

Regression models for the analysis of longitudinal Gaussian data from multiple sources

Liam M. O'Brien^{1,*},† and Garrett M. Fitzmaurice²

¹*Department of Mathematics, Colby College, 5838 Mayflower Hill, Waterville, ME 04901, U.S.A.*

²*Division of General Medicine, Brigham and Women's Hospital, 1620 Tremont Street, 3rd Floor, BC3-2Q, Boston, MA 02120, U.S.A.*

SUMMARY

We present a regression model for the joint analysis of longitudinal multiple source Gaussian data. Longitudinal multiple source data arise when repeated measurements are taken from two or more sources, and each source provides a measure of the same underlying variable and on the same scale. This type of data generally produces a relatively large number of observations per subject; thus estimation of an unstructured covariance matrix often may not be possible. We consider two methods by which parsimonious models for the covariance can be obtained for longitudinal multiple source data. The methods are illustrated with an example of multiple informant data arising from a longitudinal interventional trial in psychiatry. Copyright © 2005 John Wiley & Sons, Ltd.

KEY WORDS: covariance modelling; mixed-effects models; multiple informants; psychiatry; repeated measures

1. INTRODUCTION

Multivariate longitudinal outcomes arise in a wide variety of disciplines. The defining feature of multivariate longitudinal data is that measurements are taken repeatedly through time with multiple responses obtained at each occasion. The multiple outcomes may or may not be measured on the same scale, and may or may not be measures of the same underlying variable. When the multiple outcomes measure the same underlying variable using similar scales, we refer to them as being *commensurate*. This distinction is important because whether the outcomes are commensurate or not has an impact on the type of model for the joint outcomes that is adopted.

*Correspondence to: Liam O'Brien, Department of Mathematics, Colby College, 5838 Mayflower Hill, Waterville, ME 04901, U.S.A.

†E-mail: lobrien@colby.edu

Contract/grant sponsor: U.S. National Institutes of Health; contract/grant numbers: MH17119, MH54693, GM29745
Contract/grant sponsor: Division of Natural Sciences at Colby College

Multivariate responses that are non-commensurate are common in the health sciences. For example, the multivariate response can consist of distinct continuous and categorical outcomes that are neither measured on the same scale, nor measure the same underlying variable (see, for example, Reference [1]). On the other hand, the outcomes can be distinct but yet measured on the same scale. An example of the latter occurs in studies of AIDS, where CD4 and CD8 cell counts may be obtained longitudinally resulting in repeated measures of a bivariate outcome [2]. In this case, the two outcomes are both measured in terms of counts but are distinct markers of immune function (i.e. they are not commensurate). Another example where the outcomes are measured on the same scale, but are not commensurate, is repeated measures of systolic and diastolic blood pressures [3]. When the multiple outcomes are not commensurate, it is unlikely that the magnitudes of covariate effects on each outcome are similar, and thus separate regression models for each outcome will ordinarily be required.

With commensurate outcomes, we use the term 'multiple source' to indicate that each response comes from a different source but provides a measure of the same underlying variable (i.e. each source represents a different component of the multivariate response). Longitudinal multiple source data arise in many areas of application, e.g. environmental studies where a pollutant is measured at varying depths in a lake through time [4], or in studies of bilateral organ systems. For example, Heitjan and Sharma [5] analysed longitudinal bivariate data on intra-ocular pressure in both eyes. Studies of children are another common area where longitudinal multiple source outcomes arise. Information on children's psychopathology may be gathered from multiple informants, or multiple psychiatric instruments, repeatedly through time [6]. When the multivariate responses are commensurate, certain covariates may have similar effects on two or more of the outcomes. Due to both the commensurate nature of the responses, and the positive correlation among the source outcomes, it is advantageous to analyse the outcomes jointly rather than separately. A joint analysis will often result in an efficiency gain, but more importantly it provides a formal basis for the comparison of covariate effects across sources.

One approach to the joint analysis of multivariate longitudinal data is to simply reduce the multivariate responses to a single outcome. When the responses are continuous, this could be achieved by taking the average of the multiple responses, resulting in a single response at each occasion. As Heitjan and Sharma [5] point out, however, some information will be lost, and this approach can be problematic when there are missing data. That is, *ad hoc* methods for calculating the mean when there are missing data effectively make strong assumptions about the missingness mechanism—namely that the missingness mechanism is missing completely at random (MCAR) [7].

In this paper, we focus on joint modelling of longitudinal multiple source Gaussian data via multivariate linear regression. The development of the methods in this paper was motivated by data arising from an interventional trial in psychiatry [8]. The aim of this study was to compare two interventional strategies to prevent the 'social transmission' of affective disorders from a parent to other family members. Forty-nine families were randomized to an intervention consisting of lectures given to both parents to educate them on how to prevent such transmission from occurring. Sixty-four families randomized to the other interventional program attended clinician-facilitated counselling sessions. In order to assess whether the clinician-facilitated program was more effective than the lecture program, reports of family functioning were obtained from three sources—father, mother, and child. Family functioning was determined via a self-report measure, and was obtained from each source at five equally

spaced times over a 3-year period. Thus, each source provides a measure of the same underlying variable (family functioning) at each measurement occasion, and the analytic goal is to relate changes in family functioning to the intervention assignment. Much of the previous work on modelling multivariate longitudinal data, however, has focused on outcomes that are not commensurate (e.g. References [1, 2]). While there is some related work on longitudinal models for multiple source data, it has tended to make quite strong assumptions about the structure of the covariance among the outcomes (e.g. References [3, 5]).

Here, we present flexible regression models for longitudinal multiple source data that allow for more general covariance structures. We propose a joint model for the mean response vector. Jointly estimating the regression parameters of this model has the advantage of being more efficient than separately estimating the parameters for each source when there are common, or shared, effects across sources for one or more covariates [9]. This may often be the case with multiple source data since the outcomes are commensurate. An additional advantage of joint estimation, as mentioned earlier, is that it allows formal comparison of covariate effects across sources.

With a joint model, we also need to be concerned with specifying parsimonious models for the covariance matrix. While the model for the mean can have a large number of parameters, the proliferation of covariance parameters is of greater concern as the number of responses increases. In many longitudinal studies, the covariance parameters are considered a nuisance. Recognizing that the covariance parameters are often a nuisance, one possible approach is to utilize a 'working independence' assumption that considers the observations to be independent for the purposes of estimation, but base inference on standard errors of the regression coefficients that are estimated via the empirical (or 'sandwich') covariance estimator [10]. This avoids explicit modelling of the covariance among outcomes. For many longitudinal designs, the loss of efficiency associated with the use of the working independence estimator of the regression parameters is modest. Use of the empirical covariance estimator is not desirable in all situations though. For example, Kauermann and Carroll [11] showed that the variability of the empirical covariance estimator is larger than that of a model-based covariance estimator in certain circumstances. In particular, this occurs when the sample size is modest, and when the covariate design matrices are not identical across subjects. They also showed that confidence intervals for the regression parameters generated using the empirical covariance estimator can result in undercoverage in such cases. Additionally, use of the empirical covariance estimator can be undesirable when there are many missing responses since it relies on replications across subjects to estimate the covariance non-parametrically [12]. In practice, many studies with longitudinal multiple source data have modest sample sizes, subject-varying design matrices, and missing data, making use of the empirical covariance estimator undesirable. The use of a suitable model-based covariance estimator can avoid such problems and is preferable.

