University of Southern California
School of Medicine
Laboratory of Applied Pharmacokinetics

Technical Report 94-1

ADAPTIVE CONTROL OF DRUG DOSAGE REGIMENS: AN OVERVIEW OF BASIC CONCEPTS, RELEVANT ISSUES, AND SOME CLINICAL APPLICATIONS.

by

Roger Jelliffe, M.D.[1], Alan Schumitzky, Ph.D.[1], David Bayard, Ph.D.[1], Mark Milman, Ph.D.[1], Pascal Maire, Pharm.D.[2], Xavier Barbaut, Pharm.D.[2], Agneta Hurst, Pharm.D.[3], Bruno Charpiat, Pharm.D.[4], Valentine Breant, Pharm.D.[5], Christine Pivot-Dumarest, Pharm.D.[6], and Babak Tahani [1].

1. Laboratory of Applied Pharmacokinetics, Department of Medicine, University of Southern California School of Medicine, Los Angeles, California, USA,
2. Department of Pharmacy, Hospital Geriatrique Antoine Charial, Francheville, France.
3. University of Southern California School of Pharmacy,
4. Department of Pharmacy, Croix-Rousse Hospital, Lyon, France.
5. Department of Clinical Pediatric Pharmacology, Debrousse Hospital, Lyon, France.
6. Department of Pharmacy, Edouard Herriot Hospital, Lyon, France.

Supported by NIH grant LM05401, by French Government MRES grant 88C0573, by Bristol-Myers-Squibb, Aguettant Laboratories, the Defer Company, and by the Stella Slutzky Kunin Memorial Research Fund. Mrs. Valerie Pinela rendered expert secretarial and administrative assistance.

ADAPTIVE CONTROL OF DRUG DOSAGE REGIMENS: AN OVERVIEW OF BASIC CONCEPTS, RELEVANT ISSUES, AND SOME CLINICAL APPLICATIONS.

Roger Jelliffe, M.D.[1], Alan Schumitzky, Ph.D.[1], David Bayard, Ph.D.[1], Mark Milman, Ph.D.[1], Pascal Maire, Pharm.D.[2], Xavier Barbaut, Pharm.D.[2], Agneta Hurst, Pharm.D.[3], Bruno Charpiat, Pharm.D.[4], Valentine Breant, Pharm.D.[5], Christine Pivot-Dumarest, Pharm.D.[6], and Babak Tahani [1].

1. Laboratory of Applied Pharmacokinetics, Department of Medicine, University of Southern California School of Medicine, Los Angeles, California, USA,
2. Department of Pharmacy, Hospital Geriatrique Antoine Charial, Francheville, France.
3. University of Southern California School of Pharmacy,
4. Department of Pharmacy, Croix-Rousse Hospital, Lyon, France.
5. Department of Clinical Pediatric Pharmacology, Debrousse Hospital, Lyon, France.
6. Department of Pharmacy, Edouard Herriot Hospital, Lyon, France.

1.0 INTRODUCTION

Along with surgery and thoughtful advice to the patient, good drug therapy is the mainstay of the physician's armamentarium. Such therapy needs to be optimized, especially for many drugs which have narrow therapeutic margins of safety. Such drugs need precise and thoughtful dosage. They include not only the aminoglycoside antibiotics, vancomycin, theophylline, and cardiovascular drugs such as digoxin, digitoxin, lidocaine, and procainamide, but also many potentially toxic drugs for AIDS therapy such as AZT, trimethoprim / sulfamethoxazole (TMP/SMX), and gancyclovir. Cyclosporine, for prevention and management of rejection following transplantation, also needs such careful management.

In addition, drugs for cancer chemotherapy such as methotrexate, 6-mercaptopurine, etoposide, teniposide, suramin, and cytarabine and other drugs also need precise dosage regimens. Decision support for cancer chemotherapy using adaptive control has improved survival without increasing toxicity [1] by decreasing patient variability around the therapeutic goal of total exposure to the drug (area under the serum level curve).

Antiepileptic drugs such as phenytoin, phenobarbital, valproic acid, carbamazepine, and many other drugs used in psychiatric therapy, and therapy of Alzheimer's disease, such as tacrine, have widely varying patient responses to a given dose, and need similar careful management. Adaptive control has reduced the variability in serum concentrations of nortriptyline, imipramine, desipramine, amitriptyline, lithium, fluphenazine, and haloperidol and its decanoate [2].

In the past, management of dosage regimens of such potentially toxic drugs has been performed using several different methods.

1. Simply measuring serum drug concentrations. When used without software for modeling drug behavior and for adaptive control of dosage, therapeutic drug monitoring based only on evaluating the raw data of serum concentrations has often been ineffective, even though the therapeutic ranges for many drugs are well known [3].

2. Predictive Nomograms. These provide simple but useful guidance for the important initial regimen. However, they are not suited for subsequent adjustments based on feedback provided by measurements of response such as serum concentrations of drug.

3. Linear Least Squares Regression. [4]. This method has been in wide general use for over 20 years. It is based on linearizing a 1 compartment model by taking the logarithms of its serum concentrations. It therefore carries a significant built-in error in its fitting procedure which assumes that a serum concentration of 0.1 ug/ml, for example, is 100 times more precisely measured than a sample of 1.0 ug/ml, and 10,000 times more precisely measured than one of 10.0 ug/ml. These unrealistic assumptions, coupled with the fact that this method is wasteful of serum data, can only analyze data obtained in a single dosage interval, and can only describe at best the behavior of the traditional 1 compartment model, greatly limits this method.

4. Nonlinear Least Squares Regression. [5]. This is much better. Proper weighting of serum data can be done. Data can be analyzed over many dose intervals. However, fitting of a model having a peripheral nonserum compartment is fraught with clinical difficulties, generally limiting the clinical usefulness of the method to a 1 compartment model. In addition, one must have at least 1 serum data point for each parameter to be fitted.

5. Maximum Aposteriori Probability (MAP) Bayesian Fitting. [6,7]. This method is the best so far. It has the advantages of nonlinear least squares, above, and also can begin with only a single serum data point. Furthermore, the fitting procedure itself is more robust and less volatile that nonlinear least squares because the presence of the population parameters in the objective function restricts the parameters so they are less likely to give odd combinations which happen to fit the data "best".

For most drug therapy the MAP Bayesian method has been the best to date. Prediction (and therefore control) of subsequent drug levels has been shown to be more precise with Bayesian than with the other methods to date listed above [7-10].

However, very significant limitations still exist for this MAP method. They are present in large part because single numbers (mean and SD, for example) which describe pharmacokinetic parameters such as volume of distribution and the various rate constants are used as markers of the central tendency and of the dispersion of patient population data, or of the uncertainties in parameter values in models fitted to an individual patient's serum level data.

Since the patient can only receive a single regimen of any drug at any time, what may be a correct regimen for one perceived version of the patient (the mean or median, for example) will be incorrect for any other versions or possible parameter values that the patient actually may have. The current MAP Bayesian adaptive control procedure simply

controls the most likely version of the patient exactly, never considering the uncertainties that exist about those single-point parameter values. This will be discussed further toward the end of this chapter.

2.0 THE CLINICAL SCENARIO OF MAP BAYESIAN ADAPTIVE CONTROL.

The Reverend Thomas Baye's described in 1760 how we learn from life. He set forth the quantitative relationships between prior probabilities of certain events, the acquisition of new information about these events, and the revised or posterior probabilities of these events after the new information has been taken into account. When one wakes up in the morning one already has certain expectations. For example, he thinks that objects will probably continue to fall down, that the grass will probably still be green and the sky blue. For many years we had the expectation that life would probably be better for our children than it was for us. Then, during the day, for example, we learn new things - that rain clouds are forming, that the stock market has behaved in an unexpected manner, or that the worldwide recession has revised our expectations of life for our children. Because of this new information, our expectations are revised, and our plans for the future are altered.

This is standard feedback control. An initial plan has been revised, based on subsequent information. What was unique about Bayes' theorem was that these relationships were stated quantitatively.

When Bayes' theorem is applied to feedback control of drug dosage regimens, we find that the prior probabilities of pharmacokinetic parameter values for a certain drug are stored in a population pharmacokinetic model. This is the way we get and store our past experience with a drug so we can apply it usefully to the care of the next patient who comes to us and who seems to belong to a certain type of patient population. In this way, the past incidence of certain pharmacokinetic parameter values in a population now becomes (by our faith that the future will probably repeat the past) our prior expectation or probability that this new patient will have certain parameter values.

Since we have no idea of just how the new patient will respond to his dosage regimen, either in terms of his serum levels (pharmacokinetics) or his clinical response (pharmacodynamics), the most likely values of these parameters are contained in a model made from a relevant population of patients which resemble him. Currently, we tend to use the mean values obtained from such a population as our best guide to plan our regimen. We will see later on that this practice probably needs revision.

2.1 THE MAXIMUM APOSTERIORI PROBABILITY (MAP) BAYESIAN OBJECTIVE FUNCTION.

In the scenario of the widely used MAP Bayesian adaptive control [6-10], we balance the information obtained from our population model against the data of the serum levels we obtain from the patient. In addition, we balance the credibility of the population data against that of the patient's serum data. All this is done as we find the parameter values which minimize the overall value of the MAP Bayesian objective function shown below, which was introduced to the pharmacokinetic culture by Sheiner and Beal [6].

$$\left[\text{SUM}\ \frac{(\text{Ppop} - \text{Ppt})^2}{\text{SD}^2\ \text{Ppop}} + \text{SUM}\ \frac{(\text{Cobs} - \text{Cpt})^2}{\text{SD}^2\ \text{Cobs}}\right]$$

where Ppop and Ppt represent the parameter values of the population pharmacokinetic model and of the patient's individualized model respectively, Cobs and Cpt represent the observed (measured) serum drug concentrations and the estimates of those concentrations made with the patient's own individualized pharmacokinetic model respectively, and SD Ppop and SD Cobs are the standard deviations of the various population parameter values and of the various observed serum concentrations respectively.

In this objective function, the credibility of each population pharmacokinetic parameter value is determined by the reciprocal of the variance (the square of the SD) which it has been found to have. Thus the SD of each population parameter value, when squared and its reciprocal is then taken, provides the correct index of credibility for each population parameter value.

The same is true for the data of the measured serum drug concentrations. However, the usual practice of most clinical laboratories is to make sure that the SD's of each assay are within some selected acceptable limits for that laboratory. Once this is done, the actual error is usually ignored for purposes of therapeutic drug monitoring, and is not reported along with the serum concentration itself. The result of this is that the SD's with which serum drug concentrations are measured are usually not properly considered as a practical matter in the routine fitting of serum drug concentration data. We will discuss this more later on.

The usual way in which these Bayesian posterior parameter values are obtained is simply by trying on a set of values (for example, the population values), computing the overall value of the objective function, and then trying on another set of values to see if it is any better or not. This is currently done by the Nelder-Mead simplex procedure [11,12] in our hands, or some other function minimizer.

It is like going into a clothing store to buy a coat. One may have the expectation, based on past experience (analogous to the population model) that the correct size is a 42 long. When we put on this size, though, we find that while the shoulders fit pretty well, the sleeves are too short. We try on a similar coat, but with longer sleeves. We keep on doing this until the coat fits correctly not only in the shoulders, but also in the sleeves, and also in the back, and not only simply for standing, but also for moving and bending.

The Nelder-Mead algorithm [11,12] starts with an initial set of estimates of the parameter values (those of the population model, for example). It also sets up some competing sets of values. The total number of all sets of values is one more than the number of parameters in the model. Thus, for a simple one compartment model having only 2 parameters, it would set up 2 other sets of parameter values, for a total of 3. A simplex is a figure having one more vertex then the number of parameters in the model.For this simple model, one of the other sets might have the population value for the volume of

distribution (V) but the elimination rate constant (K) might be given a value almost zero. The other set might have the population value of K, but a V close to zero.

These 3 sets of parameter values are then used to simulate the behavior of the pharmacokinetic system. The value of the objective function is then found for each set. The set having the worst (here the largest) value of the objective function is discarded, and another set of values is selected, which is in the direction away from the worst set, just as one moved away from the coat sleeves that were too short, above. The algorithm has a built - in set of movements with which it then explores the ranges of all the parameters (the parameter space), and finally finds (finds, not computes) the set of parameter values which minimizes the value of the objective function [11]. A useful nonmathematical article describing this algorithm also appeared in BYTE magazine [12].

2.2 INDIVIDUALIZING DRUG DOSAGE REGIMENS USING BAYESIAN ADAPTIVE CONTROL.

This process involves five basic steps:

2.21 Setting an individualized explicit therapeutic goal for each patient, based on his/her need for the drug.

We have had a strong tendency, in the biomedical culture, to classify the world we see around us. We classify plants, animals, and diseases. We classify hearts as being of normal size or enlarged, rather than trying to estimate the size weight, a wall thickness.

