

Analysis of Repeated Measurements Using Linear Mixed Models

- Consider the situation in which a normally-distributed outcome variable is measured repeatedly for each subject
- Classical methodology is based on:
 - univariate repeated measures ANOVA
 - multivariate linear models with general covariance structure
- In practice, however, longitudinal studies are characterized by:
 - variation among individuals with respect to the number and timing of observations
 - missing data
 - time-dependent covariates

Analysis of Repeated Measurements Using Linear Mixed Models

- Such features make standard multivariate procedures difficult (or impossible) to apply
- An alternative approach based on the linear mixed model:
 - views the analysis of repeated measurements as a univariate regression analysis of responses with correlated errors
 - accommodates the complexities of typical longitudinal data sets
 - permits specification of models determined by subject matter considerations, rather than by limitations of the statistical methodology
 - allows explicit modelling and analysis of between- and within-individual variation

Usual Linear Model

- $y = X\beta + \epsilon$
 - y is a $n \times 1$ vector of independent observations
 - β is a $p \times 1$ vector of unknown parameters
 - X is a $n \times p$ model matrix
 - ϵ is a $n \times 1$ vector of independent errors
 - ϵ_i has mean 0 and constant variance σ^2
- The focus is to model the mean of y in terms of the unknown parameters β , which are estimated using ordinary least squares
- Straightforward generalizations include:
 - Weighted least squares: ϵ_i has variance σ_i^2
 - Generalized least squares: ϵ has mean 0 and covariance matrix Σ

Linear Mixed Model

- $y = X\beta + Z\gamma + \epsilon$
 - y is a $n \times 1$ vector of observations
 - X and Z are given $n \times p$ and $n \times q$ matrices, respectively
 - β is a $p \times 1$ vector of unknown parameters
 - γ and ϵ are unobservable random vectors of dimensions $q \times 1$ and $n \times 1$, respectively
 - $E(\gamma) = 0, \text{Var}(\gamma) = B$
 - $E(\epsilon) = 0, \text{Var}(\epsilon) = W$
 - γ and ϵ are assumed to be uncorrelated
- In this case, both the mean and variance of y can be modelled
 - $V = \text{Var}(y) = ZBZ' + W$

Linear Mixed Model

- The elements of B and W are known functions of an unknown parameter vector $\theta = (\theta_1, \dots, \theta_m)'$
- The parameter space for the model is taken to be $\{(\beta, \theta): \theta \in \Omega\}$, where Ω is the set of θ values for which W (hence V) is positive definite
- When $W = \sigma^2 I$ and $Z = 0$, the mixed model reduces to the standard linear model
- Other special cases include:
 - Multiple regression with AR(1) errors
 $\gamma \equiv 0, \Omega = \{(\sigma^2, \rho): \sigma^2 > 0, |\rho| < 1\}$
 - Mixed and random ANOVA models
 (variance-components models)

Parameter Estimation

- Estimation of the random effects (variance components) has been a longstanding problem
- Common practice for balanced ANOVA was to equate mean squares to their expectations
- Henderson (1953, *Biometrics*) developed analogous techniques for unbalanced data
- Maximum likelihood estimation of the fixed effects and variance components was not used widely, due to computational difficulties

Requires the numerical solution of a constrained nonlinear optimization problem

- Harville (1977, *JASA*) reviewed previous work, unified the methodology, and described iterative ML algorithms

Parameter Estimation

- In addition to computational difficulties, MLEs of variance components are biased downward
- Patterson and Thompson (1971, *Biometrika*) proposed the alternative REML approach
 - applies ML technique to likelihood function associated with a set of “error contrasts” rather than to that associated with y
 - accounts for the loss of degrees of freedom resulting from estimation of the fixed effects
 - gives less biased estimates of the variance components
 - yields the standard ANOVA-based estimates in balanced random and mixed ANOVA models (unlike maximum likelihood)

Example

- Consider the estimation of σ^2 in the usual linear model $y = X\beta + \epsilon$
 - y is the $n \times 1$ data vector
 - X is a $n \times p$ model matrix of rank p
 - $\epsilon \sim N(0, \sigma^2 I_n)$

- The MLE of σ^2 is

$$\tilde{\sigma}^2 = \frac{(y - X\hat{\beta})'(y - X\hat{\beta})}{n}$$

and the UMVUE is

$$\hat{\sigma}^2 = \frac{(y - X\hat{\beta})'(y - X\hat{\beta})}{n - p},$$

where $\hat{\beta} = (X'X)^{-1}X'y$

- The bias of the MLE is

$$E(\tilde{\sigma}^2 - \hat{\sigma}^2) = -\sigma^2 \left(\frac{p}{n} \right)$$

which is negative and worsens as p increases

Error Contrasts

Definition:

- A linear combination $a'y$ is an *error contrast* if $E(a'y) = 0$ for any β , i.e., if $a'X = 0'_p$

Example:

- Let $S = I_n - P_X$, where $P_X = X(X'X)^{-1}X'$
- $E(Sy) = (I_n - P_X)X\beta = X\beta - X\beta = 0$
- Each element of Sy is an error contrast
- While S is $n \times n$, its rank is $n - p$
- Thus, there are some redundancies among the elements of Sy

Question:

- How many essentially different error contrasts can be included in a single set?

Linearly Independent Error Contrasts

Definition: The error contrasts $a_1'y, \dots, a_k'y$ are *linearly independent* if a_1, \dots, a_k are linearly independent vectors

Theorem: Any set of error contrasts contains at most $n - p$ linearly independent error contrasts

Let A be a $n \times (n - p)$ matrix such that $A'A = I_{n-p}$ and $AA' = I_n - P_X$.

Theorem: $w = A'y$ is a vector of $n - p$ linearly independent error contrasts

(It is not the only such vector, however)

REML Estimation

- The REML approach applies ML to $w = A'y$
- Under the assumed model:

$$y \sim N_n(X\beta, V), \text{ where } V = ZBZ' + W$$

$$w \sim N_{n-p}(0, A'VA)$$

Question:

Would $\hat{\theta}$ obtained by maximizing $f_w(w; \theta)$ be the same as that obtained by maximizing the likelihood function associated with any other vector of $n - p$ linearly independent error contrasts?

Theorem:

Let $u = C'y$ be any vector of $n - p$ linearly independent error contrasts. The likelihood function associated with u is a scalar multiple of $f_w(w; \theta)$ that does not depend on θ

REML Estimation

- The log-likelihood function $L_R(\theta; y)$ associated with any vector of $n-p$ linearly independent error contrasts is (apart from an additive constant)

$$-\frac{1}{2} \left[\log |V| + \log |X'V^{-1}X| + (y - X\hat{\beta})'V^{-1}(y - X\hat{\beta}) \right]$$

where $\hat{\beta} = (X'V^{-1}X)^{-1}X'V^{-1}y$

- In comparison, the log-likelihood function for y is

$$L_M(\theta; y) = -\frac{1}{2} \log |V| - \frac{1}{2} (y - X\hat{\beta})'V^{-1}(y - X\hat{\beta})$$

- The only difference is that $L_R(\theta; y)$ has the additional term $-\frac{1}{2} \log |X'V^{-1}X|$
- $\hat{\theta}$ is a REML estimate of θ if $L_R(\theta; y)$ attains its maximum value at $\theta = \hat{\theta}$

