UCLA

Department of Statistics Papers

Title

A Stochastic Model for Analysis of Longitudinal AIDS Data

Permalink

<https://escholarship.org/uc/item/3m91542r>

Authors

Jeremy M. G. Taylor

W.G. Cumberland

J.P. Sy

Publication Date

2011-10-25

A Stochastic Model for Analysis of Longitudinal AIDS Data

J. M. G. Taylor; W. G. Cumberland; J. P. Sy

Journal of the American Statistical Association, Vol. 89, No. 427. (Sep., 1994), pp. 727-736.

Stable URL:

<http://links.jstor.org/sici?sici=0162-1459%28199409%2989%3A427%3C727%3AASMFAO%3E2.0.CO%3B2-1>

Journal of the American Statistical Association is currently published by American Statistical Association.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/astata.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

The JSTOR Archive is a trusted digital repository providing for long-term preservation and access to leading academic journals and scholarly literature from around the world. The Archive is supported by libraries, scholarly societies, publishers, and foundations. It is an initiative of JSTOR, a not-for-profit organization with a mission to help the scholarly community take advantage of advances in technology. For more information regarding JSTOR, please contact support@jstor.org.

A Stochastic Model for Analysis of Longitudinal AIDS Data

J. M. G. TAYLOR, W. G. CUMBERLAND, and J. P. SY*

In this paper we analyze serial CD4 T-cell measurements from the Los Angeles portion of the Multicenter AIDS Cohort Study. Our emphasis is on developing a plausible and parsimonious model to describe the stochastic process underlying the patterns of CD4 measurements. The stochastic process that we use enables us to investigate the concept of derivative tracking, for which it is assumed that the rank order of the individual's slopes is maintained over time. A general model for the analysis of longitudinal repeated measures data is

$$Y_i(t_{ij}) = X(t_{ij})\alpha + Z(t_{ij})b_i + W_i(t_{ij}) + \varepsilon_{ij}$$

where $Y_i(t_{ij})$ is the measurement of subject i at time t_{ij} , $X(t_{ij})\alpha$ represents fixed effect terms, $Z(t_{ij})b_i$ represents random effect terms, $W_i(t_{ij})$ is a stochastic process allowing correlation between measurements, and ε_{ij} is measurement error. In the simplest case, $X(t_{ij})$ and $Z(t_{ij})$ contain the times of measurements. For $W_i(t_{ij})$, we use a two-parameter integrated Ornstein–Uhlenbeck (OU) process. The OU process is the continuous mean zero Gaussian Markov process, which includes Brownian motion and white noise as special limiting cases. This model is a continuous-time version of an AR(1) process for the deviations of the derivative of y from the expected derivative of y with respect to t . This approach is flexible and tractable as the covariance structure has a closed-form expression. The model allows unequally spaced observations and can be generalized to multivariate responses. This model enables one to assess whether individuals maintain their trajectories; that is, whether their slope of Y tracks. We find no evidence in the data that the slopes of the CD4 values track.

KEY WORDS: AIDS; Ornstein–Uhlenbeck process; Repeated measures; Tracking.

1. INTRODUCTION

The concept of tracking for longitudinal data is popular with epidemiologists and has been discussed in the statistical literature. Foulkes and Davis (1981) and McMahan (1981) developed indices of tracking based on the concept of the maintenance over time of the relative or rank position of the response variable among the group of subjects. Ware and Wu (1981) regarded tracking as the ability to predict future observations for each individual, leading to a definition that a population tracks if for each individual, the expected values of serial measurements are given by a polynomial function of time. These two concepts are closely related, but not identical. In this article we adopt the first concept of tracking and apply it to the first derivative of the response variable. Thus we regard a population as having derivative tracking if the relative ranks of the rate of change of the outcome variable are maintained over time.

The typical structure of longitudinal data is numerous measurements of a possibly multivariate response variable on each subject. The measurements will usually be at unequally spaced time intervals, and the number of measurements per subject may differ between subjects. There could also be covariates, possibly time varying, that influence the response variable. The aim in the analysis of such data is to understand the changes in the mean structure of the response variable with time, to understand the effect of the covariates on the response variable, and to understand the within-subject correlation structure.

There are a number of approaches to the analysis of such data. One uses a mixed or random effects model (Laird and Ware 1982). Let $Y_i(t_{ij})$ be the response of person i at time t_{ij} ; then in a random effects model, it is assumed that

$$Y_i(t_{ij}) = X(t_{ij})\alpha + Z(t_{ij})b_i + \varepsilon_{ij},$$

where $X(t_{ij})\alpha$ represents fixed effects, $Z(t_{ij})b_i$ represents random effects, and ε_{ij} is an independent, identically distributed (iid) measurement error.

A second approach is to specify a structure for the within-subject covariance based possibly on a stochastic process. For example, it might be assumed that $Y_i(t_{ij}) = X(t_{ij})\alpha + W_i(t_{ij})$, where $W_i(t_{ij})$ is a realization of a stochastic process.

Other authors have suggested a combination of these two approaches (Chi and Reinsel 1989; Cullis and McGilchrist 1990; Diggle 1988; Jones and Bodi-Boateng 1991). The computational aspects of fitting these models to data are nontrivial and have been discussed by Jennrich and Schluchter (1986) and Lindstrom and Bates (1988).

In this article we focus on a model consisting of a combination of fixed effects, random effects, a specific stochastic process, and measurement error. In particular, we assume that

$$Y_i(t_{ij}) = X(t_{ij})\alpha + Z(t_{ij})b_i + W_i(t_{ij}) + \varepsilon_{ij}.$$

Our emphasis will be on interpreting the stochastic process with particular regard to the concept of derivative tracking. We propose using an integrated Ornstein–Uhlenbeck (IOU) process for the stochastic process. The motivation for this choice of stochastic process arose from an AIDS study of the natural history of CD4 T-cell numbers. CD4 numbers are a critical aspect of the immune system, with low values indicating more severe immune deficiency. An immunologic theory is that the rate of change of CD4 cell numbers “tracks”; that is, an immunologically weak individual who has an initial fast rate of decline of CD4 relative to other HIV-infected people will persist with a more rapid rate of

* J. M. G. Taylor is Associate Professor, W. G. Cumberland is Professor, and J. P. Sy is a graduate student, Department of Biostatistics, University of California, Los Angeles, CA 90024-1772. This work was partially supported by National Institutes of Health Grants AI29196 and AI72631.

decline of CD4 than will the others. Similarly, an immunologically strong individual with an initial slow rate of decline of CD4 will have a slower rate of decline than will others for a long period. An alternative hypothesis is that the slopes of an individual's decline in CD4 are constantly changing with time. Our approach in this article is to use a stochastic process in which both of the foregoing hypotheses are special cases. That is, a flexible family of parametric models for the covariance structure is used, and we let the data choose which of the specific submodels is most appropriate. In this sense, the data are determining the degree of tracking of the derivative rather than fixing the tracking by a particular assumption about the covariance structure.

