ratios and analogous higher order moments leads to greater efficiency than some other approach to estimating V_i. As the authors point out, the proposed likelihood estimator is a GEE estimator with a particular choice of variance matrix to weight the residuals. If the assumed variance matrix is close to the true variance matrix, the estimator will be nearly efficient; the degree of inefficiency is a simple function of the degree of misweighting as detailed in the paper.

In Figure 1, FLR present the degree of inefficiency that results from assuming the correlation is constant across individuals when it is not. First note that the conditional log odds ratio ω ranges from 0 to 10 in this illustration so that the odds ratios range between 1 and 22,026. Also, the third order term is fixed at $\kappa = 3$ so that the pairwise correlations are substantial at every value of ω. For example, the correlation between the first two observations ρ_{12} ranges as a function of the x values between 0.43 and 0.55 when $\omega = 0$ and between 0.60 and 0.98 when $\omega = 6$. Note in Figure 1 that assuming constant correlation with $\omega = 0$ (true correlations between 0.43 to 0.55) gives a nearly efficient estimate. This is because the correlations do not vary substantially and so the working variance matrix is nearly correct. Assuming the correlations are constant when they vary as a function of x between 0.6 and 0.98 ($\omega = 6$) leads to inefficient estimates. This should be no surprise. When correlations are this high and vary so dramatically with x, they must be modelled as a function of x to get reasonably efficient inferences as has been done in Liang, Zeger and Qaqish (1992) and Carey, Zeger and Diggle (1993).

To produce the high degree of correlation and dependence on x, we believe an unrealistic dependence structure has been assumed. FLR have set the third order term $\kappa = 3$. This means that when $\omega = 0$, $\mathrm{OR}(y_{i1}, y_{i2} \mid y_{i3} = 0) = 1$ and $\mathrm{OR}(y_{i1}, y_{i2} \mid y_{i3} = 1) = \exp(3) = 20$. Hence there is no association between the first two observations if the third value is 0 but an enormous positive association if the third value is 1. Is this realistic?

The challenge given to the FLR estimator which assumes constant conditional odds ratios is to assume constant correlations that range from 0 to 0.45. Note that the entire x-axis in Figure 2 corresponds to correlations that are smaller than the left most point of Figure 1 ($\omega = 0$). To illustrate the potential inefficiency of the FLR estimator, we must let the conditional odds ratios vary with the xs and use an estimator that assumes they are constant. An arbitrary degree of inefficiency can be produced in this way.

To recap our comments on efficiency, the FLR likelihood estimator is a special case of the GEE approach where the variance matrix has been specified in terms of conditional moments in such a way that the resulting equation is the score equation for a log-linear model. As a GEE estimator, it will be efficient when the assumed covariance matrix is close to the truth and inefficient when not. The same is true for any GEE estimator regardless of the approach to specifying the weighting matrix.

We once again congratulate FLR for their interesting and important paper. We look forward to the opportunity to use their methodology to analyze balanced data sets in problems where the regression parameters are the focus. Clinical trials is an area of application where this approach can be particularly important. We also concur with them that ignoring correlation when it is substantial is problematic even if robust variances are estimated. Their subsection 3.1 shows that grossly misspecifying the weighting matrix when using GEE can lead to inefficient estimates. We look forward to additional efficiency studies based upon more realistic data sets. Finally, while we have not addressed the missing data issue, we are aware of interesting recent work by one of the authors (Rotnitzky) and coworkers on handling missing at random data in the general GEE framework.

Rejoinder

Garrett M. Fitzmaurice, Nan M. Laird and Andrea G. Rotnitzky

We thank all of the discussants for their contributions. We will restrict most of our comments to four issues.

MARGINAL REGRESSION MODELS WITH STOCHASTIC TIME-VARYING COVARIATES

We are in complete agreement with the comments on the role of marginal models made by Drum and McCullagh. A related issue is the role of covariates in a longitudinal study. Our paper focused on nonstochastic covariates and the discussants' comments relate to settings where the covariates are time-stationary. However, when the covariates are both time-varying and stochastic, new issues arise regarding the interpretation and the estimation of the parameters of marginal models. These parameters may not have the implied