REML Estimation

- In the simple Gauss-Markov linear model, the REML equations have a unique solution which coincides with the UMVUE

$$\hat{\sigma}_{\text{REML}}^2 = \frac{(y - X\hat{\beta})'(y - X\hat{\beta})}{n - p}$$

- In balanced mixed and random ANOVA models, the REML equations have an explicit unique solution coinciding with the ANOVA estimate
- In general, however, the problem of obtaining a REML estimate of θ requires iterative methods of maximizing the nonlinear function $L_R(\theta; y)$ subject to the constraint $\theta \in \Omega$
- Algorithms such as Newton-Raphson and the method of scoring can be used

Mixed Model for Repeated Measurements

- The general model for longitudinal data is

$$y_i = X_i\beta + Z_i\gamma_i + \epsilon_i, \quad i = 1, \dots, n$$

- y_i is the $t_i \times 1$ vector of responses for subject i
- X_i is a $t_i \times b$ design matrix for subject i
- β is a $b \times 1$ vector of regression coefficients
- γ_i is a $g \times 1$ vector of random effects for subject i
 - the γ_i vectors are independent $N(0, B)$
- Z_i is a $t_i \times g$ design matrix for the random effects
- ϵ_i is a $t_i \times 1$ vector of within-subject errors
 - the ϵ_i vectors are independent $N(0, W_i)$
- γ_i and ϵ_i are independent

Mixed Model for Repeated Measurements

- y_1, \dots, y_n are independent $N_{t_i}(X_i\beta, V_i)$, where

$$V_i = Z_i B Z_i' + W_i$$

- The matrices X_i , Z_i , and W_i are subject specific

- This model is very general, since:

subjects can have varying numbers of observations

observation times can differ among subjects

- The within-subject covariance matrix W_i depends on i only through its dimension t_i

unknown parameters in W_i do not depend on i

- A wide variety of covariance structures for γ_i and ϵ_i can be considered

Comments

- The general mixed model for repeated measures has been studied by several authors
- While essentially similar, the various approaches differ in terms of:
 - motivation and notation
 - assumptions concerning the random effects
 - estimation method
- Some of the main references are:
 - Laird and Ware (1982, *Biometrics*)
 - Jennrich and Schluchter (1986, *Biometrics*)
 - Laird, Lange, and Stram (1987, *JASA*)
 - Diggle (1988, *Biometrics*)
 - Lindstrom and Bates (1988, *JASA*)
 - Jones and Boadi-Boateng (1991, *Biometrics*)

Laird and Ware (1982)

- Two-stage random-effects model:
 1. For each subject, $y_i = X_i\beta + Z_i\gamma_i + \epsilon_i$
 $\epsilon_1, \dots, \epsilon_n$ are independent $N(0, W_i)$
 β and γ_i are considered to be fixed
 2. $\gamma_1, \dots, \gamma_n$ are independent $N(0, B)$
 γ_i and ϵ_i are assumed independent
- Conditional independence model if $W_i = \sigma^2 I_{t_i}$
- Laird and Ware discuss Bayesian and non-Bayesian formulations of the model
- The EM algorithm is used to obtain ML and REML parameter estimates
 - (unobservable random parameters are estimated, *not* missing data)

Jennrich and Schluchter (1986)

Overview of the Paper:

- Considers the problem of how to analyze unbalanced or incomplete repeated measures
- Uses a general linear model for expected responses and arbitrary structural models for the within-subject covariances
- Describes Newton-Raphson and Fisher scoring algorithms for obtaining ML estimates
- Describes a generalized EM (GEM) algorithm for computing REML estimates
 - the likelihood is increased (rather than maximized) at each M step
- Gives an example in which several models are fit to a set of growth data

Jennrich and Schluchter (1986)

- The model is $y_i = X_i\beta + \epsilon_i$, where $\epsilon_1, \dots, \epsilon_n$ are independent $N_{t_i}(0, \Sigma_i)$
 - $\Sigma_i = \Sigma_i(\theta)$, where θ is a vector of q unknown covariance parameters
- Motivation:
 - Although estimation of β is of primary interest, efficiency may be improved by modelling Σ_i parsimoniously
 - This is especially important when sample sizes are small and the data are unbalanced
 - The ability to model Σ_i allows examination of alternative covariance structures
- $\Sigma_i = Z_i B Z_i' + W_i$ yields the mixed model as a special case

Incomplete Data Model

- Useful when a fixed number t of measurements are to be obtained from each subject, but not all responses are observed
- Each Σ_i is a submatrix of a $t \times t$ matrix $\Sigma = \Sigma(\theta)$
- Some possible incomplete data models for Σ :

Structure	q	Description
Independent observations	1	$\Sigma = \sigma^2 I_t$
Compound symmetry	2	$\Sigma = \sigma^2 I_t + \sigma_b^2 1_t 1_t'$
Random effects (g effects)	$1+g(g+1)/2$	$\Sigma = ZBZ' + \sigma^2 I_t$ $Z_{(t \times g)}$ is known
First-order autoregressive	2	$\sigma_{ij} = \sigma^2 \rho^{ i-j }$
Toeplitz (banded)	t	$\sigma_{ij} = \theta_k,$ $k = i - j + 1$
Unstructured	$t(t+1)/2$	$\Sigma_{ij} = \Sigma_{ji}$

Comments on Computational Algorithms

- Newton-Raphson uses the score vector and Hessian matrix to iteratively compute new estimates of β and θ from current values
- Fisher scoring replaces the Hessian matrix by its expectation
 - often more robust to poor starting values than Newton-Raphson
- The hybrid EM scoring algorithm:
 - is restricted to the incomplete data model
 - has the advantage of being able to fit covariance matrices with large numbers of parameters, e.g., large unstructured Σ

Computation of Standard Errors

- Standard error estimates can be computed from the inverse of the:
 - Fisher information matrix (Fisher scoring)
 - Empirical information matrix (NR)
- The standard error estimates from the empirical information matrix are preferable when the data are incomplete
- Although the hybrid EM algorithm does not produce standard errors for the elements of θ , these can be obtained by taking a single Newton-Raphson or Fisher scoring step after convergence

Comparison of Computational Algorithms

- NR algorithm has a quadratic convergence rate and generally converges in a small number of iterations (but with a higher cost per iteration)
- Hybrid EM algorithm has lowest cost/iteration, but may require a large number of iterations
- Fisher scoring algorithm is intermediate in terms of cost/iteration and number of iterations (cost/iteration often not much less than NR; can require a much higher number of iterations)
- When q is small, Newton-Raphson is preferred:
 - not restricted to the incomplete data model
 - convergence is generally clean and fast
- With large q , as when fitting an unstructured covariance matrix to more than 10 time points, only the EM scoring algorithm is feasible

Laird, Lange, and Stram (1987)

- Study the use of the EM algorithm for ML and REML estimation in the model $y_i = X_i\beta + Z_i\gamma_i + \epsilon_i$
- Consider two models:
 1. *Incomplete Data* (Jennrich & Schluchter, 1986)
 - EM algorithm requires iterative M step within each iteration (or use of GEM)
 - covariates for both observed and missing observations must be specified (choice affects rate, but not point, of convergence)
 2. *Random Effects* (Laird and Ware, 1982)
 - total data set consists of observed data plus unobservable random parameters
 - there is no missing data in the traditional statistical sense