In this paper, we specify a model for the mean response that allows for joint modelling and estimation of the regression parameters for all sources. We also present models for the within-subject covariance matrix that aim to achieve a balance between goodness-of-fit and parsimony. In Section 2, we present the regression model for the mean response and discuss two methods by which parsimonious covariance structures can be obtained. In Section 3, we illustrate these methods through the analysis of multiple informant data from a longitudinal clinical trial in psychiatry. The results of a simulation study to investigate the performance of the empirical variance estimator relative to a model-based estimator are presented in

Section 4. We conclude in Section 5 with a discussion of the benefits and drawbacks of the models presented.

2. REGRESSION MODEL FOR LONGITUDINAL MULTIPLE SOURCE DATA

2.1. Modelling the mean response

To discuss modelling of the mean response, some notation needs to be introduced. Let Y_{ijk} be the response on subject i from source j at time k ($i = 1, \dots, N; j = 1, \dots, I; k = 1, \dots, T$). We stack the responses for a subject into a vector, \mathbf{Y}_i , as $\mathbf{Y}_i = (Y_{i11}, Y_{i12}, \dots, Y_{i21}, \dots, Y_{iT})'$. We also have a matrix of subject-specific covariates, \mathbf{X}_i , which can include time- and source-varying covariates. The covariates are related to the responses through the general linear regression model,

$$\mathbf{Y}_i = \mathbf{X}_i \boldsymbol{\beta} + \boldsymbol{\epsilon}_i \quad (1)$$

where $\boldsymbol{\beta}$ is a vector of regression coefficients, and $\boldsymbol{\epsilon}_i = (\epsilon_{i11}, \epsilon_{i12}, \dots, \epsilon_{i21}, \dots, \epsilon_{iT})'$ is a vector of correlated random errors with mean zero.

While (1) gives the general formulation relating the mean response vector to the covariates, when considering longitudinal multiple source data the primary questions of scientific interest relate to whether the effects of specific covariates on the mean responses vary across sources. For example, consider the simplest possible case where observations are obtained from two sources at two time points (e.g. baseline and post-baseline) and it is intended to relate the mean of \mathbf{Y}_i to a single dichotomous covariate, x_i . A fully saturated model for the mean (with IT responses per subject) is given by

$$E(Y_{ijk}) = \beta_0 + \beta_1 x_i + \beta_2 j + \beta_3 \text{time}_k + \beta_4 x_i \times j + \beta_5 x_i \times \text{time}_k + \beta_6 j \times \text{time}_k + \beta_7 x_i \times j \times \text{time}_k \quad (2)$$

where $j = 0$ if the response is from source 1 and $j = 1$ if the response is from source 2; $\text{time}_k = 0$ for the baseline measurement and $\text{time}_k = 1$ for the post-baseline measurement. In the context of longitudinal multiple source data, the interpretation of the regression parameters is more readily apparent when the model is expressed for each source separately. For the first source ($j = 0$),

$$E(Y_{i0k}) = \beta_0 + \beta_1 x_i + \beta_3 \text{time}_k + \beta_5 x_i \times \text{time}_k$$

and for the second source ($j = 1$),

$$E(Y_{i1k}) = (\beta_0 + \beta_2) + (\beta_1 + \beta_4) x_i + (\beta_3 + \beta_6) \text{time}_k + (\beta_5 + \beta_7) x_i \times \text{time}_k$$

The coefficients β_2 , β_4 , β_6 , and β_7 allow the effects of covariates on the mean response to vary by source. For example, if $\beta_7 \neq 0$ then the covariate-by-time interaction is not the same for each source. Recall that in most longitudinal studies the covariate-by-time interaction is the effect of primary interest, since it expresses how longitudinal patterns of change differ according to levels of the covariate. This is true regardless of whether the data were obtained from an observational study or from a randomized experiment. Note that in an observational

study if there is no significant group-by-time interaction, then the test of main effect of group represents a comparison of groups in terms of their baseline response, reflecting existing differences between groups prior to the start of the study. Whereas in a randomized trial, if there is no significant group-by-time interaction, then there is no interest in the group effect (since if the groups have the same patterns of change over time, and by design do not differ at baseline, their mean response profiles must necessarily coincide). As such, the test of group effect is subsumed in the test of group-by-time interaction.

Other hypotheses may be of interest, however. Note that if $\beta_2 = \beta_4 = \beta_6 = \beta_7 = 0$, the regression models for each source do not differ, and we essentially model the average of the multiple source outcomes via a longitudinal model (although the average is more appropriately weighted by the inverse of the covariances among all source outcomes). Thus, the model for the mean response given by (1) incorporates the two extreme cases of (a) all sources requiring a unique set of regression coefficients, and (b) all sources sharing the same set of regression coefficients.

2.2. Modelling the covariance

As mentioned earlier, a major concern when modelling longitudinal multiple source data is the proliferation of covariance parameters. Given that we have IT responses per subject, we would need to estimate $\frac{1}{2}IT(IT + 1)$ covariance parameters if the within-subject covariance matrix is allowed to be unstructured. Clearly, there are many practical cases where the sample size is not large enough to provide stable estimates for so many model parameters. For example, with measurements taken from three sources at five occasions (15 observations per subject), there are 120 parameters in an unstructured covariance matrix, not counting the additional parameters in the model for the mean.

We must reduce the number of covariance parameters in some way in order to obtain stable estimates of the model parameters. We will consider two methods by which this can be accomplished. Since it is easier to place constraints upon the variances and correlations, rather than the variances and covariances, we focus on these parameters. The pairwise correlations can be separated into three subgroups. Let the two sources under consideration be denoted by j and j' , and the two times under consideration be denoted by k and k' . We denote correlations among measurements from the same source at two different times (intra-source) by $\alpha_{jkk'} = \text{Corr}(Y_{ijk}, Y_{ijk'})$; the correlations corresponding to measurements taken at the same time but from different sources (inter-source) by $\phi_{jj'k} = \text{Corr}(Y_{ijk}, Y_{ij'k})$; and the cross-correlations describing the relationship between different source responses at different times by $\tau_{jj'kk'} = \text{Corr}(Y_{ijk}, Y_{ij'k'})$. These three pairwise relationships are illustrated in Figure 1. In addition to these three types of correlations, we also have a variance parameter for each source-time combination given by $\sigma_{jk}^2 = \text{Var}(Y_{ijk})$.

Note that $\alpha_{jkk'}$ can depend on the particular source under consideration (j) and the two time points at which the reports have been taken (k and k'); $\phi_{jj'k}$ can depend on the two sources under consideration (j and j') and the time point at which the reports were taken (k); $\tau_{jj'kk'}$ can depend on the two sources (j and j') and two time points (k and k'). The variances may depend on the source (j) and time point (k) being considered.

Next, we focus on two ways in which the number of covariance parameters that need to be estimated can be reduced. The first method involves specifying two separate covariance structures: one for the association among the sources that does not depend on the




Figure 1. Graphical representation of correlations among sources over time.

measurement occasions (Σ_1), and the other for the measurement occasions that does not depend on the sources (Σ_2). Then, the overall covariance structure is obtained by taking the Kronecker product of these two separate covariance structures, $\Sigma = \Sigma_1 \otimes \Sigma_2$. This procedure was discussed first by Martin [13] who provided a theoretical rationale in the context of bi-dimensional spatial problems, and later discussed more generally by Galecki [14]. The second method that will be examined utilizes random effects, which also induce a parsimonious structure on the covariances while remaining quite flexible for modelling the relationships among longitudinal multiple source data.

