ADVANCES IN FOOD AND NUTRITION RESEARCH, VOL. 40

Chapter 17

ESTIMATION AND USE OF KINETIC PARAMETER DISTRIBUTIONS IN METABOLISM AND NUTRITION

WILLIAM F. BELTZ

Department of Medicine
University of California, San Diego
La Jolla, California 92093

I. INTRODUCTION

In the field of pharmacokinetics, there has been much recent work on developing methods for estimating interindividual variation in kinetic model parameters, particularly in sparse data situations where there are

sometimes fewer observations for each subject than there are parameters in the model. The underlying idea behind these methods is that fewer parameters need to be estimated if they are "population" parameters, which describe the distribution of the model parameters within a population of interest, rather than the kinetic model parameters for each of the individual studies. For example, a two-compartment model has 4 estimable parameters. If studies are performed in 100 subjects, there are 400 individual kinetic parameters to estimate. However, if a multivariate normal distribution is assumed, there are only 14 population parameters: 4 means, 4 variances, and 6 covariances. Several methods for "population kinetic analysis" have been developed. These differ primarily in the assumptions they make about the underlying distribution. Early methods assumed a multivariate normal distribution and estimated the population parameters listed in the example above. Recent work has concentrated on semiparametric and nonparametric approaches which make fewer assumptions and allow for more general distributions. These are particularly useful for detecting multimodality. Although most tracer studies in metabolism and experimental nutrition are designed to allow estimation of the individual kinetic parameters, the population methods are still useful in several situations, such as when only limited data may be obtained from a single animal or when there are missing data. The population methods discussed in this paper avoid making assumptions about missing data by simultaneously considering all the experiments in a rigorous and consistent way.

For the interested reader, there are several reviews of the theory and applications to pharmacokinetics and pharmacodynamics (Steimer *et al.*, 1985; Sheiner and Ludden, 1992), but so far, the methods have not been applied to tracer kinetic problems in metabolism and nutrition. The goal of this paper is to provide an introductory review of the theory, applications, and available software, with particular attention to how they relate to problems in metabolism and nutrition.

II. DEFINITIONS AND THEORY

We consider the situation in which a number of similar (but not necessarily identical) experiments are performed on a number of different subjects randomly selected from a larger population of individuals. Examples include a metabolic study in which isotopically labeled lipoprotein turnover is measured in a number of patients with heart disease, a nutritional study in which vitamin turnovers are examined in some randomly selected graduate students, or an agricultural study in which calcium kinetics are studied in a few dairy cattle randomly selected from a school's experimental herd. In

the first example, the studied patients are assumed to be representative of all heart disease patients. In the second example, the graduate students are assumed to be representative of the public at large. In the agricultural study, the experimental cattle are assumed to be representative of dairy cattle in general. In each experiment, the goal is to make some inference about the larger population. Although it might appear that the most frequent goal of mathematical modeling is to estimate the parameters for the experimental subjects, this is usually being done as the first step toward estimating the population parameters.

A. SYSTEM MODEL

The data obtained from the experiments are to be evaluated in terms of a mathematical model. Although this "system model" is only a hypothesis about how the system under investigation is thought to function, it must be mathematically well-defined so that it may be used to calculate quantitative model predictions. The system model will frequently take the form of a compartment model (Jacquez, 1985), but the ideas discussed here are applicable to other types of models as well. A very important aspect of the model is its parameters. We define the model response as $\eta(\beta,t,x)$, where β is a vector of system parameters, t is time, x is a vector of independent variables (other than time t), and the details of the model are embedded in the function η. For compartmental models, the parameters are usually the intercompartmental fractional transfer coefficients and perhaps a volume of distribution. The parameters of the system model are unknown constants that may take different values for each of the experiments. Since the experimental subjects are randomly selected from the population, the model parameters are random variables and a multivariate probability distribution exists for them. We will use $h(\theta,\beta,x)$ to denote the probability density for the system parameters, where θ is a vector of "population parameters," and x is again a vector of independent variables. The population parameters define the probability distribution h for the values of the system parameters. If h is a multivariate normal distribution, then the population parameters are the means, variances, and covariances of the distribution. We can now write the conditional probability for the values of the system parameters β given values for the population parameters θ and the independent variable x as $P(\beta|\theta,x) = h(\theta,\beta,x)$.

