ACHIEVING CONCENTRATION GOALS USING PARAMETRIC PHARMACOKINETIC MODELS - A CLINICAL REVIEW OF THE CURRENT UNIMODAL GAUSSIAN BAYESIAN APPROACH.

Roger Jelliffe, Laboratory of Applied Pharmacokinetics, USC School of Medicine.

This approach is the standard one when one uses parametric compartmental pharmacokinetic (PK) models. The usual parameter values are either the means or medians as the measures of the central tendency, and the standard deviations (SD's) as the measures of dispersion. The usual distribution is assumed to be the common Gaussian bell-shaped curve, or a lognormal distribution. Usually the mean is used as the measure of the central tendency, and the distribution is assumed to be symmetrical about it. This approach was introduced to the pharmacokinetic community by Sheiner [1], and is one of his group's most significant contributions to the field.

The (single) most likely values for each parameter (volume of distribution, rate constants, clearance, etc.) are then used to compute the dosage regimen to achieve the desired response (usually a selected target serum concentration), which is best individualized for each patient according to his/her perceived need for the drug and the risk of toxicity which is felt to be acceptable in order to obtain the most benefit from the drug. The regimen to achieve and maintain the target goal is computed and the future concentrations are predicted using these parameter values.

As data of feedback, usually in the form of measured serum concentrations, is obtained, the parameter values are revised using Bayes' theorem. The following objective function is minimized:

$$\sum \frac{(Cobs - C\ mod)^2}{Var\ (Cobs)} + \sum \frac{(Ppop - Pmod)^2}{Var(Ppop)} \qquad (1)$$

where C_{obs} is each observed serum concentration, C_{mod} is the concentration in the Bayesian fitted model, Var(Cobs) is the SD^2 of each observed concentration, P_{pop} is each population (mean) parameter value, P_{mod} is each fitted maximum aposteriori probability (MAP) Bayesian posterior parameter value in the model, and Var(Ppop) is the SD^2 of each population parameter value. This procedure is a special example of weighted nonlinear least squares fitting (see below) in which two types of data, the serum data and the population parameter values, are placed

together in the objective function. There are two sets of data points: one or more measured serum concentrations, along with the set of population parameter values. Each data point has its own SD. The fitting procedure has no information as to whether the data point is a parameter value or a serum concentration. All it sees are the various data points and their respective variances. Note that the SD's of the serum concentration and parameter value data points do the very important job of determining their relative credibility, and of determining just how much the fitting procedure proceeds toward the measured serum data or hangs back toward the population parameter values.

The MAP Bayesian fitting procedure has been shown to be somewhat better in predicting future serum concentrations than the method of weighted nonlinear least squares, which does not have the right-hand term of the above objective function, but only the serum data. MAP Bayesian fitting is also significantly better than the earlier traditional but now obsolete method of linear regression on the logarithms of the concentrations (see below).

The MAP Bayesian fitted model is then used to simulate the system to reconstruct the past behavior of the drug in that patient, using the patient's individualized pharmacokinetic model. Usually it is possible to do this over the patient's entire dosage history, especially when using a pharmacokinetic model that can accommodate changes in volume of distribution and the rate constant for elimination from dose to dose. The plot is compared with the patient's clinical behavior during the same time, thus permitting an evaluation of the patient's individual clinical sensitivity to the drug. One can re-evaluate the appropriateness of the original target goal, and can choose a different one if needed. After selecting the goal, the regimen to achieve it is computed, and the system behavior in the future is predicted using the fitted model.

Comparison with other Methods: Nonlinear and Linear Least Squares

The conventional weighted least squares procedure is not quite so smart, as its objective function has only the left hand side of the MAP Bayesian objective function, as shown below.

$$\sum \frac{(\text{Cobs} - \text{C mod})^2}{\text{Var (Cobs)}} \qquad (2)$$

Because of this, only the patient's serum data are considered in the fitting procedure, and this information is not supplemented by the additional population parameter values which represent general information of how the drug has behaved in other similar people in the past.

Because of this, fitted models made using weighted nonlinear least squares have been shown to predict future serum concentrations slightly less well than those made using MAP Bayesian fitting [2].

Weighted Nonlinear least Squares Regression

Like the MAP Bayesian procedure, this method can fit the model to data of doses and serum levels acquired over many dose intervals. There is no longer any reason to do the traditional "single dose" pharmacokinetic study. Studies can be done on the actual patients being treated, as they are receiving their therapy. The algorithm of Nelder and Mead [3] is a good one for fitting the data in both the least squares and the MAP Bayesian fitting procedures. A useful nonmathematical description of this method has appeared in BYTE magazine [4].

Secondly, like the MAP Bayesian method, weighted nonlinear least squares can provide correct weighting of serum level data according to its credibility or Fisher information [5]. It thus has the potential for obtaining good estimates of the pharmacokinetic parameter values.

However, this method, has a weakness. It cannot take into account population information that is generally known about how that drug usually behaves in patients like the individual under consideration. As the procedure moves from the starting population parameter values to others which fit the data better, it discards the general information used to begin the fitting procedure instead of supplementing it with the individual patient's data. Since no fitting procedure ever explains the entire relationship between doses given and levels found, discarding the general population information is a suboptimal feature. It may well be because of this feature that the nonlinear least-squares method, while "fitting" serum level data "best", has been shown to be a slightly poorer predictor of subsequent serum levels [2]. This method, like linear least squares, below, requires at least one serum level for each parameter to be fitted, or at least two serum levels in the models considered here, as will be discussed further below. The MAP Bayesian method, in contrast, can fit using as few as a single serum concentration data point. This is because the MAP Bayesian procedure already has one data point for each parameter. They are the collection of population parameter values themselves. The MAP Bayesian procedure therefore can start to fit with only a single serum concentration.

Linear Least Squares Regression

Another method used to fit serum concentrations has been the old traditional but now obsolete method of linear regression on the logarithms of the serum concentrations (see below). This method was the traditional one in which a pharmacokinetic model (restricted to only a single compartment) was fitted to data obtained only during a single dose interval, and specifically to the logarithms of the serum concentrations. No weighting was used. It was simple, and was widely implemented on hand calculators. It was generally the community standard for monitoring serum gentamicin levels ever since Sawchuk and Zaske showed its utility to individualize aminoglycoside dosage regimens [6].

The method requires at least 2 serum levels. It cannot handle anything more than a 1-compartment model. It takes advantage of the fact that one can linearize the solution of a first-order linear differential equation for such a model if one transforms the serum level values to their logarithms. However, the method has three important weaknesses.

First, the method can only fit serum level data acquired during a single dose interval. It discards all previous serum data (and all previous information about the patient) whenever a new set of serum levels is obtained. There is therefore a loss of continuity each time new serum data are analyzed. This method is the most wasteful of any in its use of serum levels, as the useful life span of a serum level value is shorter than with any of the other methods which do not have to discard older data, but can integrate it with more recent data from other dose intervals as nonlinear least squares and the MAP Bayesian procedure can do.