We also classify the ranges of serum drug levels in which most patients do well. This has been based on looking at the relationship between serum drug levels and the incidence, first, of effects we like to see, and which we therefore call therapeutic. Our eye usually gets drawn to a bend in the line showing that above a certain serum level, the incidence of therapeutic effects becomes appreciable or "significant". This is then taken as the beginning of the therapeutic range. Similarly, with increasing serum drug levels, the incidence of toxic effects also becomes appreciable, and this is taken as the beginning of the toxic range of serum levels.

Traditional approaches to therapeutic drug monitoring have been designed for use only in steady state situations. They develop dosage regimens only for such situations, and have been oriented to keeping serum levels within some generally accepted therapeutic range. Such approaches make 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 clinical moments” must be understood and a dosage regimen developed to achieve and maintain a desired clinical goal following a sudden change in clinical status, as in the patient receiving digoxin described later on, for example.

Individualized drug therapy cannot begin properly without first setting an explicit and individualized therapeutic goal for each patient. It is not enough to say that the serum levels should be within some so-called therapeutic range where most patients do well. That is only a general range, and is only to be considered in a general way. If that is all we do, we have not considered how to be gentle in our approach to a patient who dose not

need the drug very much, moderate to a patient who needs a moderate touch, or very aggressive to a patient who really needs his drug dosage "pushed".

The most important relationship to keep in mind, we feel, is the one between the serum level (or the area under the serum level curve, or the desired concentration at a desired time in a nonserum peripheral drug compartment, for example) and the incidence of adverse reactions. This is what turns us on or off in our approach to therapy, and what limits the dosage we are willing to give.

Clinical behavior (efficacy and toxicity) at the bottom of a wide therapeutic range is not the same as at the top. Further, using such a wide range means that the patient is not considered as a unique individual, with his or her individual need for the drug under consideration, but only as a member of a population.

Therapeutic ranges of serum concentrations are usually derived from two sets of data, as shown in Figure 1. First, there is a relationship between the serum concentration or area under the curve (AUC) reflecting the patient's exposure to the drug, and the incidence of therapeutic effects. The beginning of the therapeutic range is usually said to occur when the incidence of therapeutic effects becomes "significant". This significance is usually not explicitly defined or quantified. At a higher serum concentration (or AUC), there is a similar relationship to the incidence of adverse or toxic effects.

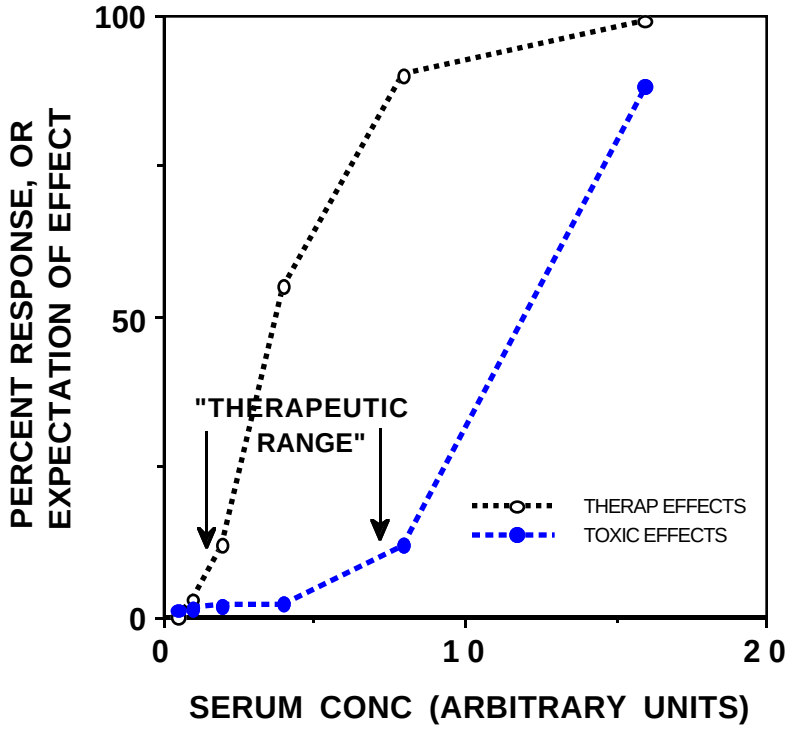

Figure 1 - *General relationships usually found between serum drug concentrations and the incidence of therapeutic and toxic effects. The eye is drawn to the bends in the curves, and the therapeutic range is found in relation to these bends. This procedure discards the*

important quantitative data of the incidence of toxicity versus serum concentration. See text for discussion.

One might choose a goal which maximizes the probability that the serum concentrations will be within some chosen therapeutic range or window. However, one must also consider the probability that the patient's serum concentration may be above or below that window. The risks of being below the window are associated with lack of efficacy - those of being above it with risks of toxicity. Simply maximizing the probability of the levels being within a certain range does not face the significant clinical risks associated with the probability of them being outside it.

Generally, one wishes to give the patient as much drug as possible, to obtain the maximum possible benefit from it. In contrast to choosing a wide therapeutic range, we feel that one should choose an explicit serum concentration peak, trough, average, or profile as a therapeutic goal, based on each patient's individual need for the drug. If the patient's need is small, one chooses a low therapeutic goal associated with a low incidence of adverse reactions. This will result in a gentle dosage regimen. If the patient's need is greater and/or if a previous dosage regimen has not brought about the desired clinical response, one can choose a higher serum concentration as the therapeutic goal, accept a greater risk of adverse reactions, and develop an appropriately higher dosage regimen to achieve that goal. In this way one can adjust the goal, not to a wide population therapeutic range, but to each patient's individual need for the drug, always holding the risk of adverse reactions only to that justified by each patient's need.

Two points are important. The first is that a knowledge of the overall quantitative relationship between serum level (or AUC) and the incidence of adverse reactions is more useful than a general therapeutic range. That relationship usually constitutes the data base from which the therapeutic range was derived in the first place. That data base permits the clinician to choose an explicit therapeutic pharmacokinetic or pharmacodynamic goal for each individual patient - gentle, moderate, or aggressive - moving appropriately up or down that overall population quantitative relationship of concentration or AUC to incidence of toxicity, based upon a careful assessment of each patient's individual need for the drug and upon a similar assessment of the risk of that particular patient having an adverse reaction which one is willing to accept (acting as the patient's advocate) in order to obtain the needed benefits of the drug. Setting such an explicit therapeutic goal is the essential first step in the individualization of drug dosage regimens.

Second, the risks and benefits of the resulting serum levels being only slightly above such a specific chosen goal are not very different from those of being only slightly below it. They can be regarded as being approximately equal. This is another argument for choosing an explicit goal based on each patient's need (and its justifiable risks) and then to achieve it as precisely as possible. It both simplifies and individualizes the decision-making process compared to the greater complexity (and lesser precision of the analysis of probabilities and utilities) associated with the possibility of being either below or above some overall wider population therapeutic range.

2.22 Achieving the desired goal with the greatest possible precision.

Successful individualization of drug dosage regimens thus requires:

1. The appropriate choice of a specific individualized therapeutic goal for each patient.
2. The precise achievement of it.

Both components of therapy are important to patient outcome. Either a poorly chosen goal or imprecise achievement of a well-chosen goal can lead to ineffective therapy or to unacceptable adverse reactions. Choosing the goal is a specific clinical decision.

Achieving the goal as precisely as possible requires a pharmacokinetic or pharmacodynamic model and software to analyze its behavior in the context of altered and/or changing renal and/or cardiac function that often occurs, and to calculate dosage regimens to best achieve the selected therapeutic goal.

2.23 Monitoring the patient and using Bayesian adaptive control to adjust the dosage regimen.

The role of Bayesian fitting is to make an individualized patient - specific model, based on data of the doses, descriptors, and serum levels. Assay errors must be carefully quantified, to balance the credibility of the serum level data correctly against the data of the population parameter values. The individualized model is plotted and used to examine relationships between its behavior and that particular patient's clinical behavior. An individualized therapeutic goal is selected once again and a new regimen developed to achieve it most precisely. These constitute the basic components of Bayesian adaptive control of pharmacokinetic and dynamic systems.

2.24 Reconstructing events of the past, using the Bayesian fitted model.

Here, one uses the Bayesian fitted model parameter values to simulate and to reconstruct the events of past drug therapy. Here is where clinical acumen comes in, evaluating and comparing the patient's past clinical behavior with the behavior of the fitted model. Here is where the sensitivity of the patient to the serum levels, and to the concentrations in his peripheral (nonserum) compartment is seen. Here is where (and how) a bioavailability problem can be recognized.

2.25 Selecting a new specific goal, if needed, and adjusting the dosage regimen to best achieve it.

Here is where another specific individualized therapeutic goal can be selected for the patient if needed. At any rate, an adjustment in the dosage regimen can be developed to best achieve the goal, as described earlier. The cycle then repeats itself each time a new set of serum levels becomes available to add to those already obtained, to re-evaluate the fitted model and the patient's clinical response.

2.26 Useful Types of Pharmacokinetic Models

Useful pharmacokinetic models have consisted of an absorptive compartment (gut or IM injection site) and a central (serum level) compartment. The parameters are the apparent volume of distribution of the central compartment (Vc) in L/kg, the descriptor-linked (usually the renal) component (Kslope) of the elimination rate constant (Kel), showing the increment of the Kel added for each unit of the descriptor of elimination, the non-descriptor-linked (usually the nonrenal) component (Kint) of the Kel, and the rate constant for absorption (Kabs). In this model, Kel = Kint + Kslope times the value of the descriptor of elimination or metabolism (usually CCr). Thus in the USC*PACK software used by us, both Vc and Kel can change as the patient's weight and CCr, for example, may change from dose to dose [13].

Parameter values for a similar model containing an additional peripheral compartment exchanging with the central compartment can also be used. They consist of the rate constants Kcp and Kpc for exchange from central to peripheral compartment and from peripheral back to central compartment respectively.

3.0 DOES BAYESIAN ADAPTIVE CONTROL WORK?

3.1 GENTAMICIN THERAPY.

With Bayesian adaptive control, the Kslope pharmacokinetic model and the Bayesian fitting procedure resulted in significantly better prediction of future serum concentrations (see Figure 2) than those made using the linear regression (Sawchuk-Zaske) method [14, and see 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.

It is perhaps for this reason that the predictions of future serum levels made using the Kslope model and Bayesian fitting procedure in that study [14] were not quite significantly better than those made with the population model used, without fitting to any serum level data at all. There was still visible room for improvement in Bayesian prediction with these very unstable patients. The scatterplots of predicted versus measured serum levels made with the Bayesian fitted 1 compartment model, and the population 1 compartment Kslope model [4, 14] are shown in Figure 2, while those obtained using linear regression are shown in Figure 3.

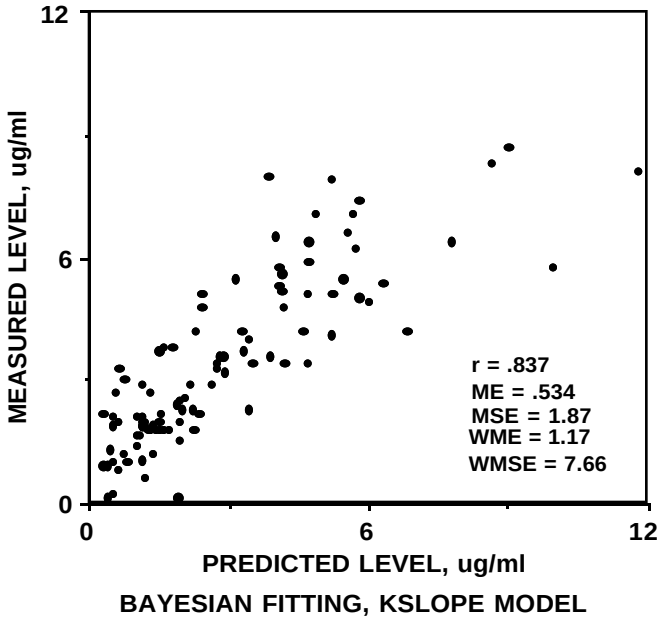

Figure 2 - *Predicted versus measured serum Gentamicin levels found with M.A.P. Bayesian fitting and the Kslope model. r = correlation coefficient, ME = mean error, MSE = mean squared error. WME = mean weighted error. WMSE = weighted mean squared error. See text for discussion.*

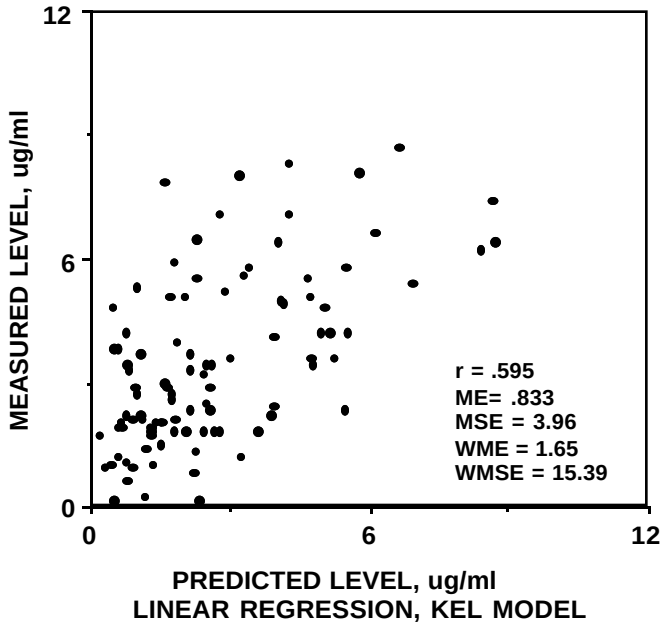

Figure 3 - *Predicted versus measured serum levels found with nonlinear least squares fitting and the Kslope model. Other symbols as in Figure 2.*

Because these programs are designed to operate in the presence of significant changes in renal function from dose to dose, they have also been useful in the analysis and management of aminoglycoside therapy for patients who must undergo periodic hemodialysis. Specific techniques for their use in such patients have been described [15].