As mentioned above, the system parameters are unknown and constant. One of the primary goals of modeling is to estimate the system parameters. The fact that they are unknown distinguishes them from independent variables or covariates. Note that by this definition, a variable in a model equation may be a parameter in one analysis and an independent variable

in another. In one study for example, a volume of distribution may be estimated while in an otherwise similar study it may be known from independent measurements. The fact that the parameters must be constant will sometimes require reparameterization of a model. For example, consider a compartmental system in which the mass in the central compartment C is changing with time (nonsteady state) and the elimination rate from this compartment, $R_{0c}(t)$, is nonlinear. Specifically, let the elimination rate follow Michaelis–Menten kinetics and be $R_{0c}(t) = V_{\max}Q_c(t)/(Q_c(t) + K_m)$. Now the fractional transfer coefficient is $V_{\max}/(Q_c(t) + K_m)$ and the system parameters are $V_{\max}$ and K_m.

The distribution estimated by population analysis is often called the "prior" distribution because it is known (or assumed) before an experiment is performed. After an experiment, the new results may be incorporated to estimate a new distribution, the "posterior" distribution, which then becomes the prior for the next experiment!

B. EXPERIMENTS AND THE ERROR MODEL

We now consider a series of experiments performed on the model and the observations obtained during those experiments. The statistical model is the usual one

$$y_{ij} = \eta(\beta_j, t_{ij}, x_j) + \varepsilon_{ij},$$

where y_{ij} is the ith observation for the jth experiment, $\eta(\beta_j, t_{ij}, x_j)$ is the model-predicted value for y_{ij}, and ε_{ij} is the residual error for this observation. The values for y_{ij}, t_{ij}, and x_j are known and values for β_j are to be estimated. The residual errors are random variables and we assume they are distributed according to a probability distribution of known form and quantified by perhaps unknown parameter values. This distribution comprises the "error model" and its parameters are the "error model parameters." The usual assumption is that the errors are normally distributed with mean 0 and variance σ_{ij}^2. In general, however, other forms may be used.

C. LIKELIHOOD

To estimate the system parameters β, we need a quantity that measures how well our model fits the data. This is provided by the likelihood function. For experiment j, the likelihood l is defined as the probability of obtaining the actual observations, given the model, the model parameters, and the known independent variables. In this form, the likelihood is a function of the system parameters

$$l_j(\beta) = P(y_{ij}, \ldots, y_{nj}|\beta).$$

This likelihood for the jth experiment may be calculated by multiplying the probabilities for the individual observations

$$l_j(\beta) = \prod_{i=1}^{n} P(y_{ij}|\beta),$$

where the individual probabilities can be calculated from the error model. The model parameters may be estimated by finding the values that maximize this function. If the errors are normally distributed, maximizing this likelihood function is equivalent to minimizing the weighted least-squares function.

D. POPULATION KINETIC ANALYSIS

To make sense of sparse, "routine" pharmacokinetic data, Sheiner *et al.* (1977) recognized that fewer parameters need to be estimated if one estimates the distribution of parameters rather than the individual parameter values for each experiment. The fundamental idea behind the population methods is that if one knows the parameter distribution $h(\theta)$, then a likelihood L for an experiment may be calculated which does not depend on actual parameter estimates for the experiment

$$L_j(\theta) = \int l_j(\beta)h(\beta,\theta)d\beta.$$

This likelihood is calculated by weighting the parameter-dependent likelihood for each parameter value by the probability of the system parameters taking on that value. This likelihood is no longer a function of the system parameters but is a function of the population parameters. A likelihood for a several experiments ($L(\theta)$) may be calculated as a product of the likelihoods for the individual experiments. The job of the various population kinetic analysis algorithms is to estimate $h(\beta,\theta)$ by finding the θ that maximizes $L(\theta)$.