Second, linear regression contains the assumption that the assay error is a constant percent of the measured concentrations. The lower the level, the more accurately it is assumed to be known. Because of this, if the assay has any other error pattern over its working range (and it almost always does!), this method greatly overestimates the credibility of low serum levels over high ones. This can be seen if one considers two serum levels, one of 8.0 ug/ml for example, and one of 1.5 ug/ml, as shown in Figure 1. One usually wishes to attach approximately equal credibility (weight) to these data points. One might thus assume that their laboratory error is approximately equal. Since the Fisher information (an index of credibility) of a data point having a normally distributed error is proportional to the reciprocal of the variance of that data point [5], the relative weights given by linear regression to serum levels of 8.0 and 1.5 ug/ml would be proportional to the reciprocal of their squares [5]. Because of this, the method of linear least squares, which assumes that the error bars are equal on the log scale, arbitrarily gives the value of

1.5 ug/ml a weight of $8^2/1.5^2 = 64/2.25 = 28.4$ times the weight of the level of 8.0 ug/ml. A level of 0.1 has 100 times the weight of a level of 1.0, and 1000 times the weight of a level of 10.0. Because of this assumption, the error pattern is often quite unrealistic, and results in parameter values that are significantly different from those obtained by other methods [2].

Third, this method ignores all population data, and therefore all past experience, concerning the behavior of the drug.

Comparison of the Methods

The MAP Bayesian method [1] appears to be the best of these three [2]. As with nonlinear least squares, it can provide correct weighting of serum level data according to the known laboratory assay error, and it can analyze such data over many dose intervals. In addition, it supplements population data (general knowledge) with specific information about each patient, instead of discarding it. Because of this, the method has been a slightly better predictor of future serum levels [2]. Lastly, the method requires only a single serum level to begin the analysis, no matter how many parameters are present in the population pharmacokinetic model. As more serum levels are obtained, the fitted model gradually becomes less of a population model and more of a patient-specific model. Both general and patient-specific data are combined intelligently in the M.A.P. Bayesian procedure to provide the most probable single-point estimates of the parameter values given both types of data and their respective standard deviations.

Finally, one other fitting procedure, now coming on the scene, holds promise of doing better than the MAP Bayesian method. This is the "Multiple Model" method of dosage design [7]. It is a stochastic rather than a deterministic method, and is based on nonparametric population and individualized pharmacokinetic models. It will be discussed more fully in another paper in this collection.

Examples of MAP Bayesian Model - Based Approaches

<u>Gentamicin therapy</u>

With a 1-compartment pharmacokinetic model in which the elimination rate constant (Kel) was composed of a nonrenal component (Knr) and a renal component having a slope (Kslope) relationship to CCr so that Kel = Knr + Kslope x CCr, the MAP Bayesian procedure resulted in significantly better prediction of future serum concentrations (see Figure 2) than those made using

linear regression (Figure 3). In contrast to most patients in the literature, who may have either normal or reduced renal function but whose renal function is stable, many patients in the above study were highly unstable and had changing renal function, to a quite significant degree, during their therapy [2].

Because the software used in that study [2] was designed to operate in the presence of significant changes in renal function from dose to dose, it has also been useful in the analysis and management of aminoglycoside therapy for patients who must undergo periodic hemodialysis.

Amikacin Therapy

MAP Bayesian adaptive control has been used to manage amikacin therapy in geriatric patients, often for extended periods, by Maire et al [8]. In their patients, whose renal function was often quite reduced but who were generally stable, visibly better prediction (and therefore control) of serum levels was seen with MAP Bayesian analysis than with their unfitted population model, in contrast to the more unstable patients receiving gentamicin described above [2]. These results are shown in Figure 4. They are better than those found in the gentamicin patients with unstable renal function [2] shown in Figure 2 above. Further, Figure 5 shows the poorer predictions based simply on the population model for Amikacin, without any fitting to the serum data.

Vancomycin Therapy

Vancomycin therapy was evaluated by Hurst et al [9] using a Kslope 2 compartment (central plus peripheral compartment) model. Using traditional linear regression, extremely poor prediction was found, as shown in Figure 6. In contrast, the 2 compartment model, coupled with Bayesian fitting, led to significantly better prediction of future serum levels than did the linear regression method, as shown in Figure 7.

Digoxin Therapy

The digoxin population model used in the USC*PACK MAP Bayesian software [10] is based on that described by Reuning, Sams, and Notari [11]. That two - compartment model uses both a central (serum) and a peripheral (nonserum) compartment. Computed concentrations of drug in the peripheral compartment correlate much better with inotropic effect than do serum levels [11]. The USC*PACK digoxin software not only uses this model, but also develops dosage

regimens to control either the peak peripheral compartment or the central serum compartment concentration.

Use of this MAP Bayesian software to manage digoxin therapy is illustrated by the following example. A 58 year old man developed rapid atrial fibrillation at another center, after missing his usual daily dose of 0.25 mg. He was clinically titrated with several intravenous doses of digoxin, and converted to sinus rhythm. He was then placed back on his original oral maintenance dosage. After a day, atrial fibrillation recurred, showing that his digoxin requirements had changed. He again was titrated with several doses of intravenous digoxin and again converted to sinus rhythm. Once again, he was placed on his original oral maintenance dosage, and once again, after about two days, atrial fibrillation recurred. For a third time he was titrated with several intravenous doses of digoxin, and for a third time he converted to sinus rhythm. A week of hospital time had been consumed during this phase of his care.

At this point the MAP Bayesian digoxin software was used to analyze his situation. Data of three serum levels, all taken during the post-distributional phase after a dose, showed almost no correlation with the patient's clinical behavior. As shown in Figure 8, he was in atrial fibrillation when the first serum level of 1.0 ng/ml was obtained, and was in sinus rhythm when the second and third serum levels of 1.0 and 1.2 ng/ml were obtained. However, when the 2 - compartment digoxin population model was fitted to the data of his various doses and these serum levels, the resulting fitted model, shown in Figure 8, was very informative.

Relating this fitted model to the patient's clinical behavior, sinus rhythm was present whenever peripheral concentrations were 10.0 to 13.0 ug/kg. Based on this, a therapeutic goal of 11.5 ug/kg was chosen for the desired peripheral compartment peak body concentration. The resulting regimen was 0.25 mg for the first day, and then averaged 0.57 mg/day. He was placed on a maintenance regimen of 0.5 and 0.625 mg on alternating days. On this regimen he was able to leave the hospital in sinus rhythm, and was still in sinus rhythm without evidence of toxicity when seen in the clinic 2 weeks later.

Why We Really Monitor Serum Levels: for Model-based, Goal-oriented Drug Therapy

Traditional approaches to therapeutic drug monitoring have been designed for use only in steady state situations, and usually have employed only 1 - compartment models. They have developed dosage regimens only for such situations, and have been oriented to keeping serum levels within a generally accepted therapeutic range rather than to achieving a specific target

goal. Such approaches have made it impossible to deal with patients in their most important clinical moments, as, for example, during changing renal function or dialysis, or when certain “golden moments” must be understood and a dosage regimen developed to achieve and maintain a desired clinical goal immediately, as in the case of the above patient receiving digoxin.

The above patient on digoxin also shows how truly individualized drug therapy begins with clinical selection of an explicit therapeutic goal for each patient, based on that individual patient’s need for the drug. One then should achieve that goal with the greatest possible precision, without any zone of indifference about it. The approach was highly cost-effective, when compared to the fact that an entire week of hosoital time was spent in the previous attempts at dosage adjustment without the aid of a model - based, goal - oriented method.

