2}\bar{\sigma}_i t_r; 0, \theta)}{\phi(\bar{b}_i + \sqrt{2}\bar{\sigma}_i t_r; \bar{b}_i, \bar{\sigma}_i^2)} \\ &= \sqrt{2}\bar{\sigma}_i \sum_{r=1}^R \omega_r e^{t_r^2} L_i^c(\bar{b}_i + \sqrt{2}\bar{\sigma}_i t_r)\phi(\bar{b}_i + \sqrt{2}\bar{\sigma}_i t_r; 0, \theta). \end{aligned} \quad (6)$$

The Laplace approximation is given by (6) for $R = 1$, $\omega_1 = 1$ and $t_1 = 0$.

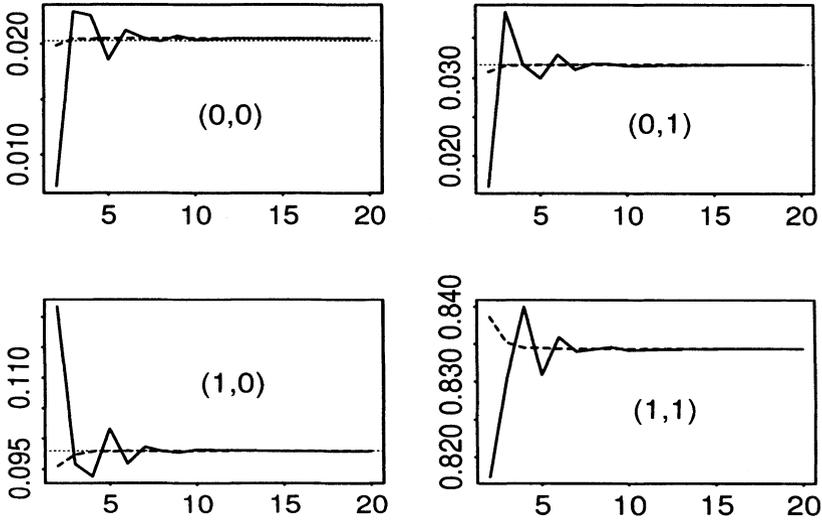


FIGURE 1. Marginal probabilities (ordinates) of each of four possible outcomes with matched pairs of binary outcome data estimated using standard (solid line) and adaptive (dashed line) Gauss-Hermite quadrature. The outcome vector is shown in the center of each panel. The abscissae show the number of quadrature points.

Use of standard and adaptive Gauss-Hermite quadrature to approximate the integrated likelihood is illustrated in Figure 1 for the special case of

matched pairs of binary outcome data with a single binary covariate that varies within clusters. Defining $\text{expit}(x) = [1 + e^{-x}]^{-1}$, this model has $\mu_{i1}^b = \text{expit}(\alpha_0 + b_i)$, $\mu_{i2}^b = \text{expit}(\alpha_0 + \alpha_1 + b_i)$ and $v(\mu) = \mu(1 - \mu)$. The parameter values used were $\alpha_0 = 4$, $\alpha_1 = -2$ and $\theta = 4$. The four panels of the figure show plots of the approximated marginal probabilities of the four possible outcomes for (y_{i1}, y_{i2}) , namely (0,0), (0,1), (1,0), and (1,1), as functions of the number R of quadrature points. Using the adaptive procedure, the integral approximations converged to the fourth decimal place with $R = 5$. By contrast, with the standard procedure, they continued to oscillate for R well beyond 5 .

Pinheiro and Bates [36] adapted this methodology for multidimensional integrals and their methods have been incorporated into PROC NL MIXED in SAS.[45] In the sequel we compare results using NL MIXED and GLIM-MIX, *i.e.*, using ML and PQL, with two well studied sets of data.

4 Meta Analysis of Clinical Trials Data

Our first example involves a series of 2×2 tables of counts of “successes” and “failures” among 293 patients distributed in treatment and control groups in eight clinical centers (Table 1). Introduced to statisticians by Beitler and

TABLE 1. Clinical trial of topical cream for infection
Center Treatment Response Total Success
Success Failure patients Rate (%)

Center	Treatment	Success	Failure	patients	Success Rate (%)
1	Drug	11	25	36	30.6
	Control	10	27	37	27.0
2	Drug	16	4	20	80.0
	Control	22	10	32	68.8
3	Drug	14	5	19	73.7
	Control	7	12	19	36.8
4	Drug	2	14	16	12.5
	Control	1	16	17	5.9
5	Drug	6	11	17	35.3
	Control	0	12	12	0.0
6	Drug	1	10	11	9.1
	Control	0	10	10	0.0
7	Drug	1	4	5	20.0
	Control	1	8	9	11.1
8	Drug	4	2	6	66.7
	Control	6	1	7	85.7
Total	Drug	55	75	130	42.3
	Control	47	96	143	32.9