The program for MAP Bayesian adaptive control has also been used to achieve more aggressive gentamicin therapeutic goals felt by Haug et.al. to be necessary and useful in critically ill surgical ICU patients [16].

3.2 AMIKACIN THERAPY.

MAP Bayesian adaptive control has been used to manage amikacin therapy in geriatric patients, often for extended periods, by Maire et al [17]. In those 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 [14]. These results are shown in Figure 4. They are clearly better than those found in the gentamicin patients with unstable renal function [14] shown in Figure 2 above. In addition, long-term control of amikacin levels in their patients has permitted successful therapy of patients with osteomyelitis, for example, without significant evidence of ototoxicity or nephrotoxicity. In contrast, Figure 5 shows the predictions based on the population model for Amikacin, without any fitting to the serum data [17].

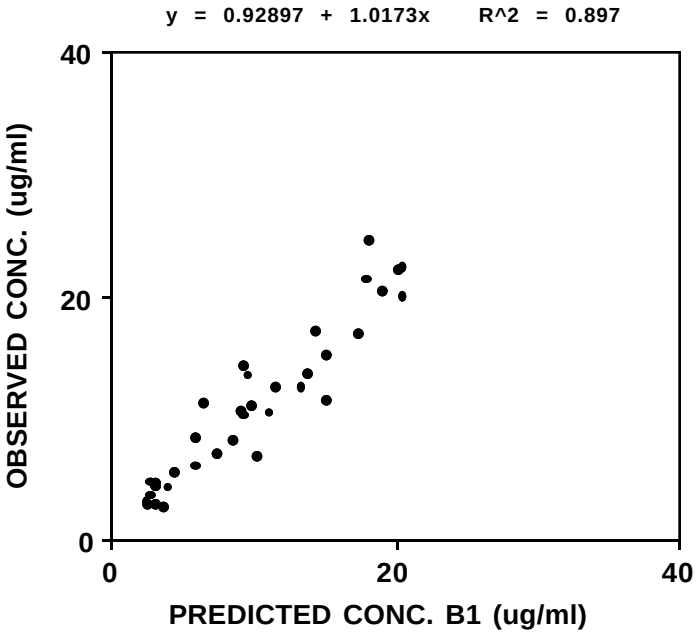

Figure 4 - *Predicted versus measured serum Amikacin levels found with M.A.P. Bayesian fitting, 1 compartment Kslope model (B1) .*

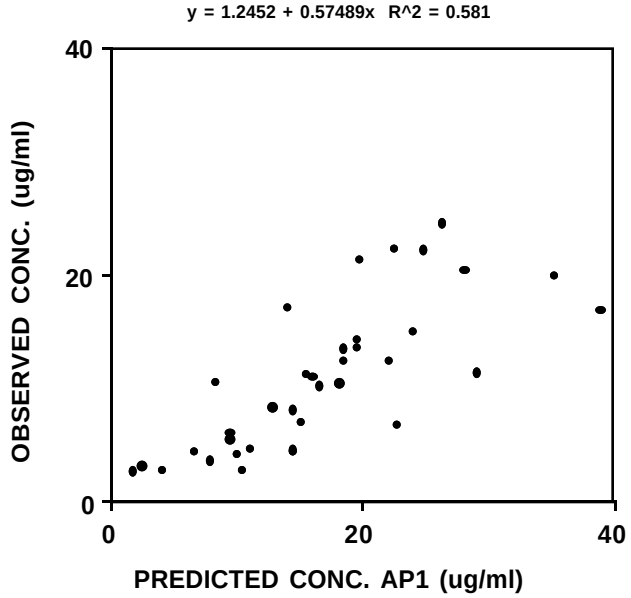

Figure 5 - *Predicted versus measured serum Amikacin levels found with A Priori population 1 compartment Kslope model (AP1) .*

3.3 VANCOMYCIN THERAPY.

Vancomycin therapy was evaluated by Hurst et al [18] 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.

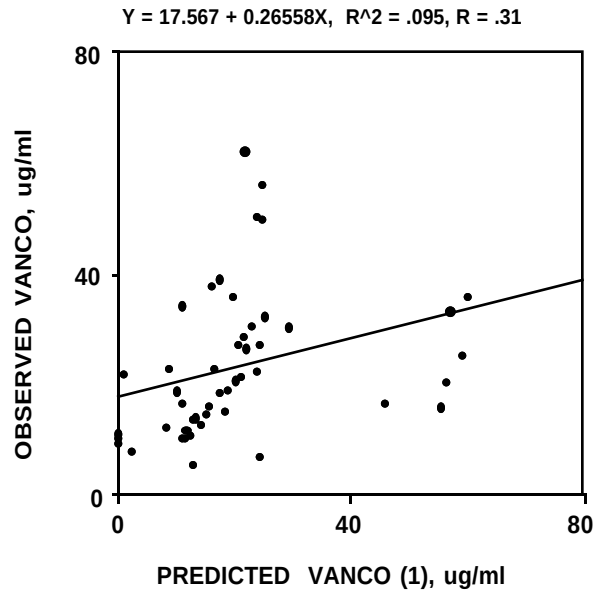

Figure 6 - *Predicted versus measured serum Vancomycin levels found with Linear regression (1)* .

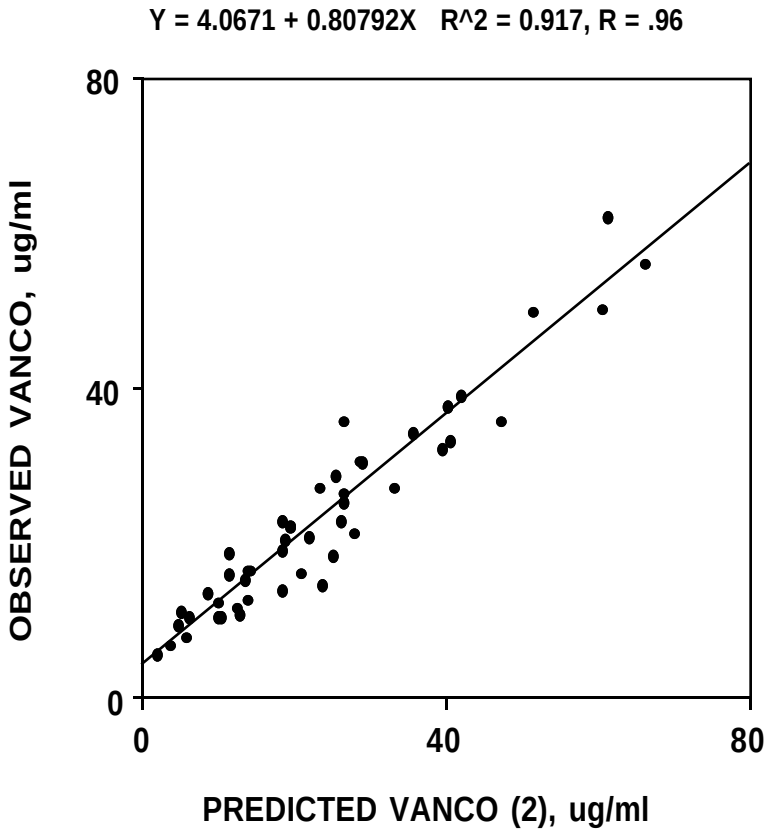

Figure 7 - *Predicted versus measured serum Vancomycin levels found with a 2 compartment Kslope model and Bayesian fitting (2) .*

3.4 CANCER CHEMOTHERAPY.

In cancer chemotherapy, Rodman, and Evans and his group, using the USC software modified to control AUC rather than serum level goals, have shown that the outcome of treatment of children with acute lymphocytic leukemia with high-dose methotrexate correlated with total systemic exposure [1]. Ratain, Schilsky, Vogelzang et al [19] showed that individualized dosing of etoposide permits safe increase in dose-intensity.

4.0 ILLUSTRATIVE CLINICAL CASE EXAMPLES.

4.1 LIDOCAINE THERAPY

Using pharmacokinetically oriented lidocaine infusion regimens developed by a previous version the USC program [20], clinically effective serum levels of lidocaine were reached earlier than with conventional therapy [21]. In a community hospital audit

comparing pharmacokinetically designed versus conventional regimens, significantly fewer breakthrough arrhythmias were noted (2 of 78 versus 33 of 78), with a reduction in the incidence of ventricular fibrillation to 2 of 78 versus 8 of 78 patients [21]. The program also interfaces with "smart" infusion apparatus for automated and reliable delivery of the stepwise infusion regimens often useful with lidocaine, and for smooth transition from loading to maintenance infusions.

4.2 DIGOXIN THERAPY

Very early versions of these programs for digitalis therapy significantly reduced the incidence of digitalis toxicity without reducing efficacy [22], even before serum level assays had become available.

The digoxin population model now used in our Bayesian adaptive control software is based on that described by Reuning, Sams, and Notari [23]. This 2 - compartment model describes the behavior of digoxin in both its serum and its peripheral (nonserum) compartment. Computed concentrations of drug in that peripheral compartment correlate much better with inotropic effect than do serum levels [23]. The USC*PACK MB program for digoxin uses this model and develops oral dosage regimens to control the peak peripheral compartment (not the central or serum level compartment) at a chosen total body concentration.

In using the MB program for most drugs, volumes of distribution are in liters per kg, and therefore in liters when adjusted for body weight. Since doses are usually expressed in milligrams, serum concentrations are therefore expressed in mg/L or ug/ml. When using digoxin or digitoxin, however, doses must be entered in micrograms (not mg) so the resulting serum concentrations can be in ug/L or ng/ml. The words do not change, as MB is a very general program. Entering digoxin or digitoxin doses in ug results in serum concentrations in ng/ml. Similarly, the peripheral body concentrations become ug/kg instead of mg/kg.

Use of the MB program 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 received several intravenous doses of digoxin, as shown in Table 1, and converted to sinus rhythm. He was placed on his original oral maintenance dosage. After 1 day, atrial fibrillation recurred. He again was given several doses of intravenous digoxin and again converted to sinus rhythm. Once again, he was placed on his original oral maintenance dosage. Once again, after about 2 days, atrial fibrillation recurred.

For a third time he received several intravenous doses of digoxin, and for a third time he converted to sinus rhythm. At this point the MB program was used to analyze his situation.

Patient: digoxin pt bill nicholson's Chart #: 123
Room: 456 Age (Years) : 58.00 Sex: M Height(Inches): 68.0

Therapy Day #1 was June 11, 1987 Filename = billspt.MB

Press Shift-PrtSc to print the screen, and/or enter to continue.

Dose #	Route	Rx Day	Time of Day (HH:MM)	Time Into Regimen (HOURS)	Weight (KG)	CCR (ML/MIN/ 1.73MSQ)	Infuse Time (HOURS)	Dose Interv (HOURS)	Dose Amount (MG)	Dose Rate (MG/HR)
19	IV	19	18:55	442.92	75.00	99.19	0.10	0.75	250.00	2500.00
20	IV	19	19:40	443.67	75.00	99.19	0.10	0.50	250.00	2500.00
21	IV	19	20:10	444.17	75.00	99.19	0.10	3.33	250.00	2500.00
22	IV	19	23:30	447.50	75.00	99.19	0.10	9.00	250.00	2500.00
23	PO	20	08:30	456.50	75.00	99.19	0.00	9.50	0.00	0.00
24	PO	20	18:00	466.00	75.00	99.19	0.00	18.17	250.00	0.00
25	IV	21	12:10	484.17	75.00	99.19	0.10	3.33	250.00	2500.00
26	IV	21	15:30	487.50	75.00	99.19	0.10	26.50	250.00	2500.00
27	PO	22	18:00	514.00	75.00	99.19	0.00	24.00	250.00	0.00
28	PO	23	18:00	538.00	75.00	99.19	0.00	24.00	250.00	0.00
29	PO	24	18:00	562.00	75.00	99.19	0.00	15.25	250.00	0.00
30	IV	25	09:15	577.25	75.00	99.19	0.10	3.42	250.00	2500.00
31	IV	25	12:40	580.67	75.00	99.19	0.10	3.58	250.00	2500.00
32	IV	25	16:15	584.25	75.00	99.19	0.10	5.00	250.00	2500.00
33	IV	25	21:15	589.25	75.00	99.19	0.10	8.25	250.00	2500.00
34	PO	26	05:30	597.50	75.00	99.19	0.00	7.42	0.00	0.00
35	IV	26	12:55	604.92	75.00	99.19	0.10	19.08	250.00	2500.00

Press Shift-PrtSc to print the screen, and/or enter to continue.

