Published in final edited form as:

Stat Med. 2017 March 15; 36(6): 958–970. doi:10.1002/sim.7199.

On the analysis of very small samples of Gaussian repeated measurements: An alternative approach

Philip M. Westgate¹ and Woodrow W. Burchett²

¹Department of Biostatistics, College of Public Health, University of Kentucky, Lexington, KY 40536, U.S.A.

²Department of Statistics, College of Arts and Sciences, University of Kentucky, Lexington, KY 40536, U.S.A.

SUMMARY

The analysis of very small samples of Gaussian repeated measurements can be challenging. First, due to a very small number of independent subjects contributing outcomes over time, statistical power can be quite small. Second, nuisance covariance parameters must be appropriately accounted for in the analysis in order to maintain the nominal test size. However, available statistical strategies that ensure valid statistical inference may lack power, whereas more powerful methods may have the potential for inflated test sizes. Therefore, we explore an alternative approach to the analysis of very small samples of Gaussian repeated measurements, with the goal of maintaining valid inference while also improving statistical power relative to other valid methods. This approach uses generalized estimating equations with a bias-corrected empirical covariance matrix that accounts for all small-sample aspects of nuisance correlation parameter estimation in order to maintain valid inference. Furthermore, the approach utilizes correlation selection strategies with the goal of choosing the working structure that will result in the greatest power. In our study, we show that when accurate modeling of the nuisance correlation structure impacts the efficiency of regression parameter estimation, this method can improve power relative to existing methods that yield valid inference.

Keywords

correlation selection; generalized estimating equations; multivariate Gaussian linear model; power; test size

1 Introduction

Studies that produce very small samples of Gaussian repeated measurements can arise in a variety of research settings. Examples can be found in, for instance, Skene and Kenward [1]. Although multiple approaches have been proposed for the analysis of such data, it can be challenging to maintain the nominal type I error rate without a loss in power in small-sample settings.

If two or more trial arms are to be compared over time, and data are complete and balanced, very popular analysis methods are repeated measures ANOVA and repeated measures MANOVA [2]. An advantage of these methods is that they do not require small sample

approximations [3] in order to maintain valid inference, although corrections to repeated measures ANOVA may be required to account for potential violations in the assumed covariance structure [4, 5, 6, 7, 8, 9]. However, when there are missing data or a saturated model is not fit, the analysis of very small samples of Gaussian repeated measurements can be challenging. Although power can be quite small due to very few subjects contributing repeated measurements, power can be further hampered by the need to choose an analysis method that ensures the nominal Type I error rate is maintained, which is often the need that is emphasized in fundamental research [10].

A very popular method is the use of a multivariate Gaussian linear model with likelihood-based inference. With this approach, both the choice and estimation of the working covariance structure can be very influential both with respect to test size and power, as explained in detail by Skene and Kenward [1, 3] and summarized in the next section. In short, a major reason the analysis of such data can prove difficult is due to the lack of data to help model a working covariance structure and to estimate corresponding nuisance parameters [1, 3]. As a result, available statistical strategies using a multivariate Gaussian linear model often have the potential to result in inflated test sizes, and the use of an unstructured working covariance matrix with the hope of ensuring valid inference may result in a lack of power and is not actually guaranteed to maintain the nominal test size.

Skene and Kenward [1, 3] addressed these small-sample pitfalls of likelihood-based inference with a multivariate Gaussian linear model in two papers. First, they proposed the use of generalized estimating equations (GEE) [11] with a working independence structure, or ordinary least squares, for regression parameter estimation, which allows the analyst to avoid having to choose and estimate a working correlation structure [1]. They showed that this approach, when coupled with the use of the empirical sandwich estimator implementing the small-sample bias correction proposed by Mancl and DeRouen [12] and the degrees of freedom approximation approach of Pan and Wall [13], results in valid inference, albeit with the potential for a large loss in power. Therefore, they next proposed a modification to Box's correction [4, 5] that only requires the estimation of the true covariance structure of the data for use in the approximation of the approximate sampling distribution of the independence-based test statistic [3]. They showed that their modified Box correction results in valid inference and improved power.

Both approaches proposed by Skene and Kenward [1, 3] assume a working independence structure in order to avoid covariance parameter estimation. However, accurate modeling of the covariance structure can result in more efficient regression parameter estimation, depending on the study design, thus decreasing standard errors (SEs) and increasing power [14, 15, 16]. Therefore, as an alternative approach to the analysis of very small samples of Gaussian repeated measurements, we explore the use of methodology developed for improved inference with GEE in small-sample settings. In short, we study empirically driven methods for computing unbiased SE estimates and approximating degrees of freedom, which have been shown to maintain the desired Type I error rate for any working correlation structure [17, 18, 19]. Furthermore, we use empirically driven correlation selection strategies, along with this robust SE estimation approach, that is designed to select the working structure that results in the smallest SEs and therefore the greatest power [20, 18,

19]. We note that this approach is driven by statistical power, rather than by the need to either model the correct correlation structure or use a working independence structure in order to ensure the Type I error rate is maintained.

In Section 2 we review in more detail the multivariate Gaussian linear model along with corresponding likelihood-based approaches to inference, in addition to the methods proposed by Skene and Kenward [1, 3]. We then give details on our alternative GEE-based approach. In Section 3 we conduct a simulation study comparing the performances, in terms of test size and power, of these methods, which are then compared via an application example in Section 4. Concluding remarks are given in Section 5.

2 Methods

2.1 Multivariate Gaussian Linear Model

Let $\mathbf{Y}_i = [Y_{i1}, \dots, Y_{in_i}]$, $i = 1, \dots, N$, denote the repeated measurements from the i th subject. We assume $\mathbf{Y}_i \sim MVN(\mathbf{X}_i\boldsymbol{\beta}, \boldsymbol{\Sigma}_i)$, $i = 1, \dots, N$, where $\mathbf{X}_i$ is the i th subject's matrix of covariate values, and $\boldsymbol{\beta} = [\beta_0, \beta_1, \dots, \beta_{p-1}]^T$. Parameters from the multivariate Gaussian linear model are often estimated using restricted maximum likelihood (REML) [21, 22]. Denoting $\hat{\boldsymbol{\Sigma}}_i$ as

the REML estimate for $\boldsymbol{\Sigma}_i$, we have $\hat{\boldsymbol{\beta}} = (\sum_{i=1}^N \mathbf{X}_i^T \hat{\boldsymbol{\Sigma}}_i^{-1} \mathbf{X}_i)^{-1} \sum_{i=1}^N \mathbf{X}_i^T \hat{\boldsymbol{\Sigma}}_i^{-1} \mathbf{Y}_i$. Assuming $\boldsymbol{\Sigma}_i$, $i = 1, \dots, N$, are correctly specified, the Kenward and Roger [23, 24] method is routinely utilized in small-sample settings as it approximates degrees of freedom and corrects for two sources of small-sample bias in the typical model-based covariance matrix,

$$\widehat{\text{Cov}}_{MB}(\hat{\boldsymbol{\beta}}) = \left(\sum_{i=1}^N \mathbf{X}_i^T \hat{\boldsymbol{\Sigma}}_i^{-1} \mathbf{X}_i \right)^{-1}. \text{ In short,}$$

$$E\left[\left(\sum_{i=1}^N \mathbf{X}_i^T \hat{\boldsymbol{\Sigma}}_i^{-1} \mathbf{X}_i \right)^{-1} \right] \neq \left(\sum_{i=1}^N \mathbf{X}_i^T \boldsymbol{\Sigma}_i^{-1} \mathbf{X}_i \right)^{-1}, \text{ and due to the estimation of nuisance}$$

covariance parameters, $\text{Cov}(\hat{\boldsymbol{\beta}}) = \left(\sum_{i=1}^N \mathbf{X}_i^T \boldsymbol{\Sigma}_i^{-1} \mathbf{X}_i \right)^{-1} + \boldsymbol{\Lambda}$ [25], where $\boldsymbol{\Lambda}$ is defined in