III. USES FOR PRIOR PARAMETER DISTRIBUTIONS

A. BAYESIAN ESTIMATION

Bayesian estimation is used to estimate the model parameters, which would otherwise be unidentifiable, by taking into account the prior distri-

bution and using this information to determine the most probable value for each parameter. The estimation is performed by finding values for system parameters for the jth experiment (β_j) that maximize $P(y_j|\beta_j,x_j) \times h(\beta_j,\theta,x_j)$, where y_j is the vector of observations for the jth experiment and x_j is the vector of independent variables for the jth experiment. In a data-rich situation, $P(y_j|\beta_j,x_j)$ dominates this calculation while in a data-poor situation, $h(\beta_j,\theta,x_j)$ dominates.

B. ADAPTIVE CONTROL

A frequent use of the prior parameter distribution in pharmacokinetics is for adaptive control (Schumitzky, 1986). The idea is to provide patients with individualized dosing to reach some therapeutic goal. Frequently, the goal is to maintain the blood level of a drug within some range but many other goals are possible. When a patient is treated for the first time, the appropriate kinetic parameters are unknown. The most likely parameter value for the subject may be determined from the prior distribution and the appropriate dosage estimated and applied. The "adaptive" part of the method occurs after as few as one drug level has been observed for this subject. New parameter estimates may then be obtained by Bayesian estimation and the dosage adjusted as needed. At the start, dosing is based on the best estimates from previous subjects but as therapy progresses, treatment is based more on estimates of the subject's own parameter values.

C. EXPERIMENTAL DESIGN

Experimental design is the optimization of experimental "controls" so as to maximize the likelihood of achieving a desired result. There are several goals that a researcher might wish to achieve, including minimizing the uncertainties for the estimated parameters, maximizing the ability to control specific responses of the system, and maximizing the ability to discriminate between two or more models. For any of these goals, accounting for parameter variability will provide better results or at least make the researcher aware of the range of results that might be expected.

D. MONTE CARLO SIMULATIONS

By definition, a Monte Carlo analysis is the generation of several randomly generated simulations, usually with a digital computer. There are several reasons for doing this, including evaluating statistical procedures, estimating the power of a statistical test, and verifying predictions made by other means. Clearly, to produce a set of random experiments, some

aspect of each experiment must be drawn from a probability distribution. If an estimate for the parameter distribution within the population is available, then this can be done easily. Parameter values are selected randomly from the prior distributions for each simulated experiment. The model response is then determined and random error added to the predicted values consistent with the error model. These "observed" data may then be fit using the same or a different model, depending on the goal of the analysis. After many simulations, the distribution of results (parameter estimates, response variable values, or both) may be evaluated in light of the "known" true system.

IV. APPLICATIONS TO METABOLISM AND NUTRITION

Although the mathematics is the same, kinetic modeling in metabolism and experimental nutrition differs from pharmacokinetics in at least two ways. First, while the modern methods of population kinetic analysis were developed specifically for the sparse data problem, in metabolism and nutrition we usually have adequate data for each experiment. Second, while pharmacokinetic analysis tends to be control-oriented, the ultimate goal being optimal dosing, metabolic modeling tends to be structure-oriented, the goal being to determine the true structural model. Nevertheless, as we will see, population kinetic analysis can be appropriate in some relatively common situations.

A. MISSING VALUES

In the metabolic world, we are not usually dealing with the sparse data encountered by pharmacokineticists. We usually have the two-stage methods (see below) in the backs of our minds and design turnover studies that allow good identification of all parameter values for each experiment. Nevertheless, even the best planned studies sometimes go afoul. Consider the situation where the protocol calls for a final blood sample at the very end of a study. This observation might be crucial for accurately estimating the terminal slope of the plasma curve. If this sample is missed or lost, it can render the model poorly identifiable. What is frequently done in practice is that some assumption is made about the parameter values for this individual. For our example, this might be the slope of the terminal slope or the relative size of a peripheral compartment or exchange pool. The assumed value is based upon the other experiments. The strategy is therefore to model all the experiments with complete data, fix a parameter or parameters in the incomplete experiment to the mean of those obtained for the other

