Nonlinear Mixed-Effects Modeling: Individualization and Prediction

Erik Olofsen, David F. Dinges, and
Hans P.A. Van Dongen

Olofsen E, Dinges DF, Van Dongen HPA. ***Nonlinear mixed-effects modeling: individualization and prediction.*** **Aviat Space Environ Med 2004; 75(3, Suppl.):A134–40.**

The development of biomathematical models for the prediction of fatigue and performance relies on statistical techniques to analyze experimental data and model simulations. Statistical models of empirical data have adjustable parameters with a priori unknown values. Inter-individual variability in estimates of those values requires a form of smoothing. This traditionally consists of averaging observations across subjects, or fitting a model to the data of individual subjects first and subsequently averaging the parameter estimates. However, the standard errors of the parameter estimates are assessed inaccurately by such averaging methods. The reason is that intra- and inter-individual variabilities are intertwined. They can be separated by mixed-effects modeling in which model predictions are not only determined by fixed effects (usually constant parameters or functions of time) but also by random effects, describing the sampling of subject-specific parameter values from probability distributions. By estimating the parameters of the distributions of the random effects, mixed-effects models can describe experimental observations involving multiple subjects properly (i.e., yielding correct estimates of the standard errors) and parsimoniously (i.e., estimating no more parameters than necessary). Using a Bayesian approach, mixed-effects models can be "individualized" as observations are acquired that capture the unique characteristics of the individual at hand. Mixed-effects models, therefore, have unique advantages in research on human neurobehavioral functions, which frequently show large inter-individual differences. To illustrate this we analyzed laboratory neurobehavioral performance data acquired during sleep deprivation, using a nonlinear mixed-effects model. The results serve to demonstrate the usefulness of mixed-effects modeling for data-driven development of individualized predictive models of fatigue and performance.

Keywords: mixed-effects models, random effects, fixed effects, between-subjects variance, within-subjects variance, individualization, prediction, Bayes posterior distribution estimation, neurobehavioral performance, sleep deprivation.

EMPIRICAL DATA SETS in human fatigue and performance research often involve multiple subjects, each contributing a series of data points obtained over time. Such data sets contain two distinct sources of variance: between-subjects variance (due to differences among individuals) and within-subjects variance (due to changes over time and/or noise). When describing the data with a statistical model, these two sources of variance must be separated in order to get optimal estimates of the parameters of the model and their standard errors. This can be done by means of mixed-effects modeling. To explain the mixed-effects modeling approach we will first discuss the analysis of a very simple model in such a degree of detail that all essential steps of the procedure are presented without having to go into technicalities. For more details and discussions we refer to the literature (5,6,9,11,14,16). Implementations of nonlinear mixed-effects modeling can be found in the program NONMEM (an acronym for "NONlinear Mixed-Effects Modeling") (6), in module NLME in S-PLUS (7), and in PROC NLMIXED in SAS (15). For all analyses described in this paper we used the program NONMEM.

METHODS

Within-Subjects and Between-Subjects Variance

Suppose that y_{ij}, the j-th experimental observation ($j = 1, \ldots, M$) from the i-th subject ($i = 1, \ldots, N$) in a study, can be modeled by:

$$\hat{y}_{ij} = a_i + \epsilon_{ij}. \qquad \text{Eq. 1}$$

The a_i are unknown constants and the subscript i reflects inter-individual (between-subjects) variability; the ϵ_{ij} are independent and normally distributed noise with mean zero and variance σ^2, and reflect intra-individual (within-subjects) variability. Suppose that $a_i = a + \eta_i$, where a is a constant and the η_i are independent and normally distributed with mean zero and variance ω^2, so that:

$$\hat{y}_{ij} = a + \eta_i + \epsilon_{ij}. \qquad \text{Eq. 2}$$

This is a simple mixed-effects model: a is the fixed effect and η_i is the random inter-individual effect for the vector $\hat{y}_i$. In more elaborate mixed-effects models for time series, time is a fixed effect, the model parameters are functions of one or more ηs, and stationarity is reflected in the assumptions about ϵ.

To estimate a, the observations could first be averaged over all subjects so that:

$$\hat{y}_j = a + \frac{1}{N}\sum_{i=1}^{N}\eta_i + \frac{1}{N}\sum_{i=1}^{N}\epsilon_{ij} = a + \bar{\eta} + \bar{\epsilon}_j \qquad \text{Eq. 3}$$

From the Department of Anesthesiology, Leiden University Medical Center, Leiden, The Netherlands (E. Olofsen); and the Division of Sleep and Chronobiology, Unit for Experimental Psychiatry, University of Pennsylvania School of Medicine, Philadelphia, PA (D. F. Dinges and H. P. A. Van Dongen).