Laird, Lange, and Stram (1987)

- Show that the random effects model is more general and includes the incomplete data model as a special case
 - avoids specification of covariates for missing observations
 - eliminates need for GEM or iterations within each M step (in a broad class of models)
- Provide computing formulas for ML and REML estimation using the EM algorithm
- Discuss the choice of starting values for the EM iterations (and give several possibilities)
- Give two methods of speeding convergence of the EM algorithm

Diggle (1988)

- Model is $y_i = X_i\beta + \epsilon_i$, where $\epsilon_i \sim N(0, \Sigma_i)$ and

$$\Sigma_i = \tau^2 I + \nu^2 J + \sigma^2 R(t_i)$$

- J is a square matrix with elements of 1
- $t_i = (t_{i1}, \dots, t_{in_i})'$ is vector of measurement times for subject i
- $R(t)$ is a symmetric matrix with (k, l) th element $\exp(-\alpha|t_k - t_l|^c)$, where $c = 1$ or 2
- Thus, the within-subject covariance structure has four parameters: τ^2 , ν^2 , σ^2 , and α
- Parameters are estimated using ML and REML
- Empirical semi-variogram of residuals is used to suggest an appropriate correlation structure

Lindstrom and Bates (1988)

- Consider the special case with $W_i = \sigma^2 I$
- Develop efficient implementations of NR and EM algorithms for ML estimation
 - Make four improvements to Jennrich and Schluchter's (1986) ML algorithm to speed convergence and improve behavior
- Compare NR, EM, and EM with Aitken's acceleration in fitting 3 models to 2 data sets:
 - 11 subjects, average of 28 observations/subject
 - 74 subjects, average of 11 observations/subject
- Conclude that NR is generally preferable, based on number of iterations and average time/iteration

Jones and Boadi-Boateng (1991)

- They consider the model $y_i = X_i\beta + Z_i\gamma_i + \epsilon_i$ when each subject is observed at different and unequally spaced time points
- Observations within subjects are assumed to be either uncorrelated or to have a continuous-time AR-1 structure
- An alternative method of estimation is proposed:
 - state space representation of the model is used
 - likelihood is calculated using the Kalman filter
 - ML estimates are obtained using a nonlinear optimization program
- Advantage is that likelihood can be calculated recursively without using large $(t_i \times t_i)$ matrices

Example

- A study conducted in two groups of children (16 boys and 11 girls)
- At ages 8, 10, 12, and 14, the distance (mm) from the center of the pituitary gland to the pteryomaxillary fissure was measured
- Let y_{hij} denote the response at time j for the i th subject in group h , where $h = 1$ for boys and $h = 2$ for girls
- Jennrich and Schluchter (1986) use maximum likelihood to fit eight models to these data


Reference

Potthoff, R. F. and Roy, S. N. (1964). A generalized multivariate analysis of variance model useful especially for growth curve problems. *Biometrika* **51**, 313–326.

Dental Measurements

Group	ID	Age 8	Age 10	Age 12	Age 14
Boys	1	26.0	25.0	29.0	31.0
	2	21.5	22.5	23.0	26.5
	3	23.0	22.5	24.0	27.5
	4	25.5	27.5	26.5	27.0
	5	20.0	23.5	22.5	26.0
	6	24.5	25.5	27.0	28.5
	7	22.0	22.0	24.5	26.5
	8	24.0	21.5	24.5	25.5
	9	23.0	20.5	31.0	26.0
	10	27.5	28.0	31.0	31.5
	11	23.0	23.0	23.5	25.0
	12	21.5	23.5	24.0	28.0
	13	17.0	24.5	26.0	29.5
	14	22.5	25.5	25.5	26.0
	15	23.0	24.5	26.0	30.0
	16	22.0	21.5	23.5	25.0
	Mean	22.9	23.8	25.7	27.5
Girls	1	21.0	20.0	21.5	23.0
	2	21.0	21.5	24.0	25.5
	3	20.5	24.0	24.5	26.0
	4	23.5	24.5	25.0	26.5
	5	21.5	23.0	22.5	23.5
	6	20.0	21.0	21.0	22.5
	7	21.5	22.5	23.0	25.0
	8	23.0	23.0	23.5	24.0
	9	20.0	21.0	22.0	21.5
	10	16.5	19.0	19.0	19.5
	11	24.5	25.0	28.0	28.0
	Mean	21.2	22.2	23.1	24.1

Mean Dental Measurements



Potential Models

1. “Classical” multivariate methods
2. Models with fixed effects for group and time, and arbitrary covariance structures, as in Jennrich and Schluchter (1986)
3. Random intercept model
 - $E(y_{hij}) = \alpha_{hi} + \beta_h j$, where $\alpha_{hi} \sim N(\alpha_h, \sigma_\alpha^2)$
 - Equivalent to the fixed effects linear model with compound symmetry covariance structure
4. Random intercept and slope model
 - $E(y_{hij}) = \alpha_{hi} + \beta_{hi} j$, where

$$\begin{pmatrix} \alpha_{hi} \\ \beta_{hi} \end{pmatrix} \sim N_2\left(\begin{pmatrix} \alpha_h \\ \beta_h \end{pmatrix}, \begin{pmatrix} \sigma_\alpha^2 & \sigma_{\alpha\beta} \\ \sigma_{\alpha\beta} & \sigma_\beta^2 \end{pmatrix}\right)$$
 - Appealing due to its simplicity, but much more flexible than the random intercept model

Creation of SAS Data Set

```

data a;
keep id sex age distance;
input sex id d8 d10 d12 d14;
if sex=2 then id=id+16;
age=8;    distance=d8;    output;
age=10;   distance=d10;   output;
age=12;   distance=d12;   output;
age=14;   distance=d14;   output;
cards;
1  1 26.0 25.0 29.0 31.0
      . . .
1 16 22.0 21.5 23.5 25.0
2  1 21.0 20.0 21.5 23.0
      . . .
2 11 24.5 25.0 28.0 28.0
;
data b; set a;
m08=(sex=1 & age=8);
m10=(sex=1 & age=10);
m12=(sex=1 & age=12);
m14=(sex=1 & age=14);
f08=(sex=2 & age=8);
f10=(sex=2 & age=10);
f12=(sex=2 & age=12);
f14=(sex=2 & age=14);

```

Model 1

- A separate mean for each (h, j) : $y_{hij} = \beta_{hj} + \epsilon_{hij}$
and a common unstructured W_{hi}

```
proc mixed method=ml;
class id;
model distance=m08 m10 m12 m14
           f08 f10 f12 f14 / noint s;
repeated / type=un subject=id r;
```

- The `s` option of the `model` statement prints the estimates of the fixed effects
- Options of the `repeated` statement:
 - `type=un` specifies unstructured covariance matrix for within-subject errors
 - `subject=id` subject identifier is required
 - `r` prints the first block of the W_{hi} matrix of within-subject errors

Alternate Parameterization of Model 1

- The previous parameterization was chosen to match Jennrich and Schluchter (1986)
- A factorial model could also be considered
- In this case, the covariates must be specified in the `class` statement

```
proc mixed method=ml;  
  class id age sex;  
  model distance=age sex age*sex / s;  
  repeated / type=un subject=id r;
```

- A “corner-point” parameterization is used for class variables, with the highest category serving as reference
- Although the two models are equivalent, the tests of the fixed effects are interpreted differently