experiments, and identify the remaining parameters. Although this is a common and useful procedure, there are some problems with it. First, the effect of the assumption is seldom considered, let alone examined in any detail. Second, it does not consider the uncertainty in the assumed value. Since the value of this parameter is fixed, other parameters for this experiment will often have lower coefficients of variation (fractional standard deviations) than the same parameters from experiments without the constraint. Thus values from this experiment appear to be identified with more certainty, when in fact they should be considered more suspect since they are based upon an additional assumption. Finally, if another experiment has missing data for a different part of the curve and a different parameter is to be assigned an assumed value, one can get into a loop where results from experiment A are used to fix a parameter for experiment B and results from experiment B are used to fix a parameter for experiment A. Although this procedure will probably converge, this is exactly the type of situation for which the population analysis methods were developed. In fact, they were created to specifically handle the extreme case, that in which all experiments have missing values!

B. SPARSE DATA

There are good examples from the metabolic literature of studies in which the number of data observed for a single subject is limited. An example is measuring the mean residence time of low density lipoprotein in the rabbit aortic wall (Schwenke and Carew, 1989). In experiments such as these, samples of the aorta may only be obtained once, at the end of the experiment. Thus there is only one datum for each tracer used. Schwenke and Carew (1989) used two iodine tracers, administered at different times, but the compartmental model they used has four parameters. Clearly, parameter values cannot be estimated for each animal without using information from experiments in other animals.

C. COVARIATES

The x in $h(\beta,\theta,x)$ is sometimes called a "covariate" (Wade *et al.*, 1994) and the estimation of covariate effects is one of important goals in population kinetic analysis (Sheiner and Ludden, 1992). In pharmacokinetics, a common goal of covariate analysis is to get reliable predictors of drug clearance. The value should be obvious. If one knows how the residence time of a drug in the body depends on age, gender, body weight, renal function, and so on, then one has a good starting point from which to

individualize dosing. In metabolism and nutrition, one might similarly be interested in how absorption rates, fractional catabolic rates, and other system parameters depend not only on age, gender, and body weight, but on dietary status, genotype, and even species. Since the covariate acts directly on the parameter distribution, the two must be estimated simultaneously. This has the additional advantage that all sources of variation are included in the estimation procedure.

V. ESTIMATION OF PRIOR PARAMETER DISTRIBUTIONS

A. NAIVE DATA POOLING

In naive data pooling, observations from different subjects are pooled and treated as though they were obtained from a single experiment. This technique is fairly common but as Steimer *et al.* (1985) point out, the results obtained can be misleading. Averaging of data can either average out phenomena in the data that may be of interest or can suggest complexity that is really due to interindividual variation. No special software is required for naive data pooling. Any of several spreadsheets or statistical packages may be used to average the data and any of several nonlinear regression programs may be used to model the mean values.

B. TWO-STAGE METHODS

In the so-called two-stage methods, each experiment is first modeled individually to obtain parameter estimates and these estimates are then used to make inferences about the population distribution. There are two recognized two-stage methods. The standard two-stage method (STS) ignores uncertainties in the individual estimates while the global two-stage method takes them into account.

1. Standard Two-Stage Method

This is perhaps the most common method of estimating a parameter distribution, though some researchers may not be aware that this is what they are doing. It refers to modeling each experiment individually to obtain point estimates, then treating these estimates as error-free and performing standard statistical analyses on them. It does not consider that there are uncertainties in the modeling results. In particular, it doesn't allow for differences between experiments in their relative impact on the results. In the STS method, covariate effects are determined by regression and

correlation analysis on the point estimates. No special software is required for the standard two-stage method. Any of several nonlinear regression programs may be used to model the individual experiments and any of even more statistical programs may be used to manipulate the point estimates obtained.