[23]. Denoting $\widehat{\text{Cov}}_{KR}(\hat{\boldsymbol{\beta}})$ as the Kenward and Roger [23, 24] bias-corrected model-based matrix, inference can be carried out using the scaled F-statistic,

$$F = \lambda l^{-1} (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})^T \mathbf{L}^T (\mathbf{L} \widehat{\text{Cov}}_{KR}(\hat{\boldsymbol{\beta}}) \mathbf{L}^T)^{-1} \mathbf{L} (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}), \text{ such that under } H_0: \mathbf{L}\boldsymbol{\beta} = \mathbf{0} \text{ it is assumed } F \sim F_{l, \nu}, \text{ where } \lambda \text{ and } \nu \text{ are obtained via moment matching [23, 24, 1].}$$

When the true covariance structure is not complicated, i.e. not many nuisance covariance parameters to estimate, and is correctly specified by the data analyst, then to our knowledge the Kenward and Roger [23, 24] method works quite well at maintaining the nominal test size, although some exceptions exist. See, for example, our simulation study and [26, 1, 27]. However, due to the lack of information in small samples, it is very difficult to consistently and correctly specify the true covariance structure [1, 3, 28, 29], assuming the true structure is even considered. Therefore, methods that are commonly used to select a structure, such as the Akaike [30] and Bayesian [31] information criteria (AIC, BIC), can result in inflated test sizes due to the use of model-based SEs that assume correct covariance specification [1, 3, 26, 32, 28].

Another popular approach, when the timing of observations is balanced for all subjects, is to use an unstructured working covariance matrix (which could also be considered for selection as described above) due to its flexibility and thus increased potential for modeling the true structure. However, even when this structure does embed the true structure, if the sample size is small the Kenward and Roger [23, 24] method may result in inflated test sizes due to the multiple covariance parameters that must be estimated in addition to the need to approximate the degrees of freedom. For instance, see [1, 3, 27, 33] and our simulation study below. Furthermore, even when the Kenward and Roger [23, 24] method maintains the nominal test size when using an unstructured working covariance matrix, the need to estimate all of these nuisance parameters can result in a loss in power relative to fitting a simpler structure, assuming the simpler structure is the true structure [1, 3]. The reason for this is that estimating these additional nuisance parameters will result in a loss in degrees of freedom and, depending on the setting, possibly increase Λ .

In short, when utilizing a model-based covariance matrix for inference, to our knowledge the Kenward and Roger [23, 24] method is preferred. However, this method is not guaranteed to maintain the nominal test size. Specifically, if the true covariance structure is not modeled, or if multiple nuisance covariance parameters must be estimated in order to accurately model the true structure, then the actual test size can be larger than its nominal value.

2.2 Methods of Skene and Kenward

Skene and Kenward [1, 3] proposed moving away from likelihood-based approaches due to the need to correctly specify Σ_{β} , $i = 1, \dots, N$, along with having to deal with the issues that arise from estimating the corresponding nuisance parameters. Specifically, estimating these nuisance parameters influence $\hat{\beta}$ and therefore $\text{Cov}(\hat{\beta})$, and they are needed in order to obtain $\widehat{\text{Cov}}_{KR}(\hat{\beta})$. In their first paper, Skene and Kenward [1] proposed a generalized estimating equations (GEE) [11] type approach that uses a working independence correlation structure, thus ignoring the need to estimate nuisance parameters when obtaining estimates for β .

In this first approach, $\hat{\beta}_{GEE} = \left(\sum_{i=1}^N \mathbf{X}_i^T \mathbf{V}_i^{-1} \mathbf{X}_i \right)^{-1} \sum_{i=1}^N \mathbf{X}_i^T \mathbf{V}_i^{-1} \mathbf{Y}_i$, where

$\mathbf{V}_i = \mathbf{A}_i^{1/2} \mathbf{R}_i \mathbf{A}_i^{1/2}$, $\mathbf{A}_i$ is a diagonal matrix of working marginal variances, and $\mathbf{R}_i$ is the working correlation matrix for $\mathbf{Y}_i$, $i = 1, \dots, N$, that is assumed to be independence, i.e. $\mathbf{R}_i$ is an identity matrix. The well-known consistent estimator for $\text{Cov}(\hat{\beta}_{GEE})$ is given by

$$\widehat{\text{Cov}}(\hat{\beta}_{GEE}) = \left(\sum_{i=1}^N \mathbf{X}_i^T \mathbf{V}_i^{-1} \mathbf{X}_i \right)^{-1} \left(\sum_{i=1}^N \mathbf{X}_i^T \mathbf{V}_i^{-1} \widehat{\text{Cov}}(\mathbf{Y}_i) \mathbf{V}_i^{-1} \mathbf{X}_i \right) \left(\sum_{i=1}^N \mathbf{X}_i^T \mathbf{V}_i^{-1} \mathbf{X}_i \right)^{-1}$$

, where $\widehat{\text{Cov}}(\mathbf{Y}_i) = (\mathbf{Y}_i - \hat{\boldsymbol{\mu}}_i)(\mathbf{Y}_i - \hat{\boldsymbol{\mu}}_i)^T$. When the number of independent clusters is

small, use of $\widehat{\text{Cov}}(\mathbf{Y}_i) = (\mathbf{Y}_i - \hat{\boldsymbol{\mu}}_i)(\mathbf{Y}_i - \hat{\boldsymbol{\mu}}_i)^T$ within this empirical sandwich covariance estimator results in negative bias. Therefore, Skene and Kenward [1] proposed using the Mancl and DeRouen [12] correction, which incorporates

$$\widehat{\text{Cov}}(\mathbf{Y}_i) = (\mathbf{I}_{n_i} - \mathbf{H}_i)^{-1} (\mathbf{Y}_i - \hat{\boldsymbol{\mu}}_i)(\mathbf{Y}_i - \hat{\boldsymbol{\mu}}_i)^T (\mathbf{I}_{n_i} - \mathbf{H}_i^T)^{-1}$$

within the empirical sandwich estimator, where $\mathbf{H}_i = \mathbf{X}_i \left(\sum_{j=1}^N \mathbf{X}_j^T \mathbf{V}_j^{-1} \mathbf{X}_j \right)^{-1} \mathbf{X}_i^T \mathbf{V}_i^{-1}$. Furthermore, they utilized the

method proposed by Pan and Wall [13] to approximate degrees of freedom to use with t-tests and F-tests. In short, this approach accounts for the estimation variability of the empirical sandwich covariance estimator in order to approximate degrees of freedom, and we refer the reader to [13] for details.

This particular GEE approach was found to maintain valid inference but at the cost of power. Therefore, in a second manuscript Skene and Kenward [3] proposed the use of a modification to Box's [4, 5] correction. Their motivation for a modified Box correction is that it does not require the working covariance matrix to be estimated in order to obtain estimated regression parameters or their estimated covariance matrix. Although we direct the reader to [3] for details, we note that their proposed test statistic does not involve a working covariance structure and therefore inherently assumes an independence structure. However, attaining the approximate distribution for the test statistic under the null hypothesis does require the use of Σ_i , $i = 1, \dots, N$, which were assumed to be unstructured and were estimated using REML.