A popular model for analyzing longitudinal data is one in which X and Z are the times of the measurements and there are random effects terms but no stochastic process. Such a model is a strong derivative tracking model. The short-term behavior of such linear random effects models seems reasonable for CD4 T-cell numbers. But long-term behavior of such random effects assumptions in which individuals maintain the same linear path seems unlikely for something as complex as the measurement of a person's immune system. The model we propose has the same short-term behavior as a linear random effects model, but its long-term behavior is less restrictive; in this sense we believe the model is more biologically plausible for CD4 T-cell numbers.

A number of authors have undertaken modeling of serial CD4 measurements (Berman 1990; DeGruttola, Lange, and Dafni 1991; Kanazawa 1990; Lange, Carlin, and Gelfand 1992; Longini et al. 1989; Munõz et al. 1988). DeGruttola, Lange, and Dafni (1991) applied a linear random effects plus measurement error model to a San Francisco cohort with 2-year follow-up and five serial measurements. They imputed the unknown date of infection from a specified distribution and analyzed CD4 numbers on a square root scale. They assumed that the population mean declines linearly with time and that each individual's derivatives "track" over his entire HIV-infected lifespan. Our results in Section 3 suggest that there is very little correlation between the rate of change of CD4 at two time points that are more than about 6 months apart. Other more ad hoc methods, which just consider the correlation between pairs of nonoverlapping differences in CD4, also show no evidence of any tracking of the slopes (Taylor, Tan, Detels, and Giorgi 1991). Lange et al. (1992) extended the analysis in DeGruttola's article to an updated version of the same data using a similar random effects model and performed a Bayesian analysis incorporating prior information.

Berman (1990) assumed that the CD4 decline can be explained on the log scale by an unknown negative drift and a stationary Gaussian process. Kanazawa (1990) assumed that the CD4 decline can be explained by the sum of an unknown negative drift and Brownian motion. These authors were concerned mainly with estimating the unknown date of infection. Although these two authors' aims are different from ours, the models they chose are similar to the best-fitting one to our data, although it should be noted that our model also includes the important measurement error term and a random intercept.

Section 2 of this article describes the IOU process. Section 3 describes the analysis of CD4 numbers in three AIDS-related data sets; and Section 4 contains a simulation study designed to investigate some of the issues raised in the data analysis. Section 5 contains a discussion, including a multivariate extension of the stochastic process.

2. INTEGRATED ORNSTEIN-UHLENBECK STOCHASTIC PROCESS

As a motivation for our model, consider the following discrete time model in which the derivatives are allowed to "track" for short periods. Let $\beta_{t,i}$ be the slope for person i at time t and let $E(\beta_{t,i})$ be the expected slope for that person; then an AR(1) process for the derivatives is

$$\beta_{t+\Delta t,i} - E(\beta_{t+\Delta t,i}) = e^{-\alpha\Delta t}(\beta_{t,i} - E(\beta_{t,i})) + e,$$

where $\text{var}(e) = \sigma^2\Delta t$.

In this model there is an elastic pull back toward each person's mean slope. The strength of the pull depends on the magnitude of α . The term e represents perturbation of this elastic pull, with the magnitude of the perturbation determined by σ^2 . A model in which the derivatives "track" for long periods can be obtained with a small α and σ^2 ; alternatively, a model in which both α and σ^2 are large has the current derivative essentially unrelated to the earlier derivative—that is, the derivatives do not "track."

The limit of the aforementioned autoregressive process for the derivatives as $\Delta t \rightarrow 0$ is an OU process (Cumberland and Sykes 1982). We are primarily interested in the process itself, not its derivative. If the derivative is an OU process, then the process itself is an IOU process. Let $Y(t)$ be a mean 0 OU process; then $Y(t)$ is Gaussian and Markov, and for large t , $Y(t)$ is stationary with

$$\text{cov}(Y(s), Y(t)) = \frac{\sigma^2}{2\alpha} e^{-\alpha|t-s|}. \quad (1)$$

Let $W(t) = \int_0^t Y(u) du$ be an IOU process; then $W(t)$ is Gaussian and

$$\begin{aligned} \text{cov}(W(s), W(t)) \\ = \frac{\sigma^2}{2\alpha^3} [2\alpha \min(s, t) + e^{-\alpha t} + e^{-\alpha s} - 1 - e^{-\alpha|t-s|}]; \end{aligned} \quad (2)$$

hence $\text{var}(W(t)) = (\sigma^2/\alpha^3)[\alpha t + e^{-\alpha t} - 1]$. It is interesting to note that scaled Brownian motion is a special case of $W(t)$ in which α is infinitely large and σ^2/α^2 is constant.

The model we propose for the observed response $Y_i(t_{ij})$ at time t_{ij} is

$$Y_i(t_{ij}) = X(t_{ij})\alpha + Z(t_{ij})b_i + W_i(t_{ij}) + \varepsilon_{ij}, \quad (3)$$

where α are fixed effects coefficients, b_i are random effects coefficients, $W_i(t_{ij})$ is an IOU process, and ε_{ij} is measurement error.

To be more specific, consider the situation where the only coefficients in X and Z are the intercept, t , and t^2 . Then (3) can be rewritten (with a slight change in notation) as

$$Y_{it} = a_i + b_it + c_it^2 + W_i(t) + \varepsilon_{it}, \quad (4)$$

where

$$\begin{bmatrix} a_i \\ b_i \\ c_i \end{bmatrix} \sim N \left[\begin{bmatrix} a \\ b \\ c \end{bmatrix}, \begin{bmatrix} \sigma_a^2 & \sigma_{ab} & \sigma_{ac} \\ \sigma_{ab} & \sigma_b^2 & \sigma_{bc} \\ \sigma_{ac} & \sigma_{bc} & \sigma_c^2 \end{bmatrix} \right],$$

$\varepsilon_{it} \sim N(0, \sigma_\varepsilon^2)$ independent.