Address reprint requests to: Erik Olofsen, M.Sc., who is a research associate, Department of Anesthesiology P5Q, Leiden University Medical Center, PO Box 9600, 2300 RC Leiden, The Netherlands; e.olofsen@lumc.nl

where $\bar{\eta}$ and $\bar{\epsilon}_j$ have variance ω^2/N and σ^2/N, respectively. Least-squares analysis of the averaged observations $\bar{y}_j$, in this case equivalent with averaging again, yields $\hat{a} = a + \bar{\eta} + \bar{\epsilon}$. This estimate has a possibly severely underestimated variance $VAR\{\hat{a}\} = \sigma^2/(N \cdot M)$ because information in the variability η_j was lost, as $\bar{\eta}$ is a constant.

An alternative approach, generally called the two-stage approach, entails first estimating a_i for each individual:

$$\hat{a}_i = a + \eta_i + \frac{1}{M}\sum_{j=1}^{M} \epsilon_{ij} = a + \eta_i + \bar{\epsilon}_i \quad \text{with VAR}\{\hat{a}_i\} = \frac{\sigma^2}{M} \qquad \text{Eq. 4}$$

Subsequently the $\hat{a}_i$ are averaged which yields:

$$\hat{a} = \frac{1}{N}\sum_{i=1}^{N} \hat{a}_i, \quad \text{with VAR}\{\hat{a}\} = \frac{\omega^2}{N} + \frac{\sigma^2}{N \cdot M} \qquad \text{Eq. 5}$$

This procedure is suboptimal because the separate estimates $\hat{a}_i$ may need to be obtained from small M. This usually yields biased estimates of model parameters and their variance, in particular in situations where the model is nonlinear, where M is not equal for every subject, where the model parameters depend on covariates (such as subjects' age), and/or where two or more parameters need to be estimated that may be correlated. Furthermore, the magnitudes of the separate intra- and inter-individual variability terms (σ^2 and ω^2) remain unknown, so even though the above expression for VAR{â} is correct, we have no means of evaluating it.

Maximum-Likelihood Estimation and Mixed-Effects Modeling

Parameters of mixed-effects models can be estimated by application of the maximum likelihood principle. According to this principle, the best estimates for a, σ^2, and ω^2 are those that would yield the greatest likelihood that the samples y_{ij} have the observed values. This requires knowledge of the type of probability distribution functions for the inter-individual variability of the parameters, which need to be postulated in practice (but were given for the analytical example of the previous section). For each subject, the likelihood l_i of observing the set of samples y_i is then given by ($\propto$ stands for proportionality):

$$l_i(y_i; a, \sigma, \eta_i) \propto \prod_{j=1}^{M} p_N(y_{ij}; a + \eta_i, \sigma) \qquad \text{Eq. 6}$$

where $p_N(x; \mu, \sigma)$ denotes the normal density function at x with mean μ and variance σ^2. Although η_i is unknown, it can be integrated out, which gives the marginal likelihood:

$$L_i(y_i; a, \sigma, \omega) \propto \int_{\eta_i=-\infty}^{\infty} l_i(y_i; a, \sigma, \eta_i) \cdot p_N(\eta_i; 0, \omega) \cdot d\eta_i \qquad \text{Eq. 7}$$

The likelihood of observing the entire set of samples y_{ij} is then given by:

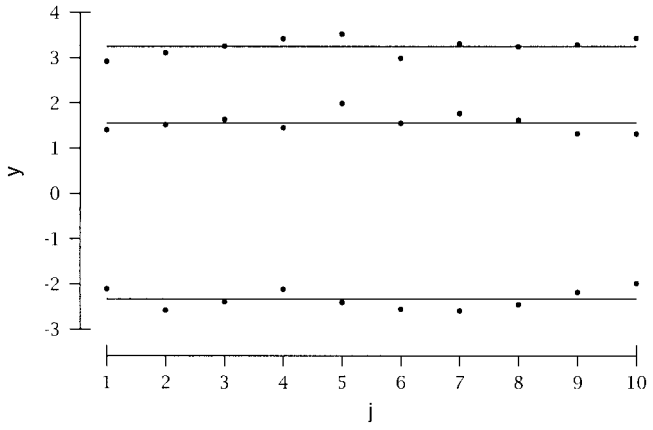

Fig. 1. Three sets of observations of y_{ij} (j = 1, . . . , 10; arbitrary i) from the analytical model given by Eq. 2 with N = 20, M = 10, a = 1, σ^2 = 0.04, and ω^2 = 2.25 (dots); and individual estimates of a = â + η_i (lines).

$$L(y; a, \sigma, \omega) = \prod_{i=1}^{N} L_i(y_i; a, \sigma, \omega) \qquad \text{Eq. 8}$$

By maximizing L, estimates of a, σ and ω and their variances can be obtained. For more complicated models and distributions of the variability terms, the integral in Eq. 7 will usually be intractable and approximations are required. A variety of numerical approximation algorithms is available for this purpose in the mixed-effects modeling software routines (6,7,15).