Model 2

- Linear relationship between distance and age, with separate lines for boys and girls:

$$y_{hij} = \alpha_h + \beta_h t_j + \epsilon_{hij},$$

where $t_j=8,10,12,14$ for $j=1,2,3,4$, respectively

- W_{hi} is unstructured (as in Model 1)

```
proc mixed method=ml;
  class id sex;
  model distance=sex age*sex / noint s;
  repeated / type=un subject=id r;
```

- Equivalent to Potthoff-Roy growth curve

model with $G = S$

- LR test comparing Model 2 to Model 1 is

$$419.48 - 416.51 = 2.97, \quad \text{df} = 18 - 14 = 4, \quad p = .56$$

- Relationship appears to be linear for each sex

Alternate Parameterizations of Model 2

- Model 2 was parameterized with a separate intercept and slope for each sex
- One alternative is a model with an intercept, a sex increment, and separate slopes for each sex:

```
proc mixed method=ml;  
class id sex;  
model distance=sex age*sex / s;  
repeated / type=un subject=id r;
```

- The INTERCEPT is now the intercept for females
- The SEX effect is now the increment (for males) to the overall intercept (intercept for females)
- In this case, a test of the intercept difference is automatically produced

Alternate Parameterizations of Model 2

- Another alternative is a model with an overall intercept and slope, plus an intercept increment and a slope increment:

```
proc mixed method=ml;  
class id sex;  
model distance=sex age age*sex / s;  
repeated / type=un subject=id r;
```

- The INTERCEPT effect is the intercept for females
- The AGE effect is the slope for females
- The SEX and AGE*SEX effects are the intercept and slope increments for males
- Tests of the intercept difference and slope difference are automatically produced

Model 3

- Linear relationship between distance and age, with a common slope for boys and girls:

$$y_{hij} = \alpha_h + \beta t_j + \epsilon_{hij}$$

- Unstructured W_{hi} (as in Models 1 and 2)

```
proc mixed method=ml;
class id sex;
model distance=sex age / noint s;
repeated / type=un subject=id r;
```

- LR test comparing Model 3 to Model 2 is

$$426.15 - 419.48 = 6.68, \quad df = 14 - 13 = 1, \quad p = .01$$

- Slopes for boys and girls differ significantly
- Models 4–8 will use the same model for the means as Model 2 and will examine the adequacy of restricted covariance structures

Model 4

- $y_{hij} = \alpha_h + \beta_h t_j + \epsilon_{hij}$
- Banded covariance matrix with 4×4 blocks

$$\begin{pmatrix} \theta_1 & \theta_2 & \theta_3 & \theta_4 \\ \theta_2 & \theta_1 & \theta_2 & \theta_3 \\ \theta_3 & \theta_2 & \theta_1 & \theta_2 \\ \theta_4 & \theta_3 & \theta_2 & \theta_1 \end{pmatrix}$$

```
proc mixed method=ml;
class id sex;
model distance=sex age*sex / noint s;
repeated / type=toep subject=id r;
```

- LR test comparing Model 4 to Model 2 is

$$424.64 - 419.48 = 5.17, \quad \text{df} = 14 - 8 = 6, \quad p = .40$$

- The banded covariance structure provides an adequate fit relative to the unstructured covariance matrix

Model 5

- $y_{hij} = \alpha_h + \beta_h t_j + \epsilon_{hij}$
- First-order autoregressive covariance structure with $w_{ij} = \sigma^2 \rho^{|i-j|}$

```
proc mixed method=ml;
class id sex;
model distance=sex age*sex / noint s;
repeated / type=ar(1) subject=id r;
```

- The estimated covariance parameters are

$$\hat{\sigma}^2 = 4.89099772, \quad \hat{\rho} = 0.6071465$$

- LR test comparing Model 5 to Model 2 is

$$440.68 - 419.48 = 21.2, \quad \text{df} = 14 - 6 = 8, \quad p = .007$$

- The AR(1) structure does not provide an adequate fit

Model 6

- Random coefficients model $y_{hij} = a_{hi} + b_{hi}t_j + \epsilon_{hij}$
 - $(a_{hi}, b_{hi})'$ are independent $N_2((\alpha_h, \beta_h)', B)$
 - ϵ_{hij} are independent $N(0, \sigma^2)$

```
proc mixed method=ml;
class id sex;
model distance=sex age*sex / noint s;
random intercept age / g subject=id type=un;
```

- `random` statement defines random effects constituting the γ vector in the mixed model
- `g` option prints estimated covariance matrix of γ
- LR test comparing Model 6 to Model 2 is
$$427.81 - 419.48 = 8.33, \quad \text{df} = 14 - 8 = 6, \quad p = .14$$
- Only $\hat{\sigma}^2$ is significantly different from zero

Comments on Model 6

- Note that Model 6 did not require use of the repeated statement
- This model could have equivalently been fit using the statements:

```
proc mixed method=ml;  
class id sex;  
model distance=sex age*sex / noint s;  
random intercept age  
          / g subject=id type=un;  
repeated / type=simple  
          subject=id r;
```

- Although the covariance parameter information is reported somewhat differently, the results are otherwise unchanged

Random Intercept and Slope Model

- Although often used, this model has a potential shortcoming, in that the covariance matrix of $y_{hi} = (y_{hi1}, \dots, y_{hik})'$ is nonstationary
- For equally-spaced time points $j = 1, 2, \dots$, general expressions are:

$$\text{Var}(y_{hij}) = \sigma_{\alpha}^2 + 2j\sigma_{\alpha\beta} + j^2\sigma_{\beta}^2 + \sigma^2$$

$$\text{Cov}(y_{ij}, y_{ij'}) = \sigma_{\alpha}^2 + (j + j')\sigma_{\alpha\beta} + jj'\sigma_{\beta}^2$$

- If $\sigma_{\alpha\beta} > -\sigma_{\beta}^2$, then $\text{Var}(y_{hij})$ will increase monotonically over time
- In terms of the correlation coefficient

$$\rho = \frac{\sigma_{\alpha\beta}}{\sigma_{\alpha}\sigma_{\beta}},$$

this condition is

$$\rho > -\frac{\sigma_{\beta}}{\sigma_{\alpha}}$$

Random Intercept and Slope Model

- For arbitrary time point j , the general relationships are:

$$j > -\frac{\sigma_{\alpha\beta}}{\sigma_{\beta}^2} \quad \text{variances increase after time } j$$

$$j < -\frac{\sigma_{\alpha\beta}}{\sigma_{\beta}^2} \quad \text{variances decrease up to time } j$$

where $-\frac{\sigma_{\alpha\beta}}{\sigma_{\beta}^2}$ can be replaced by $-\rho \frac{\sigma_{\alpha}}{\sigma_{\beta}}$

- Only if $\sigma_{\alpha\beta} = -0.5(2j + 1)\sigma_{\beta}^2$ are the j th and $(j + 1)$ st variances equal
(in this case, all subsequent variances increase over time and all previous variances decrease over time)
- These consequences of the random intercept and slope model do not appear to be widely known

Random Intercept and Slope Model

- The features observed in this example appear to occur commonly:
 - $\sigma_{\alpha\beta}$ is negative
 - σ_{β}^2 is close to zero