This patient’s case emphasizes the fact that one does not use serum levels simply to see whether or not they are in some general "therapeutic range", nor even to correlate them with the patient’s clinical behavior, although that is often possible, but significantly not so in this patient. This patient clearly shows that the real reason for monitoring serum levels is rather to find out how each patient actually handles the drug, how the drug (and its model) really behaves in that individual patient, especially in non-steady-state situations, and to correlate the behavior of patient’s fitted model with his own clinical behavior. Only then can one optimally evaluate each patient's clinical sensitivity to, and specific need for, a drug. MAP Bayesian adaptive control, in the context of model based, goal-oriented individualized drug therapy, brings a precision and capability to drug dosage which is not possible with older obsolete approaches based on linear regression or simply on the raw data of serum levels alone.

Entering Initial Conditions: Changing Population Models during the Fitting Procedure.

Most pharmacokinetic analyses deal with patients, and their pharmacokinetic models, who have stable values for their various parameters such as volume of distribution, rate constants, clearances, etc.. However, this is not always so, even though one can express a rate constant as an intercept plus a slope times a descriptor of elimination such as creatinine clearance or cardiac output [12], so that renal function can change from dose to dose during therapy, and the patient's drug model can keep up with these changes as they take place.

Probably the most serious problem in analyzing pharmacokinetic data in patients is caused by sudden significant changes in a patient's volume of distribution (Vd) of the central (serum concentration) compartment. It is generally known, for example, that patients in an ICU setting

have larger values for the Vd of gentamicin and other aminoglycosides than do general medical patients. Indeed, young very healthy people who suddenly require an aminoglycoside for a perforated or gangrenous appendix often have even smaller values for Vd [10]. It is interesting that each patient, himself, also goes through such transitions as his clinical status changes.

An Aminoglycoside Patient with sudden change in Clinical Status and Volume of Distribution

An interesting example of such an individual change was a 54 year old man in Christchurch, New Zealand, seen through the courtesy of Dr. Evan Begg in the fall of 1991. He was 69 in tall, weighed 80 kg, and his serum creatinine on admission was 0.7 mg/dL. He had a pyelonephritis, was receiving tobramycin 80 mg approximately every 8 hours. He had a measured peak serum concentration of 4.6 and a trough of 0.4 ug/ml respectively, and had been felt by all to be having a satisfactory clinical response. During this time, his Vd was 0.18 l/kg, based on those two serum samples. However, on about the 6th day, he suddenly and unexpectedly relapsed and went into clear-cut septic shock. This patient's antibiotic therapy was discussed in another paper in this collection. The present paper concerns how the analysis was able to proceed from one population model to the other, switching from that of a general medical patient at first to that of an ICU patient, and finally back to that of a general medical patient again. At the time of the patient's care, the appropriate software was not available. It was only later, in retrospect, that the analysis described below could be done and the patient's case could be more fully understood.

Following this initial phase and his surprising relapse on therapy, he was aggressively treated with much larger doses. He received 300 mg every 12 hours during this time. His serum concentrations rose to peaks of 10.1 ug/ml. During this period of sudden septic shock, his serum creatinine rose from 0.7 to 3.7 and his CCr fell to about 18 ml/min/1.73m^2. After about another 10 days he improved. At that time his serum concentrations rose to a peak of 16, and it was necessary to sharply reduce his dose to 140 mg about every 12 to 24 hours. His serum creatinine fell to 1.1 to 1.3 mg/dL, and his CCr rose to 57 ml/min/1.73m^2.

When one tries retrospectively to fit the entire data set, it was simply not possible to get a good fit to all the serum data. Most data points were obtained during the second, his sickest phase, and they dominated the fit. The ones at the beginning, prior to his sepsis, and at the end, after his recovery from it, were not at all well fitted.

Because of this, the data was divided into three parts - an initial one before his relapse into sepsis, a second one when he was septic, and a third one following his recovery, but before it was felt safe to discontinue his therapy. Each data set was was fitted separately, using the USC*PACK programs [10].

During the first data set, the first 6 days, when his behavior was that of a general medical patient, not seriously ill, his Vd was 0.18 L/kg as described above. The problem then was to pass on the ending values of his serum and peripheral compartment concentrations as initial conditions for the fitting process for the second data set. This was done, using that feature of the USC*PACK software.

A major change in his Vd was then seen when fitting the data obtained during his second, septic, phase. The Vd rose from 0.18 in his previous phase to 0.51 L/kg, and his Kslope, the increment of elimination rate constant per unit of CCr, fell to zero. However, his Kcp, the rate constant from serum to peripheral compartment, rose to 0.255 hr^{-1}, suggesting that he was "third-spacing" his tobramycin somewhere. The ending concentrations in his central (serum) compartment for this data set were 2.09 ug/ml, and for his peripheral compartment were a very high 44.1 ug/kg.

These ending values were then passed on to the third part of his data set, that of his recovery. During this time his peaks were 16 and 12 ug/ml, and his dose was reduced to 140 mg every 12-24 hours. His Vd during this third phase, that of recovery, when he was no longer seriously ill, had fallen greatly to 0.15 L/kg, close to his previous value as a general medical patient.

The use of initial conditions and of changing population models during the overall fitting procedure permitted the intelligent analysis of this patient's data, especially as quite significant concentrations were present not only in his central (serum) compartment, but also in his peripheral compartment, during the transition from his second to his third, the recovery, phase.

At the Cleveland Clinic, Drs Peter Slugg and Marcus Haug [13] have spoken of "Vd collapse", as the Vd would drop from a larger to a smaller value, as have described it as an indicator of incipient recovery of the patient. The present patient not only demonstrated such Vd collapse, but also its opposite, Vd expansion, as he made the earlier transition from being a general medical patient with a pyelonephritis to a seriously ill ICU patient with life-threatening septic shock. Thus not only do different populations of aminoglycoside patients have different

values of Vd, but it appears that each individual patient goes through these transitions, as demonstrated by this patient. The analysis of this patient's data was greatly facilitated, and indeed was only possible, by breaking his dosage history up into several parts. Each part was then analyzed, and the ending concentrations from one part were passed on to the next data set as initial conditions or concentrations of drug present prior to the first dose given in the next data set, with the appropriate population model, if needed, as well.

A Patient on Digoxin when Quinidine was Added.

Another example of the utility of using initial conditions is the example, provided through the courtesy of Dr. Marcus Haug, of a 72 year old woman, 4 ft 10 in tall, weighing only 75 pounds. She was admitted to a hospital with congestive heart failure and atrial fibrillation. Her creatinine clearance (CCr) was 38 ml/min/1.73m^2, falling to 23 after admission. She had been receiving 0.25 mg of digoxin daily. This was continued afgter admission to the hospital. A serum digoxin level was 1.8 ng/ml on admission.

Following this, her serum creatinine rose to 1.8 mg/dL, and her digoxin level after 5 days rose to 2.5 ng/ml. Her digoxin was stopped, even though she had no clinical manifestations of toxicity. The next day her serum level had fallen to 2.0, and the next day it was down to 1.4 ng/ml.

At this point her ventricular rate with her atrial fibrillation had risen beyond a reasonable resting rate, and she was begun again on digoxin, again at 0.25 mg/day, to control it. However, quinidine was now also added to her regimen. Her CCr was 22 ml/min/1.73m^2. Five days later her serum digoxin level was measured and found to be 7.6 ng/ml for a trough and 10.0 ng/ml two hours after the next dose was given.

What was going on here? She again had no clinical evidence of toxicity. Was all of this due to the digoxin - quinidine reaction? Was it a problem of digoxin - like material appearing in the assay as a result of her poor renal function? Was there something else in addition?