A disadvantage of the IOU process is that it is not stationary; hence it is necessary to have a natural time 0 for each person. This nonstationary property can be seen from the form of $\text{cov}(W(s), W(t))$ in Equation (2), which depends on s and t and not just on their difference. In some applications, it may be that there is no natural time 0, or that time 0 is not exactly known. Our method of overcoming this problem with $\text{cov}(W(s), W(t))$ is to analyze differences. Let Y_{iE_i} be the first measurement on person i at time E_i . Let $D_{it} = Y_{it} - Y_{iE_i}$, for $t > E_i$. Then

$$D_{it} = b_i(t - E_i) + c_i(t^2 - E_i^2) + W_{it} - W_{iE_i} + \varepsilon_{it} - \varepsilon_{iE_i} \quad (5)$$

and

$$\text{cov}(D_{it_1}, D_{it_2}) = A + B + C,$$

where

$$A = (t_1 - E_i)(t_2 - E_i) \text{var } b_i + (t_1^2 - E_i^2)(t_2^2 - E_i^2) \text{var } c_i \\ + \{(t_1 - E_i)(t_2^2 - E_i^2) + (t_1^2 - E_i^2)(t_2 - E_i)\} \\ \times \text{cov}(b_i, c_i),$$

$$C = \sigma_\varepsilon^2(1 + I(t_1 = t_2)),$$

and

$$B = \text{cov}(W_{it_1} - W_{iE_i}, W_{it_2} - W_{iE_i}) \\ = \frac{\sigma^2}{2\alpha^3} (2\alpha(t_1 - E_i) - 1 - e^{-\alpha(t_2 - t_1)} + e^{-\alpha(t_2 - E_i)} \\ + e^{-\alpha(t_1 - E_i)}) \quad \text{for } t_1 \leq t_2.$$

Notice that B depends only on differences in times, so it avoids the need to define natural time zero. An alternative to analyzing the differences is to regard the starting time as a random effect, as has been done by others; for example, DeGruttola, Lange, and Dafni (1991).

A second advantage of analyzing differences is that it is no longer necessary to make any distributional assumptions concerning the intercept term a_i , nor is it necessary to specify the part of the correlation structure concerning a_i .

An appealing aspect of the IOU model is that the limiting case as $\alpha \rightarrow 0$ is a special case of a random effects model. In particular, it can be shown that the covariance structures of the two models,

$$Y_{it} = a_i + b_i t + \varepsilon_{it},$$

and

$$Y_{it} = a_i + bt + W_{it} + \varepsilon_{it},$$

are the same provided that $\sigma_{ab} = 0$ and $\sigma_b^2 = \sigma^2/2\alpha$.

Equations (3) and (5) specify the mean and covariance structure of the multivariate normal vector of observations.

Parameter estimates can be obtained by maximizing the likelihood, and standard errors can be obtained in the usual way from the second derivative of the log-likelihood (Jennrich and Schluchter 1986).

3. AIDS APPLICATION

We consider data from a cohort study of 1,637 homosexual and bisexual men recruited in Los Angeles in 1984 and 1985 and followed semi-annually until the present time. Details of the Multicenter AIDS Cohort study are presented elsewhere (Kaslow et al. 1987). The response variable was the CD4 T-cell number. All the observations were transformed by a fourth-root power to achieve homogeneity of within-subject variance (Taylor, Tan, Detels, and Giorgi 1991). This transformation also has the effect of making distributions of the observations appear more Gaussian, as is assumed in the models. Three separate data sets within the study were considered: the seroconverters cohort, the seroprevalent cohort, and the seronegative cohort.

3.1 Seroconverters

This cohort consisted of 87 people who were observed to change from HIV antibody negative to HIV antibody positive. The time interval between the last HIV-negative measurement and the first HIV-positive measurement was usually less than 9 months and always less than 15 months. The midpoint of the interval was taken to be the date of infection. All observations within 6 months following this date of infection were excluded, because the CD4 number is known to decrease abruptly on seroconversion (Lang et al. 1989; Margolick et al. 1993). The total number of CD4 measurements considered was 722. Each individual had between 1 and 18 measurements, spaced at a minimum of 3-month intervals. Observations after December 1989 were excluded to reduce the influence of AZT treatments on these data.

We fit a variety of models of the general form given by Equation (4),

$$Y_{it} = a_i + b_i t + c_i t^2 + W_i(t) + \varepsilon_{it},$$

in which $t > 0$ is the time from HIV infection.

Besides the general model, specific submodels were considered; in particular, (a) the quadratic term was excluded, (b) the IOU term was excluded, (c) the IOU term was excluded and the measurement error term was replaced by an AR(1) measurement error term with covariance structure given by Equation (1), and (d) the random effects terms b_i and c_i were assumed to have zero variance. The results of the models are described in Table 1. They indicate that a population quadratic term is not necessary. Furthermore, the models that best fit the covariance structure require either a random quadratic term or the IOU or Brownian motion term. In terms of a large likelihood and small number of parameters, the best model appears to be a random intercept, linear population decline, Brownian motion plus measurement error. The likelihood values clearly indicate that the simple random intercept plus random slope plus independent measurement error model (model 1) is not a good description of the data. This supports our contention that simple linear

Table 1. Parameter Estimates from Analysis of Seroconverter Data Set

Model*	Intercept $\hat{a}$	Linear $\hat{b}$	Quadratic $\hat{c}$	IOU process		ME	-Log-likelihood	No. of parameters
	SE($\hat{a}$)	SE($\hat{b}$)	SE($\hat{c}$)	$\hat{\alpha}$	$(\hat{\sigma}/\hat{\alpha})^2$	$\hat{\sigma}_e$		
1. RE(a, b) + ME	5.159 (.058)	-.228 (.029)				.264	255.4	6
2. RE(a, b) + FE(c) + ME	5.196 (.070)	-.268 (.051)	.009 (.009)			.264	255.0	7
3. RE(a, b, c) + ME	5.215 (.074)	-.289 (.070)	.014 (.016)			.248	238.6	10
4. RE(a) + FE(b) + IOU + ME	5.172 (.054)	-.221 (.026)		2.5	.13	.236	238.1	6
5. RE(a) + FE(b) + BM + ME	5.170 (.054)	-.220 (.025)		∞	.11	.212	238.3	5
6. RE(a, b) + IOU + ME	5.170	-.223		20.7 $\uparrow$	.09	.221	237.8	8
7. RE(a, b) + BM + ME	5.170 (.055)	-.233 (.027)		∞	.09	.217	237.8	7
8. RE(a, b) + OUME	5.162 (.057)	-.228 (.028)					244.6	7

* Model is $Y_{it} = a_i + b_i t + c_i t^2 + W_i(t) + \varepsilon_{it}$ or reduced versions of this; see equation (4). RE denotes random effects, FE denotes fixed effects, IOU denotes integrated OU stochastic process, BM denotes Brownian motion, ME denotes independent measurement error, and OUME denotes correlated measurement errors.

random effects models are not biologically reasonable for these data. The fact that the quadratic random effects model has a much higher likelihood than the linear random effects model is also strong evidence against individuals being on straight line paths, and further suggests that CD4 numbers do not derivative track. The quadratic random effects model is itself biologically unsatisfactory for interpreting individual paths and is not recommended for individual prediction or extrapolation outside the range of the observations.