- The resulting estimated covariance matrix of $y_{hi} = (y_{hi1}, \dots, y_{hi4})'$ is:

$$\begin{pmatrix} 4.6216 & 2.8891 & 2.8727 & 2.8563 \\ 2.8891 & 4.6839 & 3.0464 & 3.1251 \\ 2.8727 & 3.0464 & 4.9363 & 3.3938 \\ 2.8563 & 3.1251 & 3.3938 & 5.3787 \end{pmatrix}$$

- In comparison, the pooled covariance matrix of $y_{hi} = (y_{hi1}, \dots, y_{hi4})'$ is estimated as:

$$\begin{pmatrix} 5.4155 & 2.7168 & 3.9102 & 2.7102 \\ 2.7168 & 4.1848 & 2.9272 & 3.3172 \\ 3.9102 & 2.9272 & 6.4557 & 4.1307 \\ 2.7102 & 3.3172 & 4.1307 & 4.9857 \end{pmatrix}$$

- While intuitively appealing, the random intercept and slope model is not the best fitting one

Model 7

- Random intercept model $y_{hij} = a_{hi} + \beta_h t_j + \epsilon_{hij}$
 - a_{hi} are independent $N(\alpha_h, \sigma_a^2)$
 - ϵ_{hij} are independent $N(0, \sigma^2)$

```
proc mixed method=ml;
class id sex;
model distance=sex age*sex / noint s;
random sex / g subject=id type=simple;
repeated / type=simple subject=id r;
```

- The `type=simple` option specifies the simple covariance structure $\Sigma = \sigma^2 I$

- LR test comparing Model 7 to Model 6 is

$$428.64 - 427.81 = 0.83, \quad \text{df} = 8 - 6 = 2, \quad p = .66$$

- LR test comparing Model 7 to Model 2 is

$$428.64 - 419.48 = 9.16, \quad \text{df} = 14 - 6 = 8, \quad p = .33$$

Comments on Model 7

- The repeated statement is not required
(output not displayed)

```
proc mixed method=ml;  
class id sex;  
model distance=sex age*sex  
      / noint s;  
random sex  
      / g subject=id type=simple;
```

- The intercept keyword could instead be
used in the random statement

```
proc mixed method=ml;  
class id sex;  
model distance=sex age*sex  
      / noint s;  
random intercept  
      / g subject=id type=simple;
```

- This changes only the covariance parameter
portion of the output

Model 7 (Alternate Formulation)

- The random intercept model is equivalent to the model

$$y_{hij} = \alpha_h + \beta_h t_j + \epsilon_{hij},$$

where $\epsilon_{hi} = (\epsilon_{hi1}, \dots, \epsilon_{hi4})'$ are independent $N_4(0, \Sigma)$ and

$$\Sigma = \begin{pmatrix} \sigma^2 & \rho\sigma^2 & \rho\sigma^2 & \rho\sigma^2 \\ \rho\sigma^2 & \sigma^2 & \rho\sigma^2 & \rho\sigma^2 \\ \rho\sigma^2 & \rho\sigma^2 & \sigma^2 & \rho\sigma^2 \\ \rho\sigma^2 & \rho\sigma^2 & \rho\sigma^2 & \sigma^2 \end{pmatrix}$$

has compound symmetry structure

```
proc mixed method=ml;
class id sex;
model distance=sex age*sex / noint s;
repeated / type=cs subject=id r;
```

- The `type=cs` option specifies compound symmetry

Model 8

- The only simpler model is $y_{hij} = \alpha_h + \beta_h t_j + \epsilon_{hij}$, where ϵ_{hij} are independent $N(0, \sigma^2)$

```
proc mixed method=ml;
class id sex;
model distance=sex age*sex / noint s;
repeated / type=simple subject=id r;
```

- LR test comparing Model 8 to Model 7 is

$$478.24 - 428.64 = 49.6, \quad df = 6 - 5 = 1, \quad p < .001$$

- LR test comparing Model 8 to Model 2 is

$$478.24 - 419.48 = 58.8, \quad df = 14 - 5 = 9, \quad p < .001$$

- While Models 2, 4, 6, and 7 all fit the data, compound symmetry (Model 7) requires only six parameters and is most parsimonious

Choosing the “Best” Covariance Structure

- The likelihood ratio (LR) test can be used when:
 - two models are fit to the same data by ML or REML
 - one model is a constrained version of the other
- Potential problems of using the LR test to compare covariance models:
 - you may be dealing with parameters on the boundary of the parameter space
 - the models being compared may not be nested
- Two other model selection criteria are:
 - Akaike’s (1973) Information Criterion (AIC)
 - Schwarz’s (1978) Bayesian Information Criterion (BIC)
- Both penalize the log-likelihood for the number of parameters and/or number of observations

Akaike's Information Criterion (AIC)

- As implemented in PROC MIXED, the AIC can be used to compare models with the same fixed effects, but different covariance structures
- $AIC = \ell - q$
 - ℓ is the log likelihood
 - q is the effective number of covariance parameters (those not estimated to be on a boundary constraint)
- The model with the largest AIC is deemed best
- Most statistical references give the formula

$$AIC = -2\ell + 2p,$$

where p is the number of model parameters

Reference

Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. *Second International Symposium on Information Theory*. B.N. Petrov and F. Csaki (eds), 267–281. Budapest: Akademiai Kiado.

Bayesian Information Criterion (BIC)

- In PROC MIXED, $BIC = \ell - (q/2) \log(n^*)$
 - ℓ is the log likelihood and q is the effective number of covariance parameters
 - In version 6, $n^* = n$ for ML and $n^* = n - p$ for REML, where p is the rank of X
 - In version 7, n^* is the number of “effective subjects” as displayed in the dimensions table
- BIC has an increased penalty for overfitting compared to AIC (and the two criteria may not agree as to which covariance model is best)
- Most statistical references give the formula

$$BIC = -2\ell + p \log(n),$$

with p (n) the no. of parameters (observations)

Reference

Schwarz, G. (1978). Estimating the dimension of a model. *Ann. Statist.*, 6, 461–464.

Comparison of Models using AIC and BIC

- Models 2 and 4–8 all fit a linear relationship between distance and age, with separate lines for boys and girls
- The alternative covariance structures can be compared using AIC and BIC

Covariance Model	AIC	BIC
Unstructured	−219.739	−233.149
Banded	−216.322	−221.686
AR-1	−222.341	−225.023
Random coefficients	−217.903	−223.267
Random intercept	−216.320	−219.002
Simple	−240.121	−241.462

- The “best” model by both criteria is the random intercept model
(compound symmetry covariance structure)


Example

- A randomized, double-blind, parallel group, placebo-controlled study in patients with postoperative pain
- Two active treatments and a placebo group

Group	Description	No. of Patients
A	Treatment 1	41
B	Treatment 2	41
C	Placebo	40

- Patients received a single dose of their assigned treatment when they reported moderate to severe postoperative pain
- The primary outcome variable was the amount of rescue medication used (recorded at hourly intervals for 24 hours after dosing)

Mean Rescue Use, by Treatment Group




Protocol-Specified Primary Efficacy Analysis

“Compare the three groups using a repeated measures ANOVA for the amount of rescue medication taken over the 24-hour interval, using time as the repeated factor”

Comments:

- The protocol specified that the specific comparisons of interest were A versus placebo, B versus placebo, and A versus B
- A secondary analysis to be carried out after combining the data over prespecified time intervals was also specified
0–3 hours, 4–6 hours, 7–12 hours, 13–24 hours
- The choice of covariance structure was not specified in the protocol