To illustrate the mixed-effects modeling procedure, data from the model given by Eq. 2 with a = 1, σ^2 = 0.04, and ω^2 = 2.25 were generated for N = 20 individuals with M = 10 assessments per individual, using computer simulation. Data of three individuals are plotted in **Fig. 1.** Mixed-effects modeling as implemented in NONMEM was used to estimate the parameters, resulting in â = 1.3 ± 0.3, $\hat{\sigma}^2$ = 0.034 ± 0.004, $\hat{\omega}^2$ = 2.1 ± 0.7 (estimate ± SE), in good agreement with their exact values. The procedure was also used to estimate a_i for each individual, which is represented by the lines in Fig. 1. In the next section, we explain how the mixed-effects modeling framework is used to calculate these individual estimates.

Bayesian Estimation for Predictive Modeling

Subject-specific parameters are functions of the inter-individual variability terms η_i. Although these are unknown, they can be estimated by the mode of the Bayes posterior densities; they are those η_i that maximize:

$$p(y_i; a, \sigma, \omega) = \frac{l_i(y_i; a, \sigma, \eta_i) \cdot p_N(\eta_i; 0, \omega)}{L_i(y_i; a, \sigma, \omega)} \qquad \text{Eq. 9}$$

This equation specifies that the probability density for y_i, conditional on the knowledge of the properties of the population, equals the likelihood of observing y_i multiplied by the prior (population) density of the inter-individual variability term η_i. The denominator serves as a normalization to yield a proper probability density function. In practice, population estimates are substituted for a, σ, and ω since with experimental data their true values are unknown. Notice that when more data are obtained for an individual, the likelihood of

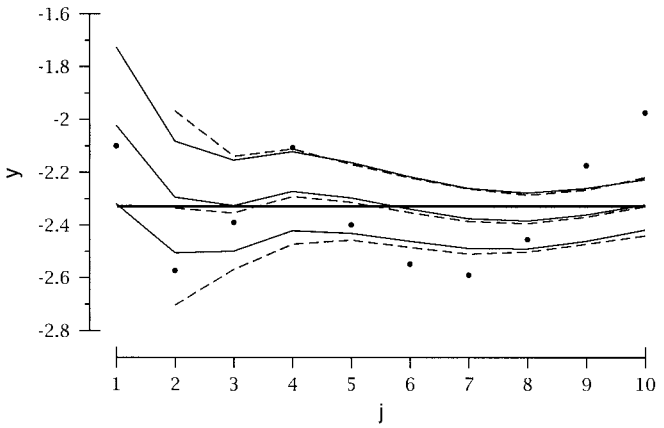

Fig. 2. Iteratively estimated Bayesian values of a together with their 90% probability intervals for individual 20 (thin lines), based on availability of only the first data point, the first two data points, etc. The exact value of a in this individual (thick line) and the subject's observations (dots) are shown as well. The panel further shows iterative estimates of a with 90% probability intervals based solely on this individual's data (dashed lines). When no knowledge from other individuals is available, at least two observations are necessary to estimate this individual's a and its variance, and many more observations are required when dealing with more complex models and less informative data (e.g. 16).

observing their data is an increasingly localized function of η_i (M increases in Eq. 6); and η_i, therefore, converges from the population value of 0 to the value that corresponds to the exact value for this individual. Consequently, this analysis does not critically depend on the assumptions made about the distributions of inter-individual variabilities.

To show the usefulness of the Bayes posterior distribution estimates, we first considered only individuals 1–19, which resulted in slightly different mixed-effects modeling estimates of a, σ, and ω. These parameter estimates provide information that can be used to predict the value of $a_{20} = a + \eta_{20}$ when incomplete data for the 20th subject become available post hoc. To demonstrate this, we iteratively determined the Bayes posterior density (Eq. 9) of y_{20} by incorporating observations y_{20j} starting first with $y_{20,1}$, then adding $y_{20,2}$, etc. For reference, y_{20} and a_{20} correspond to the individual shown at the bottom in Fig. 1. With each added data point, a more accurate estimate of a_{20} is achieved, as shown in **Fig. 2.**

The information about individuals 1–19 helps to estimate a_{20}. Even with only one (or in fact, without any) data point for the 20th individual, an estimate and confidence interval for a_{20} can be given. As data points are added, the estimate for a_{20} reaches the exact value more rapidly than when no information about individuals 1–19 were available (cf. solid vs. dashed lines in Fig. 2). The distribution* of $\hat{a}_i$ becomes narrower (i.e., the confidence interval becomes smaller) as data points are added; it converges from the population distribution [approximately $p_N(a, \omega)$] to the distribution for case i [approximately $p_N(a_i, \sigma / \sqrt{M})$]. Thus, the Bayesian estimation technique in the mixed-effects modeling framework can be extremely valuable in predictive modeling of longitudinal data for individual subjects.