Their study showed that the modified Box correction results in nominal test size. Furthermore, they showed that it is more powerful than their previously described GEE approach in [1], the use of likelihood-based inference using an unstructured working covariance matrix and the Kenward and Roger [23, 24] method (when valid), and Box's correction [4, 5].

2.3 Alternative GEE-Based Approach

The alternative approach to the analysis of very small samples of gaussian repeated measurements that we explore in this manuscript is based on our previous work with GEE [17, 20, 18, 34, 19]. In order to maintain valid inference when using GEE, for any working structure employed by GEE and not just the independence structure as in Skene and Kenward [1], we use modifications to the empirical sandwich covariance estimator. Furthermore, one of these modifications can increase degrees of freedom, and hence power. We also utilize a correlation structure selection approach that can increase power in settings in which the working structure influences estimation efficiency.

The general form for the modified empirical sandwich covariance estimator is given by

$$\widehat{\text{Cov}}_{BCE}(\hat{\beta}) = (I_p + \hat{G}) \hat{\Sigma}_M \left(\sum_{i=1}^N D_i^T V_i^{-1} \widehat{\text{Cov}}(Y_i) V_i^{-1} D_i \right) \hat{\Sigma}_M (I_p + \hat{G})^T, \quad (1)$$

where $G = (G_0, G_1, \dots, G_{p-1})$,

$$G_r = - \hat{\Sigma}_M \sum_{i=1}^N X_i^T A_i^{-1/2} R_i^{-1} \frac{\partial R_i(\hat{\alpha}(\beta))}{\partial \beta_r} R_i^{-1} A_i^{-1/2} (Y_i - \mu_i(\beta)),$$

$$\widehat{\text{Cov}}(\mathbf{Y}_i) = (\mathbf{I}_{n_i} - \tilde{\mathbf{H}}_i)^{-c} (\mathbf{Y}_i - \hat{\boldsymbol{\mu}}_i) (\mathbf{Y}_i - \hat{\boldsymbol{\mu}}_i)^T (\mathbf{I}_{n_i} - \tilde{\mathbf{H}}_i^T)^{-c}, \tilde{\mathbf{H}}_i = \mathbf{X}_i (\mathbf{I}_p + \mathbf{G}) \left(\sum_{j=1}^N \mathbf{X}_j^T \mathbf{V}_j^{-1} \mathbf{X}_j \right)^{-1} \mathbf{X}_i^T \mathbf{V}_i^{-1}$$

and $\hat{\boldsymbol{\Sigma}}_M = \left(\sum_{i=1}^N \mathbf{X}_i^T \mathbf{V}_i^{-1} \mathbf{X}_i \right)^{-1}$ [17, 18, 19]. Here, $c = 1$ ($c = 1/2$) for the Mancl and DeRouen [12] (Kauermann and Carroll [35]) correction. We consider the Kauermann and Carroll [35] correction, rather than just the Mancl and DeRouen [12] correction as in the first Skene and Kenward [1] method, because it has been shown to work well and is less conservative [36, 18, 34, 19]. Furthermore, $\mathbf{G}$ is utilized in Equation (1) to account for the inflation of $\text{Cov}(\hat{\boldsymbol{\beta}})$ due to the estimation of nuisance correlation parameters [17, 18]. To attain valid inference, the Pan and Wall [13] degrees of freedom approach can be utilized.

To potentially increase power, Westgate and Burchett [19] proposed using $\widehat{\text{Cov}}_{WB}(\mathbf{Y}_i) =$

$$\mathbf{A}_i^{1/2} \left(\frac{1}{N} \sum_{j=1}^N \mathbf{A}_j^{-1/2} (\mathbf{I}_{n_j} - \tilde{\mathbf{H}}_j)^{-c} (\mathbf{Y}_j - \hat{\boldsymbol{\mu}}_j) (\mathbf{Y}_j - \hat{\boldsymbol{\mu}}_j)^T (\mathbf{I}_{n_j} - \tilde{\mathbf{H}}_j^T)^{-c} \mathbf{A}_j^{-1/2} \right) \mathbf{A}_i^{1/2}, \quad (2)$$

rather than $\widehat{\text{Cov}}(\mathbf{Y}_i)$, in Equation (1). If a common marginal variance is assumed, such that $\mathbf{A}_j = \sigma^2 \mathbf{I}$, then $\widehat{\text{Cov}}_{WB}(\mathbf{Y}_i)$ is a bias-corrected empirical estimate of an unstructured covariance matrix which is assumed to be equivalent for each subject. We note that this formula can be modified to handle missing data [19]. The idea here of assuming a common structure [37] is to reduce the variability of the empirical sandwich covariance matrix and thus increase degrees of freedom and ultimately power. Such an increase in power will arise when conducting tests on regression parameters corresponding to covariates that have distributions of values that vary across subjects, and calculation of degrees of freedom is described by Westgate and Burchett [19].

We note that thus far, the proposed alternative GEE approach is similar to the first Skene and Kenward [1] method. However, the small-sample corrections just described have the potential to increase power. Furthermore, the use of these particular small-sample corrections do not limit the analyst to an independence working structure. The importance of this is that the working independence structure can be very inefficient in longitudinal studies and can therefore result in much less power than if a more realistic working correlation structure is utilized [14, 15, 16]. As we will demonstrate in our simulation study, the working structure does not have to be the true structure in order to improve power. However, accurate modeling of this structure can improve estimation efficiency, and thus power, even further.

In practice, the true correlation structure will be unknown, and the data analyst will have to decide on a structure to use. If the analyst is considering more than one structure, a criterion can be utilized to select the best of the competing structures; i.e., the structure that results in the smallest standard errors and thus greatest power. Multiple criteria for selecting a working correlation structure have been proposed [20]. For instance, Westgate [20] found the 'trace of the empirical covariance matrix' (TECM) criterion and 'correlation information criterion' (CIC) [38] both work well. Therefore, our alternative GEE-based approach utilizes one of

these criteria that incorporate Equation (1). For instance, the TECM for any given structure is given by $\text{tr}[\widehat{\text{Cov}}_{BCE}(\hat{\beta}_i)]$, and the structure that results in the smallest value is chosen. We note that selection accuracy can increase when Equation (2) is used within Equation (1) [19]. Furthermore, in order to attain valid inference, a degree of freedom penalty may need to be employed such that for each additional structure under consideration, one degree of freedom is lost. For example, if exchangeable and AR-1 are under consideration for selection, a 1 degree of freedom penalty would be applied to the distribution of the test statistic under the null hypothesis. This type of penalty was demonstrated in both longitudinal study and cluster randomized trial settings [19, 39].

3 Simulation Study

3.1 Study Description

We now conduct a simulation study to assess how our alternative GEE approach compares with the methods of Skene and Kenward [1, 3] in terms of maintaining the nominal 5% significance level and resulting power. We further present results from the use of a multivariate Gaussian linear model utilizing either the true (unknown in practice), unstructured, or a chosen covariance structure with the Kenward and Roger [23, 24] adjustment. Use of the true structure gives an ideal power, although this structure will likely be unknown in practice. Therefore, for each simulated dataset, we allow a variety of structures (independence, exchangeable, AR-1, unstructured) to be considered for selection by the AIC and BIC. We note that in practice the true structure may not even be considered for selection, with the possible exception of the unstructured covariance matrix. However, results from the sole use of the unstructured matrix will demonstrate problems resulting from the need to estimate numerous nuisance parameters. Furthermore, even though we allow the true structure to be selected by the AIC and BIC in our simulation study, our results will show that inflated test sizes can still result.