In the model where the estimate of α in the IOU process is 2.5 (model 4), the likelihood surface is very flat for $\alpha > 2.5$ and the 95% profile likelihood confidence interval is $(.8, \infty)$, with $\alpha = \infty$ corresponding to Brownian motion. This high estimate of α indicates that the derivatives track for only very short periods. The correlation between derivatives at time Δt apart is given by $e^{-\alpha \Delta t}$; with $\alpha = 2.5$, the estimated correlation between slopes 3 months apart is .54, between slopes 6 months apart .28, and between slopes 1 year apart .08. This feature of a decreasing correlation with time is an attractive aspect of the IOU model that we used.

Figures 1 and 2 show some comparisons between the data and the model predictions for models 1 and 5. In these graphs the data are first grouped into 8-week windows of time, and then the means, standard deviations, and correlations are calculated on this grouped $(CD4)^{1/4}$ data. Figure 1 shows the population standard deviation; based on this graph, both models appear adequate, and neither model appears better than the other despite the fact that one has a much smaller likelihood. Figure 1 also clearly indicates that any model that is stationary will be inappropriate for these data, because the standard deviation is increasing with time. It is interesting to note that beyond the range of the data, the predicted standard deviations rapidly differ. It is a feature of random effects models that their predicted variances grow rapidly outside the range of the observations; this may make such models questionable for extrapolation purposes.

The correlation structure of the observations was evaluated in two ways. In a correlogram (not shown), the correlation

between Y_{is} and Y_{it} was plotted against $s - t$ ($s > t$) for various choices of t . Again, we found that the difference in the predicted correlations between models 1 and 5 is small and negligible compared to the variation in empirical correlations calculated from the data.

Figure 2 shows the correlation between $Y_{it_1} - Y_{i(t_1+\delta)}$ and $Y_{it_2} - Y_{i(t_2+\delta)}$ plotted against δ , for $t_2 > t_1 + \delta$. The graph indicates that model 5 fits well and gives a better description of this aspect of the covariance structure than model 1. In calculating the empirical correlations in Figure 2, some individuals contribute more than one set of four observations, and also there is a wide range of values of $t_2 - (t_1 + \delta)$. If one just considers a single quadruplet of observations per person and a restricted range of $t_2 - (t_1 + \delta)$, the graph looks similar, although more noisy, and the conclusion is the same.

The estimates and standard errors of a , b , and c all appear reasonable and are very similar irrespective of the assumed covariance structure. Thus if the main interest is on the fixed effects, then the exact specification of the covariance structure is not vital. The values of $\hat{a}$ and $\hat{b}$ in the best model are such that the population median of CD4 10 years after infection would be approximately 78. Clinical AIDS develops when the CD4 value is around 100 and the median time to AIDS is approximately 10 years. Hence the rate of decline of CD4 as determined by $\hat{b}$ is reasonable.

In this article our focus has been on developing a stochastic model for the process of CD4 decline. The IOU model that we use manifests itself through the covariance structure. We do not intend to suggest that the mean structure is unimportant, and we demonstrate how the effect of covariates on the mean structure can be examined within the IOU model framework. The effect of age at seroconversion on the slope and intercept of the CD4 T-cell number was investigated. Three models of the form

$$Y_{it} = a_i + b_i t + \gamma_1 \text{age}_i + \gamma_2 \text{age}_i t + W_i(t) + \varepsilon_{it}$$

were fit to the data, where a_i and b_i are random effects as before and $W_i(t)$ is either Brownian motion or an IOU pro-

Figure 1. Standard Deviation of $(CD4)^{1/4}$ Versus Time From Seroconversion. * data, discretized into 8-week windows; —, random effects model (Model 1 in Table 1); --- Brownian motion (Model 5 in Table 1).

cess. The results are presented in Table 2. Model 1 is the standard random effects model with $W_i(t) = 0$. In Model 2 the slope is not random and $W_i(t)$ is an IOU process, and in Model 3 the slope is not random and $W_i(t)$ is Brownian motion. We see again that the IOU and Brownian motion models give a better fit to the data than the standard random effects model, and that the estimate of α in the IOU model is not near 0. The magnitude of γ_2 suggests a faster average rate of decline for older subjects, although not statistically significant at the 5% level. The estimated rate of decline is .16 $CD4^{1/4}$ /year for a 24-year-old at HIV infection (10th percentile of age distribution) and .28 $CD4^{1/4}$ /year for a 42-year-old (90th percentile). This result is supportive of other findings (Taylor and Chon 1992) indicating a shorter incubation period for older men.

3.2 Seroprevalent Cohort

There were 809 subjects who were infected with HIV prior to enrollment. Each individual had up to 13 measurements, for a total of 4,636 CD4 values. All observations after December 1987 were excluded, to minimize the influence of AZT on our conclusions.

For all of the subjects, the date of infection, which is the natural time 0, is unknown. But we do know the approximate distribution of the dates of infection; the mean date of in-

fection is about mid-1982, and the standard deviation is about 1 year (Taylor and Chon 1994). The notation describing the various time scales is given in Figure 3.

The most general model we attempt to fit is given by equation (4):

$$Y_{it} = a_i + b_i\tau + c_i\tau^2 + W_i(\tau) + \varepsilon_{it}$$

or, in terms of differences,

$$D_{it} = Y_{it} - Y_{iE_i} = b_i(\tau - E_i) + c_i(\tau^2 - E_i^2) + W_i(\tau) - W_i(E_i) + \varepsilon_{it} - \varepsilon_{iE_i},$$

where τ is the unknown time from HIV infection. This can be written in terms of S_i and known quantities t and V_i as

$$D_{it} = Y_{it} - Y_{iV_i} = (b_i - 2c_iS_i)(t - V_i) + c_i(t^2 - V_i^2) + W_i(t - S_i) - W_i(V_i - S_i) + \varepsilon_{it} - \varepsilon_{iV_i}.$$

In addition to $ES_i = 0$, $\text{var}(S_i) = 1$, we assume that S_i is independent of (b_i, c_i) . We further assume that joint distribution of the new random effect terms can be approximated by a bivariate normal,

$$\begin{pmatrix} b_i - 2c_iS_i \\ c_i \end{pmatrix} \sim N \left[\begin{pmatrix} b \\ c \end{pmatrix}, \begin{pmatrix} \sigma_b^2 + 4(c^2 + \sigma_c^2) & \sigma_{bc} \\ \sigma_{bc} & \sigma_c^2 \end{pmatrix} \right].$$

Figure 2. Correlation $(Y_{it_1} - Y_{it_1+\delta}, Y_{it_2} - Y_{it_2+\delta})$ Versus δ , Where $Y = (CD4)^{1/4}$. * data, discretized into 8-week windows (The number of sets of four measurements contributing to the 13 correlations are, from left to right, 343, 1128, 889, 452, 643, 317, 245, 260, 85, 120, 49, 28, and 11); — random effects model (Model 1 in Table 1); --- Brownian motion (Model 5 in Table 1).