2.2.1. Kronecker product covariances. Martin [13] introduced the use of the Kronecker product to construct covariance matrices for bi-dimensional spatial data. Such covariance structures have also been used for data that have two repeated factors, such as days nested within menstrual cycles [15]. Kronecker product covariance matrices are also useful for multivariate data that have more than one level of complexity (e.g. multivariate outcomes, with each outcome measured repeatedly) since each level can be modelled separately. For example, this method has been used in environmental statistics to model the covariance among spatio-temporal data, whereby two covariance matrices were specified: one for the spatial factor that does not depend upon the measurement occasions; and one for the measurement occasions that does not depend on the spatial factor [4]. Because each of these two covariance matrices does not depend upon the levels of the other factor, we refer to these two separate matrices as ‘marginal’ covariance matrices.

In the context of longitudinal multiple source data, a natural analogue is to allow the ‘marginal’ covariance for sources to be modelled separately from the ‘marginal’ covariance for the measurement occasions. Taking the Kronecker product of these two marginal covariances to obtain the overall covariance matrix greatly reduces the number of parameters that need to be estimated. This results in a reduction in the number of variance parameters that need to be estimated. However, the most significant reduction comes about by eliminating the need for estimating the cross-correlations since they are simply products of the marginal correlations. That is, the overall correlation matrix is the Kronecker product of the correlation matrices corresponding to Σ_1 and Σ_2 (i.e. $\tau_{jj'kk'} = \phi_{jj'}\alpha_{kk'}$), which is shown in the appendix. Note

that because $\phi_{jj'k}$ is time invariant in the Kronecker product setting, we denote the inter-source correlations by $\phi_{jj'}$. Similarly, since $\alpha_{jkk'}$ is source invariant, we denote the intra-source correlations by $\alpha_{kk'}$. These implicit assumptions are also shown in the appendix.

For example, if we specify an unstructured marginal covariance for both the sources and measurement occasions, which is the least parsimonious of the Kronecker product structures, we only need to estimate $\frac{1}{2}[I(I+1)+T(T+1)]-1$ covariance parameters. This results in a net reduction of the estimation of the $2\binom{I}{2}\binom{T}{2}$ cross-correlations if the overall covariance matrix is left unstructured. It also reduces the number of estimated variances from IT to $I+T-1$. Note that the marginal covariance structures are scaled such that one of the marginal variance terms is set to 1; and an overall scale factor is multiplied through the resulting covariance matrix. This ensures that the overall covariance matrix is uniquely estimable. In the case where three sources provide reports at six times, using the least parsimonious Kronecker product structure reduces the number of covariance parameters from 171 to 26. With more parsimonious models there will be an even greater decrease in the number of covariance parameters that need to be estimated.

The marginal covariance patterns that could be utilized may include many different structures. The marginal covariance for the measurement occasions could be unstructured or could assume any of the patterned covariance structures adopted from the time series literature, e.g. Toeplitz, exponential, autoregressive, etc. However, when considering the marginal covariance for the sources, there is less inherent structure. Ordinarily, only the unstructured and compound symmetry matrices would be appropriate, although it may be possible to model relationships among the sources. For example, if we had information from a father, mother, and child in a multiple informant setting, we may wish to consider the parents as exchangeable but not the child.

In general, we would expect that the cross-correlations comparing responses between two different sources at two different times should be weaker than either the intra-source or inter-source correlations; this is implicitly assumed in the Kronecker product structures since the cross-correlations are products of the marginal correlations. Additionally, the variance of the observations provided by source j at occasion k is simply the product of the marginal variance of the responses from source j and the marginal variance of the responses at occasion k (as shown in the appendix). Thus, if there is a pattern in the marginal variances (e.g. increasing variances over time) this pattern will be reflected in the diagonal elements of the overall covariance matrix for \mathbf{Y}_i .

Martin [13] introduced the Kronecker product covariance structures for cases when a two-dimensional lattice process is 'separable', thereby providing a theoretical basis for this structure. Here, we present some additional motivation for the assumptions about the correlations embodied in these structures. Consider a response from source j at occasion k , along with another response from the same source at occasion k' (say, $k' > k$), as well as a report from source j' at occasion k . The correlation between $Y_{ij'k}$ and $Y_{ijk'}$, conditioning on Y_{ijk} (i.e. the partial correlation between Y_{ijk} and $Y_{ij'k'}$, given $Y_{ij'k}$), can be expressed in terms of the three pairwise correlations, $\phi_{jj'k}$, $\alpha_{jkk'}$, and $\tau_{jj'kk'}$ given in Figure 1. Specifically, the partial correlation is given by

$$\frac{\tau_{jj'kk'} - \phi_{jjk'}\alpha_{jkk'}}{\sqrt{1 - \alpha_{jkk'}^2}\sqrt{1 - \phi_{jj'k}^2}} \quad (3)$$

Under the Kronecker product correlation, $\phi_{jj'k} = \phi_{jj'}$, $\alpha_{jkk'} = \alpha_{kk'}$, and $\tau_{jj'kk'} = \phi_{jj'}\alpha_{kk'}$; thus (3) is equal to 0 for the Kronecker product structures. Therefore it makes a strong conditional independence assumption that, if we already had information from source j at time k , then additional information from any source j' at time k would not help to predict the response from source j at time k' . The appropriateness of the Kronecker product covariances for any particular application will depend to a large extent on whether this conditional independence assumption is tenable. It also depends on the appropriateness of the requirement of the time invariance of the inter-source correlations, and the source invariance of intra-source correlations.

2.2.2. Random effects covariance. While the Kronecker product structures are a useful way at arriving at parsimonious models for the covariance, we next examine how the introduction of random effects induces structure on the three types of pairwise relationships—the inter-source correlations ($\phi_{jj'k}$), the intra-source correlations ($\alpha_{jkk'}$), and the cross-correlations ($\tau_{jj'kk'}$). We also consider how the variances are modified.

Let us consider a simple mixed-effects model for longitudinal multiple source data with only a random subject effect. The model implies that all subjects track along the same trajectory through time, but that some may have consistently higher or lower responses than the population average. In addition, we could include random source effects. The random source effect could be drawn from a single distribution, or each source effect could be drawn from a separate distribution. If drawn from separate distributions, the random effects may or may not be correlated. Additionally, the random effects could consist of both random intercepts and random trajectories of the underlying response over time. For simplicity, we first illustrate the case in which the random source effects are drawn from a single distribution. Thus, we have source random effects nested within a subject random effect,

$$Y_{ijk} = \mu_{ijk}(\boldsymbol{\beta}) + \gamma_i + \eta_{ij} + \varepsilon_{ijk} \quad (4)$$

Here i indexes the subjects, j indexes the sources, and k indexes the repeated measures. The fixed effects for the response from source j at time k on subject i are given by $\mu_{ijk}(\boldsymbol{\beta}) = \mathbf{X}_{ijk}\boldsymbol{\beta}$, the random subject effect is given by the γ_i term, and the random source effect is given by the η_{ij} term. In this illustration, the random effects and error terms are assumed to be independent of each other and distributed as follows:

$$\gamma_i \sim N(0, \sigma_\gamma^2), \quad \eta_{ij} \sim N(0, \sigma_\eta^2), \quad \varepsilon_{ijk} \sim N(0, \sigma_\varepsilon^2)$$

This mixed-effects model results in the following expressions for the marginal correlations and variances:

$$\alpha = \frac{\sigma_\gamma^2 + \sigma_\eta^2}{\sigma^2}, \quad \phi = \tau = \frac{\sigma_\gamma^2}{\sigma^2}$$

where $\sigma^2 = \sigma_\gamma^2 + \sigma_\eta^2 + \sigma_\varepsilon^2$.