We present results from three scenarios that vary in terms of the efficiency gain that can occur due to accurate modeling of the working correlation structure. The first scenario represents a situation in which accurate modeling of the working correlation structure can notably improve estimation efficiency. Specifically, we generate multivariate normal outcomes from the marginal mean structure given by $\mu_{ij} = \beta_0 + \beta_1 x_{1ij} + \beta_2 x_{2ij}$, $j = 1, 2, 3, 4$, based on other studies [40, 38, 20, 18, 19, 34]. Here, $\text{Var}(\epsilon_{ij}) = 1$, $j = 1, 2, 3, 4$; $i = 1, \dots, N$, $\beta = [0, 0, 0]^T$ to assess test size, and $\beta = [0, 0.3, 1.0]^T$ to assess power.

For the second and third scenarios, MVN data are generated from the marginal mean structure given by $\mu_{ij} = \beta_0 + \beta_1 \text{Time}_j + \beta_2 (\text{Time}_j \times \text{Intervention}_i)$; $j = 1, 2, 3, 4$; $\text{Time}_j = j - 1$; based on a general randomized trial. Here, Intervention_i is an indicator for the i th subject receiving intervention and is equal to 1 for half of the subjects, $\text{Var}(\epsilon_{ij}) = 1$, $j = 1, 2, 3, 4$; $i = 1, \dots, N$, $\beta = [0, 0, 0]^T$ to assess test size, and $\beta = [0, 0, 0.5]^T$ to assess power relative to the interaction. In Scenario 2, the model for inference is the true model with three regression parameters, representing a situation in which accurate modeling of the working correlation structure can improve estimation efficiency, but not to a point such that power can be as greatly improved as in Scenario 1. In Scenario 3, a saturated model is fit. In this

situation, with no missing data, repeated measures ANOVA and repeated measures MANOVA [2] are popular. However, based on a simulation design of Skene and Kenward [1] allowing for missing data, we allow one subject to drop out, completely at random, at a random time (1, 2, or 3) in each trial arm. We note that if there were no missing data, the choice of a working correlation structure would have no impact on estimation efficiency in Scenario 3. This is similarly the case with this missing data scenario. Furthermore, the method of Westgate and Burchett [19], using $\widehat{\text{COV}}_{WB}(\mathbf{Y}_i)$ in Equation (1), will not improve power in Scenarios 2 and 3 because the distribution of covariate values for time are equivalent for subjects within the same trial arm, with the exception of possible missing data points.

Table 1 presents empirical test sizes and powers resulting from the methods under comparison, and Table 2 presents the number of times each structure under consideration is selected. In Table 2, we only present results from settings used to obtain empirical test sizes, as selection frequencies are similar in the settings used to obtain empirical powers. Results for each setting are based on the analyses of 10,000 generated datasets. We note that in order to account for random simulation error in terms of judging the validity of inference, empirical test sizes ranging from approximately 0.046 to 0.054 would have corresponding confidence intervals that cover the nominal 0.05 level. Each small-sample setting is comprised of either 10 or 20 subjects. The true correlation structure is AR-1 with correlation parameter value of 0.8, although we set this value to 0.5 in the settings of Scenarios 2 and 3 with 10 subjects to avoid lack of convergence. We note that with low values for the correlation, differences between the methods being compared will be small due to small differences in efficiency. For instance, we found this with a correlation value of 0.2, and therefore results are not presented. Furthermore, results from settings in which exchangeable is the true structure were similar to the presented results from a true AR-1 structure, with the exception that use of a working exchangeable structure would be preferred over a working AR-1 structure, and are therefore not shown.

We note that we present results from four different versions of our alternative GEE-based approach. Specifically, either the Mancl and DeRouen [12] or Kauermann and Carroll [35] correction is used, and a common unstructured covariance matrix in the center of the empirical sandwich covariance estimator is either assumed or not assumed. When not assumed, i.e. use of $\widehat{\text{COV}}(\mathbf{Y}_i) = (\mathbf{I}_{n_i} - \tilde{\mathbf{H}}_i)^{-c} (\mathbf{Y}_i - \hat{\boldsymbol{\mu}}_i)(\mathbf{Y}_i - \hat{\boldsymbol{\mu}}_i)^T (\mathbf{I}_{n_i} - \tilde{\mathbf{H}}_i^T)^{-c}$ in Equation (1), then we present results from the use of Pan and Wall [13] degrees of freedom. When a common structure is assumed, i.e. use of Equation (2) in Equation (1), we utilize the degrees of freedom of Westgate and Burchett [19]. Additionally, our GEE approach utilizes correlation selection with a degree of freedom penalty. Specifically, we use the TECM to select either a working exchangeable or AR-1 structure, as these are popular structures that may not lack efficiency relative to independence and do not require the estimation of multiple nuisance parameters as with the unstructured matrix [34]. We note that similar results were obtained when utilizing the CIC rather than the TECM, and therefore these results are not presented. Because the AR-1 structure is the truth in our simulation study and is allowed to be selected by our alternative GEE-based approach, empirical powers resulting from this approach may be thought to be unfairly high, although they do demonstrate the

potential advantage of correlation selection. Therefore, we also present results from the sole use of a working exchangeable structure to demonstrate that notable power gains, relative to the use of the Skene and Kenward methods, can still be observed without correlation selection and when a misspecified working structure is utilized.

3.2 Study Results

3.2.1 Validity of Inference—Use of the true covariance structure with the multivariate Gaussian linear model and the Kenward and Roger [23, 24] adjustment maintained the nominal 0.05 level, albeit typically on the conservative side. Alternatively, even the use of the Kenward and Roger [23, 24] adjustment could not maintain the nominal 0.05 level when the unstructured covariance matrix was used, with the exception of Scenario 3 in which a saturated mean model is used and $N = 20$. Furthermore, also as expected, use of the AIC and BIC resulted in inflated test sizes in most settings, particularly when $N = 10$, because the true AR-1 structure was not consistently selected. We note that in practice, the true structure may not even be considered, in which case test sizes may be inflated even further.

The first Skene and Kenward [1] approach did well at maintaining the nominal 0.05 level, except in Scenario 3 with $N = 10$. Alternatively, test sizes were inflated with their modified Box correction in Scenarios 2 and 3. The proposed alternative approach utilizing GEE, whether solely using a working exchangeable structure or selecting between the exchangeable and AR-1 structures, did well in terms of maintaining the 0.05 level for settings with 20 subjects. However, for Scenarios 1 and 2 settings with 10 subjects and Scenario 3 with 20 subjects, the Mancl and DeRouen [12] correction was needed to ensure the validity of inference. In Scenario 3, both the Kauermann and Carroll [35] and Mancl and DeRouen [12] corrections were liberal for 10 subjects.