The assumptions concerning the mean and variance of S_i are reasonable from epidemiologic information concerning the growth of the HIV epidemic in Los Angeles. The independence assumption between S_i and (b_i, c_i) is reasonable, because prior to 1988 it is unlikely that the time of HIV infection is related to the subsequent disease progression. In addition to this general model, a similar set of submodels to those applied to the seroconverter cohort was considered, although in this case the differences D_{it} rather than the original observations Y_{it} were analyzed.

As shown in Table 3, a significant finding from the various fits to the data was that the data would not support both random effects terms and either the IOU or Brownian motion stochastic process, and would always choose the stochastic process by forcing the variance of the random effects terms to 0 (see model 4).

The results in Table 3 indicate that Brownian motion plus quadratic drift and measurement error is the best description of these data. The estimates of α and σ^2 in the IOU model are both very large and tended to infinity when the software

Table 2. Effect of Age at Seroconversion on CD4 Decline

Model*	Intercept		Linear		IOU process	ME	-Log-likelihood
	Constant ($\hat{a}$)	Age ($\hat{\gamma}_1$)	Constant ($\hat{b}$)	Age ($\hat{\gamma}_2$)	$\hat{\alpha}$	$\hat{\sigma}_e$	
1. RE(a, b) + Age + ME	5.056 (.295)	.003 (.009)	-.002 (.146)	-.007 (.004)		.264	254.0
2. RE(a) + FE(b) + Age + IOU + ME	5.118 (.273)	.002 (.008)	-.007 (.136)	-.007 (.004)	3.0	.235	236.6
3. RE(a) + FE(b) + Age + BM + ME	5.101 (.275)	.002 (.008)	-.000 (.008)	-.007 (.004)	∞	.213	236.7

* Model is $Y_{it} = a_i + b_i t + \gamma_1 \text{ age}_i + \gamma_2 \text{ age}_i t + W_i(t) + e_{it}$ or reduced versions of this. RE denotes random effects, FE denotes fixed effects, IOU denotes integrated OU stochastic process, BM denotes Brownian motion, and ME denotes independent measurement error.

Figure 3. Time Scales for Seroprevalent Cohort. W , t , and V_i are known. E_i , τ , and S_i are unknown. The mean and variance of S_i are assumed known, $ES_i = 0$ and $\text{var } S_i = 1$.

indicated convergence. This suggests that the derivatives track for at most only very short periods. The likelihood surface is very flat for $\alpha > 5$ in models that contain the IOU term, and the 95% profile likelihood confidence interval is approximately $(3.0, \infty)$ for both models 2 and 9, with $\alpha = \infty$ corresponding to Brownian motion.

As a sensitivity analysis, the assumed mean (ES) and variance (var S) of the distribution of dates of infection were varied, but this had very little effect on either the parameter estimates or the conclusions regarding which model is preferred.

3.3 Seronegative Cohort

This cohort consisted of 726 subjects who remained HIV seronegative throughout their follow-up time. All data after December 1987 were excluded, giving a total of 4,652 observations, with a maximum of 13 measurements per person. The analysis was performed on the differences from baseline. The results are given in Table 4. Somewhat surprisingly, the

simple, appealing stationary model $Y_{it} = a_i + \varepsilon_{it}$ could be substantially improved on in terms of likelihood. The best models included either random coefficients for linear and quadratic terms or Brownian motion in addition to measurement error. But comparison with the results of the seroprevalent and seroconverter data sets reveals that the magnitude of the variance of the random coefficient or the variance of the Brownian motion is much smaller, indicating a more stable population. In addition, as expected, the fixed effects coefficients for the linear and quadratic terms were not significant.

For the AIDS application and in all the simulation work described in the next section, a Fisher scoring algorithm (Jennrich and Schluchter 1986), written in SAS-IML, was used to maximize the likelihood. We found that some reparameterization improved the convergence properties of the algorithm; in particular, we treated α and σ^2/α^2 as the parameters for the IOU model and reparameterized correlations for the random effects terms as $\text{TANH}^{-1}(\rho)$. We experienced some convergence problems with the scoring algorithm, particularly as α became increasingly large or small or as σ_b approached 0. Although we did not reparameterize α , it may be advantageous to use $1/\alpha^2$ or $\log \alpha$ if α is expected to be large. The starting values for the algorithm were obtained from the results of BMDP program P5V, applied to a discretized version of the data.

4. SIMULATION STUDY

The analyses described in Section 3 raised a number of issues, in particular:

- Is it valid to interpret the estimate of α , which is a parameter in the covariance structure?

Table 3. Parameter Estimates from Analysis of Seroprevalent Data Set

Model*	Linear $\hat{b}$	Quadratic $\hat{c}$	IOU process		ME	-Log-likelihood	No. of parameters
	SE($\hat{b}$)	SE($\hat{c}$)	$\hat{\alpha}$	$(\hat{\sigma}/\hat{\alpha})^2$	$\hat{\sigma}_e$		
1. RE(b) + ME	-.180 (.010)				.292	1,882.5	3
2. FE(b) + IOU + ME	-.176 (.010)		23↑	.15	.220	1,803.1	4
3. FE(b) + BM + ME	-.176 (.010)		∞	.15	.213	1,803.1	3
4. RE(b) + BM + ME	-.176 (.010)		∞	.15	.212	1,803.1	4
5. RE(b) + OUME	-.175 (.010)					1,822.6	4
6. FE(b, c) + ME	-.095 (.044)	-.009 (.006)			.364	2,221.6	3
7. RE(b) + FE(c) + ME	-.004 (.038)	-.024 (.005)			.291	1,871.2	4
8. RE(b, c) + ME	-.038 (.047)	-.019 (.006)			.274	1,833.0	6
9. FE(b, c) + IOU + ME	-.069 (.047)	-.015 (.006)	21↑	.15	.221	1,800.5	5
10. FE(b, c) + BM + ME	-.069 (.047)	-.015 (.006)	∞	.15	.214	1,800.5	4
11. RE(b) + FE(c) + OUME	-.036 (.044)	-.019 (.006)				1,817.5	5

* Model is $Y_{it} = a_i + b_i t + c_i t^2 + W_i(t) + \varepsilon_{it}$ or reduced versions of this, and differences are analyzed; see equation (5). RE denotes random effects, FE denotes fixed effects, IOU denotes integrated OU stochastic process, BM denotes Brownian motion, ME denotes independent measurement error, and OUME denotes correlated measurement errors.