2. *Global Two-Stage Methods*

The global two-stage method attempts to account for the uncertainty in the individual parameter estimates and considers correlations among the various model parameters. In essence, it bases the population mean on a weighted average of the individual values. The most readily available software for a global two-stage analysis is the Extended Multiple Studies Analysis (EMSA), available in the SAAM/CONSAM package (Berman and Weiss, 1978). Details of EMSA are provided by Lyne *et al.* (1992). EMSA assumes that the distribution is normal and, as a result, EMSA will not detect multimodal distributions. As with other two-stage analyses, each experiment is first studied individually. During this stage, SAAM writes the estimated parameter values, the estimated standard deviations for the parameter values, and the parameter correlation matrix to a file. The results for the several experiments are collated into a single file which serves as input for the second step. During this stage, means and variances for the population distribution are estimated and printed. Since the parameter values and correlations for the first stage are generated by SAAM, a wide variety of compartmental and algebraic models are available for analysis. SAAM is distributed in executable binary form and is available for VAX/VMS, Unix, PC/DOS, and Macintosh. It is available from Loren Zech at the National Institutes of Health (Building 10, Room 6B-13, National Institutes of Health, Bethesda, MD 20892).

C. POPULATION KINETIC ANALYSIS

The population analysis methods use all the available data to estimate the population. The best estimates for the parameters of an individual study are only obtained after the population distribution has been estimated by Bayesian estimation. Essentially, the various methods estimate the population parameters θ in $h(\beta,\theta,x)$. The methods differ primarily in the form that $h(\beta,\theta,x)$ is assumed to have. Despite the fact that all arrive at a quantitative description of $h(\beta,\theta,x)$, the different forms have been divided into parametric, semiparametric, and nonparametric. Each of these will be described.

1. *Parametric Methods*

The parametric methods assume that $h(\beta,\theta,x)$ is multivariate normal. The population parameters (θ) therefore correspond to means, variances, and covariances. The most used software package for population analysis is NONMEM (NONMEM Project Group, 1992), which stands for "nonlinear mixed effects models" and uses a parametric method. The theory behind NONMEM is described by Beal and Sheiner (1982). NONMEM is distributed as Fortran code and will therefore run on a wide variety of computers though the user must be prepared to deal with editing the code and compiling the program. Although NONMEM includes a library of the most common pharmacokinetic models, for any but the simplest of metabolic models, the user will probably have to write a subroutine to calculate the model response. NONMEM is available from the NONMEM Project Group at the University of California, San Francisco (mail stop C255, San Francisco, CA 94143).

2. *Semiparametric Methods*

One limitation to the parametric approach is that any multimodal distribution will be approximated by the unimodal normal curve and an important characteristic of the population could be overlooked. Several alternative forms of h have therefore been suggested. The class of so-called semiparametric methods allows for forms of $h(\beta,\theta,x)$ more general than the normal distribution but still places some limitations on the structure of this function. The only readily available program for semiparametric analysis appears to be NLMIX (Davidian and Gallant, 1992), which also stands for "nonlinear mixed effects models." The method used is described as smooth nonparametric estimation (Davidian and Gallant, 1992). The class of density distributions actually allowed by the method essentially consists of a multivariate normal distribution multiplied by a polynomial, allowing multimodal, fat-tailed, and skewed densities, none of which are allowed by the parametric methods. The population parameters (θ) correspond to the means, variances, and covariances of the normal distribution, as well as the coefficients of the polynomial. Like NONMEM, NLMIX is distributed as Fortran code. To use the package, the user must be prepared to do some Fortran programming, more than with NONMEM. The user must again supply a subroutine to calculate the model response, but there is no model library included with NLMIX. The source code and documentation are available by anonymous ftp from the StatLib statistical software collection at Carnegie-Mellon University. The easiest access is to make a world wide

web (www) connection to lib.stat.cmu.edu and look in the "general" archive.

3. *Nonparametric Methods*

The nonparametric methods (Mallet *et al.*, 1988) make no assumptions about the form of $h(\beta,\theta,x)$. The estimated probability density function takes the form of a finite number of allowable parameter vectors, the "points of support," and a probability assigned to each point of support. Although there are an unlimited number of parameter values the estimation procedure can choose from, the final result is really a discrete multivariate probability distribution. The population parameters (θ) correspond to the parameter values at the points of support and the probability assigned to each. Thus the "nonparametric" distribution is quantitative and described by a finite number of variables. After estimation, any of various smoothing techniques can be used to make the function more presentable, but calculations using the distribution are based on the unsmoothed form. Calculations can actually be faster than with other distributions because rather than integrating over a continuous function, you only need to sum over the points of support.