Mean Rescue Use, by Treatment Group (Collapsed Hourly Measurements)




The Dilemma

- With 24 repeated measurements per subject, one is limited in the types of covariance structures that can be considered
- It is not possible to compare the adequacy of a reduced covariance model relative to the “full” (unstructured) model
a 24×24 covariance matrix, with 300 parameters
- The company stated that they wrote an analysis plan prior to breaking the blind, which specified the use of an AR-1 covariance structure
- The FDA considered several alternative covariance models
- The analysis results differed depending on the covariance model selected

Variances and Covariances, by Hour



Variances and Covariances, by Lag



Potential Covariance Models

The covariance matrix of $y_i = (y_{i1}, \dots, y_{i,24})'$ is

$$V_i = Z_i B Z_i' + W_i$$

- Sphericity (compound symmetry as special case)

$B = 0$ and $W_i = W$ has 2 parameters (diagonal elements are σ^2 , off-diagonal elements are $\rho\sigma^2$)

- AR-1 (2 parameters)

$$V_{ij} = \sigma^2 \rho^{|i-j|}$$

- Random intercept + AR-1 (3 parameters)

$$V_{ij} = \sigma^2 \rho^{|i-j|} + \sigma_e^2$$

- Random intercept and slope (4 parameters)


$$B = \begin{pmatrix} \sigma_\alpha^2 & \sigma_{\alpha\beta} \\ \sigma_{\alpha\beta} & \sigma_\beta^2 \end{pmatrix}, \quad W_i = \text{diag}(\sigma^2)$$

Analysis Results


- Preliminary models showed that the:
 - the effect of time was nonlinear
 - profiles for the treatment groups were parallel
- Thus, the final model included main effects for treatment group, linear time, and quadratic time

	Covariance Model		
	AR-1	AR-1 + Rand. Int.	Rand. Int. and Slope
Treatment effect (<i>p</i> -value)	0.010	0.189	0.198
A – Placebo			
Estimate	−0.329	−0.339	−0.304
S.E.	0.138	0.243	0.229
<i>p</i> -value	0.018	0.163	0.185
B – Placebo			
Estimate	−0.399	−0.419	−0.395
S.E.	0.138	0.243	0.229
<i>p</i> -value	0.004	0.085	0.085