The clinical problem was analyzed as follows. First, her original dosage history on digoxin alone was fitted to her serum levels, using the 2-compartment population model for digoxin made from the work of Reuning, Sams, and Notari [11]. This included the three measured serum levels. At the end of that part of her history, just before her first dose of quinidine was added, her fitted and predicted central compartment (serum) concentration was 1.19 ng/ml, and her peripheral

(nonserum) compartment concentration was 7.58 ug/kg. These relationships are shown in Figure 9.

These two ending values from this first phase of her analysis were passed on as initial concentrations of drug already present in those compartments of her pharmacokinetic model at the time her digoxin was restarted, but now with quinidine as well. The population model for digoxin with quinidine [10] was used. This model was not fitted to her subsequent serum levels, but merely used to supply predictions of those high measured levels. If the prediction was good, the interpretation would be that the interaction would quantitatively account for the measured levels found. If not, then another explanation would have to be devised.

As shown in Figure 10, the predicted concentration of 7.2 ng/ml closely matched the measured one of 7.6 ng/ml. In addition, the measured level of 10.0 ng/ml was predicted as 9.9 ng/ml. Because of these good predictions, it was felt that the digoxin-quinidine interaction explained the measured levels well, and that no other alternative explanation was needed. This is a good example of how pharmacokinetic analyses can be used to evaluate experiences with drugs, and can provide strong evidence for or against a particular question or issue, much more than a clinical opinion made without the aid of such a model. The use of initial conditions was the key to being able to change from one population model to another in the middle of a patient's history. In the same way, one can make the transition from regular theophylline to long-acting preparation, for example. With the use of initial conditions, one can thus follow the patient as he goes from one situation to another, passing on the data from one set to another.

Limitations of MAP Bayesian Adaptive Control

The MAP Bayesian approach to adaptive control and dosage individualization is straightforward and robust. However, it does not represent an optimal approach to dosage individualization, and it has two significant drawbacks from optimality.

The first drawback is that the parameter values used to describe the behavior of the drug are assumed to be either normally or log-normally distributed. This is often not so, as many drugs, for example, have rapid and slow metabolizers within the population, and have parameter distributions for the elimination rate constant which may be multimodal. Furthermore, the volume of diatribution for drugs such as the aminoglycosides is affected by the patient's clinical state as a general medical patient or a patient in an intensive care unit, for example. Because of this, parameter distributions are often not either normal or lognormally distributed, and are not

optimally described by mean, median, or mode values. This point reflects the problems associated with making parametric population models. It is largely overcome by making nonparametric population models which describe the entire joint density within the population, with virtually one support point (set of parameter values, and its estimated probability) for each subject studied in the population [14,15].

The second drawback is that there is no tool in the MAP Bayesian strategy for evaluating the precision with which a desired dosage regimen developed to hit a desired target goal actually will do so. The method lacks a vital performance criterion.

The separation or heuristic certainty equivalence principle is well known among the stochastic control community, but less so amobg the pharmacokinetic community. It states [16] that when the task of controlling the behavior of a system is separated into the steps of:

1. Obtain the best single point parameter values in the model describing the behavior of the system, and then,
2. Using these single point values to design the inputs to control the system, that the task is inevitably performed supoptimally. Yet this is exactly what the MAP Bayesian, and all methods which use single point parameter values, do.

There is no performance criterion in the MAP Bayesian strategy (estimated precision with which the desired target will be hit, for example) as there is only one set of parameter values, and the target is assumed to be hit exactly. There is no tool to evaluate the predicted precision with which the regimen will, or will fail, to hit the desired target. These two limitations are overcome by the combination of the new nonparametric population models [14,15] and the "multiple model" design of dosage regimens [7]. These approaches will be discussed more fully in other papers in this collection.

ACKNOWLEDGMENTS

Supported by US Government grants LM 05401, RR 01629, RR 11526, and by the Stella Slutzky Kunin Research Fund.

References

1. Sheiner LB, Beal S, Rosenberg B, and Marathe V: et al.: Forecasting Individual Pharmacokinetics. Clin. Pharmacol. Ther., 31: 294-305, 1979.

2. Jelliffe R, Iglesias T, Hurst A, Foo K, and Rodriguez J: Individualizing Gentamicin Dosage Regimens: A Comparative Review of Selected Models, Data Fitting Methods, and Monitoring Strategies. Clin. Pharmacokinet. 21: 461-478, 1991.

3. Nelder JA and Mead R: A Simplex Method for Function Minimization. Computer Journal 7: 308-313, 1965.

4. Caceci MS and Cacheris WP: Fitting Curves to Data: The Simplex Algorithm is the Answer. BYTE Magazine, May 1984, pp. 340-362.

5. De Groot MH: Probability and Statistics, Second Edition, Addison-Wesley Publishing Co., Reading, MA, 1989, pp. 422-423.

6. Sawchuk R and Zaske D: Pharmacokinetics of Dosing Regimens which Utilize Multiple Intravenous Infusions: Gentamicin in Burn Patients. J. Pharmacokin. Biopharm 4: 183-195, 1976.

7. Jelliffe R, Schumitzky A, Bayard D, Milman M, Van Guilder M, Wang X, Jiang F, Barbaut X, and Maire P: Model-Based, Goal-Oriented, Individualised Drug Therapy: Linkage of Population Modelling, New "Multiple Model" Dosage Design, Bayesian Feedback, and Individualized Target Goals. Clin. Pharmacokinet. 34: 57-77, 1998.

8. Maire P, Jelliffe R, Dumarest C, Roux D, Breant V, Charpiat B, Vermeulen E, Brazier J, and Courpron P: Controle Adaptatif Optimal des Posologies: Experience des Aminosides en Geriatrie. in Information et Medicaments. Comptes Rendus du Colloque AIM-IF et IRT, Paris, December 1989, ed. by Venot A and Degoulet P, Volume 2 of Informatique et Sante, directed by Degoulet P, Springer Verlag, Paris, 154-169, 1989.

9. Hurst A, Yoshinaga M, Mitani G, Foo K, Jelliffe R, and Harrison E.: Application of a Bayesian Method to Monitor and Adjust Vancomycin Dosage Regimens. Antimicrob. Agents Chemother., 34; 1165-1171, 1990.

10. Jelliffe R, Schumitzky A, Van Guilder M, and Jiang F: User Manual for Version 10.7 of the USC*PACK Collection of PC Programs. December 1, 1995. Laboratory of Applied Pharmacokinetics, University of Southern California School of Medicine, Los Angeles, CA.

11. Reuning R, Sams r, and Notari r: Role of Pharmacokinetics in Drug Dosage Adjustment. 1. Pharmacologic Effects, Kinetics, and Apparent Volume of Distribution of Digoxin. J. Clin. Pharmacol. 13: 127-141, 1973.

12. Jelliffe R: A Mathematical Analysis of Digitalis Kinetics in Patients with Normal and Reduced Renal Function. Math. Biosci. 1: 305-325, 1967.

13. Marcus Haug, Pharm.D., and Peter Slugg, M.D., personal communication.

14. Schumitzky A: Nonparametric EM Algorithms for Estimating Prior Distributions. App. Math. and Computation. 45: 143-157, 1991.

15. Mallet A: A Maximum Likelihood Estimation Method for Random Coefficient Regression Models. Biometrika. 73: 645-656, 1986.

16. Bertsekas D: Dynamic Programming: deterministic and stochastic models. Englewood Cliffs (NJ): Prentice-Hall, pp. 144-146, 1987.