The most readily available nonparametric software is in the USC*PACK collection of pharmacokinetic computer programs for population analysis, Bayesian estimation, and individualized dosing (Jelliffe, 1991). The rationale for the package as a whole is described in detail by Jelliffe (1986). Included in the package is the NPEM2 program, which performs nonparametric population analysis on a three-compartment model using the EM algorithm (Schumitzky, 1991). The USC*PACK programs are distributed in executable binary form for PC/DOS. The programs are available from the Laboratory of Applied Pharmacokinetics at the University of Southern California (USC School of Medicine, CSC 134-B, 2250 Alcazar Street, Los Angeles, CA 90033).

VI. IDENTIFIABILITY ISSUES

Unfortunately, there has been minimal work on identifiability issues with respect to population kinetic analysis. The current work on identifiability of kinetic models (Jacquez, 1985; Cobelli and DiStefano, 1980; Cobelli and Saccomani, 1990) focuses on estimation for an individual experiment. With regard to population analysis and Bayesian estimation of parameters for an individual from the population, it is the population analysis step for which identifiability issues need to be considered. Once a prior distribution

is available, Bayesian estimates can be calculated for any individual, even if there are no data from that subject. In this case, the values will be the modes of the prior and identifiability is not a problem. It is the estimation of the first prior that could be troublesome. It is also this estimation that the various procedures discussed in this paper attempt to do.

Despite the lack of published theory, a couple reasonable points can be made about the identifiability of the multivariate parameter distribution. First, the system model must be identifiable for a single data-rich experiment. This means that if all the data were presumed to be from a single experiment, then the parameters for that experiment must be uniquely identifiable. Second, there must be some "information" about each parameter contained within the data set. This means that all the data together must cover the range required for an individual experiment to be identifiable.

Inclusion of the covariates in the estimation step introduces some additional complications. In a recent report, Wade *et al.* (1994) demonstrated with simulated data that if significant covariates were not included in the analysis, then what was in reality a monoexponential response could be interpreted as a multiexponential response. Variation between experiments was interpreted as variation within the structural model. In a sense, this is an example of using the wrong model. Nevertheless, it demonstrates some of the dangers in population analysis of questionably identifiable models and that the presence of the covariates in the population model makes the situation more complicated from an identifiability standpoint.

In metabolism and nutrition, where each experiment has been designed to be complete and population analysis is used to fill in missing values and to incorporate relative uncertainties into the estimation, a good procedure would be to first examine in detail those individual studies which are the most complete. This will familiarize the user with the behavior of the model, produce initial estimates for the system parameters, provide a chance to verify that these values are reasonable, and allow the use of tools for the identifiability of individual experiments (Jacquez and Perry, 1990). After this exercise has been complete, all experiments, including those that are incomplete, can be pooled for population analysis and testing the effects of covariates. If required, the final step would be to use the estimated distributions to obtain Bayesian parameter estimates for the individual experiments. This procedure should yield the most appropriate estimates for the incomplete experiments.

Finally, it must be pointed out that it is ultimately the responsibility of the user to be aware of identifiability issues. The ability of a computer program to provide parameter estimates should not always be interpreted as proving that the model is identifiable (Cobelli and Saccomani, 1990).

VII. CONCLUSIONS

The population estimation methods developed by pharmacokineticists to handle sparse data have potential use in the relatively data-rich studies encountered in metabolism and experimental nutrition. They provide a consistent and logical method to combine information from several experiments, accounting for interindividual variation in a theoretically sound manner. Anyone working in kinetic analysis should be aware of these tools, their advantages, and their limitations.

REFERENCES

Beal, S. L., and Sheiner, L. B. (1982). Estimating population kinetics. *CRC Crit. Rev. Biomed. Eng.* **9,** 195–222.