RESULTS

Application to Experimental Data

As part of a larger study, N = 13 subjects (27.7 ± 5.4 yrs old) spent 20 d inside a laboratory (12). Informed consent was obtained from each subject. After 3 baseline days with 8 h time-in-bed (23:30–07:30), subjects' sleep was restricted to 4 h time-in-bed (03:30–07:30) for 14 d. Neurobehavioral performance was tested every 2 h during wakefulness, and included a psychomotor vigilance test (PVT; 4). Daily averages (09:30–23:30) were computed for PVT lapses (reaction times ≥ 500 ms).

The PVT lapse data y were modeled by the following equation:

$$y(t) = \alpha + \beta t^{\theta} + \epsilon(t) \qquad \text{Eq.10}$$

where the fixed effect t is time in days (t = 0, . . . , 14), the parameters (affected by random effects) α, β, and θ denote baseline performance†, trend gain (i.e., rate of performance impairment), and a nonlinearity exponent, respectively, and ϵ denotes within-subjects variability.

The stochastic variables α and β were assigned log-normal distributions, which is a reasonable description for parameters that show large inter-individual variability but should always be positive. This was implemented by multiplying their typical (i.e., the median of the distribution when it holds exactly) values with factors $\exp(\eta)$, where η is normally distributed with mean zero and variance ω^2. Variable θ was assigned a normal distribution. Variable ϵ was assumed to be normally distributed with mean zero and variance σ^2 for each individual (which was a reasonable assumption according to visual inspection of the residuals as a function of time). Covariances between the random effects variables were assumed to be zero (and are usually not well-estimable from small populations).

Table I presents NONMEM's estimates of the parameters and their SE for this data set. We compare these to the results from a two-stage analysis (also using NONMEM for the first stage) as shown in **Table II.** Notice that the means of the individual estimates of α, β, and

* The determined confidence intervals are approximations to prediction intervals for the set of observations, since population estimates of *a*, σ, and ω were used in their construction (see Eq. 9). The bootstrap may be used to deal with this issue (2).

† Note that when t = 0, $y(0) = \alpha + \epsilon(0)$; normalizing y by dividing by y(0) to eliminate α would cause the uncertainty of $\epsilon(0)$ to propagate to the estimates of the remaining parameters.

TABLE I. MIXED-EFFECTS ANALYSIS OF THE EXPERIMENTAL DATA: ESTIMATES OF MODEL PARAMETERS (α, β, θ) AND RESIDUAL NOISE VARIANCE (σ^2), WITH SE, AND INTER-INDIVIDUAL VARIABILITY TERMS ω^2 OF THE PARAMETERS, WITH SE ω^2.

Parameter	Estimate	SE	ω^2	SE ω^2
α	1.74	1.09	1.41	0.727
β	1.18	0.789	1.72	1.24
θ	0.833	0.184	0.179	0.144
σ^2	8.17	1.74		

TABLE II. TWO-STAGE ANALYSIS OF THE EXPERIMENTAL DATA: MEAN AND SD OF PARAMETER ESTIMATES (α, β, θ) AND RESIDUAL NOISE VARIANCE (σ^2), AS WELL AS THEIR MEDIANS AND MEDIAN ABSOLUTE DEVIATIONS (MAD).

Parameter	Mean	SD	SEM*	Median	MAD†
α	3.18	4.37	1.21	2.29	2.33
β	2.33	2.26	0.628	2.20	2.17
θ	1.33	1.11	0.308	0.806	0.286
σ^2	6.36	4.54	1.26	6.25	3.79

*SEM were obtained by dividing the SD by the square root of the number of individuals (N = 13).
†The median absolute deviations give an indication of the stability of the SD as the former are robust against outliers; asymptotically, MAD equals SD times 0.675. The SE of the median equals SEM times $\sqrt{(\pi/2)}$ (ref. 8).