Note that none of the expressions for the correlations depends on the sources or time points under consideration (i.e. $\phi_{jj'k} = \phi$, $\alpha_{jkk'} = \alpha$, and $\tau_{jj'kk'} = \tau$). This structure also constrains the inter-source (ϕ) and cross-correlations (τ) to be equal. While this structure is remarkably parsimonious (only three covariance parameters need to be estimated), it seems quite unlikely that it would provide a reasonable fit to longitudinal multiple source data. In order for it

to hold, the variances must be constant for all source/time combinations; as well as the requirement that $\tau = \phi$ is constant, regardless of the time separation.

Assumptions about the random effects can be altered such that the resulting correlations and variances are not so constrained. For example, we can allow either the subjects, or sources, or both, to have random trajectories over time in addition to random intercepts. This will result in the correlations and variances being parametric functions of time. Consider the following model:

$$Y_{ijk} = \mu_{ijk}(\boldsymbol{\beta}) + \gamma_{i0} + \gamma_{i1}t_k + \eta_{ij0} + \eta_{ij1}t_k + \varepsilon_{ijk} \tag{5}$$

with,

$$\begin{pmatrix} \gamma_{i0} \\ \gamma_{i1} \end{pmatrix} \sim N \left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{\gamma 0}^2 & \sigma_{\gamma 01} \\ \sigma_{\gamma 01} & \sigma_{\gamma 1}^2 \end{pmatrix} \right]$$

$$\begin{pmatrix} \eta_{ij0} \\ \eta_{ij1} \end{pmatrix} \sim N \left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{\eta 0}^2 & \sigma_{\eta 01} \\ \sigma_{\eta 01} & \sigma_{\eta 1}^2 \end{pmatrix} \right]$$

$$\varepsilon_{ijk} \sim N(0, \sigma_{\varepsilon}^2)$$

The resulting correlations and variances are now functions of time, but note that they do not vary by source. They are given by,

$$\alpha(t_k, t_{k'}) = \frac{\sigma_{\gamma 0}^2 + (t_k + t_{k'})\sigma_{\gamma 01} + t_k t_{k'} \sigma_{\gamma 1}^2 + \sigma_{\eta 0}^2 + (t_k + t_{k'})\sigma_{\eta 01} + t_k t_{k'} \sigma_{\eta 1}^2}{\sigma(t_k)\sigma(t_{k'})}$$

$$\phi(t_k) = \frac{\sigma_{\gamma 0}^2 + 2t_k \sigma_{\gamma 01} + t_k^2 \sigma_{\gamma 1}^2}{\sigma^2(t_k)}, \quad \tau(t_k, t_{k'}) = \frac{\sigma_{\gamma 0}^2 + (t_k + t_{k'})\sigma_{\gamma 01} + t_k t_{k'} \sigma_{\gamma 1}^2}{\sigma(t_k)\sigma(t_{k'})}$$

where $\sigma^2(t_k) = \sigma_{\gamma 0}^2 + 2t_k \sigma_{\gamma 01} + t_k^2 \sigma_{\gamma 1}^2 + \sigma_{\eta 0}^2 + 2t_k \sigma_{\eta 01} + t_k^2 \sigma_{\eta 1}^2 + \sigma_{\varepsilon}^2$. While including random intercepts and trajectories over time for both the subjects and sources may seem unnecessary, it does allow us to determine the relative magnitude of the variability due to subject and source components.

To introduce a source dependence among the variances and correlations, we can allow the source random effects to be drawn from separate distributions. This introduces a source dependence in the variance and intra-source correlation terms. If, in addition, we allow the source random effects to be correlated, the inter-source and cross-correlations become source dependent. Thus, the introduction of random effects allows us to fit a flexible class of models to account for the covariance among longitudinal multiple source outcomes. Using the commonly adopted notation for mixed-effects models for longitudinal data [16], we can consider models for the $IT \times 1$ vectors of outcomes on the i th subject given by

$$\mathbf{Y}_i = \mathbf{X}_i \boldsymbol{\beta} + \mathbf{Z}_i \mathbf{b}_i + \boldsymbol{\varepsilon}_i \tag{6}$$

where $\mathbf{b}_i = (\mathbf{b}_{i1}, \mathbf{b}_{i2}, \dots, \mathbf{b}_{il})' \sim \text{MVN}(\mathbf{0}, \mathbf{G})$, and $\varepsilon_{ijk} \sim N(0, \sigma_j^2)$.

The matrix, \mathbf{G} , can be unstructured, allowing the variances and intra-source correlations to differ by source. This would also allow the inter-source and cross-correlations to be source dependent. Note that we could also allow the variance of the random measurement error to differ depending on source.

Next, we discuss estimation of the model parameters. Since the focus of this paper is on Gaussian responses, we utilize maximum likelihood (ML) estimation and briefly discuss some computational issues.

2.3. Estimation

The general model is given by

$$\mathbf{Y}_i = \mathbf{X}_i\boldsymbol{\beta} + \boldsymbol{\epsilon}_i$$

where \mathbf{Y}_i is the vector of responses for subject i . Due to possibly missing observations, \mathbf{Y}_i is of dimension $t_i \times 1$ with $t_i \leq IT$. Also, \mathbf{X}_i is a $t_i \times p$ known design matrix, $\boldsymbol{\beta}$ is a $p \times 1$ vector of unknown regression parameters, and the $\boldsymbol{\epsilon}_i$ are independently distributed as $N(\mathbf{0}, \Sigma_i)$, with Σ_i a function of a $q \times 1$ vector of unknown covariance parameters, $\boldsymbol{\theta}$. The parameters in $\boldsymbol{\theta}$ comprise either the marginal covariance parameters from the Kronecker product structures or the covariance parameters from a random effects model. The log-likelihood is given by

$$l = \text{Const} - \frac{1}{2} \sum_{i=1}^n \left[\log |\Sigma_i| - \frac{1}{2} (\mathbf{Y}_i - \mathbf{X}_i\boldsymbol{\beta})' \Sigma_i^{-1} (\mathbf{Y}_i - \mathbf{X}_i\boldsymbol{\beta}) \right] \quad (7)$$

For the random effects model given by (6), $\boldsymbol{\beta}$ and $\boldsymbol{\theta}$ can be obtained using standard statistical software packages (e.g. Proc Mixed in SAS; SAS Institute, Cary, NC). However, many of the covariance patterns introduced via the Kronecker product are not readily available in the standard statistical software packages. For the latter, we implemented the modified Fisher scoring algorithm given by Jennrich and Schlutter [17] in S-plus 2000 (MathSoft Engineering and Education Inc., Cambridge, MA).

To illustrate the model for longitudinal multiple source data, and the two approaches for obtaining the covariances, we present an example arising from a longitudinal clinical trial in psychiatry.

3. EXAMPLE

To illustrate the application of the linear regression model presented in the previous section, we utilize the data described in Section 1 arising from a longitudinal interventional study conducted at The Judge Baker's Children's Center. The data arise from 113 families recruited through a health maintenance organization in Boston, MA. Each family has at least one parent identified as having an affective disorder within the 12 months prior to contact. Each family also had at least one child between the ages of 8 and 15 years. Families that had children with a prior affective disorder or other psychiatric diagnosis were excluded, as were families in which a parent was diagnosed with any prior psychosis. Also, families that had experienced severe strain for other reasons, such as the death of another family member, were excluded.