Berman, M., and Weiss, M. F. (1978). "SAAM Manual," DHEW Publ. No. (NIH) 78-180 U.S. Govt. Printing Office, Washington, DC.

Cobelli, C., and DiStefano, J. J., III (1980). Parameter and structural identifiability concepts: A critical review and analysis. *Am. J. Physiol.* **239,** R7–R24.

Cobelli, C., and Saccomani, M. P. (1990). Unappreciation of a priori identifiability in software packages causes ambiguities in numerical estimates. *Am. J. Physiol.* **258,** E1058–E1059.

Davidian, M., and Gallant, A. R. (1992). Smooth nonparametric maximum likelihood estimation for population pharmacokinetics, with application to quinidine. *J. Pharmacokinet. Biopharm.* **20,** 529–556.

Jacquez, J. A. (1985). "Compartmental Analysis in Biology and Medicine." Univ. of Michigan Press, Ann Arbor.

Jacquez, J. A., and Perry, T. (1990). Parameter estimation: Local identifiability of parameters. *Am. J. Physiol.* **258,** E727–E736.

Jelliffe, R. W. (1986). Clinical applications of pharmacokinetics and control theory: Planning, monitoring, and adjusting dosage regimens of aminoglycosides, lidocaine, digitoxin, and digoxin. *In* "Topics in Clinical Pharmacology and Therapeutics" (R. F. Maronde, ed.), pp. 26–82. Springer-Verlag, New York.

Jelliffe, R. W. (1991). The USC*PACK PC programs for population pharmacokinetic modeling, modeling of large kinetic/dynamic systems, and adaptive control of dosage regimens. *In* "Fifteenth Annual Symposium on Computer Applications in Medical Care" (P. D. Clayton, ed.), pp. 922–924. McGraw-Hill, New York.

Lyne, A., Boston, R., Pettigrew, K., and Zech, L. (1992). EMSA: A SAAM service for the estimation of population parameters based on model fits to identically replicated experiments. *Comput. Methods Programs Biomed.* **38,** 117–151.

Mallet, A., Mentré, F., Steimer, J.-L., and Lokiec, F. (1988). Nonparametric maximum likelihood estimation for population pharmacokinetics, with application to cyclosporine. *J. Pharmacokinet. Biopharm.* **16,** 311–327.

NONMEM Project Group (1992). "NONMEM Users Guides." University of California, San Francisco.

Schumitzky, A. (1986). Stochastic control of pharmacokinetic systems. *In* "Topics in Clinical Pharmacology and Therapeutics" (R. F. Maronde, ed.), pp. 13–25, Springer-Verlag, New York.

Schumitzky, A. (1991). Nonparametric EM algorithms for estimating prior distributions. *Appl. Math. Comput.* **45,** 143–157.

Schwenke, D. C., and Carew, T. E. (1989). Initiation of atherosclerotic lesions in cholesterol-fed rabbits. II. Selective retention of LDL vs. selective increases in LDL permeability in susceptible sites of arteries. *Arteriosclerosis (Dallas)* **9,** 908–918.

Sheiner, L. B., and Ludden, T. M. (1992). Population pharmacokinetics/dynamics. *Annu. Rev. Pharmacol. Toxicol.* **32,** 185–209.

Sheiner, L. B., Rosenberg, B., and Marathe, V. V. (1977). Estimation of population characteristics of pharmacokinetic parameters from routine clinical data. *J. Pharmacokinet. Biopharm.* **5,** 445–479.

Steimer, J.-L., Mallet, A., and Mentré, F. (1985). Estimating interindividual pharmacokinetic variability. *In* "Variability in Drug Therapy: Description, Estimation, and Control" (M. Rowland, L. B. Sheiner, and J.-L. Steimer, eds.), pp. 65-111, Raven Press, New York.

Wade, J. R., Beal, S. L., and Sambol, N. C. (1994). Interaction between structural, statistical, and covariate models in population pharmacokinetics analysis. *J. Pharmacokinet. Biopharm.* **22,** 165–177.