θ are higher than their medians, suggesting that the distributions of the two-stage estimators are right-skewed. However, this does not necessarily imply that the inter-individual variability distributions are right-skewed as well. The median α and β are higher than the estimates obtained by the mixed-effects analysis, indicating a bias in either the mixed-effects estimates, the two-stage estimates, or both (the true values are obviously unknown, in contrast to the Monte Carlo simulations below). The difference between the two-stage and mixed-effects estimates of intra-individual variability (σ^2) may be explained by the fact that these are based on only 15 data points for the former and 195 (15 times 13) for the latter analysis. The residual variance of the two-stage estimate of σ^2 will, therefore, be biased with a factor of approximately 12/15 (where 12 is the number of data points minus the number of parameters to be estimated) which corresponds well with the ratio 6.36/8.17. However, in general the degree of bias of an estimator is unknown and NONMEM cannot routinely apply bias correction. Inter-individual variability is usually poorly quantified in a two-stage analysis, as the squares of the SD (possibly transformed for non-normal distributions) of the two-stage estimates are only equivalent to the inter-individual variability terms ω^2 in Table I when the value of σ^2 is negligible.

Monte Carlo Simulation

Monte Carlo simulations were performed to demonstrate and explain the differences in results from a mixed-effects approach versus a two-stage approach. Monte Carlo simulation is a technique tailored to the study of the properties of estimators. Data sets are obtained from a model with specified model parameters, and random-numbers generators provide values for the random variables.‡ For each data set, model parameters are estimated using the technique under investigation (e.g., mixed-effects modeling), and the mean and SD of those estimates provide information about the bias and true SE of the estimators under study.

Table III presents the results of the mixed-effects modeling approach applied to 1,000 data sets, comprised of data for 13 simulated subjects, generated using Eq. 10 with population estimates of the parameters α, β, and θ, and σ and variances of their inter-individual variability terms (ω^2) obtained from Table I. There is consistency between these results and those from the original data set. Bias, defined by the expectation of an estimator minus the true value of the parameter, and estimated here by the difference between the mean of the Monte Carlo estimates and the fixed value used for the simulation, is relatively small.§

The SD of the estimated parameters (columns 2 and 6 in Table III) are a measure of their true variability (as they are obtained by repeatedly analyzing new data). They correspond well with the means of the SE (columns 3 and 7), estimated by maximum likelihood theory, demonstrating their validity. An exception is the relative disagreement between the SD of the inter-individual variability estimates of α (1.57) and the mean estimated SE (0.823) of α; the latter would seem to be underestimated. However, the median absolute deviation of the inter-individual variability estimates of α was 0.447 (not shown) so they are probably not normally distributed, causing the standard deviation to be overestimated in this case. The standard deviation of σ^2 (0.890) in Table III is considerably lower than the standard error (1.74) obtained from the experimental data (Table I). This may be due to inter-individual variability in σ^2 in the experimental data which we did not take into account.§§

The results of the two-stage approach applied to

‡ Extensive simulations like these have recently become feasible on personal computer platforms due to the considerable increases in computational power.
§ Although maximum likelihood estimators are asymptotically unbiased, they are likely to be biased with a finite data set. Precision of the estimates of bias may be obtained by dividing the standard deviations by the square root of the number of Monte Carlo trials.
§§ This could be done in NONMEM, but at the expense of a poorer approximation of the likelihood function as the Laplacian approximation is not available in this case (6).

TABLE III. MEANS (AND SD) OF PARAMETER ESTIMATES AND THEIR STANDARD ERRORS (SE) OBTAINED FROM 1,000 MONTE CARLO SIMULATIONS USING THE POPULATION MODEL IN EQ. 10 WITH PARAMETERS AS GIVEN IN TABLE I, AND ANALYZED WITH THE MIXED-EFFECTS APPROACH.

Parameter	Mean*	(SD)†	Mean SE†	(SD)	Mean* ω^2	(SD)	Mean SE ω^2	(SD)
α	2.04	(0.969)	1.01	(0.478)	1.35	(1.57)	0.823	(2.32)
β	1.35	(0.681)	0.651	(0.470)	1.70	(1.20)	0.886	(0.690)
θ	0.852	(0.203)	0.165	(0.114)	0.156	(0.112)	0.0811	(0.0823)
σ^2	8.16	(0.890)	0.788	(0.295)				

* Biases, given by the difference between the mean parameter estimates and their values in Table I, are relatively small.
† The SD of the parameters, which are a measure of their true SE, correspond well the means of NONMEM's estimated SE.

TABLE IV. MEANS (AND STANDARD DEVIATIONS) OF PARAMETER ESTIMATES AND THE STANDARD ERRORS OF THE MEAN (SEM), AS WELL AS THE MEDIANS AND MEDIAN ABSOLUTE DEVIATIONS (MAD), OBTAINED FROM 1,000 MONTE CARLO SIMULATIONS USING THE POPULATION MODEL IN EQ. 10 WITH PARAMETERS AS GIVEN IN TABLE I, AND ANALYZED WITH THE TWO-STAGE APPROACH.