- How accurate is the estimate of α likely to be?
- It is reasonable to infer that one model is preferred to another based on larger likelihood values for nonnested models?
- Are the asymptotic standard errors reasonable measures of the uncertainty of the parameter estimates for the sample sizes we used, even for misspecified models?
- How important is it to use the correct covariance structure when the primary emphasis is on inference associated with the mean structure?

To investigate these issues, we used a small Monte Carlo simulation study. The design of the study was based on the seroconverter data set. We used a balanced design with 70 subjects and 10 measurements per subject at equally spaced 4-month intervals, with the first measurement at $t = 4$ months; thus the range of t is .33 to 3.33 years. Multivariate normal data were generated according to a linear random effects (LRE) model, a quadratic random effects (QRE) model, an integrated OU (IOU) model, and a Brownian motion (BM) model; that is, models 1, 3, 4, and 5 in Table 1. For all four generating schemes, it was assumed that $a = 5.165$, $b = -.225$, and $c = 0$ and that the covariance parameters were those obtained from the seroconverter data. To distinguish the IOU model from the BM model, we assumed a lower value of $\alpha = 1$ to generate the data. For every generated data set, all four models were fit to the observations (except the QRE model was not fit to LRE data). Two hundred Monte Carlo replications were used.

Table 5 gives the median, 5th, and 95th percentiles of the estimates of α when the IOU model is fit. The results show that it is reasonable to interpret α as a measure of the degree of derivative tracking, with small α indicating a stronger degree of derivative tracking.

For the mean function parameters, we considered the regression coefficients a and b and the predicted population means at $t = 0$ (intercept), $t = 2$, $t = 4$, and $t = 6$.

Table 6 shows the coverage rates of 95% confidence intervals (estimate ± 1.96 SE) for b ; similar results were ob-

Table 5. Monte Carlo Distribution of Estimate of α in IOU Model

True model	Percentile		
	5th	Median	95th
LRE	0	.10	.60
QRE	.25	.82	2.13
IOU	.38	1.00	2.93
BM	1.53	∞	∞

tained for a and the predicted population means. These results for mean function parameters show that for this sample size and design using the correct covariance structure is not vitally important, provided that a reasonably flexible covariance structure is assumed. An alternative, although possibly less efficient, procedure that would give correct coverage rates for fixed effects is to assume a simple covariance structure and then use robust covariance estimates (Royall 1986) to obtain valid inference.

Table 7 shows the relative efficiency of the estimate of $g(t)$, the population mean, $\text{var } \hat{g}(t) (\text{fitted}) / \text{var } \hat{g}(t) (\text{correct})$. Except in the QRE model, the results for $\hat{b}$ (not shown) were similar to those for $\hat{g}(0)$ (which is $\hat{a}$). The results for relative efficiency show that only a small loss of efficiency is associated with fitting a model with the incorrect LRE covariance structure and no loss of efficiency is associated with incorrectly assuming an IOU or BM covariance structure. A larger loss of efficiency is associated with fitting the QRE model, due to its larger number of parameters; this loss is particularly severe outside the range of the data.

To investigate the interpretation of log-likelihood values, we calculated for each simulated data set the difference in log-likelihood between the fitted correct model and the fitted incorrect model. This difference was then summarized over the 200 simulations. Table 8 shows the median and percentiles of the difference in log-likelihood between fitting the correct model and fitting a different model. The table shows that both the QRE model and the IOU model are flexible

Table 4. Parameter Estimates from Analysis of Seronegative Data Set

Model*	Linear $\hat{b}$	Quadratic $\hat{c}$	IOU process		ME	-Log-likelihood	No. of parameters
	SE($\hat{b}$)	SE($\hat{c}$)	$\hat{\alpha}$	$(\hat{\alpha}/\hat{\alpha})^2$	$\hat{\sigma}_e$		
1. FE($b = 0$) + ME	0				.264	972.1	1
2. FE(b) + ME	.005 (.004)				.264	971.3	2
3. RE(b) + ME	.006 (.005)				.255	953.0	3
4. FE($b = 0$) + IOU + ME	0		100†	.014	.247	950.7	3
5. FE($b = 0$) + BM + ME	0		∞	.014	.247	950.7	2
6. FE(b) + BM + ME	.005 (.005)				.247	950.1	3
7. RE(b, c) + ME	.003 (.014)	.001 (.004)			.252	947.4	6
8. FE($b = 0$) + OUME	0				.266	964.6	2
9. RE(b) + OUME	.005 (.005)				.257	950.9	4

* Model is $Y_{it} = a_i + b_i t + c_i t^2 + W_i(t) + e_{it}$ or reduced versions of this, and differences are analysed; see equation (5). RE denotes random effects, FE denotes fixed effects, IOU denotes integrated OU stochastic process, BM denotes Brownian motion, ME denotes independent measurement error, OUME denotes correlated measurement errors.

Table 6. Monte Carlo Coverage Rates for Regression Coefficient; Nominal Level = 95%

True model	Fitted model			
	LRE	QRE	IOU	BM
LRE	.96	ND	.95	.90
QRE	.96	.95	.95	.91
IOU	.95	.95	.97	.94
BM	.95	.96	.96	.96

enough to accommodate other covariance structures. But the smaller number of parameters in the IOU model makes it more appealing. The larger differences in log-likelihood when either LRE or BM is the fitted model indicates that these models are less able to accommodate more general covariance structures.

In the data analysis in Section 3, we found that BM was usually the best fitting model, with QRE a close competitor in terms of log-likelihood. In the simulation, we found that QRE has the highest log-likelihood 98% of the time when QRE is the true model, but only 11% of the time when BM is the true model. These and the other simulation results indicate that it is reasonable to try to interpret the parameters of the covariance structure for the sample sizes in our AIDS application, thus supporting our conclusions in Section 3.