Variances and Covariances, by Lag



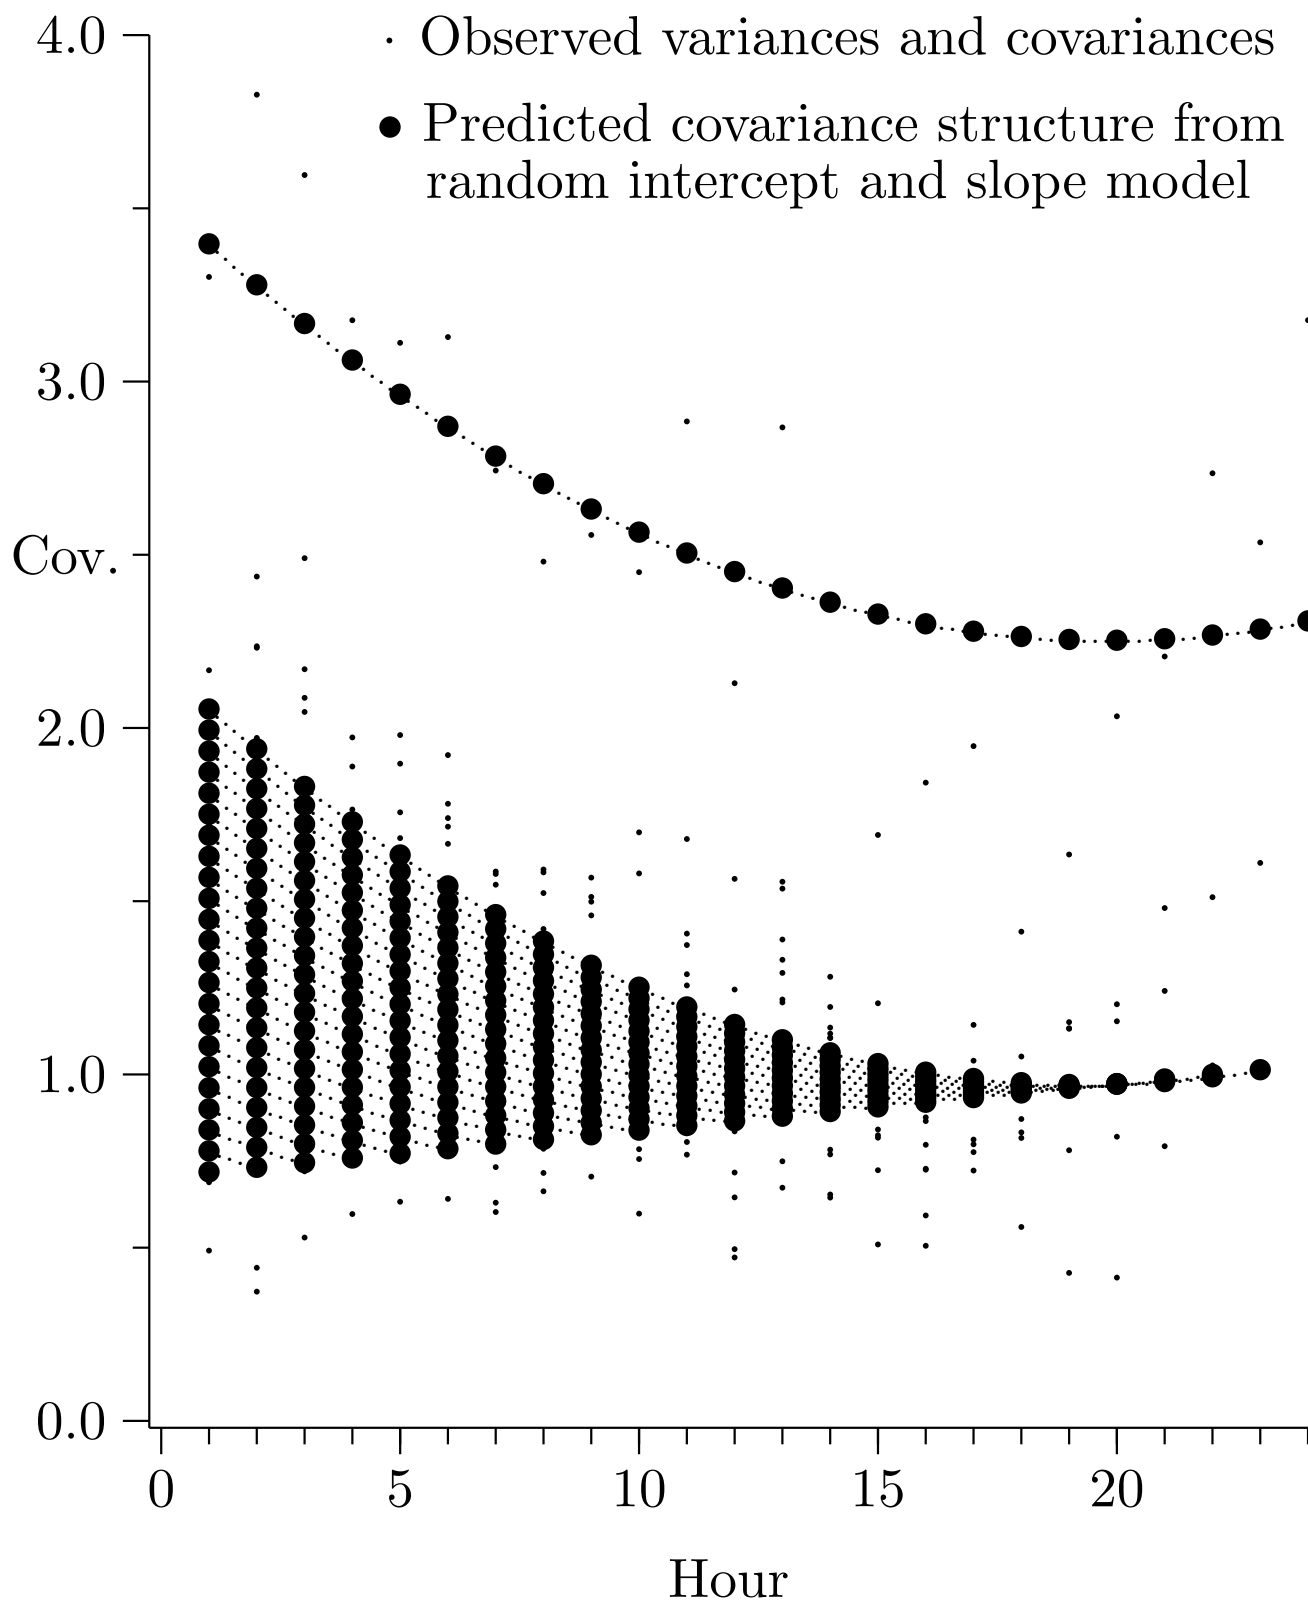
Variances and Covariances, by Lag



Variances and Covariances, by Lag



Variances and Covariances, by Hour



Covariance Parameters from the Random Intercept and Slope Model

- $B = \begin{pmatrix} \sigma_\alpha^2 & \sigma_{\alpha\beta} \\ \sigma_{\alpha\beta} & \sigma_\beta^2 \end{pmatrix}, \quad W_i = \text{diag}(\sigma^2)$

- The parameter estimates are

$$\hat{B} = \begin{pmatrix} 2.23581 & -0.06396 \\ -0.06396 & 0.00322 \end{pmatrix}, \quad \hat{\sigma}^2 = 1.282$$

- Since

$$-\frac{\hat{\sigma}_{\alpha\beta}}{\hat{\sigma}_\beta^2} = \frac{0.06396}{0.00322} = 19.84,$$

the variances decrease from hour 1 to hour 20,
and then increase to hour 24

- In general, this type of variance assumption is not sensible for repeated measurements
- It may, however, be reasonable in this example

Example

- Jones & Boadi-Boateng (1991) analyze data from a longitudinal study of 619 patients in 4 groups:

Group	Sample Size	Kidney Disease	Hypertensive
1	294	Yes	Yes
2	193	Yes	No
3	73	No	Yes
4	149	No	No

- The response variable is the reciprocal of serum creatinine; this variable ranges from 0.028 to 2.5
- The explanatory variables are group and patient age (which ranges from 18–84 years)
- The number of observations per patient ranges from 1 to 22

Example

- The spacing between observations ranges from 0.002 years (one day) to 22 years
- The median spacing between adjacent observations on the same patient is 1 year
- The spacings among adjacent observations differ among the four groups

The median spacings in groups 1–4 are 0.8, 2, 4, and 3 years, respectively

- Jones and Boadi-Boateng fit a structural model with separate lines for each of the four groups

Reference

Jones, R. H. and Boadi-Boateng, F. (1991). Unequally spaced longitudinal data with AR(1) serial correlation. *Biometrics* **47**, 161–175.

Creation of SAS Data Set

- The treatment group effect will be parameterized using three indicator variables:
group2, group3, group4

```

data scr;
input id group age scr;
agesq=age*age;
group2=(group=2);
group3=(group=3);
group4=(group=4);
cards;
    1 1 35.765    .182
    1 1 37.990    .088
    2 2 24.997  1.429
    2 2 27.441  1.111
    2 2 30.524  1.429
    ...
    616 1 32.676  1.429
    616 1 33.295  1.429
    616 1 33.867  1.429
    616 1 34.374  1.667
    617 1 61.331   .476
    618 1 27.477  1.111
    619 2 25.555  1.250
    ;

```

Analysis Strategy

- As in Jones & Boadi-Boateng (1991), we'll model SCR^{-1} (`scr`) as a function of group and age
- We'll examine alternative covariance structures and consider incorporation of nonlinear age effects
- Models of the form

```
scr=age group2 group2*age group3
      group3*age group4 group4*age;
```

will be used to fit an intercept and slope for group 1, and incremental intercept and slope parameters for groups 2–4

- Similarly, quadratic models of the form:

```
scr=age agesq group2 group2*age group2*agesq
      group3 group3*age group3*agesq group4
      group4*age group4*agesq;
```

fit a quadratic polynomial for group 1 and incremental intercept, slope, and quadratic parameters for groups 2, 3, and 4

“Naive” Approach

- If we ignore the fact that observations obtained from the same subject are correlated, standard linear model programs can be used
- The following PROC GLM statements fit the quadratic model and test the significance of the nonlinear terms:

```
proc glm;
model scr=age agesq group2 group2*age
      group2*agesq group3 group3*age
      group3*agesq group4 group4*age
      group4*agesq;
contrast 'Quadratic Effects' agesq 1,
      group2*agesq 1, group3*agesq 1,
      group4*agesq 1;
```

- The test of nonlinearity is highly significant ($F_{4,1573} = 10.25, p < 0.001$)
- Still, we'll fit the linear model with a separate intercept and slope for each group

“Naive” Approach

- The SAS statements are:

```
proc glm;
model scr=age group2 group2*age group3
      group3*age group4 group4*age;
```

- The parameter estimates are:

Parameter	Estimate	t	p-value	s.e.
INTERCEPT	1.2138	31.44	0.0001	0.0386
AGE	-0.0136	-15.91	0.0001	0.0008
GROUP2	0.3213	3.86	0.0001	0.0831
AGE*GROUP2	-0.0024	-1.11	0.2668	0.0022
GROUP3	-0.2012	-2.06	0.0397	0.0977
AGE*GROUP3	0.0137	6.54	0.0001	0.0021
GROUP4	-0.0469	-0.67	0.5002	0.0696
AGE*GROUP4	0.0124	7.33	0.0001	0.0017

- In groups 1, 2 & 4, scr decreases as age increases (the rate of decrease is greatest in group 2, while the slope is nearly zero in group 4)
- The slope in group 3 is slightly greater than zero

“Naive” Approach using PROC MIXED

- The MIXED procedure can also be used to fit this model; the statements are:

```
proc mixed noclprint;  
class id;  
model scr=age group2 group2*age group3  
        group3*age group4 group4*age / s;  
repeated / type=simple subject=id;
```

- The `noclprint` option suppresses the printing of the levels of the class variable(s)
(it is not necessary to see the list of the 619 subject identifiers)
- The `s` option of the `model` statement requests that the solution for the fixed effects be printed
- The `type=simple` option specifies the within-subject covariance structure of independent errors with common variance