Parameter	Mean*	(SD)†	SEM	(SD)	Median*	(SD)	MAD	(SD)
α	3.17	(2.27)	1.51	(1.36)	1.97	(1.17)	2.14	(0.906)
β	3.63	(1.93)	1.25	(1.10)	2.16	(1.10)	1.51	(0.732)
θ	1.03	(0.213)	0.181	(0.066)	0.943	(0.220)	0.361	(0.135)
σ^2	6.50	(0.795)	0.713	(0.193)	6.24	(0.932)	1.60	(0.625)

* Biases, given by the difference between the mean or median parameter estimates and their values in Table I, are higher than when the mixed-effects approach is applied.
† The SD of the parameters are also higher than those obtained using the mixed-effects approach (see Table III).

1,000 generated data sets are presented in **Table IV**. For each data set, the population parameters were determined by taking both the mean and the median in the second stage (as was done for Table II), and subsequently their means and standard deviations over the 1,000 simulations were calculated. Based on the means, there is considerable bias in the estimates of α (3.17 minus 1.74) and β (3.63 minus 1.18). Based on the medians, a considerable bias for β (2.16 minus 1.18) remains. A possible cause for this bias is inter-individual variability, as it disappeared when we repeated the Monte Carlo simulations without inter-individual variability on β (note that it was the highest, with $\omega^2 = 1.72$ in table 1). The fact that the difference between the (mean) mean (3.63) and the (mean) median (2.16) of β is higher than obtained from the experimental data set (2.33 and 2.20) would suggest that β is not lognormally distributed in this set (by definition it was lognormally distributed in the Monte Carlo simulations). However, mixed-effects analysis with a normal distribution for β yielded a worse fit (judged by the value of the likelihood). The difference between the median mean (3.23) and the median median (1.96) (not shown) suggests that the discrepancy is not primarily caused by non-normality of the two-stage estimator of β, but by bias.** For a two-stage analysis, the individual estimates have to be averaged and the arithmetic means are likely to be biased estimates, as the distributions of the individual estimates are often not normal. On the contrary, the medians of the first-stage estimates are unaffected by assumptions concerning those distributions.

** Correlation matrices revealed a large correlation between β and θ for two-stage analyses (for $\theta \approx 1$, $y = \alpha + \beta t^{\theta} \approx \alpha + \beta\theta t$). Furthermore, the model is by necessity not centered with respect to the fixed effect (time), which may be another source of correlation.

Individualization and Prediction

When a population model has been established (i.e., the experimental data of a previously studied group of subjects are adequately described by a set of mathematical equations with identified parameter values and their uncertainties), it can be used to predict the performance of an as yet unstudied individual by means of Bayes posterior distribution estimation. Initially, the individual could be described using the population average response to the experimental circumstances. Then, as data for the individual are being obtained, the model parameters can be adjusted from their population values to the ones describing the individual at hand—even when only one or two data points for the new individual are available as yet (as we demonstrated for the example problem in the first part of the paper). Moreover, the individualized parameters can be used to make a one-day-ahead prediction (and beyond).

Fig. 3 and 4 show the results for three subjects. For the first day, the prediction of both the model parameters and the performance outcomes are set to equal the estimates of an analysis based on the population without these three subjects. As data for these individuals are obtained, the model parameters are updated taking

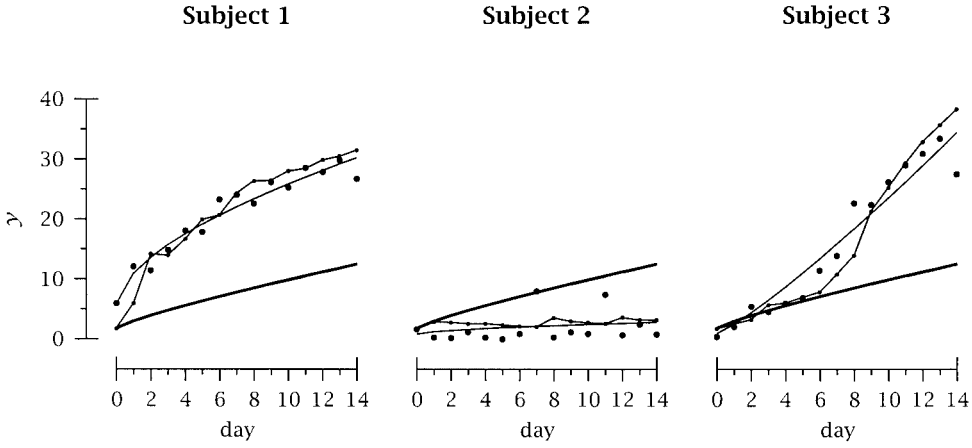

Fig. 3. Observations from three individuals (dots), model fits (thin solid lines), and population average model profiles (thick solid lines) based on the mixed-effects analysis, and one-day-ahead predictions (connected dots) adjusted each day as more data for the individual were obtained. The three panels illustrate the variety in vulnerability of these individuals to the effects of chronic partial sleep deprivation compared to the population average; subject 2 was much less vulnerable than subjects 1 and 3.