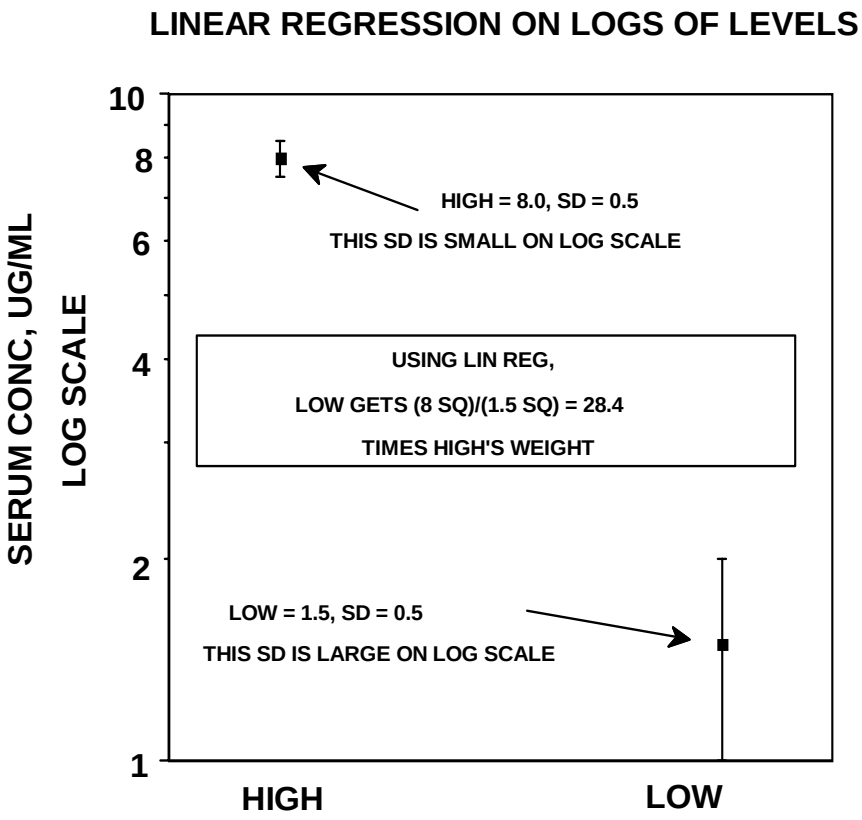

Figure 1. Error pattern assumed using fitting by linear regression on logarithms of serum levels. Note the much greater weighting given to the lower levels.

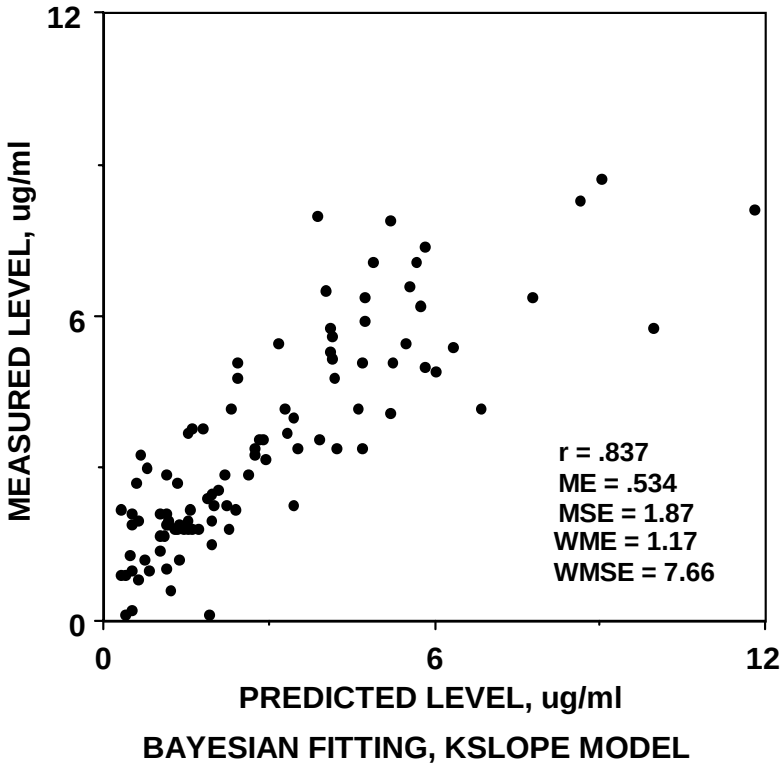

Figure 2. - *Predicted versus measured serum Gentamicin levels found with M.A.P. Bayesian fitting and the Kslope model. r = correlation coefficient, ME = mean error, MSE = mean squared error. WME = mean weighted error. WMSE = weighted mean squared error. See text for discussion.*

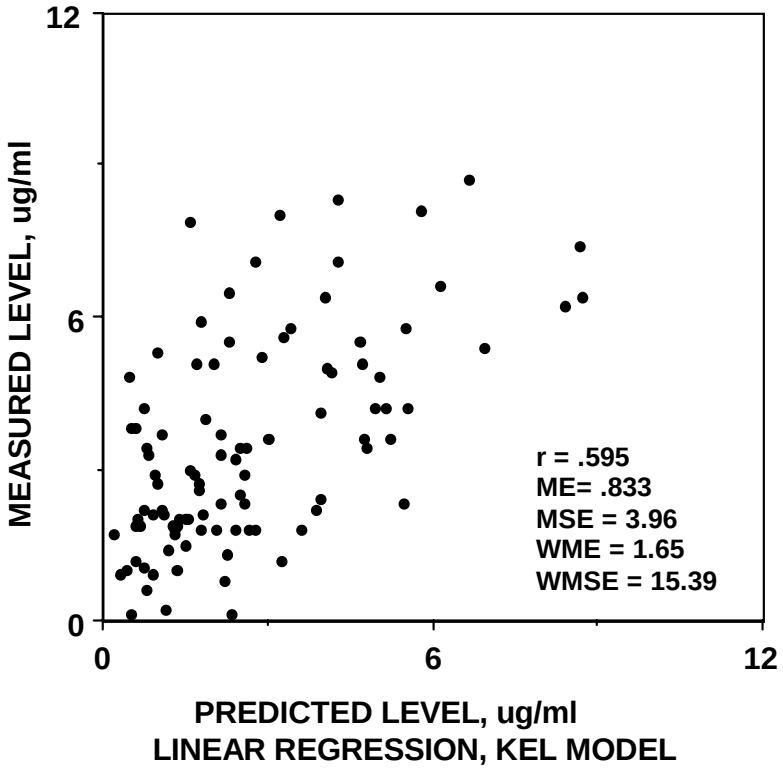

Figure 3. - *Predicted versus measured serum levels found with linear regression on the logarithms of the serum levels. Other symbols as in Figure 2.*