Each family was randomly assigned to one of two preventive treatment strategies designed to improve family functioning. One strategy consisted of two 1-h lectures given to the parents. The lectures were uniform and were given by a trained clinician. The other treatment consisted of clinician-facilitated psycho-educational sessions involving the entire family. These sessions were tailored to the individual families' educational levels and consisted of 6–10

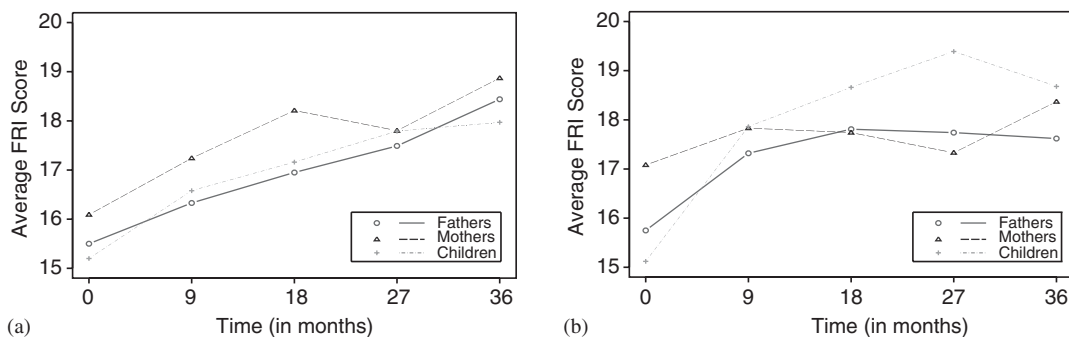


Figure 2. Mean FRI scores over time by informant, (a) clinician-facilitated group, (b) lecture group.

sessions. Sixty-four families were assigned to the clinician-facilitated group and 49 to the lecture group. Both treatment groups also received literature designed to educate them about affective disorders, and how parents afflicted with them can prevent their children from feeling the effects themselves. The families were assessed at baseline and at four approximately equally spaced times over the next 3 years. The aim of the study was to determine whether the clinician-facilitated treatment strategy was more efficacious in improving the functioning of the families over time.

In families with multiple children, one child was randomly selected for inclusion in the analysis. Thus, each family could have as many as 15 reports (three sources reporting at five times). In all, approximately 30 per cent of the responses were missing, but there was no reason to believe that the missingness was non-ignorable. Throughout we assume that the missing responses are missing at random (MAR) [7]. The psychometric measure that was used to obtain family functioning information was the family relationships index (FRI). The FRI is a self-report measure that uses three subscales—cohesion, expressiveness, and conflict—to obtain a final composite score. The conflict subscale is weighted negatively in the final score determination (see References [18, 19], for more information about the FRI). The FRI composite score is a continuous measure that was observed to be approximately normally distributed. Here, the unit of analysis is the family and the multiple 'sources' are the father, mother, and child.

Note that if an unstructured covariance matrix was assumed, we would need to estimate 120 covariance parameters in addition to the regression coefficients. Given data on only 113 independent families, estimation may be unstable. Thus, reducing the number of covariance parameters that need to be estimated is necessary. We considered covariance structures resulting from both the Kronecker product covariances and the random effects models.

Figure 2(a) shows the mean observed FRI scores for each source through time for the clinician-facilitated group, and Figure 2(b) shows the mean observed FRI scores for the lecture group. The final model for the mean was obtained via a stepwise procedure utilizing likelihood ratio tests. The initial model that was fit to the mean was the saturated model with main effect terms for source, categorical time, treatment, and all possible two- and three-way interactions. A quadratic time trend was found to provide an adequate fit to these data. A test of the null hypothesis of no treatment effect (i.e. a joint test of treatment-by-time and treatment-by-time-by-source interactions) was not significant (likelihood ratio $\chi^2_6 = 6.1$; $p = 0.41$). An additional




Figure 3. Predicted mean FRI scores over time by informant.

test of collapsibility was conducted (i.e. a joint test of treatment main effect, treatment-by-time, treatment-by-source, and treatment-by-time-by-source interactions) and was also not significant (likelihood ratio $\chi^2_9 = 7.79$; $p = 0.56$). Thus the final model consists of a source main effect, a quadratic time trend, and both source-by-time interaction terms. That is,

$$E(Y_{ijk}) = \beta_{0j} + \beta_{1j}t_{ik} + \beta_{2j}t_{ik}^2$$

Thus, no significant treatment effect was found to exist in these data, suggesting that there is no significant difference between the two interventions. The ML estimated means for each source are plotted in Figure 3. It appears that family functioning improved more rapidly at the beginning of the study, and then levelled off. The fathers' and mothers' reports tracked each other, while the children's reports displayed a somewhat different trend. Although there were no differences in the effects of the interventions, family functioning improved over time.

The covariance structure first fit to these data was of the Kronecker product class. Note that fathers and mothers show a general increase in FRI scores through time. However, the mean response for children is more erratic. Analyses suggested that the three sources (i.e. father, mother, and child) were not exchangeable and that a Toeplitz correlation structure held over time. Due to these considerations, the Kronecker product structure that seemed most appropriate for these data considered the source marginal covariance as unstructured and the covariance matrix for the measurement occasions as Toeplitz with heterogeneous variances. This results in 5 parameters for the source covariance matrix, 8 parameters for the measurement occasions covariance matrix, and an overall scale factor for a total of 14 parameters.

A random-effects model was also assumed for these data. Recall that the exchangeability of informants assumption does not seem tenable; we also wanted to allow each source to have a random trajectory within family. Thus random source effects were drawn from a multivariate normal distribution, allowing each source to have a random intercept and trajectory, and also

Table I. Marginal inter-informant correlations at the first occasion.

Informant pairs	Kronecker product	Random effects
Father–mother	0.372	0.503
Father–child	0.206	0.251
Mother–child	0.168	0.228

Table II. Marginal intra-informant correlations for the mother.

Occasions	Kronecker product	Random effects
(1,2)	0.636	0.681
(1,3)	0.553	0.636
(1,4)	0.453	0.572
(1,5)	0.382	0.495

Table III. Marginal cross-correlations for the mother at the first occasion and child at subsequent occasions.

Occasions	Kronecker product	Random effects
(1,2)	0.107	0.330
(1,3)	0.093	0.413
(1,4)	0.076	0.463
(1,5)	0.064	0.484

allowing correlation among the intercepts and trajectories for the different sources. The model can be written as

$$Y_{ijk} = \mu_{ijk}(\boldsymbol{\beta}) + \eta_{ij0} + \eta_{ij1}t_k + \varepsilon_{ijk}$$

where $(\eta_{i10}, \eta_{i11}, \eta_{i20}, \eta_{i21}, \eta_{i30}, \eta_{i31})' \sim \text{MVN}(\mathbf{0}, \mathbf{G})$, and $\varepsilon_{ijk} \sim N(0, \sigma_{\varepsilon}^2)$. Thus the multivariate normal distribution from which the random source effects were drawn was six dimensional. Note that \mathbf{G} is assumed to be unstructured allowing the random source effects and random trajectories to be correlated. Since the psychometric reports obtained from children are often more variable than those obtained from adults [20], we also allowed the random measurement error to depend on source; for these data there was no *a priori* reason for doing so.

Selected marginal variances and correlations resulting from the Kronecker product and random effects models are listed in Tables I–III. Table I shows how the two covariance structures relate with respect to the inter-source correlations at the first measurement occasion. The father–mother correlation is larger than the parent–child correlation; of note, the Kronecker product structure estimates are consistently smaller than those from the random effects model. The intra-source correlations for mothers are given in Table II. These correlations decrease as the time separation increases as might be expected. However, they do not decrease at a constant rate. Both structured covariance matrices result in similar patterns, with the Kronecker product structure providing somewhat smaller estimates. Table III gives representative values

for the cross-correlations comparing mothers and children with the time separations indicated. While the inter-source and intra-source correlations were smaller for the Kronecker product structures than for the random effects models, the cross-correlations were discernably smaller.