Level #	Therapy Day	Time of Day (HH:MM)	Time Into Regimen (HOURS)	Taken After Dose Number	Time After Start of dose (HOURS)	Level (MCG/ML)
1	21	05:15	477.25	24	11.25	1.00
2	22	05:50	501.83	26	14.33	1.00
3	26	12:00	604.00	34	6.50	1.20

Table 1. Data of doses and serum levels in a patient with atrial fibrillation who was repeatedly converted to sinus rhythm with digoxin, but who could not be maintained in it using the raw data of his serum levels as a guide.

Data of three serum levels, shown in Table 1, bottom right, all troughs before the various series of intravenous doses, showed no correlation with the patient's clinical behavior. 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, respectively, were obtained. However, when the digoxin population model was fitted to his various doses and to these serum levels, the resulting fitted model, graphically shown in Figure 8, was very useful.

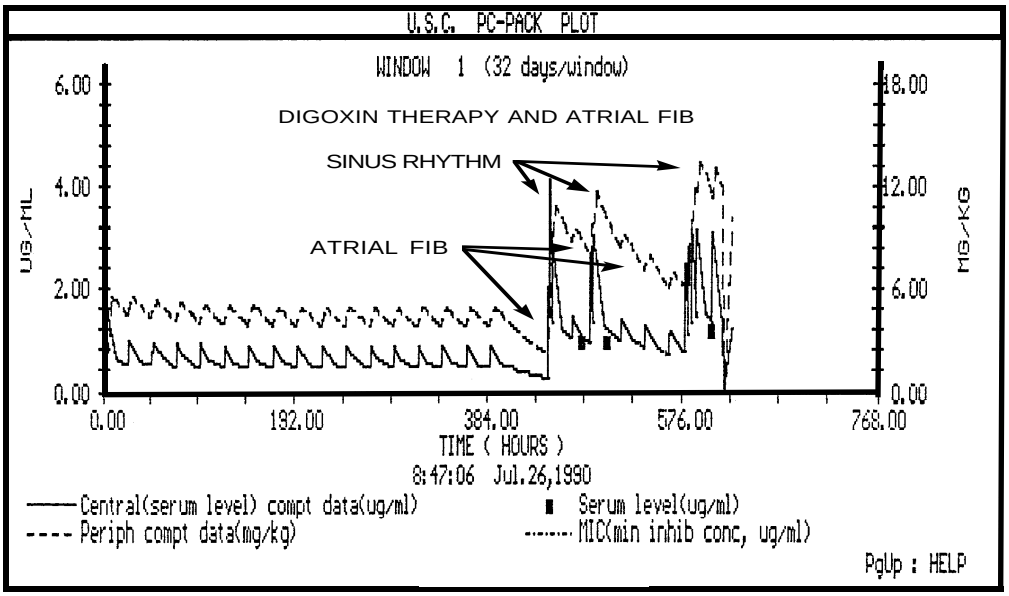

Figure 8. *Screen plot of patient with atrial fibrillation who was successfully converted to sinus rhythm with IV digoxin three separate times, but who relapsed into atrial fibrillation twice when put back on his previous maintenance dose. Sinus rhythm was consistently present when peripheral body glycoside 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 and use of the MB program led to a dosage regimen of 0.5 and 0.625 mg/day. The patient was discharged home in sinus rhythm and was still in sinus rhythm when seen in clinic 2 weeks later.*

Relating this fitted model to the patient's clinical behavior, sinus rhythm was present whenever peripheral concentrations were 10.0 to 12.0 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.

5.0 CLINICAL APPLICATIONS.

5.1 WHY WE REALLY MONITOR SERUM LEVELS.

Traditional approaches to therapeutic drug monitoring have been designed for use only in steady state situations. They develop dosage regimens only for such situations, and are oriented to keeping serum levels within a generally accepted therapeutic range. Such approaches make 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 suddenly, as in the above patient receiving digoxin.

The above patient also shows how truly individualized drug therapy begins with clinical selection of an explicit therapeutic goal for each patient, based on that patient's need for the drug. One then should achieve that goal with the greatest possible precision, without any zone of indifference.

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 patient's **fitted model** of the behavior of the drug in him with his own **clinical behavior**. Nothing less would have dealt with and resolved the clinical problem this patient presented. Bayesian adaptive control brings a precision and capability to drug therapy not possible with traditional approaches based on raw data of serum levels alone. It also led to the end of his hospital stay, whereas a week of intuitive approaches to therapy had not been able to maintain the benefit achieved twice before by the successful conversions to sinus rhythm. Bayesian adaptive control permits truly individualized therapy by a caring clinician rather than application of a "current standard of practice" simply to give standard doses or to have serum levels be kept within some generally acceptable therapeutic range.

5.2 MODELING DIFFUSION OF DRUGS INTO SIMULATED ENDOCARDIAL VEGETATIONS.

We have modeled the diffusion of drugs into a simulated porous sphere which can represent an endocardial vegetation, or perhaps even an abscess. It currently uses the diffusion coefficient of Amikacin into rabbit experimental endocardial vegetations found by Bayer, Crowell, et. al. [24].

As a result, one can now compute and plot the diffusion of substances from the serum compartment of a PK model into a vegetation of a stated diameter, with a stated diffusion coefficient, with a stated percent protein binding of drug in the serum. These plots have now been incorporated into the new release of the clinical USC*PACK programs, version 10.0.

5.3 MODELING PHASE RELATIONSHIPS (HYSTERESIS, ETC.).

Phase relationships and hysteresis loops can also be seen, with the input (usually serum compartment data) on the horizontal and output (peripheral compartment concentrations, concentrations in a vegetation, pharmacologic effect, or bacterial growth and kill dynamics, see below) on the vertical. Many clinically interesting dynamic relationships can be seen. For example, one can see visually that after drug distribution becomes complete between serum and peripheral compartments, the concentrations in

both decrease at an equal rate constant toward zero, with a straight line toward the plot origin. Similarly, the phase relationships between serum and peripheral compartments can be seen for typical dose intervals during drug therapy, as for digoxin, for example.

5.4 MODELING PHARMACOLOGIC EFFECT.

Relationships between events in serum, peripheral, or vegetation compartments and resulting pharmacological effects can now be seen and plotted, using Michaelis - Menten or Hill Models.

5.5 MODELING DYNAMICS OF BACTERIAL GROWTH AND KILL.

The dynamics of bacterial growth and kill have also been modeled, using the differential equation that the rate of increase or decrease of organisms is the unconstrained rate constant for logarithmic growth in the absence of the antibiotic, minus the effect of the antibiotic to kill it. That effect is described by a Hill model having an Emax, an E50, and a Hill constant. The growth rate constant and the Emax can often be obtained from data in the literature.

One of us (RJ) then decided to represent the patient's minimum inhibitory concentration (MIC) of the antibiotic as the rate constant for growth but with opposite sign, thus inhibiting growth but not killing the organism either. When that is done, then one of us (AS) that the E50 can be obtained and everything can be computed and plotted as the relative number of organisms in the compartment of interest. Using an assumed rate constant for growth of 0.7 hr^{-1} (the organism doubles in slightly less than an hour) and an Emax of -1.4 hr^{-1} (speaking conservatively, it would probably be a pretty poor antibiotic that did not kill maximally at twice the rate constant for unconstrained growth), and using clinical data of an MIC of 2 ug/ml, good correspondence was found in 2 cases when patients fulfilling those criteria unexpectedly relapsed and developed severe septic shock when it had not been present or anticipated before on clinical grounds.

The plot below shows one of these 2 cases, a patient seen in 1991 in Christchurch, New Zealand, through the courtesy of Dr. Evan Begg. The patient had appeared to be responding adequately to initial therapy, but unexpectedly relapsed and went into severe septic shock at the right side of the plot. His initial serum tobramycin levels, as shown, were on the low side, with peaks only about 5.0 ug/ml, levels known now often to be associated with suboptimal outcome. After apparent initial success in killing, one can see the organisms beginning to grow out in the center of the plot, and then escaping from control during a 30 hour period between the last dose shown here and the next dose, which is not shown here but which was given at about 152 hours on the time scale shown here. His further history is not shown here. He had a very stormy course but eventually made a good recovery and went home after about 3 more weeks of intensive therapy.

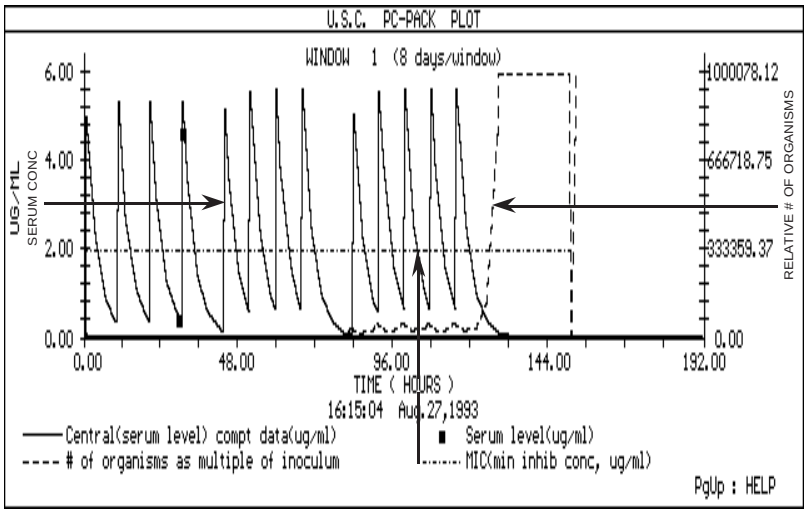

Figure 9. *Plot of serum Tobramycin levels (solid line and left scale) and computed relative number of organisms in serum (dashed line and right scale) in a patient who unexpectedly relapsed after initially appearing to do well on this therapy. The organisms grow out at the right side of the plot.*

6.0 FUNDAMENTAL ISSUES TO BE FACED IN BAYESIAN ADAPTIVE CONTROL

6.1 WHEN TO GET SERUM LEVEL SAMPLES: D-OPTIMAL STRATEGIES FOR THERAPEUTIC DRUG MONITORING

Let us now consider what are the specific optimal times at which serum levels should be obtained.

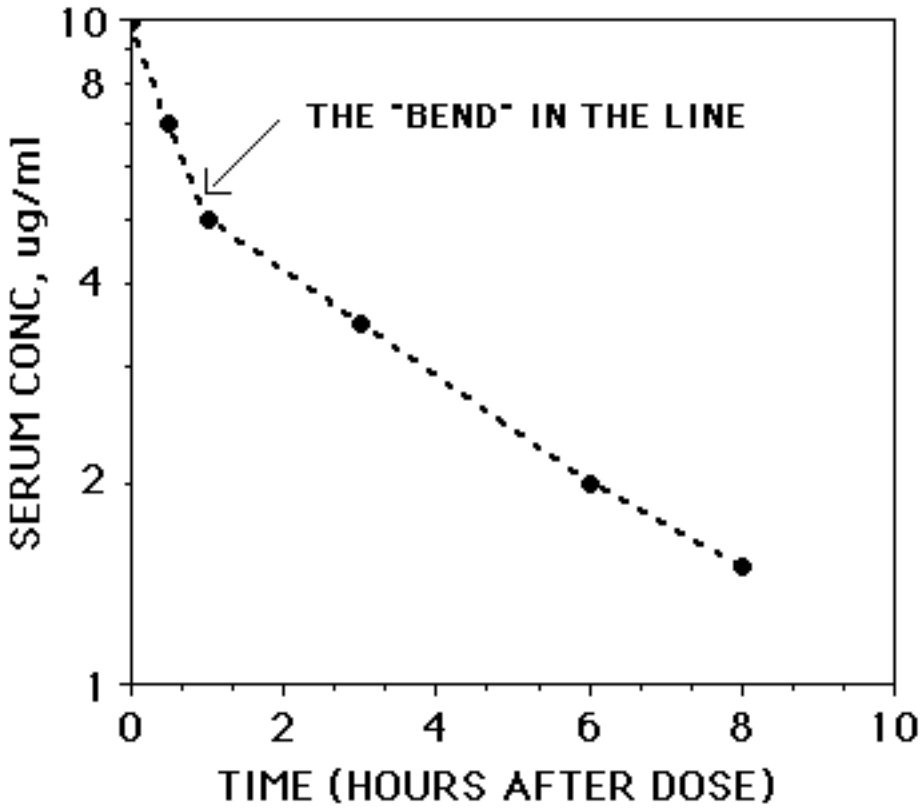

Figure 10. *Typical semilog plot of a gentamicin disappearance curve. The bend in the line takes place about 1/2 hr after the end of the infusion.*

When one examines data such as that shown in both Figures 1 and 10, the eye is drawn to the bends in the lines, and a decision is commonly made at the location of the bend. Figure 10 illustrates the typical 2-compartment disappearance of a drug with an initial early distributive phase. This is the case when such data are plotted on the

traditional semilogarithmic scale. For the aminoglycosides, it occurs at about 1/2 to 1 hour into the disappearance curve.