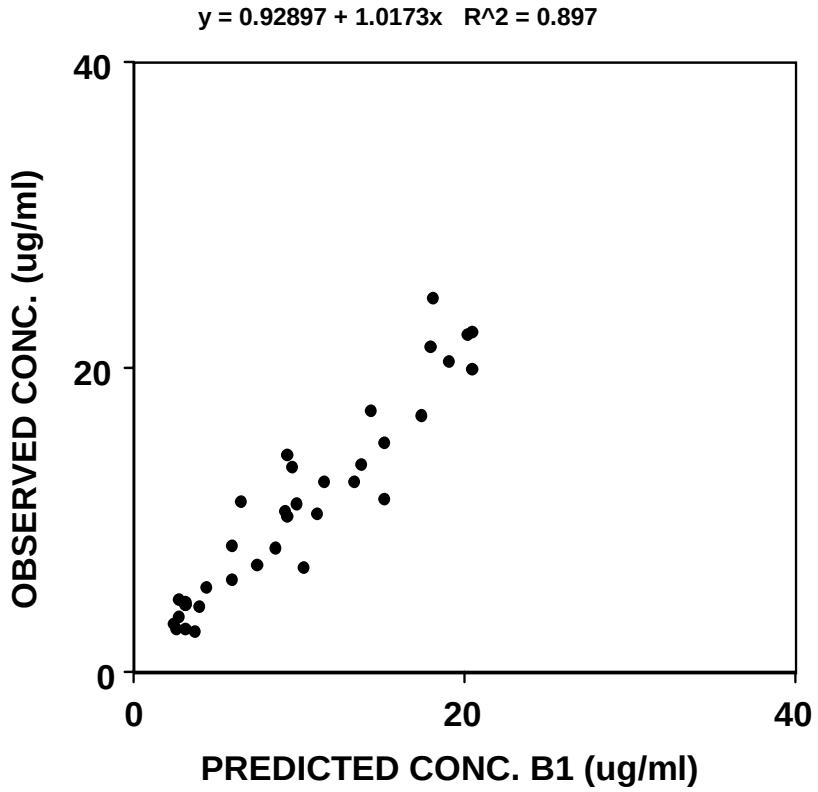

Figure 4 - *Predicted versus measured serum Amikacin levels found with M.A.P. Bayesian fitting, 1 compartment Kslope model (B1) .*

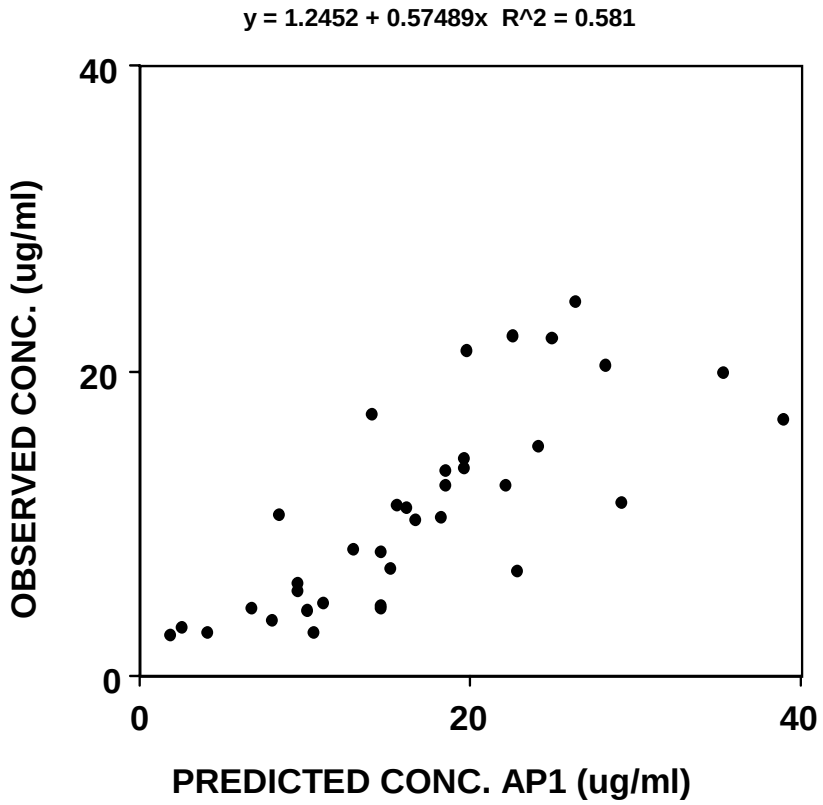

Figure 5 - *Predicted versus measured serum Amikacin levels found with A Priori population 1 compartment Kslope model (AP1)* .

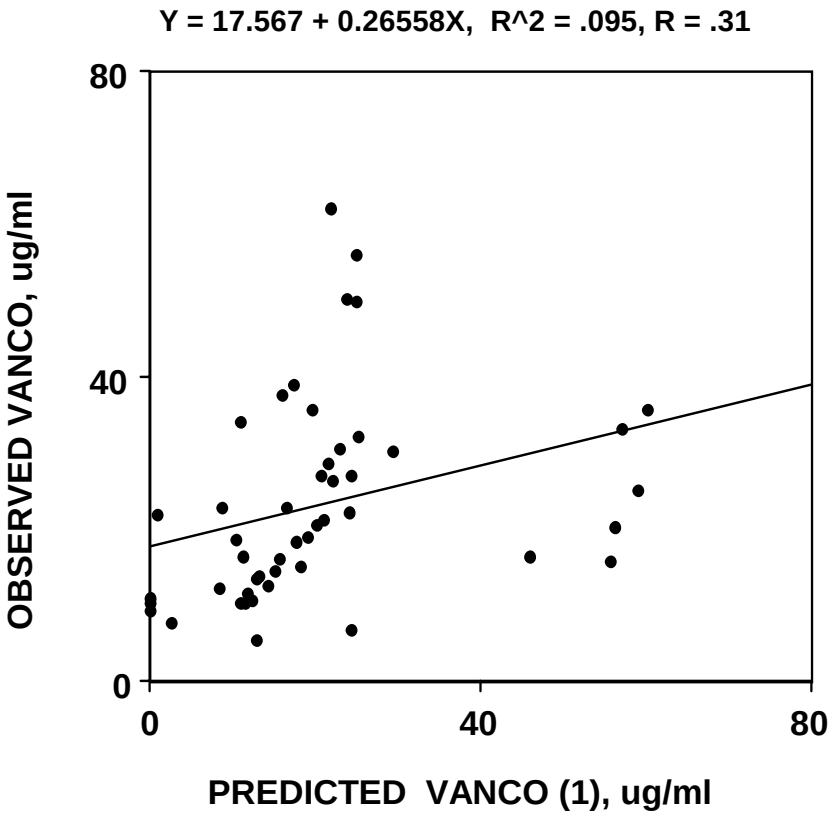

Figure 6 - *Predicted versus measured serum Vancomycin levels found with Linear regression (1)*

.

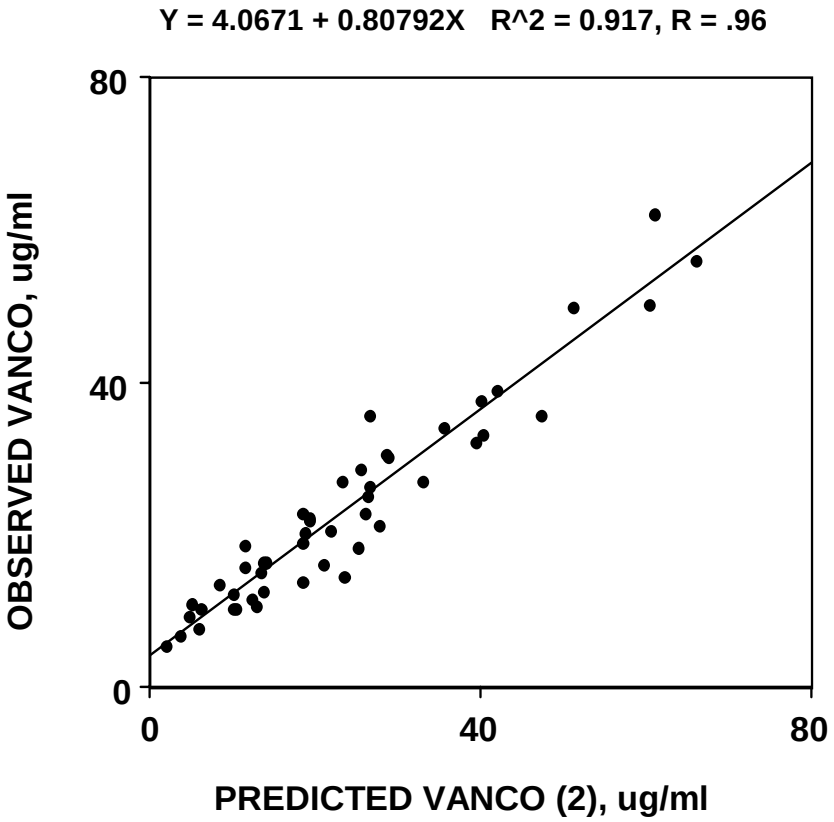

Figure 7 - *Predicted versus measured serum Vancomycin levels found with a 2 compartment Kslope model and Bayesian fitting (2)* .