3.2.2 Power—Ideally, use of the true, albeit unknown in practice, covariance structure with the multivariate Gaussian linear model and the Kenward and Roger [23, 24] adjustment will result in the greatest power. Alternatively, use of the unstructured covariance matrix resulted in lower power in the setting in which its resulting empirical size was not liberal (Scenario 3, $N = 20$). Furthermore, even in some settings in which test sizes were inflated, power was lower when using the unstructured covariance matrix due to the loss in degrees of freedom resulting from the need to estimate numerous nuisance covariance parameters. We note that these results agree with previous literature. In general, because the use of the AIC and BIC typically resulted in invalid inference, and possibly even more so in practice if the true structure is not considered, empirical powers from the use of these selection criteria should not be compared with powers resulting from the use of valid methods. However, we do note that when valid inference was attained, empirical powers were similar to those resulting from the sole use of the true AR-1 structure because this structure was selected in the majority of simulations in such settings.

As in Skene and Kenward [1, 3], in Scenario 3 we found their first method, based on the use of GEE with a working independence structure, lacks power. Alternatively, use of their modified Box correction in Scenario 3 actually resulted in empirical powers that were notably larger than the empirical powers resulting from the use of the true covariance

structure with the multivariate Gaussian linear model and the Kenward and Roger [23, 24] adjustment. However, this result is in part due to inflated test sizes, and we therefore also studied the use of half the denominator degrees of freedom with this approach. Resulting empirical sizes and powers were 0.028 and 0.206, respectively, when $N = 10$ and 0.030 and 0.775, respectively, when $N = 20$. Therefore, the use of half the denominator degrees of freedom does appear to result in valid inference and power that is similar to power resulting from the use of the true covariance structure with the multivariate Gaussian linear model. However, in Scenarios 1 and 2 in which time-dependent covariates were continuous, both Skene and Kenward [1, 3] methods, when valid, resulted in large losses in power relative to the ideal use of the true covariance structure with the multivariate Gaussian linear model.

In Scenarios 1 and 2, our alternative GEE (AGEE) approach with small-sample corrections and correlation selection greatly improved power relative to both Skene and Kenward [1, 3] methods, given that comparisons were among versions of these methods that were valid. In Scenario 1, the sampled distributions of covariate values varied across subjects, i.e. covariate values varied across subjects, and as a result the AGEE version assuming and using a common unstructured covariance matrix in the center of the empirical sandwich covariance estimator was more powerful and resulted in empirical powers that were very close to the empirical powers resulting from the use of the true covariance structure with the multivariate Gaussian linear model. Alternatively, in Scenario 2 in which the continuous time covariate did not vary in value across subjects within the same trial arm, with the exception of any dropout, assuming a common unstructured covariance matrix in the center of the empirical sandwich covariance estimator was not greatly more powerful than when not assuming this. When fitting a saturated model, Scenario 3 shows that the modified Box correction is preferred over the AGEE approach.

Table 2 shows that the TECM, no matter which of the four different versions of the working estimate for $\text{Cov}(\hat{\beta})$ is utilized, selected the true AR-1 structure the majority of the time. Furthermore, the number of times AR-1 was selected increased with N . However, even when GEE solely utilized the exchangeable working structure, greater empirical powers were observed than with either Skene and Kenward [1, 3] approach in Scenarios 1 and 2.

3.2.3 Summary—Results from the use of the multivariate Gaussian linear model agree with results found in many other studies and were provided simply to give a basis of comparison for the Skene and Kenward [1, 3] methods and the different versions of our AGEE approach. In short, sole use of the true structure, although likely not possible in practice as it will be unknown, was initially assumed to yield ideal power and valid inference. Use of the unstructured covariance matrix either lacked the validity of inference or power, and use of the AIC or BIC also tended to result in inflated test sizes. When fitting a saturated model in the settings of Scenario 3, use of the modified Box correction of Skene and Kenward [3] with half the denominator degrees of freedom resulted in valid inference and was approximately as powerful as the sole use of the true covariance structure with the multivariate Gaussian linear model and the Kenward and Roger [23, 24] adjustment. We note that these results coincide with the results of Skene and Kenward [3], with the exception of the need to use fewer denominator degrees of freedom. Furthermore, in the presence of continuous time-dependent covariates, both methods proposed by Skene and

Kenward [1, 3] had a notable lack in power, whereas the use of our AGEE approach tended to result in much greater power even when solely using the incorrect exchangeable structure. Furthermore, the validity of inference and level of power with the AGEE approach are slightly influenced by the type of standard error bias correction that is utilized, and possibly by whether or not a common unstructured covariance matrix is used in the center of the empirical sandwich covariance estimator.

4 Application

Our application example utilizes data from 20 subjects from the Prevention of Alzheimer's Disease by Vitamin E and Selenium (PREADViSE) clinical trial [41, 42]. Each year, for 4 to 5 consecutive years, Consortium to Establish a Registry for Alzheimer's Disease (CERAD) Word List Learning scores [43] were obtained as a partial measure of global cognitive status [41, 42]. Potential predictors of Word List Learning scores include time in years from study entry and baseline age ranging from 62 to 77 years, and we use the marginal model given by

$$E(\text{Score}_{ij}) = \beta_0 + \beta_1 \text{time}_j + \beta_2 \text{Age}_i; \quad j=1, \dots, 4 \text{ or } 5; \quad i=1, \dots, 20.$$

As this is a realistic application example, we do not know the true covariance structure. Therefore, we use the AIC and BIC to choose a working structure for the multivariate Gaussian linear model, and we use the TECM to select a working structure for use within GEE. The AIC and BIC select unstructured and exchangeable, respectively, whereas each of the four different versions of the TECM select the AR-1 structure. Results from the analyses of this dataset using a multivariate Gaussian linear model with the Kenward and Roger [23, 24] adjustment, both Skene and Kenward [1, 3] methods, and the four different versions of the alternative approach utilizing GEE are presented in Table 3. Specifically, regression parameter estimates, estimated standard errors, and p-values are given.

Regression parameter estimates were notably different for the multivariate Gaussian linear model with unstructured working covariance matrix, possibly demonstrating the small-sample instability that can arise when estimating numerous nuisance covariance parameters. The first method of Skene and Kenward [1] utilizes GEE with a working independence structure, and therefore there was a small difference in regression parameter estimates between this approach and our GEE-based approach which utilizes an AR-1 structure for the analysis of this particular dataset. Similarly, there were small differences in estimates resulting from the use of the exchangeable structure in the multivariate Gaussian linear model. Finally, the modified Box correction of Skene and Kenward [3] just corresponds to a test and ignores regression parameter estimation; therefore, no estimates are presented in Table 3 for this approach. Furthermore, due to the large number of denominator degrees of freedom, p-values are similar whether the full (results presented) or half of the denominator degrees of freedom are utilized.

Table 3 shows that with respect to time (β_1), p-values were similar resulting from the use of the modified Box correction of Skene and Kenward [3] and the different GEE-based approaches. With respect to baseline age (β_2), p-values resulting from the use of the modified Box correction or the AGEE version assuming and using a common unstructured

covariance matrix in the center of the empirical sandwich covariance estimator were also similar. P-values were a little larger with the AGEE version not making this assumption, as not making this assumption may not be as powerful with respect to a covariate such as baseline age that varies in value across subjects [19]. Alternatively, the first Skene and Kenward [1] approach resulted in the largest p-values as this approach lacks power. We note that the smallest p-values occurred from the use of the multivariate Gaussian linear model with unstructured or exchangeable working covariance matrices. However, our simulation study suggests that this result from the use of the unstructured matrix could be due to the fact that use of this structure may not result in valid inference. Similarly, use of the exchangeable structure may not result in valid inference as there is no way of knowing if this is actually the true structure.