“Naive” Approach using PROC MIXED

- By default, PROC MIXED gives REML estimates
- The estimate of the error variance is labelled `DIAG` and is equal to 0.10693227
(identical to the error mean square from GLM)
- The solutions for the fixed effects are also identical to those obtained from PROC GLM
- The REML $-2 \log$ likelihood is 1015.678
- If we fit the model with linear and quadratic age effects, the REML $-2 \log$ likelihood *increases* to 1042.203
- Using REML estimation, it is not possible to compare the fit of two models using a likelihood ratio test

“Naive” Approach using PROC MIXED

Maximum Likelihood Estimation

- The same models can also be fit by maximum likelihood using the procedure statement:

```
proc mixed noclprint method=ml;
```

- This gives identical estimates of the fixed effects
- The estimate of the error variance from the linear model (0.1063925) and the standard errors of the fixed effects are slightly smaller than the REML estimates
- The likelihood ratio test statistic comparing the linear and quadratic models is

$$946.6528 - 905.8743 = 40.78$$

with 4 df, which is highly significant ($p < 0.001$)

Comments on Unstructured Correlation Model

- The most general type of within-subject covariance structure is the unspecified one:

```
proc mixed noclprint;
class id;
model scr=age group2 group2*age group3
      group3*age group4 group4*age / s;
repeated / type=un subject=id;
```
- Not useful with unequally spaced observations
- Since there are up to 22 observations/subject, this fits a 22×22 within-subject covariance matrix
- Thus, the distance between adjacent observations is not taken into account
- Even if a 22×22 covariance matrix was reasonable, there are computational difficulties:

```
REML Estimation Iteration History
Iteration Evaluations Objective Criterion
0          1          -1882.654336
Unable to make hessian positive definite
```

Jones and Boadi-Boateng Model

- Since the number & spacing of the repeated measurements vary considerably across subjects, a parsimonious covariance model is required
- JBB fit a model with:
 - Separate lines for each of the 4 groups
 - Subject-specific random effects for the intercept and slope
 - Within-subject covariance structure is a continuous time AR(1) process with observational error
 - Eight fixed effects and six random effects
- Three of the random effects parameters are from the 2×2 covariance matrix of the random intercept and slope

Jones and Boadi-Boateng Model

- The other three random effects are within-subject parameters (σ^2 , ϕ , σ_0^2)
 - Correlation function is $\rho(\tau) = e^{-\phi|\tau|}$, where τ is the time between two observations
 - Additional observational error component σ_0^2 allows observations very close together in time to be different

- The SAS statements are:

```
proc mixed noclprint;
class id;
model scr=age group2 age*group2 group3
      age*group3 group4 age*group4 / s;
random intercept age / type=un
                        subject=id g;
repeated / type=sp(exp)(age) local
          subject=id;
```

- The `local` option specifies the observational error component

Results of JBB Analysis

- 2×2 random effects covariance matrix is singular
- The estimated fixed effects are:


Group	Parameter	Estimate	<i>p</i> -value
1	intercept	1.406	< 0.001
	slope	-0.018	< 0.001
2	int. increment	0.084	0.417
	slope increment	0.003	0.264
3	int. increment	-0.359	0.002
	slope increment	0.018	< 0.001
4	int. increment	-0.178	0.039
	slope increment	0.015	< 0.001

- In all four groups, the SCR reciprocal is estimated to decrease as age increases
 - Rate of decrease is greatest in group 1 and is nearly zero in group 3
- Groups 1 and 2 are similar
- Intercept is greatest in group 2, followed by groups 1, 4, and 3


Effect of Choice of Covariance Structure

- Used the same structural model
(separate intercepts and slopes in each group)
- Considered nine “reasonable” covariance models:
 1. AR(1) with observational error
 2. AR(1)
 3. Compound symmetry
 4. Random intercept and slope + AR(1) with observational error
 5. Random intercept and slope + AR(1)
 6. Random intercept and slope + independent within-subject errors
 7. Random intercept + AR(1) with observational error
 8. Random intercept + AR(1)
 9. Random intercept + independent within-subject errors

Parameter Estimates from the Nine Models



Tests of Significance from the Nine Models



Properties of Tests

- The properties of three types of test statistics were compared:

UM	Unstructured multivariate approach using Wilks' Λ
MM(ML)	Mixed model approach using ML estimation
MM(REML)	Mixed model approach using REML estimation

- The case of two groups and four time points was considered
- The model for the mean in group h at time j was $\mu_{hj} = \beta_{h0} + \beta_{h1}j$, where


$$\beta_{20} = \beta_{10} + \delta_0, \quad \beta_{21} = \beta_{11} + \delta_1$$

- Two hypotheses were of interest:
 - equality of groups $H_0: \delta_0 = \delta_1 = 0$
 - parallelism $H_0: \delta_1 = 0$

Properties of Tests

- Data were generated from 5 correlation models:
 - AR-1 with $\rho = 0.3, 0.7$
 - compound symmetry with $\rho = 0.3, 0.7$
 - unstructured
- The mixed model analyses were carried out using four assumed correlation structures:
 - independence
 - AR-1
 - compound symmetry
 - unstructured
- Sample sizes of 15 and 25 observations/group were considered
- Several values of (δ_0, δ_1) were studied
- 1000 replications were carried out for each combination of factors

Test Size versus Sample Size



Unstructured multivariate approach: ●


Linear mixed model approach:

Assumed Correlation Structure	ML	REML
Independence	□	+
AR-1	○	×
Compound Symmetry	▷	*
Unstructured	◇	★

Test of Equality of Group Effects

- The unstructured multivariate approach tends to have test sizes closest to the nominal 5% level
- The mixed model approach tends to yield anti-conservative tests
- For a given assumed correlation structure, sizes of REML tests are smaller (closer to the nominal level) than those of ML tests
- The unstructured multivariate approach is robust to the true correlation structure
- The performance of the mixed model tests depends highly on the structures of the true and assumed correlation models
- The UM approach is preferred, even in a four-parameter model with a 4×4 covariance matrix and a sample size of only 30

Test Size versus Sample Size



Unstructured multivariate approach: ●


Linear mixed model approach:

Assumed Correlation Structure	ML	REML
Independence	□	+
AR-1	○	×
Compound Symmetry	▷	*
Unstructured	◇	★

Test of Parallelism

- The results for testing parallelism are somewhat different from those for testing group effects
- The test for parallelism is based on the differences among adjacent responses, which might be less sensitive to the correlation structure than the responses themselves
- The unstructured multivariate approach tends to have test sizes closest to the nominal 5% level
- The mixed model approach tends to yield anti-conservative tests
- For a given assumed correlation structure, sizes of REML tests are smaller (closer to the nominal level) than those of ML tests
- The mixed model ML and REML tests assuming independence are conservative

Power (test of parallelism) vs. Trt. Difference




Unstructured multivariate approach: ●

Linear mixed model approach:

Assumed Correlation Structure	ML	REML
Independence	□	+
AR-1	○	×
Compound Symmetry	▷	*
Unstructured	◇	★

Power (test of parallelism) vs. Trt. Difference



Unstructured multivariate approach: ●

Linear mixed model approach:

Assumed Correlation Structure	ML	REML
Independence	□	+
AR-1	○	×
Compound Symmetry	▷	*
Unstructured	◇	★