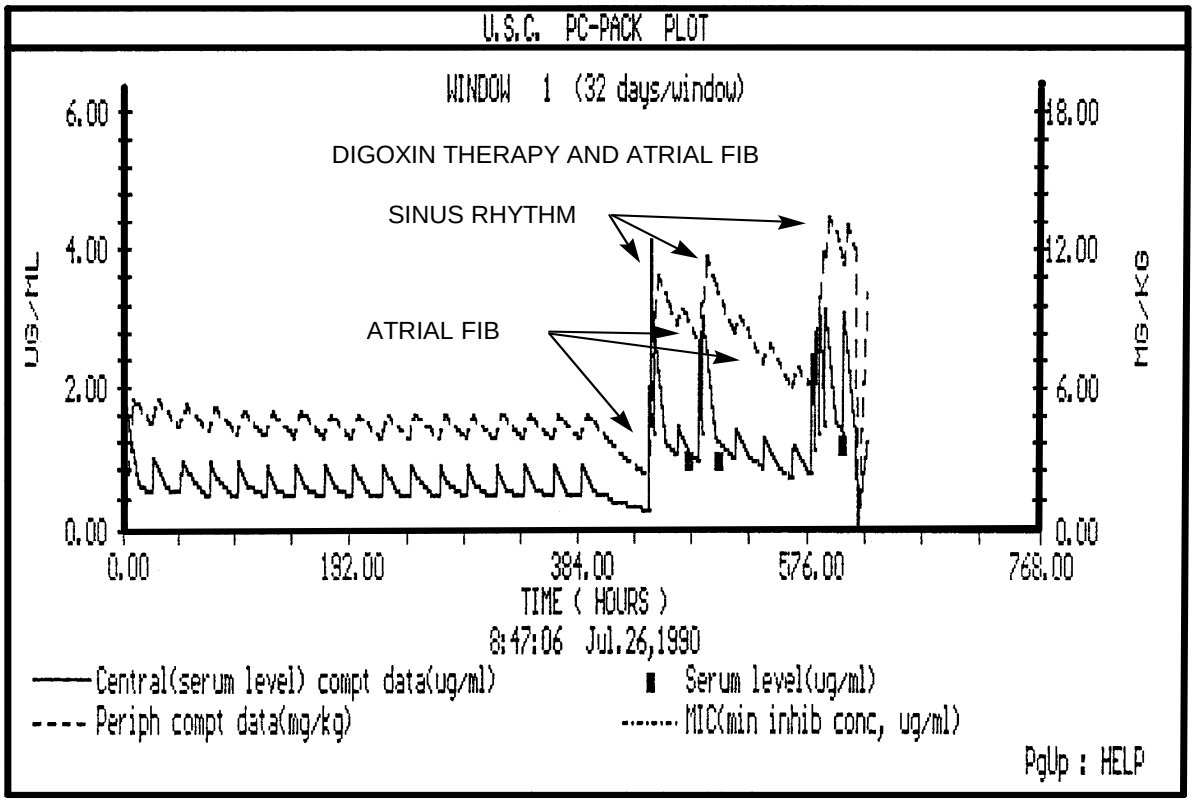

Figure 8. *Screen plot of patient with atrial fibrillation who was successfully converted to sinus rhythm with IV digoxin three separate times, but who relapsed into atrial fibrillation twice when put back on his previous maintenance dose. Sinus rhythm was consistently present when peripheral body glycoside concentrations were 10-12 ug/kg (right hand scale, and not mg/kg as labeled). Selection of a therapeutic goal of 11.5 ug/kg in the peripheral compartment led to a dosage regimen of 0.5 and 0.625 mg/day. On that regimen, the patient could be discharged home in sinus rhythm and was still in sinus rhythm when seen in clinic 2 weeks later.*