Source: Beitler and Landis [26]

Landis [2], these data have been widely used to illustrate different methods for mixed effects modeling of categorical data. The developers of GLIMMIX, for example, noted that their macro converged more consistently if one first converted the table of counts to a series of binary outcome variables and covariates.[26, p. 440] The data also featured prominently in a recent review by Agresti and Hartzel [1] of methods for meta analysis of binary outcome data. The object of many of these analyses has been to estimate the clinic specific treatment effect, expressed as an odds ratio and assumed constant over clinics, while adjusting for clinic to clinic variation in baseline success rates via random effects modeling. There has also been interest in deciding whether there is evidence for treatment by center interaction.

Let y_{ij} denote the binary outcome (1 for success, 0 for failure) for the j^{th} subject in the i^{th} clinic. Suppose the covariate x_{ij} takes values $-\frac{1}{2}$ for control and $+\frac{1}{2}$ for treatment. This coding helps to orthogonalize the design matrix and render more plausible the implicit assumption of independence between random intercept and random slope (interaction) terms in what follows. Two models of interest are I: $\text{logit } E(y_{ij}|b_i) = \alpha_0 + \alpha_1 x_{ij} + b_i^0$; and II: $\text{logit } E(y_{ij}|b_i) = \alpha_0 + \alpha_1 x_{ij} + b_i^0 + b_i^1 x_{ij}$, the first corresponding to the hypothesis of constant odds ratio. The parameter of interest α_1 represents the *within clinic* log odds ratio comparing treatment and control groups. This is assumed constant across clinics in Model I but may vary by clinic in Model II. Tables 2 and 3 compare results obtained using four procedures for fitting GLMMs, including the PQL2 procedure mentioned at the end of §2. Also shown for Model I are results for the “exact” conditional maximum likelihood (CML) analysis, based on convolutions of the non-central hypergeometric distributions that arise when one conditions on all four marginal totals in each table. [10, §2.5] The analog for Model II is the GLMM that adds a random effect to the log odds ratio parameter in each non-central hypergeometric distribution. This may be fitted by PQL using methods previously described.[5, §6.4] Some notable features of

TABLE 2. Estimates \pm standard errors for Model I

Method	α_0	α_1	θ_0
NLMIXED (ML)	-0.828 \pm 0.533	0.739 \pm 0.300	1.96 \pm 1.19
GLIMMIX (PQL)	-0.784 \pm 0.537	0.724 \pm 0.296	2.03 \pm 1.26
MLwiN (PQL)	-0.784 \pm 0.537	0.724 \pm 0.296	2.03 \pm 1.19
MLwiN (PQL2)	-0.789 \pm 0.606	0.859 \pm 0.310	2.56 \pm 1.46
Hypergeometric (CML)		0.756 \pm 0.303	

this comparison include: (i) the lack of any suggestion for a treatment by clinic interaction; (ii) the excellent agreement between the estimates and standard errors obtained by ML (adaptive quadrature) and PQL, especially for the variance component of the random intercept; and (iii) the fact that the PQL2 results are substantially different from the others. Note that the

standard errors of the variance components estimated by the GLIMMIX and MLwiN implementations of PQL differ slightly. Otherwise the results were identical.

TABLE 3. Estimates \pm standard errors for Model II

Method	α_0	α_1	θ_0	θ_1
NLMIXED (ML)	-0.830 \pm 0.535	0.746 \pm 0.323	1.97 \pm 1.20	0.02 \pm 0.32
GLIMMIX (PQL)	-0.791 \pm 0.538	0.749 \pm 0.333	2.04 \pm 1.27	0.12 \pm 0.41
MLwiN (PQL)	-0.791 \pm 0.538	0.749 \pm 0.333	2.04 \pm 1.15	0.12 \pm 0.37
MLwiN (PQL2)	-0.870 \pm 0.614	0.830 \pm 0.367	2.61 \pm 1.46	0.20 \pm 0.45
Hypergeometric (PQL)		0.793 \pm 0.352		0.16 \pm 0.48