5. DISCUSSION

In addition to the covariance structure, which has been the main focus of this article, and the mean structure, which is also important, there are other issues in modeling serial CD4 measurements. Many of the data sets are from prevalent cohort studies, in which those with AIDS are excluded from enrolling in the study. Hence there is a tendency for persons with very fast rates of CD4 decline to be truncated from the sample, giving a biased sample. A second source of bias is due to informative censoring. That is, a subject in the study whose CD4 numbers drop rapidly will develop AIDS and then drop out of the study due to illness or death. A comparison of $\hat{b}$, the linear rate of decline of $CD4^{1/4}$, between the seroconverters and the seroprevalent cohort suggests that perhaps these biases are influencing the estimates. In particular, the average rate of decline of CD4 is larger for seroconverter subjects than for seroprevalent subjects.

There are other important markers of disease progression in AIDS, such as serum neopterin and beta-2-microglobulin. It would be interesting and possibly important to model the stochastic development of these markers jointly with CD4. Both the OU and IOU processes have multivariate generalizations that could be used for the analysis of such a vector of response variables. A K -variate stationary OU process $\{Y(t); t \geq 0\}$ is defined as the k -variate, stationary, Gaussian, Markov process that is continuous in probability and satisfies $E\{Y(t)\} = 0$. Its covariance structure is given by $Ce^{B(t-s)}$, $0 \leq s \leq t$, where B and C are $K \times K$ real matrices. The IOU process, defined by $W(t) = \int_0^t Y(S) ds$, is Gaussian and has mean 0 and covariance function (Cumberland and Rohde 1977; Schach 1971)

Table 7. Monte Carlo Relative Efficiency for Predicted Population Mean $\hat{g}(t)$, $\text{Var } \hat{g}(t)$ (Fitted)/ $\text{Var } \hat{g}(t)$ (Correct), where $\hat{g}(t) = \hat{a} + \hat{b}t$ or $\hat{g}(t) = \hat{a} + \hat{b}t + \hat{c}t^2$

True model	t	Fitted model			
		LRE	QRE	IOU	BM
LRE	0	1.00	ND	1.01	1.03
LRE	2	1.00	ND	1.00	1.01
LRE	4	1.00	ND	1.00	.99
LRE	6	1.00	ND	1.00	.99
QRE	0	.67	1.00	.72	.72
QRE	2	.90	1.00	.88	.88
QRE	4	.89	1.00	.86	.86
QRE	6	.29	1.00	.29	.28
IOU	0	1.04	1.30	1.00	1.00
IOU	2	1.05	1.18	1.00	.99
IOU	4	1.05	1.15	1.00	1.00
IOU	6	1.04	3.10	1.00	1.00
BM	0	1.09	1.41	1.00	1.00
BM	2	1.08	1.27	1.00	1.00
BM	4	1.07	1.25	1.00	1.00
BM	6	1.06	4.42	1.00	1.00

$\text{cov}(W(s), W(t))$

$$= (e^{B's} - I - B's)B'^{-2}C + CB^{-2}(e^{Bt} - e^{B(t-s)} - Bs),$$

$$0 \leq s \leq t.$$

The multivariate IOU process is an attractive and flexible stochastic process for analyzing multivariate longitudinal data. But the escalation in the number of parameters is considerable for $K > 2$, and a substantial amount of data would be required to fit this model.

In this article we have suggested a flexible and tractable family of models for the analysis of longitudinal data. Our main focus in the AIDS application has been to give a parsimonious and biologically plausible model for the underlying stochastic process that manifests itself through the correlation structure. Furthermore, this model can be useful both at the population level and at the individual level. We have emphasized the concept of "derivative tracking" in which an individual's measurements have the tendency to maintain the same trajectory. Using an IOU process as part

Table 8. Differences in Log-Likelihood ($\log L(\text{CORRECT}) - \log L(\text{FITTED})$)

Fitted model	No. of parameters	Correct model	No. of parameters	Percentiles of difference in log-likelihood		
				5th	50th	95th
LRE	6	QRE	10	2.6	7.3	15.8
LRE	6	IOU	6	1.6	8.0	18.3
LRE	6	BM	5	7.1	17.9	31.4
QRE	10	IOU	6	-4.8	-4	3.7
QRE	10	BM	5	-2.4	4.3	11.2
IOU	6	LRE	6	.1	4.1	10.4
IOU	6	QRE	10	1.8	5.6	12.1
IOU	6	BM	5	-1.7	0	0
BM	5	LRE	6	5.4	13.3	23.2
BM	5	QRE	10	4.2	10.5	18.6
BM	5	IOU	6	0.3	3.9	9.9

of the covariance structure allows the amount of derivative tracking to be estimated from the data.

We find substantial evidence that the derivatives of CD4 do not track for any extended period. This implies, excluding the contribution from measurement error, that prior CD4 values are of little additional assistance in predicting future values, if the current CD4 value is known.

Two caveats regarding this conclusion are (1) it applies to a largely healthy cohort whose CD4 values are generally higher than 100, and (2) our measurements were usually 6 months apart and always at least 3 months apart, so it is possible that the derivatives may track for time scales of less than 3 months. If the derivatives do track for periods shorter than 3 months, this can be easily accommodated by an IOU model with a large value for α .

Other authors (Fusaro, Neilsen, and Scheike 1993; Lin, Fischl, and Schoenfeld 1993; Phillips et al. 1991; Taylor, Tan, Detels, and Giorgi 1991; Tsiatis et al. 1992) have investigated, either directly or indirectly, whether future disease progression is dependent on prior CD4 values if the current CD4 value is known. The methods used are very different from those described in this article, and in some cases the endpoint is clinical disease or death rather than future CD4 values. Most authors find that the latest CD4 value contains most of the information about future disease progression and that prior values add only a little information at most. These findings are consistent with our conclusion that, ignoring the effect of measurement error through which past CD4 values may help to define the current CD4 value, the prior CD4 value adds no information concerning future values given the current value.

The finding of no derivative tracking allows some very cautious substantive interpretation—in particular, that fixed factors, such as possibly age or genetic factors, which might have a continual effect on a subject's immune system, are not the sole dominant determinant of an individual's rate of decline of CD4. Rather, time-dependent factors that cause the slope of CD4 to vary with time, such as concurrent infections, are more likely candidates to explain an individual's CD4 path. This hypothesis is supported by our analysis of the effect of age at seroconversion on the slope of CD4 decline. In this analysis we showed that age has a possibly large effect on the rate of decline of CD4, but there was no significant improvement of the fit of the model by including age. The conclusion of no derivative tracking may also have implications for how CD4 values are used in a clinical trial in which the endpoint is to be based on this marker.

[Received August 1992. Revised September 1993.]