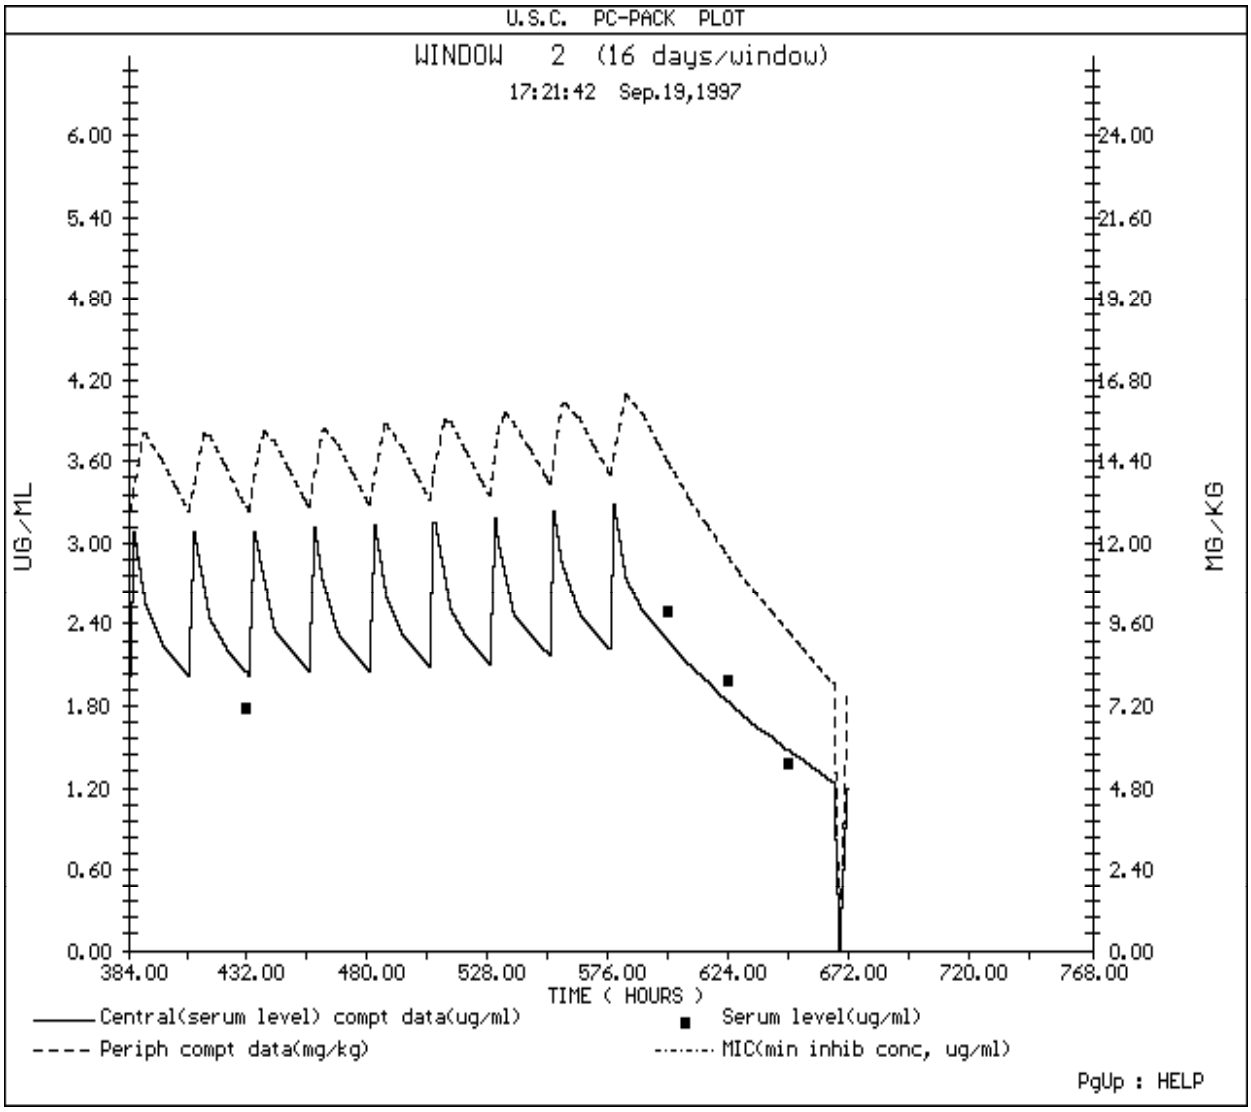

Figure 9. Plot of serum and peripheral compartment digoxin concentrations of patient admitted receiving digoxin. She was receiving 0.25 mg of digoxin daily, ands weighed only 75 lb. Solid rectangles - measured serum levels. Solid line and left hand scale - digoxin serum concentrations. Dashed line and right hand scale - digoxin peripheral (nonserum) compartment concentrations. Using the Bayesian approach, the population model for digoxin was fitted to the patient's data of doses and serum levels. Serum levels rose as her renal function worsened. Digoxin was stopped after the serum level of 2.5 ng/ml was obtained, after which her serum levels fell to 1.4, and, in the fitted model, finally to 1.19ng/ml at the end of this plot, when digoxin was begun again, but along with quinidine.

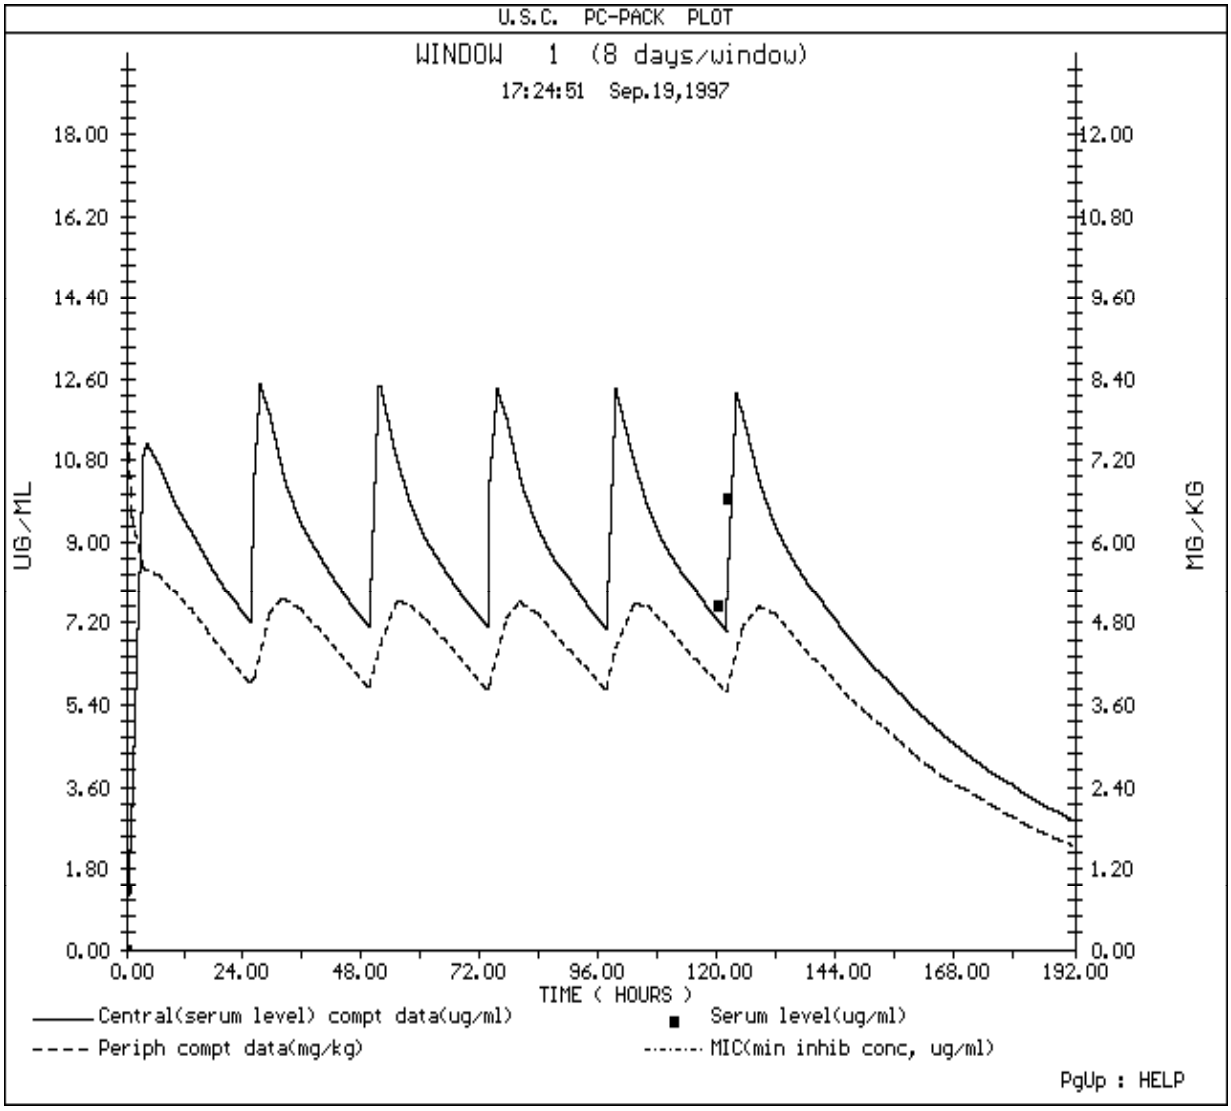

Figure 10. Plot of serum and peripheral compartment digoxin concentrations of patient admitted receiving digoxin. In this plot, digoxin was restarted at 0.25 mg/day, but along with quinidine. Solid rectangles - measured serum levels. Solid line and left hand scale - digoxin serum concentrations predicted using the population model for digoxin with quinidine [2]. Dashed line and right hand scale - predicted digoxin peripheral (nonserum) compartment concentrations. This plot begins with initial conditions equal to the final concentrations found at the end of the plot in Figure 9.