Table 4 reports results of a small simulation study designed to evaluate more systematically the performance of PQL in this setting.[4] For each of 10,000 simulations, 8 pairs of independent binomial observations $r_{ij} \sim \text{binom}(p_{ij}, n_{ij})$, $i = 1, \dots, 8$, $j = 1, 2$ were drawn with denominators n_{ij} chosen equal to those in the penultimate column of Table 1. The GLMM was specified by logit $p_{ij} = \alpha_0 + \alpha_1(2x_{ij} - 1) + b_i^0 + b_{ij}^1$ where $b_i^0 \sim \mathcal{N}(0, \theta_0)$ and $b_{ij}^1 \sim \mathcal{N}(0, \theta_1/2)$ were mutually independent sets of random effects. Thus the b_i^0 were random clinic effects, with roughly the same amount of clinic-to-clinic variation as for the data in Table 1, while the differences between b_{ij}^1 for $j = 1$ and $j = 2$ represented the variation in treatment effects (log odds ratios). Parameter settings were $\alpha_0 = 0$, $\theta_0 = 2$, $\alpha_1 = 0, 1, 2$ and $\theta_1 = \text{Var}(b_{i1}^1 - b_{i2}^1) = 0, 0.5, 1, 2$. $\bar{\alpha}_1$ and $\bar{\theta}_1$ refer to the averages of the estimates of these two parameters over the 10,000 replications. The error rates refer to the proportion of replicates for which the 95% confidence interval for α_1 excluded the true value on the left or the right side.

The simulated data were analyzed using PQL as described above for the log odds ratio GLMM based on the non-central hypergeometric distribution. As with any mixed model, there was a tendency to over-estimate slightly the small (or null) values of the variance component since negative estimates were not allowed. The systematic underestimation of variance components often observed with clustered binary data (see §6 below) was not a problem here, probably because of the relatively large denominators and mid-range values for many of the binomial observations. PQL estimates of the regression coefficient α_1 and of the larger values of the variance component were remarkably unbiased. Error rates for interval estimation were quite satisfactory. Not shown here are corresponding results for the empirical transform (ET) method, which consisted of applying ordinary LMM methods to derived outcome variables. The derived variable was the logarithm of the observed odds ratio in each table, with 0.5 added to both cells whenever any marginal total of success or failure was zero, so as to avoid infinities. Conditional on the random effects, this outcome variable was treated as normally distributed with variance equal to the in-

TABLE 4. Results of the simulation study of PQL

	True values		Estimates		Error rates	
	θ_1	α_1	$\hat{\theta}_1$	$\bar{\alpha}_1 - \alpha_1$	Left	Right
0.0		0	0.15	0.000	0.015	0.016
		1	0.16	0.015	0.012	0.017
		2	0.18	0.030	0.013	0.018
0.5		0	0.58	0.002	0.032	0.027
		1	0.58	0.013	0.029	0.033
		2	0.60	0.023	0.018	0.034
1.0		0	1.05	-0.003	0.030	0.032
		1	0.96	-0.012	0.027	0.035
		2	1.04	0.002	0.024	0.038
2.0		0	2.00	-0.016	0.026	0.031
		1	1.98	0.000	0.030	0.032
		2	1.99	-0.000	0.025	0.029

Source: Breslow, Leroux and Platt [4]

verse of the sum of reciprocals of the cell frequencies. The ET estimates of both the variance component and the regression coefficient were seriously biased towards zero, so that the random effect predictors were similarly misbehaved.[4, pp. 57-58] A similar tendency of ET to underestimate the variance component was observed for simulated Poisson observations representing spatially correlated rates when the mean rates were very small.[4, pp. 58-59] Thus the recent recommendation that ET methods be used in preference to PQL in such situations appears to be unfounded.[30, p. 283]

5 Longitudinal Series of Counts

Our second example involves a series of counts of seizures recorded by 59 patients with epilepsy for each of four two-week periods that preceded clinic visits. Introduced by Thall and Vail [47], these data also have been used by numerous statisticians to illustrate methods for analysis of longitudinal data with discrete outcomes. Figure 2 plots the patient trajectories of the log counts, augmented by 0.5 to avoid infinities. Each trajectory starts with the log of the baseline count over the eight-week period before the study, which was divided by four for comparability. Other fixed covariates of interest included a binary treatment indicator, the logarithm of age in years, and either a binary indicator for the fourth visit or the visit number j after division by ten.