In summary, based on evidence from previous studies and our simulation study, the modified Box correction of Skene and Kenward [3] or the AGEE version assuming and using a common unstructured covariance matrix in the center of the empirical sandwich covariance estimator are preferred in terms of the validity of inference and power in this particular setting. Both methods performed similarly here in terms of resulting p-values. Although AGEE did not have an apparent gain in power here relative to its primary contender, it still has an advantage in that standard output (regression parameter estimates, standard errors, confidence intervals) can be easily obtained.

5 Concluding Remarks

Use of the multivariate Gaussian linear model is very popular in practice for the analysis of Gaussian repeated measurements. However, when the number of subjects is small, it can be a challenge to ensure valid, yet powerful, inference. Therefore, Skene and Kenward [1, 3] proposed two different approaches, with a modified Box correction being preferred. These approaches ignore the correlation structure of outcomes when obtaining regression parameter estimates by assuming an independence structure. Therefore, in this manuscript we proposed the use of an alternative approach based on GEE that uses small-sample corrections and correlation selection. Although the modified Box correction was shown to be preferred when fitting a saturated mean model, this method has the potential to result in a great loss in power when fitting a model with one or more continuous time-dependent covariates. Alternatively, our alternative approach based on GEE was shown to result in much greater power in such situations, as the use of a correlation selection strategy results in more accurate correlation structure modeling and thus improved estimation efficiency for regression parameters corresponding to such types of covariates.

We looked at four different versions of our alternative approach based on GEE. Specifically, both the Kauermann and Carroll [35] and Mancl and DeRouen [12] corrections were studied. We also used the traditional approach to obtaining the empirical sandwich covariance estimator, and an approach that assumes and uses a common unstructured covariance matrix in the center of the empirical sandwich covariance estimator. As discussed in Westgate and Burchett [19], use of a common unstructured covariance matrix only helps power when the sampled distribution of covariate values varies across subjects, as was the case in Scenario 1 of our simulation study and in our application example with respect to

baseline age. Furthermore, when there are very few subjects in a study, e.g. 10 in our simulation study, the Mancl and DeRouen [12] correction may be safer to use than the Kauermann and Carroll [35] correction in order to ensure the validity of inference.

Our goal for accurate correlation structure selection with the proposed GEE approach is to choose the working structure, of those that are considered, that will result in the smallest standard errors of regression parameter estimates and thus the greatest power. We note that in practice this best structure may not be the true structure, as the analyst will not know if any of the structures considered for selection are actually the true structure. Furthermore, in previous work [34], we discuss a small-sample simulation study setting in which the true structure was unstructured, and yet use of an AR-1 working structure, although not the true structure, resulted in the smallest empirical standard deviations of regression parameter estimates, which as a result would yield the greatest power. As noted in [34], in our experience it is difficult to find a small-sample setting in which use of the working unstructured matrix will result in the smallest standard errors, even when the other structures under consideration are incorrect, due to the need to estimate multiple correlation parameters.

Therefore, we suggest that the analyst consider simple structures such as exchangeable and AR-1 for selection, as the use of a working independence structure has the potential to be very inefficient. Even if neither structure is the true structure, there can be a gain in power relative to the use of the independence structure, as demonstrated by the sole use of the exchangeable structure in our simulation study. We do note, however, that if the true correlation value is small, use of different working correlation structures can result in similar standard errors due to low impacts on efficiency, as noted in Section 3.1 for a correlation value of 0.2, in which case the accuracy of selecting the best structure is not critical. Furthermore, we would like to point out that correlation structure selection accuracy increases with the number of subjects, as can be observed in Table 2, because $Cov(\hat{\beta})$ used within selection criteria is estimated more precisely as N increases.

Our simulation results demonstrate that in small-sample settings the alternative approach based on GEE may be the preferred method for analysis when models contain only a small number of covariates, whereas the modified Box correction approach should be used when a saturated mean model is fit. In terms of the validity of inference, Scenario 3 demonstrated that the GEE approach may result in inflated test sizes in the presence of numerous covariates, whereas the modified Box correction approach works well, particularly when half the denominator degrees of freedom are utilized. Alternatively, the GEE approach resulted in greater power when there were only a small number of covariates, because the use of a more accurate working structure other than independence has the potential to notably improve power with respect to time-dependent covariates when the model is not saturated. However, we do note that these are only general suggestions for when to use these methods, as they are based on the findings from the three scenarios in our simulation study. Therefore, we encourage future studies with focus given to specific scenarios.

Future studies, focusing on scenarios similar to Scenario 3 when a saturated model is fit, could be particularly beneficial with respect to studying the modified Box correction

approach. This approach worked very well in [3], although we found that test sizes could be inflated. Therefore, we tried the use of half the denominator degrees of freedom. Although this resulted in valid, powerful inference, test sizes were conservative. Therefore, future work should consider modifications to the denominator degrees of freedom with the aim of attaining nominal sizes such that valid, yet the most powerful, inference will result.

Acknowledgments

We thank an anonymous associate editor and two reviewers for their constructive comments that resulted in the improvement of this manuscript.

We also thank Dr. Richard J. Kryscio, Dr. Frederick A. Schmitt, and Dr. Erin Abner for allowing us to use data from the PREADViSE trial, which was supported by a grant from the National Institute on Aging (R01 AG019241).

This publication was supported by the National Center for Research Resources and the National Center for Advancing Translational Sciences, National Institutes of Health, through Grant UL1TR000117. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

References

1. Skene SS, Kenward MG. The analysis of very small samples of repeated measurements I: An adjusted sandwich estimator. *Statistics in Medicine*. 2010; 29:2825–2837. [PubMed: 20839367]
2. Fitzmaurice, GM., Laird, NM., Ware, JH. *Applied Longitudinal Analysis*. 2nd. Hoboken, New Jersey: John Wiley & Sons, Inc.; 2011.
3. Skene SS, Kenward MG. The analysis of very small samples of repeated measurements II: A modified Box correction. *Statistics in Medicine*. 2010; 29:2838–2856. [PubMed: 20860066]
4. Box GEP. Some theorems on quadratic forms applied in the study of analysis of variance problems, I. Effect of inequality of variance in the one-way classification. *The Annals of Mathematical Statistics*. 1954; 25:290–302.
5. Box GEP. Some theorems on quadratic forms applied in the study of analysis of variance problems, II. Effects of inequality of variance and of correlation between errors in the two-way classification. *The Annals of Mathematical Statistics*. 1954; 25:484–498.
6. Geisser S, Greenhouse SW. An extension of Box's results on the use of the F distribution in multivariate analysis. *The Annals of Mathematical Statistics*. 1958; 29:885–891.
7. Greenhouse SW, Geisser S. On methods in the analysis of profile data. *Psychometrika*. 1959; 24:95–112.
8. Huynh H, Feldt LS. Estimation of the Box correction for degrees of freedom from sample data in randomized block and split-plot designs. *Journal of Educational and Behavioral Statistics*. 1976; 1:69–82.
9. Muller KE, Edwards LJ, Simpson SL, Taylor DJ. Statistical tests with accurate size and power for balanced linear mixed models. *Statistics in Medicine*. 2007; 26:3639–3660. [PubMed: 17394132]
10. Oberfeld D, Franke T. Evaluating the robustness of repeated measures analyses: The case of small sample sizes and nonnormal data. *Behavior Research Methods*. 2013; 45:792–812. [PubMed: 23184532]
11. Liang KY, Zeger SL. Longitudinal data analysis using generalized linear models. *Biometrika*. 1986; 73:13–22.
12. Mancl LA, DeRouen TA. A covariance estimator for GEE with improved small-sample properties. *Biometrics*. 2001; 57:126–134. [PubMed: 11252587]
13. Pan W, Wall MM. Small-sample adjustments in using the sandwich variance estimator in generalized estimating equations. *Statistics in Medicine*. 2002; 21:1429–1441. [PubMed: 12185894]
14. Fitzmaurice GM. A caveat concerning independence estimating equations with multivariate binary data. *Biometrics*. 1995; 51:309–317. [PubMed: 7766784]