In the past 10-15 years it has become increasingly common to obtain "peak" aminoglycoside levels actually drawn from 20 to 60 minutes after the end of the infusion, at about the point where the above bend occurs in Figure 10. The reason usually given for this recommendation is that the true peak is transient, and that it is "better" to be on the terminal phase of the logarithmic disappearance curve, (after the bend in the curve), in order to better capture the "terminal" Kel as perceived by such a delayed peak level and the subsequent trough level. However, no rigorous justification for this strategy has ever been given, to our knowledge. One rationale often given is that a 1-compartment model cannot accurately describe the distributive phase of the serum levels which is seen with the aminoglycosides and other drugs. To the authors' knowledge, however, there has been no objective evidence that altering such a 1-compartment model by obtaining a delayed "peak" level in this manner is of any actual benefit.

All pharmacokinetic models are imperfect. Nevertheless, <u>some</u> model, albeit imperfect, must be used to manage a patient's therapy. For whatever model one has decided to use to manage the patient's therapy, 1-compartment or multicompartment, one can calculate the times at which serum drug levels are maximally sensitive to, and changed by, any changes in the model parameter values [5,25].

In the setting of intermittent intravenous therapy and a 1-compartment model, as is commonly done with the aminoglycosides, these most informative times are found to be at the true peak (immediately after the end of the IV infusion) and at 1.44 half-times after the end of that infusion, as shown in Figure 11. The true peak level, not a delayed "peak" drawn later on, contains the most information concerning the volume of distribution (Vd), as the partial derivative of the serum level with respect to the Vd is maximal at this time. A level drawn at 1.44 half-times after the end of the infusion contains the most information concerning the elimination rate constant (Kel), as the partial derivative of the serum level with respect to the Kel is maximal at this time. The use of such D-optimal monitoring strategies to reduce the <u>overall</u> uncertainties concerning the parameter values in any given pharmacokinetic model, having any number of compartments, has been well described [5,25].

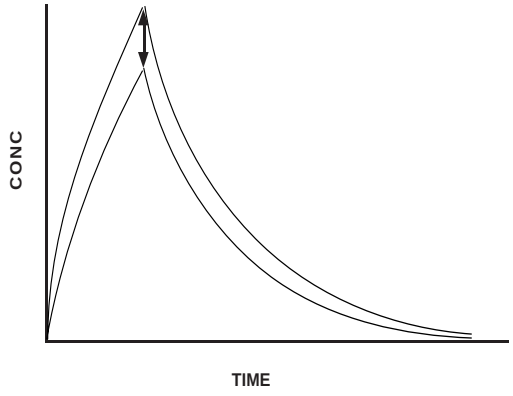

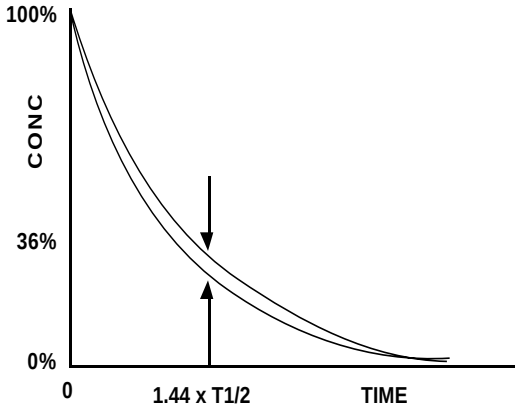

Figure 11. *Optimal times to monitor serum drug levels. A change in the volume of distribution (left) causes the greatest change in the serum level when the level is the highest (at the true peak). A change in the elimination rate constant (right) causes the greatest change in the serum level at 1.44 drug half-times after the end of the infusion, when the level has fallen to 36 % of the original true peak.*

The reasoning behind these D-optimal monitoring strategies can be seen in graphic form. Suppose, as in Figure 11, the volume of distribution (Vd) should change. For example, it might decrease by 10 per cent. The result would be that all serum levels would increase by 11/10 of their original values. The greatest change would thus take place with the highest level, at the true peak at the end of the IV infusion. High levels generally are more informative about the Vd, and a level obtained at the true peak contains the most information, offering the best opportunity to "see through" the many uncertainties always present in any clinical situation, to obtain the most precise value of the Vd. The partial derivative of the serum levels with respect to the Vd is greatest at the true peak. A level obtained at any other time contains less information concerning the Vd.

Similarly, suppose one takes two slightly different values of Kel, and computes for them (using, for example, a hand calculator that does exponentials) the values of e^{-kt} for range of times. When one compares the differences in e^{-kt}, as shown graphically in Figure 11, they are greatest at 1.44 half-times after the end of the infusion, when e^{-kt} has fallen to 36 per cent of its original value at time zero. This is the time when the partial derivative of the serum levels is greatest with respect to the Kel. When taken at this time, a serum level value will contain the most information concerning the patient's Kel.

The above pair of levels turns out to be that which minimizes the overall uncertainty concerning the parameter values for a 1 - compartment model in the clinical setting of intermittent IV therapy [5,25]. It contains more information about the traditional 1 - compartment model than the pair drawn at the delayed "peak" and at the trough . If one decides to use a 1 - compartment model, one loses information if one corrupts that model to "be like" a 2 - compartment model by getting a delayed "peak" and a trough level. This strategy only results in a 1 - compartment model being used suboptimally.

If one chooses to use a 2 - compartment or larger model, D-optimal concepts generalize to all such models, and the D-optimal times can be computed. The above pair of D-optimal times might well (and probably should) replace the current conventional ones which offer less information and are less cost-effective. In most clinical settings it is not difficult to arrange to get serum levels at these D-optimal times, especially when they can easily be known in advance. Computation of these D-optimal times is routinely done for aminoglycoside therapy by the clinical computer program [15] used in the study below.

Pharmacokinetic parameter values obtained using such D-optimal sampling strategies have been shown to be as good as those obtained using larger clusters of levels [26, 27]. In contrast, the ability to predict future serum gentamicin levels was examined in patients receiving gentamicin [14], in which the ability of their 1-compartment models, fitted to a first set of levels obtained by various monitoring strategies, to predict future levels or subsequent therapy, was evaluated. The initial cluster of 6 serum levels consisted

of a trough just before the start of the 30 minute infusion (T1); a true peak obtained from the opposite arm immediately at the end of the infusion (APK); samples at 1/2 hour (30M) and at 1.44 estimated half-times (1.44T) after the end of the infusion, a second trough immediately before the start of the next infusion (T2), and a sample 21 hours after the start of the infusion.

The monitoring strategies evaluated were the entire cluster and various subsets of it, as follows, and as shown in Figure 12:

1. The D-optimal pair: one level at the true peak, and a second level 1.44 estimated half-times later.
2. The entire cluster of 6 levels.
3. The trough before the infusion and the above delayed "peak" drawn 1/2 hour after the end of the infusion.
4. The trough before the infusion, a delayed "peak" drawn 1/2 hour after the end of the infusion, and the next trough.
5. The population (non-fitted) model.

A total of 67 subsequent serum levels were available to be predicted using the above 5 strategies. Models were fitted using the MAP Bayesian method.

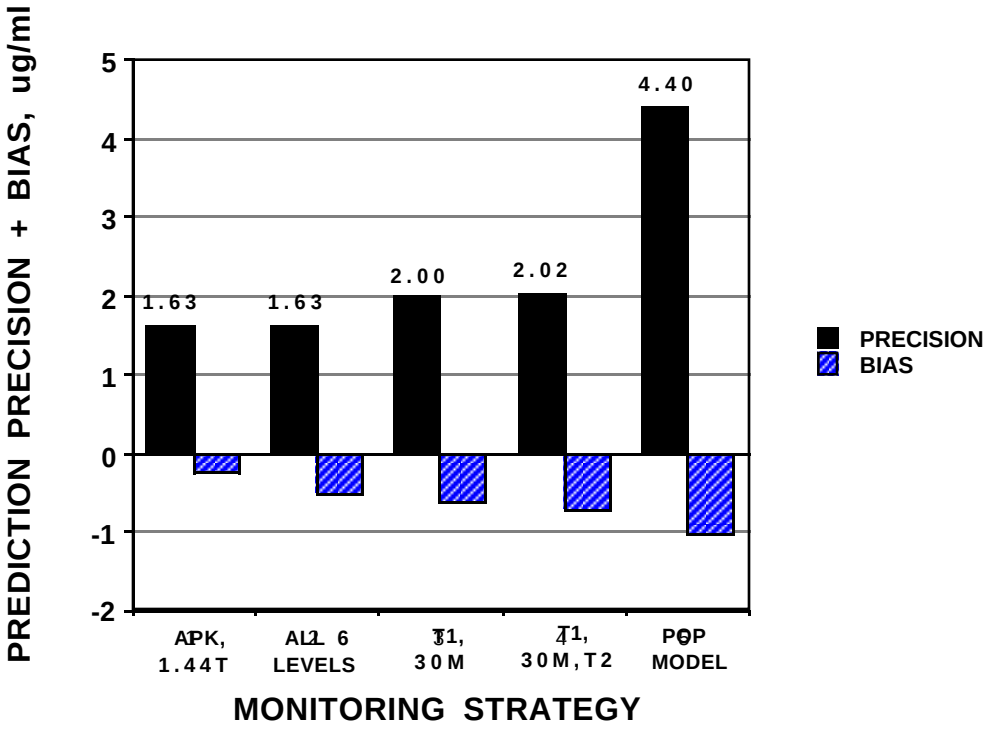

Figure 12. *Optimal strategies for monitoring serum drug concentrations. The D-optimal pair (APK + 1.44T) has less predictive bias and greater precision than all other combinations, even the full cluster. See text for abbreviations and discussion.*

The D-optimal monitoring strategy yielded predictive precision (mean squared error) equal to that seen with the full cluster of levels, as shown in Figure 12, and had a slightly smaller bias (mean error). Both those strategies gave somewhat, though not significantly, more precise and less biased predictions than did the other strategies not containing the true peak drawn at the end of the infusion. The D-optimal pair of levels was somewhat, but not significantly, more precise and less biased than the a priori population model. Predictions made with the D-optimal pair of serum levels were more cost-effective than those of any other strategy, as they gave the therapeutic precision for the number of serum samples drawn.

6.2 DESCRIBING ASSAY ERRORS OPTIMALLY

6.21 Evaluating the Credibility of Population Parameter Values and Serum Level Data.

The use of MAP Bayesian methods to make individualized pharmacokinetic models of drug behavior in patients has led to improved prediction (and therefore control) of future serum drug concentrations, especially when compared to the method of linear regression on logarithms of serum levels [14]. These Bayesian fitted patient-specific models utilize a combination of the population pharmacokinetic parameter values and their standard deviations on the one hand, and data of the patient's own measured serum concentrations on the other. The Bayesian posterior parameter values are found by minimizing the Bayesian objective function shown earlier.

The credibility of each population pharmacokinetic parameter value is thus determined by the reciprocal of its variance. Note that exactly the same is true for the data of the measured serum drug concentrations.

In contrast, the usual practice of most clinical laboratories has been to make sure that the SD's of each assay are within some selected acceptable limit for that laboratory. Once this is done, the actual error has usually been ignored for purposes of practical therapeutic drug monitoring, and has not been reported or made available along with the serum concentration itself. The result of this is that the SD's with which serum drug concentrations are measured have usually not been rigorously considered in the routine fitting of serum drug concentration data, although it is a vital part of the objective function in the Bayesian fitting procedure, as shown above. While it has been common to discuss assay errors as coefficients of variation (SD/concentration) times 100, the actual term in the objective function above is the SD. Because of this, the discussion will be in terms of the SD's.

6.22 The Practical Determination of Serum Assay Error Patterns

What is needed is a practical way to estimate the standard deviation (SD) of each serum drug concentration as it is routinely measured by the clinical laboratory. The SD usually has a nonlinear relationship to the serum concentration, being lowest in the low midrange, sometimes higher at the blank, and almost always higher at the high end.

One way to compute the probable SD with which a single determination of a serum drug concentration is measured is to do replicate measurements of several representative samples (at least in quadruplicate) and to determine the mean and SD of each sample. This can be done, for example, on a blank, a low, an intermediate, a high, and a very high sample, so that the entire working range of the assay is covered. One can fit this data with a polynomial equation, usually of second order. Using this equation, it is then easy to calculate the probable SD with which any subsequent single serum concentration is measured within that range. For example, at the suggestion of Gilman [28], the error pattern of the EMIT gentamicin assay in use at the Los Angeles County - USC Medical Center was determined. As shown in Figure 5, its polynomial equation was found to be:

$$SD = 0.56708 - 0.10563X + 0.016801X^2$$

where SD is the assay standard deviation and X is the measured serum concentration.