REFERENCES

- Berman, S. M. (1990), "A Stochastic Model for the Distribution of HIV Latency Time Based on T4 Counts," *Biometrika*, 77, 733–741.
- Chi, E. M., and Reinsel, G. C. (1989), "Models for Longitudinal Data With Random Effects and AR(1) Errors," *Journal of the American Statistical Association*, 84, 452–459.
- Cullis, B. R., and McGilchrist, C. A. (1990), "A Model for the Analysis of Growth Data From Designed Experiments," *Biometrics*, 46, 131–142.
- Cumberland, W. G., and Rohde, C. A. (1977), "A Multivariate Model for Growth of Populations," *Theoretical Population Biology*, 11, 127–139.
- Cumberland, W. G., and Sykes, Z. M. (1982), "Weak Convergence of an Autoregressive Process Used in Modeling Population Growth," *Journal of Applied Probability*, 19, 450–455.
- DeGruttola, V., Lange, N., and Dafni, U. (1991), "Modeling the Progression of HIV Infection," *Journal of the American Statistical Association*, 86, 569–577.
- Diggle, P. J. (1988), "An Approach to the Analysis of Repeated Measurements," *Biometrics*, 44, 959–971.
- Foulkes, M. A., and Davis, C. E. (1981), "An Index of Tracking for Longitudinal Data," *Biometrics*, 37, 439–446.
- Fusaro, R. E., Nielsen, J. P., and Scheike, T. H. (1993), "Marker-Dependent Hazard Estimation: An Application to AIDS," *Statistics in Medicine*, 12, 843–865.
- Jennrich, R. I., and Schluchter, M. D. (1986), "Unbalanced Repeated-Measures Model With Structured Covariance Matrices," *Biometrics*, 42, 804–820.
- Jones, R. H., and Boadi-Boateng, F. (1991), "Unequally Spaced Longitudinal Data With AR(1) Serial Correlation," *Biometrics*, 47, 161–175.
- Kanazawa, Y. (1990), "Estimating Expected Time of Infection With HIV," Technical Report 13, SIMS (Statistical Methodology for Study of the AIDS Epidemic).
- Kaslow, R. A., Ostrow, D. G., Detels, R., Phair, J. P., Polk, B. F., and Rinaldo, C. R., Jr. (1987), "The Multicenter AIDS Cohort Study: Rationale, Organization and Selected Characteristics of the Participants," *American Journal of Epidemiology*, 126, 310–318.
- Laird, N. M., and Ware, J. H. (1982), "Random Effects Models for Longitudinal Data," *Biometrics*, 38, 963–974.
- Lang, W., Perkins, H., Anderson, R. E., Royce, R., Jewell, N. P., and Winkelstein, W. Jr. (1989), "Patterns of T Lymphocyte Changes With Human Immunodeficiency Virus Infection: From Seroconversion to the Development of AIDS," *Journal of Acquired Immune Deficiency Syndrome*, 2, 63–69.
- Lange, N., Carlin, B. P., and Gelfand, A. E. (1992), "Hierarchical Bayes Models for the Progression of HIV Infection Using Longitudinal CD4 T-Cell Numbers," *Journal of the American Statistical Association*, 87, 615–632.
- Lin, D. Y., Fischl, M. A., and Schoenfeld, D. A. (1993), "Evaluating the Role of CD4 Lymphocyte Counts as Surrogate Endpoints in Human Immunodeficiency Trials," *Statistics in Medicine*, 12, 839–842.
- Lindstrom, M. J., and Bates, D. M. (1988), "Newton-Raphson and EM Algorithms for Linear Mixed Effects Models for Repeated-Measures Data," *Journal of the American Statistical Association*, 83, 1014–1022.
- Longini, I. M., Clark, W. S., Byers, R. H., Ward, J. W., Darrow, W. W., Lemp, G. F., and Hethcote, H. W. (1989), "Statistical Analysis of the Stages of HIV Infection Using a Markov Model," *Statistics in Medicine*, 8, 831–843.
- Margolick, J. B., Donnenberg, A. D., Muñoz, A., Park, L. P., Bauer, K. D., Giorgi, J. V., Ferbas, J., and Saah, A. J. (1993), "Changes in T and Non-T Lymphocytes Subsets Following Seroconversion to HIV-1: Stable CD3⁺ and Declining CD3⁺ Populations Suggest Regulatory Responses Linked to Loss of CD4 Lymphocytes," *Journal of the Acquired Immune Deficiency Syndrome*, 6, 153–161.
- McMahan, C. A. (1981), "An Index of Tracking," *Biometrics*, 37, 447–455.
- Munöz, A., Cary, V., Saah, A. J., Phair, J. P., Kingsley, L. A., Fahey, J. L., Ginzburg, H. M., and Polk, B. F. (1988), "Predictors of Decline in CD4 Lymphocytes in a Cohort of Homosexual Men Infected With Human Immunodeficiency Virus," *Journal of the Acquired Immune Deficiency Syndrome*, 1, 396–404.
- Phillips, A. N., Lee, C. A., Elford, J., Janossy, G., Timms, A., Bofill, M., and Kernoff, P. B. (1991), "Serial CD4 Lymphocyte Counts and Development of AIDS," *Lancet*, 337, 389–392.
- Royall, R. M. (1986), "Model Robust Confidence Intervals for Maximum Likelihood Estimators," *International Statistical Review*, 54, 221–226.
- Schach, S. (1971), "Weak Convergence Results for a Class of Multivariate Markov Process," *Annals of Mathematical Statistics*, 42, 451–465.
- Taylor, J. M. G., and Chon, Y. (1992), "Semi-Parametric Estimation of the Incubation Period of AIDS," in *AIDS Epidemiology: Methodological Issues*, eds. N. P. Jewell, K. Dietz, and V. T. Farewell, Boston: Birkhauser.
- Taylor, J. M. G., and Chon, Y. (1994), "Smoothing Grouped Bivariate Data to Obtain the Incubation Period Distribution of AIDS," *Statistics in Medicine*, 13, 969–981.
- Taylor, J. M. G., Tan, S.-J., Detels, R., and Giorgi, J. V. (1991), "Applications of Computer Simulation Model of the Natural History of CD4 T-Cell Number in HIV-Infected Individuals," *AIDS*, 5, 159–167.
- Tsiatis, A. A., DeGruttola, V., Strawderman, R. L., Dafni, U., Propert, K. J., and Wulfsohn, M. (1992), "The Relationship of CD4 Counts Over Time to Survival in Patients With AIDS: Is CD4 a Good Surrogate Marker?" in *AIDS Epidemiology: Methodological Issues*, eds. N. P. Jewell, K. Dietz, and V. T. Farewell, Boston: Birkhauser.
- Ware, J. H., and Wu, M. C. (1981), "Tracking: Prediction of Future Values From Serial Measurements," *Biometrics*, 37, 427–437.