15. Mancl LA, Leroux BG. Efficiency of regression estimates for clustered data. *Biometrics*. 1996; 52:500–511. [PubMed: 10766502]
16. Wang YG, Carey V. Working correlation structure misspecification, estimation and covariate design: implications for generalised estimating equations performance. *Biometrika*. 2003; 90:29–41.
17. Westgate PM. A bias correction for covariance estimators to improve inference with generalized estimating equations that use an unstructured correlation matrix. *Statistics in Medicine*. 2013; 32:2850–2858. [PubMed: 23255154]
18. Westgate PM. A covariance correction that accounts for correlation estimation to improve finite-sample inference with generalized estimating equations: a study on its applicability with structured correlation matrices. *Journal of Statistical Computation and Simulation*. 2016; 86:1891–1900. [PubMed: 27818539]
19. Westgate PM, Burchett WW. Improving power in small-sample longitudinal studies when using generalized estimating equations. *Statistics in Medicine*. 2016
20. Westgate PM. Improving the correlation structure selection approach for generalized estimating equations and balanced longitudinal data. *Statistics in Medicine*. 2014; 33:2222–2237. [PubMed: 24504841]
21. Patterson HD, Thompson R. Recovery of inter-block information when block sizes are unequal. *Biometrika*. 1971; 58:545–554.
22. Harville DA. Maximum likelihood approaches to variance component estimation and to related problems. *Journal of the American Statistical Association*. 1977; 72:320–338.
23. Kenward MG, Roger JH. Small sample inference for fixed effects from restricted maximum likelihood. *Biometrics*. 1997; 53:983–997. [PubMed: 9333350]
24. Kenward MG, Roger JH. An improved approximation to the precision of fixed effects from restricted maximum likelihood. *Computational Statistics and Data Analysis*. 2009; 53:2583–2595.
25. Kacker RN, Harville DA. Approximations for standard errors of estimators of fixed and random effects in mixed linear models. *Journal of the American Statistical Association*. 1984; 79:853–862.
26. Gomez EV, Schaalje GB, Fellingham GW. Performance of the Kenward-Roger method when the covariance structure is selected using AIC and BIC. *Communications in Statistics - Simulation and Computation*. 2005; 34:377–392.
27. Vallejo G, Fernandez MP, Livacic-Rojas PE, Tuero-Herrero E. Comparison of modern methods for analyzing repeated measures data with missing values. *Multivariate Behavioral Research*. 2011; 46:900–937. [PubMed: 26736117]
28. Gurka MJ, Edwards LJ, Muller KE. Avoiding bias in mixed model inference for fixed effects. *Statistics in Medicine*. 2011; 30:2696–2707. [PubMed: 21751227]
29. Gurka MJ. Selecting the best linear mixed model under REML. *The American Statistician*. 2006; 60:19–26.
30. Akaike H. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*. 1974; 19:716–723.
31. Schwarz G. Estimating the dimension of a model. *The Annals of Statistics*. 1978; 6:461–464.
32. Arnau J, Bono R, Vallejo G. Analyzing small samples of repeated measures data with the mixed-model adjusted F test. *Communications in Statistics - Simulation and Computation*. 2009; 38:1083–1103.
33. Schaalje GB, McBride JB, Fellingham GW. Adequacy of approximations to distributions of test statistics in complex linear models. *Journal of Agricultural, Biological and Environmental Statistics*. 2002; 7:512–524.
34. Westgate PM, Burchett WW. A comparison of correlation structure selection penalties for generalized estimating equations. *The American Statistician*. 2016
35. Kauermann G, Carroll RJ. A note on the efficiency of sandwich covariance matrix estimation. *Journal of the American Statistical Association*. 2001; 96:1387–1396.
36. Westgate PM. On small-sample inference in group randomized trials with binary outcomes and cluster-level covariates. *Biometrical Journal*. 2013; 55:789–806. [PubMed: 23852612]

37. Pan W. On the robust variance estimator in generalised estimating equations. *Biometrika*. 2001; 88:901–906.
38. Hin LY, Wang YG. Working-correlation-structure identification in generalized estimating equations. *Statistics in Medicine*. 2009; 28:642–658. [PubMed: 19065625]
39. Westgate PM. Intra-cluster correlation selection for cluster randomized trials. *Statistics in Medicine*. 2016
40. Hin LY, Carey VJ, Wang YG. Criteria for working-correlation-structure selection in GEE. *The American Statistician*. 2007; 61:360–364.
41. Caban-Holt A, Abner E, Kryscio RJ, Crowley JJ, Schmitt FA. Age-expanded normative data for the Ruff 2&7 Selective Attention Test: evaluating cognition in older males. *The Clinical Neuropsychologist*. 2012; 26:751–768. [PubMed: 22651854]
42. Kryscio RJ, Abner EL, Schmitt FA, Goodman PJ, Mendiondo M, Caban-Holt A, Dennis BC, Mathews M, Klein EA, Crowley JJ. A randomized controlled Alzheimer's disease prevention trial's evolution into an exposure trial: The PREADViSE trial. *The Journal of Nutrition, Health & Aging*. 2013; 17:72–75.
43. Chandler MJ, Lacritz LH, Hynan LS, Barnard HD, Allen G, Deschner M, Weiner MF, Cullum CM. A total score for the CERAD neuropsychological battery. *Neurology*. 2005; 65:102–106. [PubMed: 16009893]

Table 1

Empirical test sizes (nominal value of 0.05) and powers resulting from the use of the given methods for the analyses of 10,000 simulated datasets.