4. SIMULATION: EMPIRICAL VARIANCE ESTIMATOR

As mentioned in Section 1, the empirical variance estimator can have increased variability in finite samples. In this section, we demonstrate the increased variability of the empirical variance estimator, compared to the optimal model-based estimator, through a simulation study. Data were generated assuming a longitudinal multiple source setting under a mixed-effects model similar to that utilized in the previous section. The fixed effects consisted of a source main effect along with a linear time trend and a treatment effect. A treatment-by-time interaction was also included in the model as it is the parameter of main interest in a longitudinal study. Additionally, a random intercept and slope were drawn from a bivariate normal distribution for each of three sources. If we index subjects by i , sources by j , and measurement occasions by k , the true underlying mixed-effects model is

$$Y_{ijk} = \beta_0 + \beta_1 1_{j=2} + \beta_2 1_{j=3} + \beta_3 k + \beta_4 g_i + \beta_5 g_i k + \eta_{0ij} + \eta_{1ij} k + \varepsilon_{ijk}$$

where g_i indicates the group assignment for subject i with $i = 1, 2, \dots, N$; $j = 1, 2, 3$; $k = 0, 1, \dots, T$. The indicator function 1_c equals unity if condition c is true and zero otherwise. The coefficients for the source-specific random intercept and trajectory are given by $(\eta_{0ij}, \eta_{1ij})'$, and the random measurement error is given by ε_{ijk} . These random coefficients are distributed as, $(\eta_{0i1}, \eta_{1i1}, \eta_{0i2}, \eta_{1i2}, \eta_{0i3}, \eta_{1i3}) \sim \text{MVN}(\mathbf{0}, \mathbf{G})$ and $\varepsilon_{ijk} \sim N(0, \sigma^2)$, with known \mathbf{G} and σ^2 . While several different covariance structures, for different choices of \mathbf{G} and σ^2 , were considered, we present results from a model that induces an overall association structure with moderate correlations, similar to that found in the family functioning data analysed in Section 3. The inter-source correlations range from 0.15 to 0.40, the inter-source correlations range from 0.25 to 0.45, and the cross-correlations range from 0.15 to 0.35.

For each replication of the simulation, data were generated under this model and the fixed-effects (β) were estimated along with standard errors of the estimates obtained via both a model-based variance estimator and the empirical variance estimator under a correctly specified covariance model. This procedure was repeated 1000 times with samples of size $N = 100$ and 500. The medians of the model-based and empirical standard errors for the estimates of β_5 are given in Table IV for sample sizes of 100 and 500, and for $T = 3, 5$, and 8 measurement occasions (giving 9, 15, and 24 repeated measures, respectively). The true values of the standard errors of β_5 are also given for each case. As can be seen, the medians are discernably different—particularly for the smaller sample size—with the model-based medians being closer to the true values of the standard errors. The discrepancy between the medians of the model-based and empirical standard error estimates tends to be larger as the number of repeated measures increases and decreases with an increase in sample size. In general, when the repeated measures-to-independent subjects ratio increases, there are discernable differences between the medians. In all cases, the centre of the distribution of the empirical variance estimates is larger than the corresponding distribution of the model-based estimates. This was not observed for the other correlation models considered, however.

Table IV. Medians of standard errors from 1000 simulations.

	$J = 3, T = 3$	$J = 3, T = 5$	$J = 3, T = 8$
$N = 100$			
True value	0.41	0.18	0.089
Model-based median	0.42	0.21	0.17
Empirical median	0.50	0.38	0.71
$N = 500$			
True value	0.18	0.082	0.040
Model-based median	0.19	0.098	0.067
Empirical median	0.21	0.153	0.068

Table V. Quantile range (85–15 per cent) of standard errors from 1000 simulations.

	$J = 3, T = 3$	$J = 3, T = 5$	$J = 3, T = 8$
$N = 100$			
Model-based range	0.064	0.033	0.023
Empirical range	0.262	0.311	1.000
Ratio (MB/empirical)	0.24	0.10	0.23
$N = 500$			
Model-based range	0.013	0.0064	0.0045
Empirical range	0.045	0.110	0.370
Ratio (MB/empirical)	0.29	0.06	0.01

The distributions of the model-based standard errors were symmetric, however the distributions of the empirical standard errors were positively skewed. The skewness was more pronounced as the number of repeated measures per subject was increased. The top of Figure 4 illustrates the distribution of standard errors for model-based and empirical standard errors with three sources and three measurement occasions and a sample size of 100. The bottom of Figure 4 illustrates simulation results from the same mixed-effects model, but with a sample size of 500. It should be noted that data generated for the simulations used to construct Figure 4 consisted of almost independent repeated observations on each subject. It is apparent that the empirical standard errors have considerably more variability than the model-based, even when the number of subjects is relatively large compared to the number of repeated measures.

Table V gives the range of values for the standard errors between the 15th percentile and the 85th percentile. Since some distributions of the empirical standard errors were bimodal, using the interquartile range did not accurately reflect the variability of the standard error estimates in some cases. The empirical variance estimator shows poor performance, even as the number of measurements per subject increases. Increasing the sample size from 100 to 500 did not result in a discernable change in this performance.

These results indicate that the model-based estimator outperforms the empirical variance estimator with a correctly specified covariance model. This was a consistent finding across several correlation structures considered, but not reported here. Interestingly, we did not see the undercoverage reported by Kauermann and Carroll [11]. The minimum and maximum coverage probabilities observed for sample sizes of 500 when using a model-based estimator

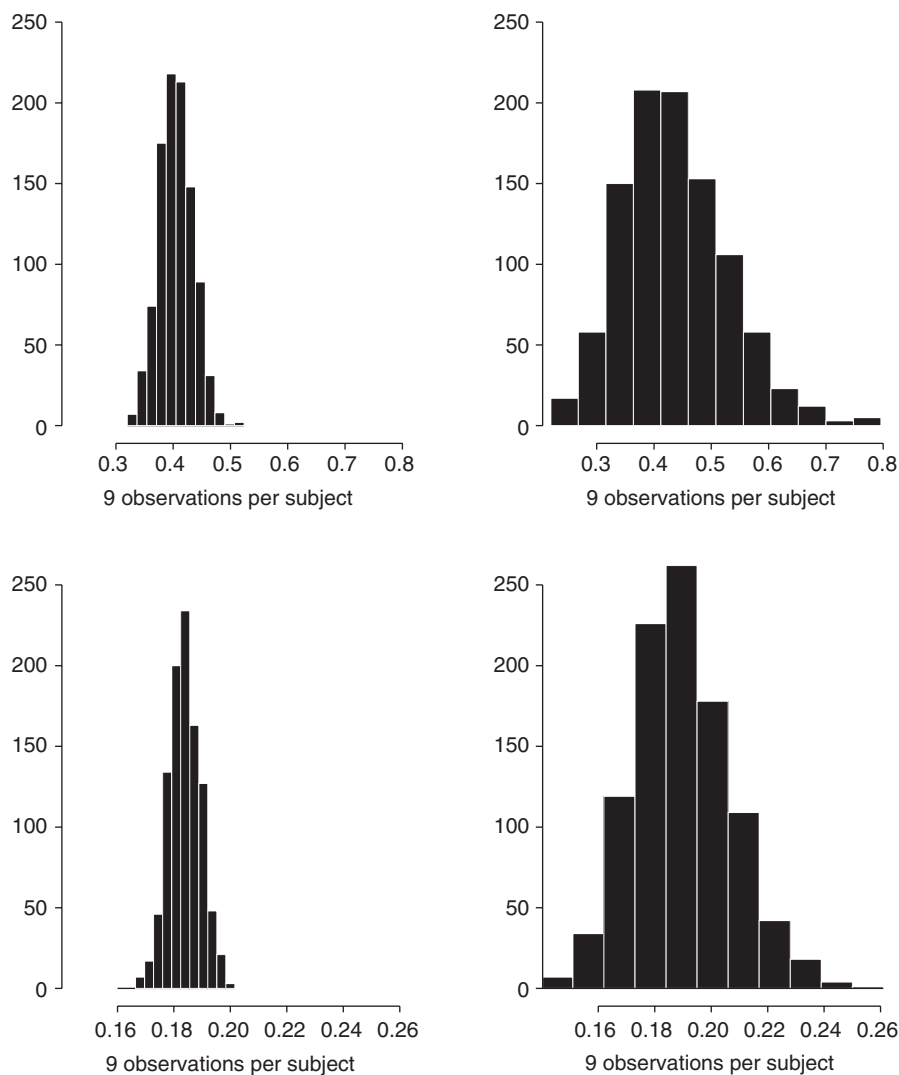


Figure 4. Distributions of model-based and empirical standard errors.