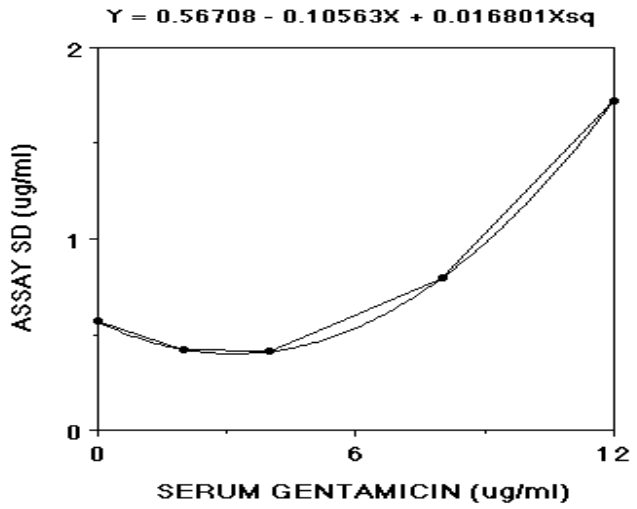

Figure 13. *Error pattern of an EMIT assay for Gentamicin and its associated polynomial equation.*

That assay had an SD of 0.57 ug/ml at 0.0 ug/ml (the blank), with variance 0.326 and weight (1/variance) of 3.07. The SD fell to 0.40 ug/ml at 3.0 ug/ml, with weight 6.25, double that of the blank. The SD then rose to 0.8 at 8.0 ug/ml and to 1.7 ug/ml at 12.0 ug/ml, when the variance was 2.89 and the weight was 0.346. The weights of these five representative points on the curve thus range over a factor of 18! The coefficients of the assay error polynomials are stored in the various USC*PACK Gentamicin programs.

Similarly, in Figure 14, the error pattern of the Abbott TDx assay for vancomycin used in our laboratory was found to be:

$$SD = 0.30752 + 0.024864X + 0.0002763X^2.$$

There was no increase in the SD as the blank was approached.

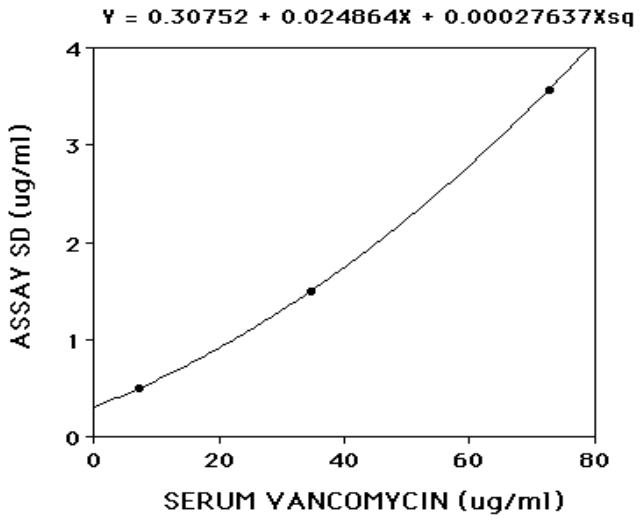

Figure 14. *Error pattern of an Abbott TDx assay for Vancomycin and its associated polynomial equation.*

6.23 An Examination of the College of American Pathologists Survey

The College of American Pathologists (CAP) sends out sample specimens containing stated drug concentrations to many clinical laboratories, which report their findings back to the College. The College then publishes the means and SD's of these findings, and the number of laboratories reporting. The results are broken down by the drug and by the type of assay used.

Two of us (RJ and BT) examined the survey results published by the College in Data Sets 1987 ZM-D, 1988 Z-D, 1989 Z-B, Z-C, and Z-D, and 1990 Z-A, Z-B, and Z-C, for Amikacin, Gentamicin, and Digoxin. We then took the published means and SD's of concentrations for the various specimens and fitted them with a polynomial, to provide a library of error patterns for the above assays. These polynomial equations can then be used to estimate the probable SD of individual serum samples for population pharmacokinetic modeling or for Bayesian fitting of pharmacokinetic models until each laboratory can determine its own assay error patterns.

6.24 Results Found

Amikacin

Fifteen sample means, ranging from 1.1 to 30.0 ug/ml, and their SD's, obtained from 339 to 725 reporting laboratories, provided the data. The following polynomial equations for the error patterns were found.

Abbott TDx SD (ug/ml) = $0.30156 + 0.0053855C + 0.0011184\ C^2$, $R^2 = 0.983$

Dupont ACA SD (ug/ml) = $0.46475 + 0.0281310C + 0.0021305C^2$, $R^2 = 0.939$

Syva Emit SD (ug/ml) = 0.23237 + 0.0470150C + 0.0016876C^2, R^2 = 0.965

All Methods SD (ug/ml) = 0.32272 + 0.0183650C + 0.0012051C^2, R^2 = 0.983

The Abbott TDx assay was the most precise. The Dupont ACA and Syva Emit assays were less so. The results found for all methods were heavily dominated by those found with the Abbott TDx assay, as so many laboratories used it.

Gentamicin

Seventeen sample means, ranging from 0.9 to 17.8 ug/ml, and their SD's, obtained from 2512 to 3600 reporting laboratories provided the data. The following polynomial equations for the error patterns were found.

Abbott TDx SD (ug/ml) = 0.02458 + 0.04948C + 0.0020318C^2, R^2 = 0.957

Dupont ACA SD (ug/ml) = 0.25719 - 0.016215C + 0.0081998C^2, R^2 = 0.982

Syva Emit SD (ug/ml) = 0.14078 - 0.002263C + 0.0184060C^2, R^2 = 0.991

All Methods SD (ug/ml) = 0.09114 - 0.043524C + 0.0045964C^2, R^2 = 0.992

The Abbott TDx assay was the most precise.

Digoxin

Seventeen sample means ranging from 0.2 to 3.0 ng/ml, and their SD's, obtained from 3160 to 4454 reporting laboratories provided the data. The following polynomial equations were found.

Abbott TDx SD (ng/ml) = 0.09211 + 0.0088626C + 0.0099406C^2, R^2 = 0.948

Baxter Stratus SD (ng/ml) = 0.144211 - 0.048708C + 0.022917C^2, R^2 = 0.911

Clinical Assays SD (ng/ml) = 0.086719 + 0.017052C + 0.011857C^2, R^2 = 0.881

Dupont ACA SD (ng/ml) = 0.15560 - 0.056293C + 0.035574C^2, R^2 = 0.562

Syva Emit SD (ng/ml) = 0.16111 + 0.051579C, R^2 = 0.451

All Methods SD (ng/ml) = 0.12312 - 0.0073104C + 0.020257C^2, R^2 = 0.951

Figure 15 shows graphically the above relationship for the Abbott TDx assay, while Figure 16 shows that found for the Syva Emit digoxin assay.

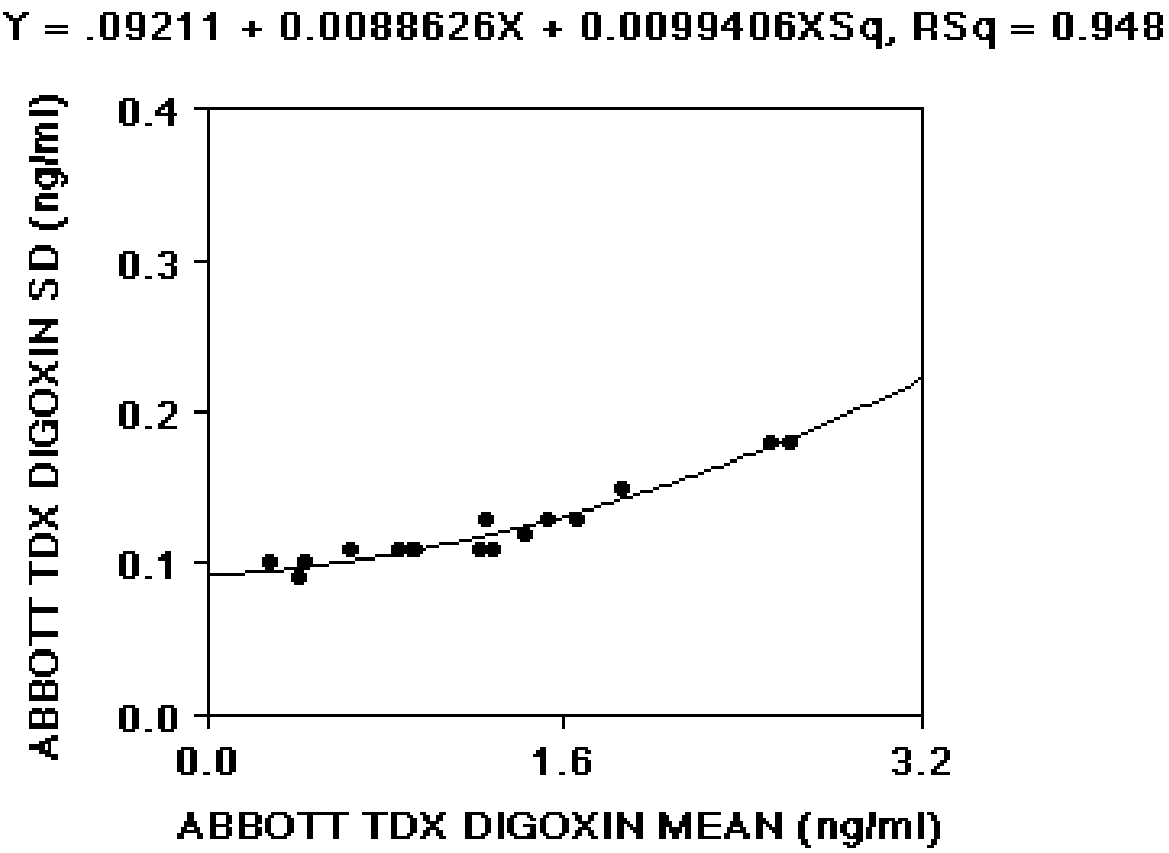

Figure 15. *Error pattern of the CAP data for the Abbott TDx assay for Digoxin, and its associated polynomial equation.*

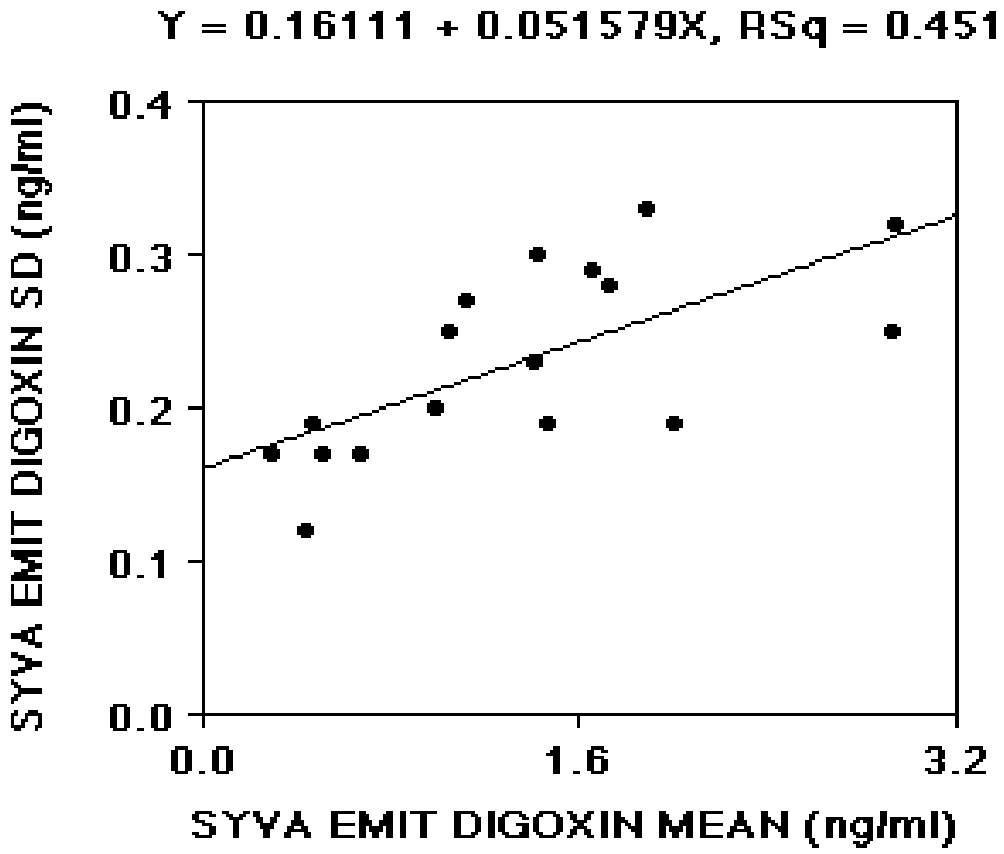

Figure 16. *Error pattern of the CAP data for the Syva EMIT assay for Digoxin, and its associated polynomial equation.*

The Abbott TDx assay was the most precise, and also had the highest R^2 (coefficient of the determination). When the Syva Emit assay findings were fitted with a second order polynomial, the curve reached a peak and then bent downward. This could yield dangerously low estimates of the SD when extrapolated beyond the range reported here (0.2 to 3.0 ng/ml). Because of this, and because the first order polynomial had almost the same value of R^2, the first order equation was used. The Syva Emit and Dupont ACA assays had the lowest values of R^2, showing that their error pattern had much more scatter around these equations. Figure 8 shows this for the Syva Emit assay. In contrast, the Abbott TDx assay, as shown in Figure 7, had a high R^2, showing that its error pattern was much less scattered around its equation. The Clinical Assays error pattern was intermediate. The Abbott TDx assay was the most precise.