N	Analysis Method*	Scenario 1			Scenario 2			Scenario 3		
		Test Size	Power	Test Size	Power	Test Size	Power	Test Size	Power	Power
		$\beta_1=0$	$\beta_1=0.3$	$\beta_2=0$	$\beta_2=1.0$	$\beta_2=0$	$\beta_2=0.5$			
10	KR (True Structure)	0.039	0.140	0.036	0.866	0.054	0.567	0.040	0.203	
	KR (Unstructured)	0.149	0.233	0.148	0.793	0.152	0.476	0.056	0.143	
	KR (AIC)	0.086	0.188	0.084	0.867	0.121	0.606	0.073	0.262	
	KR (BIC)	0.070	0.175	0.069	0.868	0.112	0.607	0.070	0.263	
	SK I	0.038	0.060	0.034	0.346	0.046	0.404	0.058	0.170	
	SK II	0.059	0.100	0.054	0.490	0.088	0.565	0.066	0.339	
	AGEE Ia	0.069	0.169	0.067	0.814	0.076	0.530	0.098	0.255	
	AGEE Ib	0.051	0.134	0.050	0.752	0.053	0.435	0.064	0.179	
	AGEE IIa	0.054	0.168	0.052	0.870	0.077	0.526	0.105	0.253	
	AGEE IIb	0.044	0.145	0.043	0.845	0.054	0.431	0.071	0.181	
20	GEE Exch Ia	0.058	0.137	0.058	0.713	0.074	0.505	0.095	0.247	
	GEE Exch Ib	0.042	0.105	0.042	0.645	0.049	0.411	0.061	0.172	
	GEE Exch IIa	0.048	0.135	0.043	0.764	0.072	0.502	0.100	0.245	
	GEE Exch IIb	0.037	0.117	0.035	0.739	0.049	0.409	0.065	0.174	
	KR (True Structure)	0.045	0.279	0.045	0.994	0.052	0.957	0.042	0.753	
	KR (Unstructured)	0.071	0.281	0.071	0.983	0.069	0.911	0.052	0.680	
	KR (AIC)	0.053	0.287	0.054	0.993	0.069	0.957	0.052	0.763	
	KR (BIC)	0.047	0.282	0.047	0.994	0.062	0.959	0.047	0.765	
	SK I	0.037	0.095	0.035	0.636	0.045	0.698	0.043	0.629	
	SK II	0.052	0.133	0.052	0.737	0.065	0.776	0.062	0.876	
	AGEE Ia	0.058	0.276	0.056	0.988	0.051	0.928	0.064	0.711	
	AGEE Ib	0.049	0.246	0.048	0.983	0.040	0.908	0.046	0.655	
	AGEE IIa	0.051	0.292	0.052	0.994	0.050	0.927	0.068	0.702	
	AGEE IIb	0.047	0.277	0.048	0.994	0.040	0.906	0.052	0.649	
	GEE Exch Ia	0.051	0.209	0.053	0.954	0.048	0.920	0.061	0.700	
	GEE Exch Ib	0.043	0.186	0.044	0.945	0.037	0.898	0.045	0.644	

N	Analysis Method*	Scenario 1			Scenario 2			Scenario 3		
		Test Size	Power	Test Size	Test Size	Power	Power	Test Size	Test Size	Power
		$\beta_1=0$	$\beta_1=0.3$	$\beta_2=0$	$\beta_2=1.0$	$\beta_2=0$	$\beta_2=0.5$			
	GEE Exch IIa	0.048	0.222	0.051	0.972	0.048	0.921	0.066	0.695	
	GEE Exch IIb	0.045	0.211	0.047	0.969	0.037	0.899	0.049	0.641	

The true structure is AR-1 with $\rho = 0.8$, with the exception that $\rho = 0.5$ in Scenarios 2 and 3 when $N=10$ subjects. Tests correspond to covariates x_1 (β_1) and x_2 (β_2) in Scenario 1, a single parameter interaction effect (β_2) in Scenario 2, and Scenario 3 corresponds to this same interaction effect but utilizes a saturated model.

* Analysis Methods are as follows:

KR - Kenward & Roger method using the following working covariance structures: the true structure, unstructured, or structure selected by the AIC or BIC

SK - The first (SK I) and second (SK II) Skene and Kenward methods using GEE or modified Box correction, respectively.

AGEE - Our alternative GEE-based approach utilizing correlation selection.

- i. Does not assume and use a common unstructured covariance matrix in the center of the empirical sandwich covariance estimator.
- ii. Assumes and uses a common unstructured covariance matrix in the center of the empirical sandwich covariance estimator.
- a. Incorporates the Kauermann and Carroll correction
- b. Incorporates the Mancl and DeRouen correction

GEE Exch - The use of GEE with the small-sample corrections above (I or II, a or b) and a working exchangeable correlation structure.

Empirical frequencies of selecting each of the given working structures in 10,000 simulated datasets

Table 2

Scenario		N	Selection Method	Selection Frequencies				
				Independence	Exchangeable	AR-1	Unstructured	
1	10	AIC		7	1127	7492	1374	
		BIC		8	1206	8028	758	
		TECM Ia			2305	7695		
		TECM Ib			2603	7397		
		TECM IIa			1467	8533		
1	20	AIC		0	459	8844	697	
		BIC		0	494	9461	45	
		TECM Ia			1254	8746		
		TECM Ib			1419	8581		
		TECM IIa			507	9493		
2	10	AIC		997	1759	5621	1623	
		BIC		1274	1815	5936	975	
		TECM Ia			4107	5893		
		TECM Ib			3917	6083		
		TECM IIa			4100	5900		
2	20	AIC		0	565	8718	717	
		BIC		0	606	9325	69	
		TECM Ia			3896	6104		
		TECM Ib			3842	6158		
		TECM IIa			3854	6146		
3	10	AIC		1293	1730	5230	1747	
		BIC		1608	1794	5531	1067	
		TECM Ia			4790	5210		
		TECM Ib						
		TECM IIa						

Selection Frequencies					
Scenario	N	Selection Method	Independence	Exchangeable	AR-1 Unstructured
3	20	TECM Ib		4517	5483
		TECM IIa		5393	4607
		TECM IIb		4626	5374
		AIC	0	589	8693
		BIC	0	638	9289
		TECM Ia		3524	6476
		TECM Ib		3338	6662
		TECM IIa		3042	6958
		TECM IIb		3560	6440
					718
					73

For simplicity, only results corresponding to test size settings are presented.

N - Number of subjects

TECM - trace of the empirical covariance matrix

- i. Does not assume and use a common unstructured covariance matrix in the center of the empirical sandwich covariance estimator.
- ii. Assumes and uses a common unstructured covariance matrix in the center of the empirical sandwich covariance estimator.
- a. Incorporates the Kauermann and Carroll correction
- b. Incorporates the Manel and DeRouen correction

Table 3

Application example analysis results.

AGEE										
	KR		KR		SK I	SK II	Ia	Ib	IIa	IIb
	Unstructured/AIC	Exchangeable/BIC	Exchangeable/BIC	Exchangeable/BIC						
Time	$\hat{\beta}_1$	0.47	0.42	0.42	0.38	-	0.41	0.41	0.41	0.41
	SE	0.22	0.16	0.16	0.22	-	0.21	0.21	0.21	0.22
	P-Value	0.05	0.01	0.01	0.11	0.07	0.08	0.09	0.08	0.09
Age	$\hat{\beta}_2$	-0.56	-0.36	-0.36	-0.37	-	-0.39	-0.39	-0.39	-0.39
	SE	0.13	0.15	0.15	0.18	-	0.16	0.17	0.14	0.15
	P-Value	0.001	0.03	0.03	0.07	0.02	0.04	0.05	0.01	0.01

Model: $E(\text{Score}_{ij}) = \beta_0 + \beta_1 \text{time}_i + \beta_2 \text{Age}_i, j = 1, \dots, 4 \text{ or } 5; i = 1, \dots, 20.$

Analysis Methods are as follows:

- KR - Kenward & Roger method using an unstructured (selected by AIC) or exchangeable (selected by BIC) working covariance structure
- SK - The first (SK I) and second (SK II) Skene and Kenward methods using GEE or modified Box correction, respectively.

AGEE - Our alternative GEE-based approach utilizing correlation selection (AR-1 always selected).

- i. Does not assume and use a common unstructured covariance matrix in the center of the empirical sandwich covariance estimator.
- ii. Assumes and uses a common unstructured covariance matrix in the center of the empirical sandwich covariance estimator.
- a. Incorporates the Kauermann and Carroll correction
- b. Incorporates the Mancl and DeRouen correction