With y_{ij} now denoting the seizure count reported at the j^{th} visit by the i^{th} patient, assumed to have a Poisson distribution after condition-

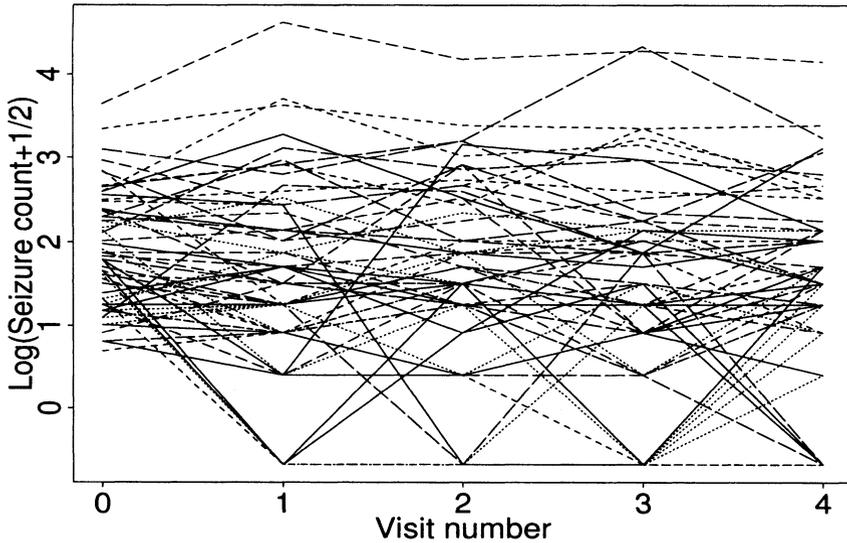


FIGURE 2. Epilepsy seizure counts at baseline and four follow-up periods.

TABLE 5. Estimates \pm standard errors for Model III

Parameter	NLMIXED (ML)	GLIMMIX (PQL)	MLwiN (PQL)	MLwiN (PQL2)
Regression coefficients				
Constant	-1.117 \pm 1.182	-1.256 \pm 1.220	-1.256 \pm 1.220	-1.335 \pm 1.239
Baseline*	0.884 \pm 0.131	0.872 \pm 0.136	0.872 \pm 0.136	0.881 \pm 0.138
Treatment	-0.933 \pm 0.401	-0.917 \pm 0.413	-0.917 \pm 0.413	-0.929 \pm 0.420
Bas* \times Trt	0.338 \pm 0.203	0.331 \pm 0.210	0.331 \pm 0.210	0.336 \pm 0.213
Age*	0.484 \pm 0.347	0.472 \pm 0.358	0.472 \pm 0.359	0.481 \pm 0.364
Visit 4	-0.161 \pm 0.055	-0.161 \pm 0.055	-0.161 \pm 0.055	-0.161 \pm 0.055
Variance component				
$\sqrt{\theta_0}$	0.503 \pm 0.059	0.524 \pm 0.062	0.524 \pm 0.059	0.529 \pm 0.060

* log transform

ing on the random effects, two models of interest were Model III: $\log E(y_{ij}|b_i) = x_i^T \alpha + b_i^0$ and Model IV: $\log E(y_{ij}|b_i) = x_i^T \alpha + b_i^0 + b_i^1 j/10$. Model IV was the more interesting in that it provided for a patient specific random slope and intercept, assumed to follow a bivariate normal distribution, to model the trends in the trajectories. Results of fitting these models using the NLMIXED and GLIMMIX procedures in SAS, and the PQL and PQL2 methods in MLwiN, are shown in Tables 5 and 6. There was remarkably good agreement in estimation of the regression coefficients and their standard errors. By contrast to the previous example, PQL2 produced regression coefficients slightly closer to those of ML than did PQL. The PQL2

estimates of the variance components, however, were slightly further from the ML estimates. The high (0.4 or so) within cluster (patient) correlation in the log epilepsy counts is reflected in the large, and highly statistically significant, estimates of variance components.

TABLE 6. Estimates \pm standard errors for Model IV

Parameter	NLMIXED (ML)	GLIMMIX (PQL)	MLwiN (PQL)	MLwiN (PQL2)
Regression coefficients				
Constant	-1.368 \pm 1.201	-1.267 \pm 1.215	-1.268 \pm 1.215	-1.361 \pm 1.241
Baseline*	0.885 \pm 0.131	0.870 \pm 0.135	0.870 \pm 0.135	0.882 \pm 0.138
Treatment	-0.929 \pm 0.402	-0.910 \pm 0.411	-0.910 \pm 0.411	-0.922 \pm 0.421
Bas* \times Trt	0.338 \pm 0.204	0.330 \pm 0.209	0.330 \pm 0.209	0.335 \pm 0.214
Age*	0.477 \pm 0.354	0.463 \pm 0.357	0.463 \pm 0.357	0.472 \pm 0.364
Visit/10	-0.266 \pm 0.165	-0.264 \pm 0.157	-0.264 \pm 0.157	-0.267 \pm 0.160
Variance components				
$\sqrt{\theta_{00}}$	0.502 \pm 0.059	0.521 \pm 0.062	0.521 \pm 0.061	0.527 \pm 0.063
θ_{01}	0.003 \pm 0.089	0.002 \pm 0.090	0.002 \pm 0.088	0.005 \pm 0.091
$\sqrt{\theta_{11}}$	0.729 \pm 0.157	0.737 \pm 0.157	0.737 \pm 0.162	0.756 \pm 0.165