were 0.937 and 0.967, respectively. When using the empirical variance estimator the minimum and maximum coverage probabilities were 0.951 and 0.976, respectively. These did not differ appreciably for samples of size 100. While the coverage probabilities are not substantially different when using the empirical variance estimator, its increased variability is undesirable in most practical longitudinal multiple source settings.

5. DISCUSSION

We have presented a multivariate linear regression model for longitudinal Gaussian outcomes obtained from multiple sources. Since the outcomes are commensurate, we model the longi-

tudinal multiple source outcomes jointly. Note that the multivariate linear regression model implies regression models for the outcomes for each source, with the same set of predictors, but potentially different regression parameters. When these models share some common regression parameters, this often results in an efficiency gain. Note that if there are no shared regression parameters (i.e. separate regression parameters relating the same set of predictors to the outcomes from each source) there is no gain in efficiency when the data are complete (i.e. no missing data), but there is a potential gain if the data are incomplete [9]. Even when there is no potential gain in efficiency, joint modelling of the outcomes allows hypotheses comparing covariate effects and time-by-covariate interaction effects across sources to be tested. For example, in psychiatric studies such comparisons can be useful when considering whether treatment effects on changes in the outcomes differ according to informant report. Furthermore, with incomplete data that are MAR, joint estimation can eliminate potential bias that might arise in separate analyses of each source. Such modelling cannot be accomplished, however, without addressing the challenge of how to model the covariance among the outcomes when the number of outcomes is relatively large in comparison to the number of replications. Provided one models the covariance correctly, using a model-based estimator is preferable to utilizing the empirical variance estimator in the longitudinal multiple source setting, as was shown through simulation results in Section 4.

There has been some previous work concerning the analysis of longitudinal multivariate data that has addressed the problem of the proliferation of covariance parameters. A model for multivariate longitudinal data was proposed by Carey and Rosner [3] that addressed this issue, achieving a parsimonious covariance structure through a set of somewhat restrictive assumptions about the relationships among outcomes. They assumed that the inter-source correlations and the variances are time invariant (however, they could differ by source). Additionally, they assumed that the intra-source and cross-correlations follow a damped autocorrelation structure, with a form that is similar to $\text{corr}(Y(t), Y(t+s)) = \gamma^{|s|^\theta}$; where γ is a time-invariant damping factor. Rochon [1] also considered models for multivariate longitudinal outcomes, with possibly incomplete data that are MCAR, using generalized estimating equations (GEEs); this approach could be extended to handle data that are MAR using the approaches of Robins *et al.* [21] and Paik [22]. The covariance structures Rochon [1] considered were of the vector autoregressive moving average (VARMA) form often used for multivariate time series data. Such structures can result in complex covariance patterns, but require that the intra-source and cross-correlations decay through time at a certain rate. Similar to the covariance structures considered by Carey and Rosner [3], these may not be useful to fit a general class of longitudinal multiple source models where the decay in the correlation levels off after a certain time lag.

Mixed-effects models have also been employed for the analysis of multivariate longitudinal data. Shah *et al.* [2] discussed models for non-commensurate outcomes in which the measurement errors for the same source at any two occasions are uncorrelated, but measurement errors between two different sources at the same occasion are correlated. In addition to correlated random measurement error, a random subject effect was included. The resulting structure assumed that outcomes at any given time were exchangeable ($\phi_{jj'k} = \text{constant}$), and that the inter-source and cross-correlations were equal. Through the introduction of a random trajectory over time for each subject, these correlations would become parametric functions of time, but would still retain the basic constraints described. Reinsel [23] specified growth curve

models for multivariate repeated measures data. The covariance structure specified by these growth models satisfied a more general multivariate sphericity condition described by Boik [24]. The structure described in Reference [23] would allow the variances and intra-source correlations to depend on the source under consideration, but are time invariant. The inter-source and cross-correlations would also depend on the two sources in question, but would also be time invariant. This assumption may not be plausible with longitudinal multiple source data. Finally, Heitjan and Sharma [5] extended the linear random effects/autoregressive linear model of Chi and Reinsel [25] to analyse longitudinal outcomes from two sources. Their model included a random subject effect and the resulting marginal covariance structure assumed the sources are exchangeable. The intra-source correlations have an autocorrelated structure, and the cross-source ($\tau_{jj'kk'}$) correlations have a damped autocorrelation structure. In all of the previous work, the proposed covariance structures are plausible in a variety of situations, but make strong assumptions about the relationships among outcomes.

The linear multivariate regression model that we have presented addresses the issue of how to achieve a parsimonious covariance structure via two methods. The Kronecker product structures are useful in that they only require that the multiple source and longitudinal covariance structures be specified separately. They have the desirable property that the cross-correlations are no longer estimated separately since they are products of the marginal intra-source and inter-source correlations. The number of variance parameters that need to be estimated is also reduced. However, their appropriateness in any application will depend on whether the underlying conditional independence assumption is tenable. Additionally, the assumptions of the source invariance of the covariance among measurement occasions, and the time invariance of the covariance among sources must also be tenable. This has been found to hold in a variety of spatial and environmental applications (see References [13, 26, 27] for examples).

We have also discussed the use of mixed-effects models that induce relatively parsimonious structures on the covariance among outcomes. These models are remarkably flexible and can accommodate any degree of complexity. To be most useful, these models will often involve either random source effects nested within random subject effects, or correlated random source effects.

Except in cases where the ratio of subjects to repeated measures is very large, it is generally advantageous to fit a structured covariance matrix. Given that there will often not be enough independent clusters for the use of the empirical variance estimator to be attractive, for reasons mentioned in Section 1, the parametric models described in this paper provide parsimonious covariance structures in the longitudinal multiple source framework. We utilized ML estimation which results in asymptotically efficient estimators and can handle data that are either MCAR or MAR. With missing data that are MAR, consistent estimates of the regression and covariance parameters are obtained via ML provided the models for both the mean and covariance are correctly specified. This underscores the importance of careful modelling of the covariance structure when the missing data are MAR. However, if the data are complete, ML estimation results in unbiased estimates for the regression parameters even if the covariance model is misspecified. Methods for the analysis of longitudinal multiple source data in the binary data setting, where specifying the joint distribution is more complex, has been done by O'Brien and Fitzmaurice [28]. In that setting, ML estimation may not be feasible requiring alternative estimation methods to be employed.