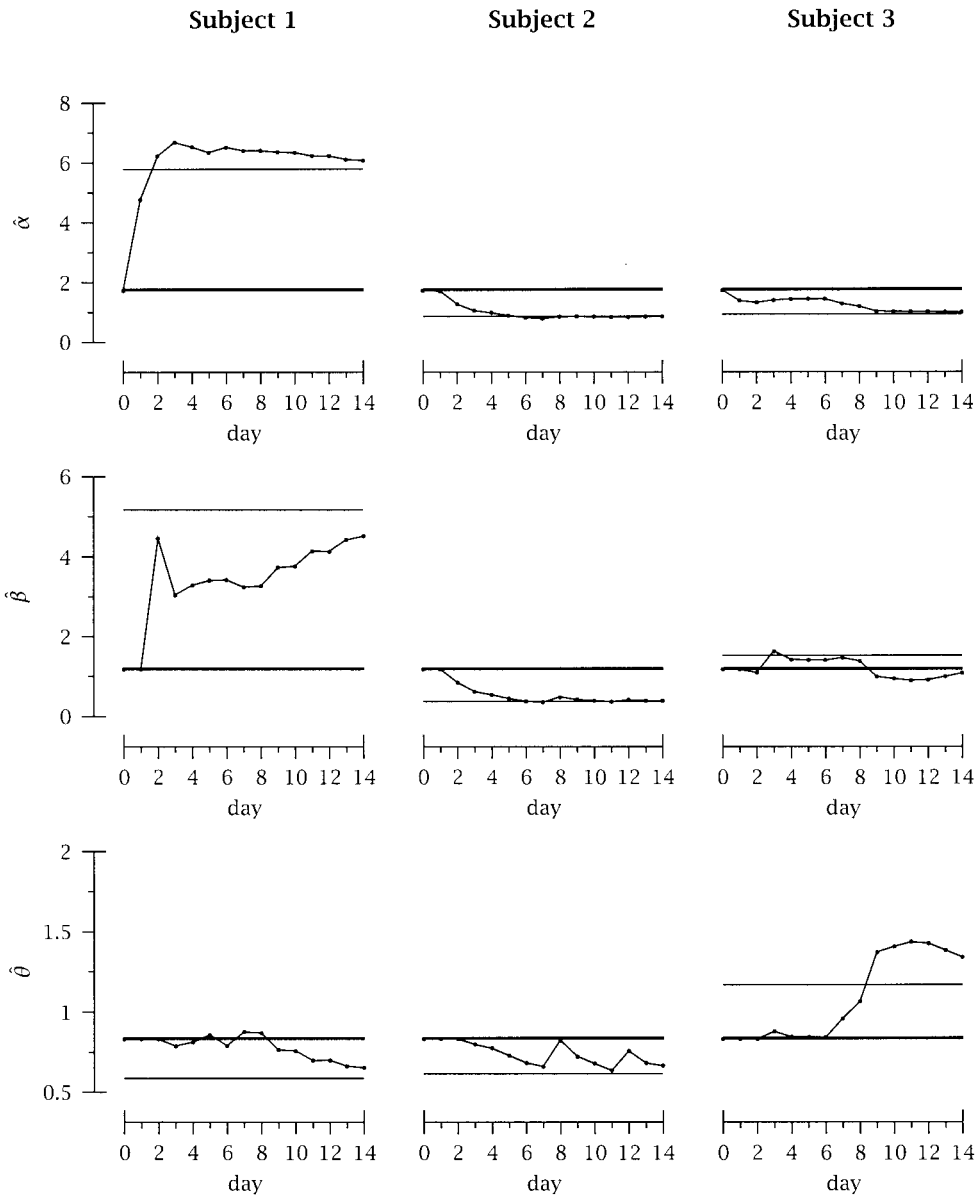

Fig. 4. One-day-ahead predictions (connected dots) for subject-specific estimates of α (upper panels), β (middle panels), and θ (bottom panels). The three panels per parameter correspond to the three individuals in Fig. 3. Also shown are the final parameter estimates for these individuals (using all 14 d of data; thin lines) and the population average estimates (thick lines). As the parameters of the model are tuned over the days and increasingly characterize the individual at hand, they allow improved prediction compared to the population average.

into account this subject-specific information. Thus, as data are obtained, the predictions converge to the values that characterize these individuals.[††]