6.25 Sources of Error

The means and standard deviations reported by the College survey are a mixture of within - run and between - run laboratory results, as well as within - laboratory and between - laboratory results. They appear to be useful for purposes of therapeutic drug monitoring and Bayesian pharmacokinetic modeling until such time as each clinical

laboratory can determine its own assay error patterns for its own drugs and develops their own polynomial equations. Fortunately, this is easily done, and at minimal cost.

One might express concern that using the above approach assumes that the entire residual variability can be expressed as the assay SD, and thus grossly overestimates the effect of the assay errors compared to other sources of error. This is not correct. However, more than previous approaches, this approach does correctly show the effect of assay error upon the parameter values obtained, using the standard objective function for Bayesian fitting described above.

Other sources of error not yet considered in this approach are the errors in recording when serum samples were obtained, when the doses were given, model mis-specification, and errors in preparation of the doses. The result of this is that the other errors mentioned above will distribute themselves as uncertainties in the parameter values obtained, and will reduce the ability of any method of adaptive control to achieve the desired therapeutic goals with optimal precision. It has been shown, in the context of aminoglycoside therapy [29], that the contribution of errors in preparation of the doses and in recording when the doses were given contribute significantly more to interfere with therapeutic precision than do errors in the assay or the phlebotomy service (recording the timing of samples drawn).

What really is needed here is to implement the pharmacokinetic models in stochastic rather than deterministic form, with a parameter in the dynamic equations for process noise to take into account the errors due to model misspecification and in preparation and administration of the doses. Such models have been implemented in prototype form [41], but not yet in clinical software. Further, another parameter is needed to model the effect of errors in timing of the serum samples. This has been implemented in the NONMEM population modeling program, but not, to our knowledge, in clinical software. All such efforts are in the future. For the present, using the current Bayesian approach to adaptive control of dosage regimens (and its objective function), it seems useful to know explicitly what the assay error is, to consider its effects correctly upon the Bayesian objective function, and to work toward the future when improved models having both process noise and measurement noise parameters will become available (see further on).

6.26 The Importance of Measuring Blanks

It is interesting that in none of the samples sent out by the College was there a blank sample. Clinical laboratories, however, usually characterize the sensitivity of their assays by choosing a value two SD's above a blank. When concentrations lower than those clearly detectable are encountered, they are often simply reported as being "less than X", where X is two SD's above the blank.

6.27 The Importance of Reporting Low Concentrations Below "Detectable Limits"

In toxicological analysis one must make a firm decision as to whether a substance is present in the body or not, without any clinical history. Because of this, minimal detectable limits of an assay are set up, usually 2 SD above the blank, and results lower than this are reported as being below this limit. However, this practice is actually an obstacle to optimal Bayesian therapeutic drug monitoring. In therapeutic drug monitoring there is usually no

question that the drug has been given. In a hospital setting, or in a reliable outpatient setting, for example, one clearly knows this from the history, the orders, and the chart. Indeed, many clinical laboratories will not measure a serum drug concentration unless the time since the last dose is clearly stated on the request slip. Since the patient never excretes the last molecule of the drug, there is no question that the drug is still present in the body. The only question is how much. Low trough aminoglycoside concentrations for example, below those clearly detectable, are needed to use, when found, for therapeutic drug monitoring and Bayesian pharmacokinetic fitting and modeling. Not reporting such results renders that measurement useless for Bayesian modeling. A vital data point has been withheld.

One might say that when compliance is unknown, the above point becomes moot. However, compliance can be evaluated by looking at the apparent volume of distribution of the drug, and to some degree quantified. If a patient's volume of distribution is twice as much as it should be, then perhaps only half the dose has been taken or absorbed. Because of this, the practice of reporting a serum concentration both as the **measured value itself** as well as the **lower detectable limit**, will let the serum concentration result be much better used to make the Bayesian fitted model, and to evaluate and perhaps to quantify compliance. In addition, it still permits the traditional toxicological report as well.

Rather than reporting a Gentamicin concentration as "less than 0.5 ug/ml" for example, the laboratory can easily report the actual value found, and can report it as "0.1 ug/ml, below secure detectable limits of 0.5 ug/ml", for example. This procedure is easily done and will answer both the needs of the toxicologists and the pharmacokineticists.

6.28 The Importance of Collecting High Serum Concentrations

The CAP Survey paid most attention to determining the laboratory errors for concentrations within the therapeutic ranges of the drugs under consideration. However, low trough concentrations, well below the usual detectable limits, are frequently found. Because of this, one might suggest that more such low concentrations, and especially blank concentrations, be included in future surveys.

It is equally important to know the errors of concentrations found well into the toxic range. Because of this, when very high concentrations are encountered, one might suggest that the laboratory run them in replicate as many times as possible, to better characterize the error of the assay at its high end, and to extend the range of the known assay error as widely as possible. When dilutions must be made to assay high levels, these can also be done in at least quadruplicate analyses of each sample to find the actual SD of the assay as performed using the dilutions.

6.29 The Importance of Improving Assay Precision at the High End

Proper weighting of serum level data, using the correct assay error pattern, is essential. Inaccurate assay error patterns or simple assumptions of a certain coefficient of variation can and do lead to grossly inaccurate model parameter values, both in individually fitted patient pharmacokinetic models and in population pharmacokinetic models.

When doing population modeling or Bayesian fitting, one can only give equal weight to various serum concentrations when they have the same SD. An assay with a constant SD over its working range is said to be homoschedastic. Such an assay will have a coefficient of variation that decreases by half as the concentration doubles. None of the assays evaluated here were homoschedastic, and such an assay is probably an unobtainable ideal.

In contrast, a heteroschedastic assay error pattern is one in which the assay SD changes over its working range. Even an assay with a constant coefficient of variation is heteroschedastic. With such an assay, as the concentration doubles, the SD also doubles, the variance quadruples, and the weight given to the assay is reduced to one fourth. If one assumes a constant coefficient of variation, a concentration of 1.0 ug/ml, for example, has a weight 100 times greater than that of a concentration of 10.0 ug/ml, and a concentration of 0.1 ug/ml has a weight 100 times that of the concentration of 1.0 ug/ml, and 1000 times that of the concentration of 10.0 ug/ml! Because of this, when a constant coefficient of variation is assumed for an assay used in Bayesian fitting, high concentrations will be relatively ignored compared to lower ones, and the fitted model will not reach out to and fit the high concentrations as closely as one perhaps might wish them to.

This is also true for the polynomial equations described above. The difference is that the polynomial equations are carefully derived from empirically measured SD's over the entire working range of the assay, and should include blank concentrations as well. Because of this, they are a more correct estimate of the assay error pattern over its working range, and the fit, while often appearing to ignore the high concentrations, is actually being correctly done by current standards, using the criterion of the Fisher information [31] as employed in the usual Bayesian objective function shown previously.

One of two things needs to be improved. Either the current Bayesian fitting procedure based on the Fisher information of the data points is incorrect, or the assays need to have their precision improved at the high end to make them, if possible, more homoschedastic. To discard the concept of Fisher information would be to overthrow several decades of carefully acquired and searchingly criticized mathematical and statistical knowledge. To improve the precision of assays at their high end is probably the most constructive thing to do. It may even be possible, for example, to alter the ratios of reagents so that the ratio of bound and unbound drug in the assay can be changed, with a resultant change in the error pattern toward homoschedasticity.

7.0 MODELING POPULATION BEHAVIOR OF DRUGS

7.1 MODELING LINEAR PHARMACOKINETIC SYSTEMS.

As we acquire experience with the clinical and pharmacokinetic behavior of a drug, it is usually optimal to store this experience in the form of a population pharmacokinetic model, and then to relate the behavior of the model to the clinical effects of the drug or to a linked pharmacodynamic model. The role of population modeling is thus to store our experience with the behavior of a drug in a certain group or population of patients or subjects.

The traditional method of Naive Pooling has been used for population modeling when experiments are performed on animals, for example, which must be sacrificed to obtain a single data point per subject. Data from all subjects is then pooled as if it came from one single subject. One can estimate pharmacokinetic parameter values, but cannot estimate the variability between the various subjects making up the population.

The Standard Two Stage (S2S) method has been used when there is enough data to obtain parameter estimates for each individual subject. That is the first stage. Then one finds the population means and standard deviations (SD's) from the various individual parameter values, and also examines their frequency distributions. An extension of this method is the Iterative Two Stage (I2S) method in which the initial population model obtained above is then used as the Bayesian prior, and each patient's individual parameter values are found once again, this time as the Bayesian posteriors. The population means and SD's are found once again, and once again each patient's Bayesian posterior parameter values are found. The method then iterates over and over again until stable individual and population parameter values are found. The Global Two Stage (G2S) method is a further refinement of the S2S and the I2S in which the covariance and correlations between the parameters are also estimated. Each of these methods, and all those below, require software to implement them.

True population modeling began with the Nonlinear Mixed Effects Model with first-order approximation (NONMEM) of Beal and Sheiner [32]. In contrast to the above methods, the overall population, even if it has only 1 data point per subject, almost always supplies enough data for this approach. This algorithm and computer program estimates means, SD's, and covariances of population parameter values. However, it has sometimes given different answers from the other above methods [33].

Mallet [34] then showed that the joint probability distribution of the parameter values in a population model is actually discrete. With his Nonparametric Maximum Likelihood (NPML) method, the joint probability density function (PDF) is optimally supported at a number of points not more than the number of subjects studied. This can be seen intuitively by saying that if we really knew each patient's true parameter values, the population PDF would be optimally represented simply by a scatterplot, for example, of the true individual parameter values for each subject in the population.

The NPML method thus estimates the entire collection of points, and their probabilities, representing the joint population PDF of the parameters, without having to make any traditional (parametric) assumptions about the shape of that distribution, such as Gaussian, etc, in which an equation (and its parameters, such as mean and standard deviation) determine the assumed shape of the PDF. In contrast to NONMEM, the NPML method also permits explicit discovery of previously unsuspected clusters of patients such as fast and slow metabolizers of a drug.

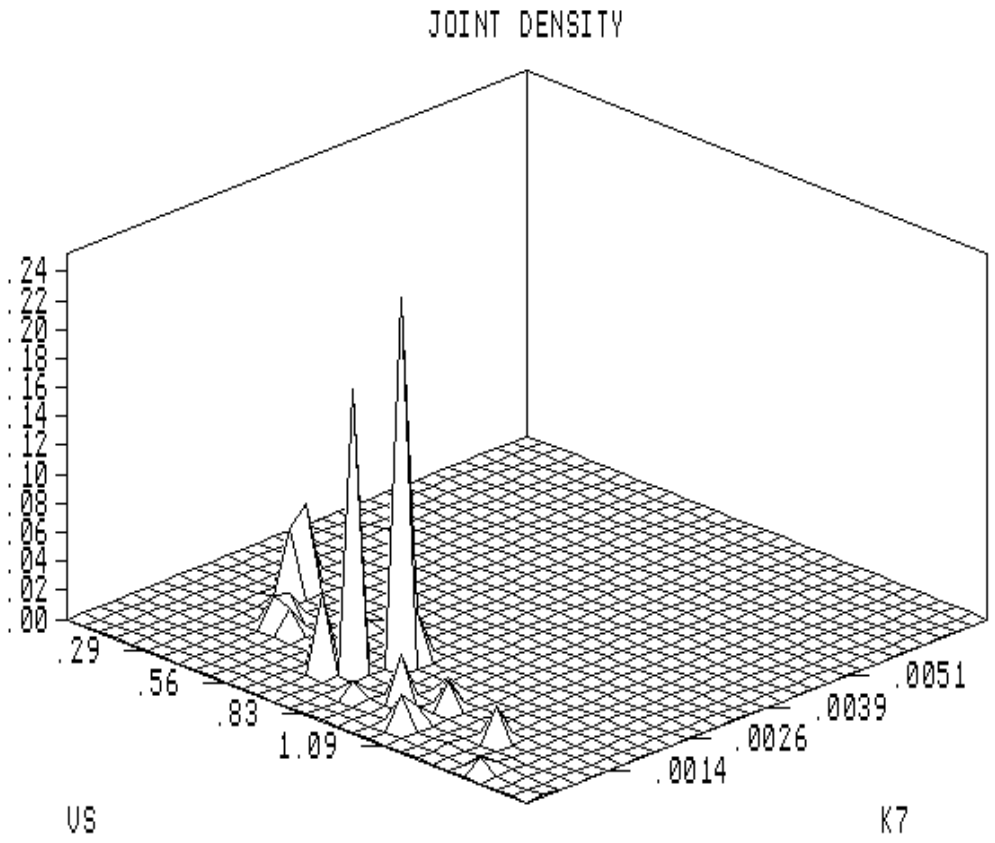

Figure 17. *3D plot of a population joint density function. Vs = slope of volume with respect to body weight. K7 = slope of Kel with respect to creatinine clearance.*

Figure 17 shows another approach to the NPML solution. This is the Nonparametric Expectation Maximization (NPEM) algorithm of Schumitzky [35]. It uses continuous PDF's which in the limit go to the discrete PDF's of Mallet.