* log transform

6 Further Simulations with Binary Outcome Data

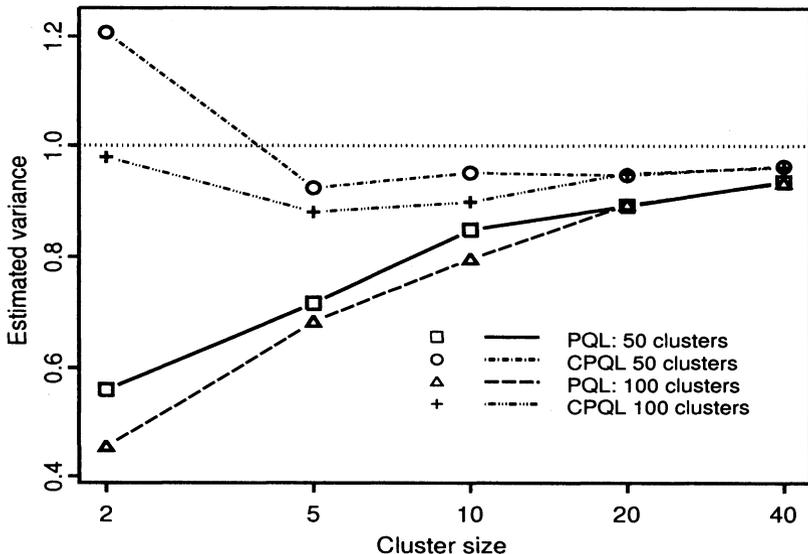
To further evaluate the bias of PQL estimates of variance components with binary outcome data, and assess the degree of correction afforded by CPQL, a new series of simulation experiments was run using a variant of a model originally proposed by Zeger and Karim [50] for clustered data. Each experiment involved K clusters of constant size n . Binary outcome variables y_{ij} for $i = 1, \dots, K$ and $j = 1, \dots, n$ were generated according to the hierarchical model

$$\text{logit}E(y_{ij}|b_i) = \alpha_0 + \alpha_1 t_{ij} + \alpha_2 x_i + \alpha_3 t_{ij} x_i + b_i,$$

where the t_{ij} were unit level covariates that were randomly generated from the uniform distribution on the interval $[-\frac{1}{2}, \frac{1}{2}]$, the x_i were subject level covariates of which the first half took the value 0 and the remainder the value 1, and the b_i were independent, normally distributed random effects with mean 0 and variance θ . The parameter values were $\alpha_0 = -0.5$, $\alpha_1 = 1$, $\alpha_2 = -1$, $\alpha_3 = 0.5$ and $\theta = 1$. The number K of clusters was 50 or 100 and the sample size n per cluster ranged between 2 and 40. Each experiment was replicated 200 times at each parameter setting.

The results in Figure 3 demonstrate the substantial bias in the PQL estimates. With matched pairs of binary outcome data, the true variance of the

FIGURE 3. Mean values of estimated variance component



subject specific effects was underestimated by about a half.[6] Even with as many as 40 observations per cluster, the variance was still underestimated by 6%. The degree of bias was affected more by cluster size than by the number of clusters. Indeed, it was worse for $K = 100$ than for $K = 50$. When $\alpha_0 = -2.5$, the bias in $\hat{\theta}_p$ with $n = 40$ was closer to 10%.

CPQL substantially reduced the bias, overcorrecting with 50 clusters of size 2. However, the slow rates at which the averages of the PQL and CPQL estimates approached the true value 1 suggests that cluster sizes might need to be quite large to eliminate entirely the bias in the variance component. As noted previously [6, 24], the bias in the regression coefficients is unimportant once the variance components have been estimated correctly.

7 Higher Order Laplace Expansions

The integral in the expression (1) for the likelihood has dimensionality equal to the number of random effects and hence, for many problems of interest, increases with the sample size n . Shun and McCullagh [42] and Shun [41] noted that the standard Laplace approximation failed to have an asymptotic ($n \uparrow \infty$) justification in such circumstances, and derived a remainder term that improved its performance. Raudenbush, Yang and Yosef [39] developed a systematic approach to higher order Laplace expansions, and provided details for two-level models involving a series of clusters