For example, subject 1, presented in the left panels of Fig. 3 and 4, has a significantly higher gain factor β than the population average, for which correction appears to be easily achieved as more data for the individual are obtained. Subject 2, displayed in the middle panels, has lower values for all three parameters compared to the population averages, but one-day-ahead predictions remain reasonable. The performance of subject 3, displayed in the right panel of Fig. 3, closely follows the population average at first, but later deteriorates more rapidly as is reflected in this individual's estimate of θ. One-day-ahead predictions for subjects were about 10 units off at maximum. They obviously depend on the quality of both the model and the data, and may be further improved by including relevant covariates. The distribution of the differences between the one-day-ahead predictions and the corresponding actual observations provides an indication of the usefulness of the predictions. Research is in progress in order to obtain confidence intervals, which is a challenge due to the fact that model selection uncertainty also has to be taken into account (1). Finally, although we assumed lognormal distributions for the inter-individual variabilities of parameters α and β, and a normal distribution for θ, their true distributions remain unknown. In a mixed-effects analysis these assumptions serve as prior information; the distributions converge to the true ones as more data are included [recall that in Eq. 9 the likelihood $l_i(y_i; a, \sigma, \eta_i)$ sharpens up].

CONCLUSIONS

We illustrated the mixed-effects modeling theory by studying a simple model for which analytical results

[††] The one-day-ahead predictions for day t = 14 are based on days t = 0, . . . , 13 and may, therefore, differ slightly from final estimates based on all 14 days.

can be derived, and by numerically analyzing experimental data with a nonlinear model. We showed that by separating within- and between-subjects variance components, model parameters can be estimated more precisely and more accurately than with a two-stage approach. Furthermore, the mixed-effects modeling serves as a foundation for the prediction of subject-specific temporal profiles, something that has been needed in performance model development (3,10) for years. As recent studies have shown trait-like inter-individual differences in vulnerability to performance impairment from sleep loss (13), there is a clear advantage of applying mixed-effects modeling approaches in the development of future biomathematical models of fatigue and performance.

ACKNOWLEDGMENTS

This work was supported by NIH grants NR04281, HL70154, and RR00040, and NASA cooperative agreements NCC 2-1077 and NCC 2-1394 with the IEPRF.

REFERENCES

1. Burnham KP, Anderson DR. Model selection and multimodel inference: A practical information-theoretic approach. 2nd ed. New York; Springer-Verlag; 2002.
2. Davison AC, Hinkley DV. Bootstrap methods and their application. Cambridge; Cambridge University Press; 1999.
3. Dinges DF, Achermann P. Future considerations for models of human neurobehavioral function. J Biol Rhythms 1999; 14:598–601.
4. Dinges DF, Powel JW. Microcomputer analyses of performance on a portable, simple visual RT task during sustained operations. Behav Res Meth Instr Comp 1985; 17:652–5.
5. Hartford A, Davidian M. Consequences of misspecifying assumptions in nonlinear mixed effects models. Comp Statist Data Anal 2000; 34:139–64.
6. NONMEM Project Group. NONMEM user's guide. San Francisco, CA: University of California at San Francisco; 1999.
7. Pinheiro JC, Bates DM. Mixed-effect models in S and S-PLUS. New York: Springer-Verlag; 2000.
8. Rice JA. Mathematical statistics and data analysis, 2nd ed. Belmont, CA: Duxbury Press; 1995.
9. Sheiner LB, Beal SL. Evaluation of methods for estimating population pharmacokinetic parameters. J Pharmacokin Biopharm 1980; 8:553–72.
10. Van Dongen HPA. Comparison of mathematical model predictions to experimental data of fatigue and performance. Aviat Space Environ Med 2004; 75(3, Suppl): A15–36.
11. Van Dongen HPA, Maislin G, Dinges DF. Dealing with inter-individual differences in the temporal dynamics of fatigue and performance. Aviat Space Environ Med 2004; 75(3, Suppl.): A147–54.
12. Van Dongen HPA, Maislin G, Mullington JM, Dinges DF. The cumulative cost of additional wakefulness: Dose-response effects on neurobehavioral functions and sleep physiology from chronic sleep restriction and total sleep deprivation. Sleep 2003; 26:117–26.
13. Van Dongen HPA, Rogers NL, Dinges DF. Sleep debt: theoretical and empirical issues. Sleep Biol Rhythms 2003; 1:5–13.
14. Wahlby U, Jonsson EN, Karlsson MO. Assessment of actual significance levels for covariate effects in NONMEM. J Pharmacokinet Pharmacodyn 2001; 28:231–52.
15. Wolfinger RD. Fitting nonlinear mixed models with the new NLMIXED procedure. Cary, NC: SAS Institute Inc.; 2000.
16. Yano I, Beal SL, Sheiner LB. The need for mixed-effects modeling with population dichotomous data. J Pharmacokin Biopharm 2001; 28:389–412.