APPENDIX: RELATIONSHIPS AMONG COVARIANCES AND CORRELATIONS
FOR KRONECKER PRODUCT STRUCTURES

Using the notation of Sections 2.1 and 2.2, we define the covariance matrix of \mathbf{Y}_i by, $\Sigma = \Sigma_1 \otimes \Sigma_2$. Let $\Sigma_1 = \mathbf{A}_1^{1/2} \mathbf{R}_1 \mathbf{A}_1^{1/2}$ where \mathbf{R}_1 is the correlation matrix corresponding to Σ_1 , and \mathbf{A}_1 is the diagonal matrix of variances corresponding to Σ_1 . Similarly, let $\Sigma_2 = \mathbf{A}_2^{1/2} \mathbf{R}_2 \mathbf{A}_2^{1/2}$ where \mathbf{R}_2 is the correlation matrix corresponding to Σ_2 and \mathbf{A}_2 is a diagonal matrix of variances corresponding to Σ_2 . Let $\Sigma = \mathbf{A}^{1/2} \mathbf{R} \mathbf{A}^{1/2}$.

Using properties of the Kronecker product [29], the covariance of \mathbf{Y}_i can then be written as

$$\begin{aligned}\Sigma &= (\mathbf{A}_1^{1/2} \mathbf{R}_1 \mathbf{A}_1^{1/2}) \otimes (\mathbf{A}_2^{1/2} \mathbf{R}_2 \mathbf{A}_2^{1/2}) \\ &= (\mathbf{A}_1^{1/2} \mathbf{R}_1 \otimes \mathbf{A}_2^{1/2} \mathbf{R}_2) (\mathbf{A}_1^{1/2} \otimes \mathbf{A}_2^{1/2}) \\ &= (\mathbf{A}_1^{1/2} \otimes \mathbf{A}_2^{1/2}) (\mathbf{R}_1 \otimes \mathbf{R}_2) (\mathbf{A}_1^{1/2} \otimes \mathbf{A}_2^{1/2})\end{aligned}$$

Thus,

$$\mathbf{A} = \mathbf{A}_1 \otimes \mathbf{A}_2 \tag{A1}$$

and

$$\mathbf{R} = \mathbf{R}_1 \otimes \mathbf{R}_2 \tag{A2}$$

It follows that the correlation between Y_{ijk} and $Y_{ij'k'}$ is obtained by taking the product of the corresponding inter-source and intra-source correlations from \mathbf{R}_1 and \mathbf{R}_2 , respectively. Similarly, the variance for the measurement obtained from source j at measurement occasion k is obtained by taking the product of the corresponding variances from \mathbf{A}_1 and \mathbf{A}_2 , respectively.

ACKNOWLEDGEMENTS

This research was supported by Grants MH17119, MH54693, and GM29745 from the U.S. National Institutes of Health. It was also supported, in part, by a grant from the Division of Natural Sciences at Colby College. We would like to thank William Beardslee and Pat Salt of The Judge Baker Children's Center for permission to use the data set on family functioning. We would also like to thank Nan Laird and two referees for their helpful comments on this manuscript.

REFERENCES

1. Rochon J. Analyzing bivariate repeated measures for discrete and continuous outcome variables. *Biometrics* 1996; **52**:740–750.
2. Shah A, Laird NM, Schoenfeld D. A random-effects model for multiple characteristics with possibly missing data. *Journal of the American Statistical Association* 1997; **92**(438):775–779.
3. Carey VJ, Rosner BA. Analysis of longitudinally observed irregularly timed multivariate outcomes: regression with focus on cross-component correlation. *Statistics in Medicine* 2001; **20**:21–31.
4. Dutilleul P, Pinel-Alloul B. A doubly multivariate model for statistical analysis of spatio-temporal environmental data. *Environmetrics* 1996; **7**:551–565.
5. Heitjan DF, Sharma D. Modelling repeated-series longitudinal data. *Statistics in Medicine* 1997; **16**:347–355.
6. Daskalakis C, Laird NM, Murphy JM. Regression analysis of multiple-source longitudinal outcomes: a 'stirling county' depression study. *American Journal of Epidemiology* 2002; **155**(1):88–94.

7. Little RJA, Rubin DB. *Statistical Analysis with Missing Data*. Wiley: New York, 1987.
8. Beardslee WR, Wright EJ, Salt P, Drezner K, Gladstone TRG, Versage EM, Rothberg PC. Examination of children's responses to two preventive intervention strategies over time. *Journal of the American Academy of Child and Adolescent Psychiatry* 1997; **36**(2):196–204.
9. Kmenta J. *Elements of Econometrics*. Macmillan: New York, 1971.
10. Huber PJ. The behaviour of maximum likelihood estimators under non-standard conditions. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, vol. 1, LeCam LM, Neyman J (eds). University of California Press: Berkeley, CA, 1967; 221–233.
11. Kauermann G, Carroll RJ. A note on the efficiency of sandwich covariance matrix estimation. *Journal of the American Statistical Association* 2001; **96**(456):1387–1396.
12. Diggle PJ, Liang K-Y, Zeger SL. *Analysis of Longitudinal Data*. Clarendon Press: Oxford, 1994.
13. Martin RJ. A subclass of lattice processes applied to a problem in planar sampling. *Biometrika* 1979; **66**(2): 209–217.
14. Galecki AT. General class of covariance structures for two or more repeated factors in longitudinal data analysis. *Communications in Statistics—Theory and Methods* 1994; **23**(11):3105–3119.
15. Park T, Lee YJ. Covariance models for nested repeated measures data: analysis of ovarian steroid secretion data. *Statistics in Medicine* 2002; **21**:143–164.
16. Laird NM, Ware JH. Random-effects models for longitudinal data. *Biometrics* 1982; **38**:963–974.
17. Jennrich RI, Schlucter MD. Unbalanced repeated-measures models with structured covariance matrices. *Biometrics* 1986; **42**:805–820.
18. Holahan CJ, Moos RH. The quality of social support: measures of family and work relationships. *British Journal of Clinical Psychology* 1983; **22**:157–162.
19. Hoge RD, Andrews DA, Faulkner P, Robinson D. The family relationship index: validity data. *Journal of Clinical Psychology* 1989; **45**(6):897–903.
20. Achenbach TM, McConaughy SH, Howell CT. Child/adolescent behavioral and emotional problems: implications of cross-informant correlations for situational specificity. *Psychological Bulletin* 1987; **101**(2):213–232.
21. Robins JM, Rotnitzky A, Zhao LP. Analysis of semiparametric regression models for repeated outcomes in the presence of missing data. *Journal of the American Statistical Association* 1995; **90**(249):106–121.
22. Paik MC. The generalized estimating equation approach when the data are not missing completely at random. *Journal of the American Statistical Association* 1997; **92**(440):1320–1329.
23. Reinsel G. Multivariate repeated-measurement or growth curve models with multivariate random-effects covariance structure. *Journal of the American Statistical Association* 1982; **77**(377):190–195.
24. Boik RJ. The mixed model for multivariate repeated measures: validity conditions and an approximate test. *Psychometrika* 1988; **53**(4):469–486.
25. Chi EM, Reinsel GC. Models for longitudinal data with ar(1) errors. *Journal of the American Statistical Association* 1989; **84**(406):452–459.
26. Verbyla AP, Cullis BR. The analysis of multistratum and spatially correlated repeated measures data. *Biometrics* 1992; **48**:1015–1032.
27. Martin RJ. The use of time-series models and methods in the analysis of agricultural field trials. *Communications in Statistics—Theory and Methods* 1990; **19**(1):55–81.
28. O'Brien LM, Fitzmaurice GM. Analysis of longitudinal multiple-source binary data using generalized estimating equations. *Applied Statistics* 2004; **53**(1):177–193.
29. Searle SR. *Matrix Algebra Useful for Statistics*. Wiley: New York, 1982.