If each subject's parameter values were truly and exactly known, and if we were examining two parameters such as volume of distribution (V or Vslope) and elimination rate constant (K or Kslope), then the joint population distribution of these parameter values would actually be the collection of each patient's individual points on a scatterplot. However, since these values can never be known exactly, but have to be estimated from data of doses given and serum concentrations found, one will find spikes, the location and height of which reflects their estimated values and their probability, as shown in Figure 17, obtained using the NPEM algorithm.

7.2 MODELING LARGE NONLINEAR KINETIC/DYNAMIC SYSTEMS.

A number of programs for modeling such as ADAPT II, MKMODEL, and PC NONLIN are available. In those programs, the differential equations describing the kinetic and dynamic behavior of the drug must be written. BOXES, a newer mouse-driven PC program for making such models [36], simplifies this process and lets one place boxes on the screen for compartments, connecting them with arrows for pathways. It then

automatically writes the differential equations and generates the Fortran source code for the model part of the program package, compiles it (one must also have IBM Professional Fortran), and then links it with object modules to make, first, an executable program for parameter identification, using the Nelder-Mead simplex algorithm and weighted nonlinear least squares, and, second, another program for model simulation. BOXES can also be used with a text editor to put the proper equations into the model portion of the ADAPT II programs. Both linear and Michaelis-Menten pathways can be made, and can be linked to clinical descriptors. Effect models (Emax, Sigmoid Emax, or Keo) can also be made, 1 effect from each compartment. Up to 3 different effects may combine to act upon a single pathway to increase or decrease its rate.

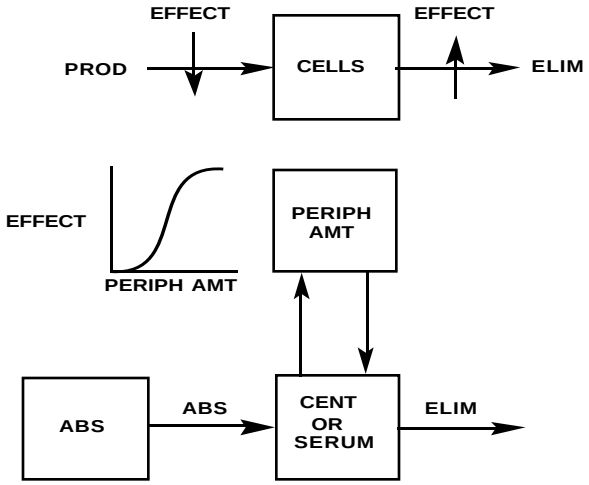

Figure 18. *Model of a drug having an absorptive, central (serum), and peripheral (nonserum) compartments, with a Hill effect model related to the amount of drug in the peripheral compartment. That effect can act either to reduce the production of a cell line or to accelerate its destruction.*

A representative effect model is shown in Figure 18. It could be used, for example, to study the pharmacokinetic behavior of a drug such as trimethoprim, and to describe its toxic effects (reduction of white blood cell count, for example) along with its kinetic behavior. A similar simulation program for the Macintosh is Stella, but apparently without a module for parameter identification.

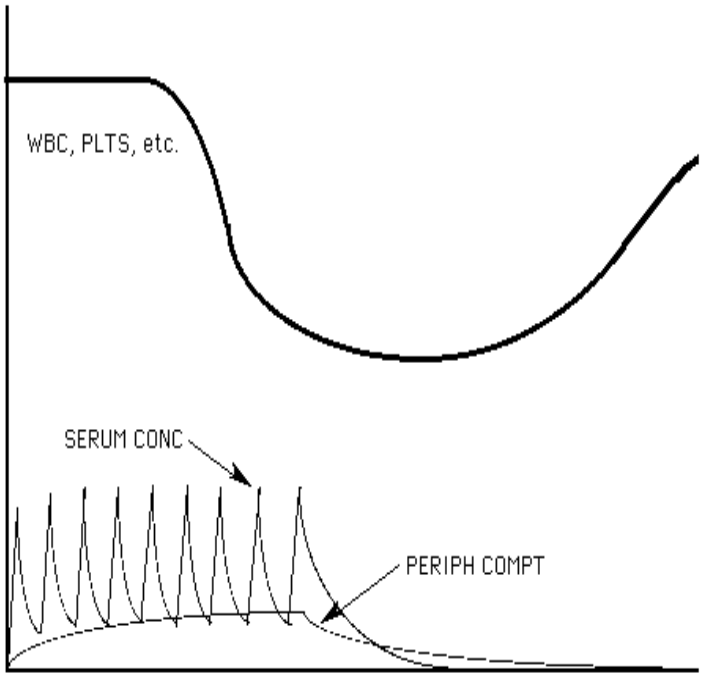

Figure 19. *Simulated data of the serum concentrations of a drug having hematologic toxicity, its slow accumulation in a peripheral compartment, and the result of the effect of the drug on the production or destruction of a cell line. The relationship between serum levels and the effect can be revised by adaptive control during the first course of therapy, not having to wait until after it has been completed to establish a correlation.*

Figure 19 shows the time course of the effect of the above drug to decrease white cell production, for example, with its delayed onset, its sharper fall, and its delayed recovery after therapy is stopped, as can be seen with trimethoprim and many other drugs. One thus can capture, in a linked manner, the delayed effect of a drug on cells, such as RBC, WBC, or platelets, for example, and the delayed recovery after therapy is stopped, with linked description of both the effect and its results on cells in a single overall model. In addition, since one can model such relationships for a single drug, one now can also do this to model combined therapy with several drugs that share a common effect, permitting quantitative descriptions of combined drug effects. BOXES can also be used to rapidly model the effects of drugs upon other drugs, and thus to yield quantitative models and descriptions of drug interactions.

In summary, population pharmacokinetic and dynamic models are used to store and describe our experiences with drug behavior in patients. They provide our initial guides from which to develop dosage regimens of drugs to achieve and maintain explicitly selected individualized therapeutic goals for each patient.

7.3 MODELING ENVIRONMENTAL NOISE FACTORS

It has been shown by Monte-Carlo methods [29], during experimental studies [37], by population pharmacokinetic studies [38], or during routine clinical drug monitoring [39,40] that nonpharmacokinetic factors in the patient's therapeutic environment can affect both the precision of the patient's estimated pharmacokinetic parameter values and the precision with which the desired serum drug levels can be achieved.

Predictions of future amikacin serum levels made during two separate studies whose major difference was the method and care with which the clinical data were collected. In one study the clinical data was collected by a trained pharmacy resident; in the other, by pharmacokinetically untrained nurses. A nonrandomized retrospective comparison of predicted serum amikacin concentrations based on Bayesian fitted pharmacokinetic models where the data collection was obtained under two different sets of circumstances [30].

7.31 Group 1: Resident Collected Data (RCD) Group:

Here, prospective data collection was done by a trained pharmacy resident specially assigned to this work.

Ten patients were studied in this group. Dosage adjustments were performed during therapy using the MAP Bayesian method [14]**.** For this group of patients, from Monday to Friday from 8:30AM to 6:30PM, and on Saturday from 8:30AM until 12:00 noon, information about drug therapy, doses given, start and stop infusion times and serum sampling times were ordered and collected by a resident specially trained in and assigned to this work.

7.32 Group 2: Nurse Collected Data (NCD) Group:

In this group, prospective data collection was done by the nurses as part of their usual nursing duties. They had no special pharmacokinetic training. This group consisted of 23 elderly patients. Dosage adjustments were performed using the same MAP Bayesian method.

For both groups, either the resident or the nurses were asked to collect four serum levels for the first set, one just before the infusion, one at 1/2 hour after the end of the 1/2 hour infusion (or at 1 hour after the dose if the dose was given intramuscularly), and one each at 3 and at 6 hours after the end of the 1/2 hour infusion. Preparation of Amikacin doses by nurses, and their administration by minibag via gravity flow, both were done under the same conditions in both studies.

For both groups, each patient had two sets of serum levels obtained. For both groups, the nurses and the residents each were asked to collect four serum levels for the first set as described above. For the second set, four serum levels were also to be obtained, usually 2 to 8 days after the first set, and at the same times as for the first set. For all patients, there was at least one sample in each set obtained 30 minutes after the end of the infusion or one hour after intra-muscular injection (the "peak" value), and one sample just before the next infusion (the "trough" value).

The MAP Bayesian method was used to fit either a one compartment model (B1) or a two compartment model (B2) to the data of the doses and the first set of serum

concentrations. Each patient's fitted model was then used to predict prospectively the subsequent serum concentrations which later were achieved on the subsequent dosage regimen he or she received. The relationship between predicted (P) and measured (M) serum concentrations was then examined using a scattergram of the data, by evaluating their correlation coefficient, by calculating the bias or mean weighted error (MWE), and by computing the precision or mean weighted squared error (MWSE). In addition, the percentage of serum levels accurately predicted (%SLAP), defined as being centered on the predicted concentration ± 20% , was determined.

7.33 Data Analysis:

In each study, the frequency distribution of the individual weighted errors was found to have a Gaussian distribution. The t-test for differences between the means of two populations was used to evaluate their differences.

However, the weighted squared errors are not normally distributed. The mean weighted squared error (MWSE) for the two groups was therefore compared using the Mann-Whitney U test. In addition, the percent of serum levels accurately predicted, within ± 20 percent, (%SLAP) for each group were compared by Chi-squared analysis. When the expected number in a cell was less than five, the Yates modified Chi-squared test was performed.

7.34 Results found

There was no significant difference between the two study groups with respect to sex, age, height, or initial estimated creatinine clearance. However, there was a significant difference between the two groups in the number of serum levels obtained for the first set of serum levels. The RCD group had 4 ± 0 samples obtained in the first cluster, while the NCD group had 3.26 ± 0.45 samples drawn, a significant difference.

Table 2 shows the correlation coefficient (R) between predicted and measured serum levels and the Bayesian predictive performance with the <u>one</u> compartment fitted model for the RCD and the NCD groups. The MWSE was significantly less in the RCD group (64 versus 271 units), and the %SLAP was significantly greater (58% versus 30%).

Table 3 shows the correlation coefficient (R) between predicted and measured serum levels and the Bayesian predictive performance with the <u>two</u> compartment fitted model for the RCD and the NCD groups. The MWSE here was also significantly less in the RCD group (82 versus 329 units), and the %SLAP was also significantly greater (72% versus 40%).

Tables 4 and 5 show for the RCD group and the NCD group the percentage of peak (PSLAP) and trough (TSLAP) serum levels accurately predicted in the interval centered on the predicted concentration ± 20%, obtained respectively with the <u>one</u> compartment (B1) and the <u>two</u> compartment (B2) models. Both PSLAP and TSLAP were greater in the RCD group, significantly so for PSLAP (80% versus 30%) for the B1 model. For the B2 model,

both PSLAP and TSLAP again were greater in the RCD group, significantly so for TSLAP (80% versus 26%).

Because there was a significant difference between the two groups in the number of serum levels obtained for the first set of serum levels, the amount of information available for each group might be different. This fact might influence the Bayesian predictive performance and thus confound the comparison, as the group with more levels and information in the first cluster might be able to predict subsequent levels more accurately because of that fact. It was clear, however, as described above, that the NCD group obtained significantly fewer serum levels than did the RCD group, and thus adhered to the desired monitoring protocol significantly less well.

Another comparison was then done, using only data of peak and trough levels for each patient in each group. The results of this second comparison basically confirmed the results obtained above [30].

Table 2: Comparison of Predicted Versus Measured Serum Levels in the RCD and NCD Groups Using the MAP Bayesian One-Compartment Fitted Model.

	RCD	NCD	p Value *
R	0.94	0.82	-
MWE	-3.6	-6.6	N.S. **
MWSE	63.9	270.6	<0.05
%SLAP	56%	27%	<0.01

*Comparing the groups
**N.S. = Not significant
R = correlation coefficient
MWE = Mean Weighted Error
MWSE = Mean Weighted Squared Error
%SLAP = percent of Serum Levels Accurately Predicted (within ± 20%).

Table 3: Comparison of Predicted versus Measured Serum Levels in the RCD and NCD Groups, Using the MAP Bayesian Two-Compartment Fitted Model.

	RCD	NCD	p Value *
R	0.96	0.81	-
MWE	2.7	-2.3	N.S. **
MWSE	82.1	328.8	<0.05
%SLAP	81%	36%	<0.001

* Comparing the groups
** N.S. = Not Significant
R = correlation coefficient

MWE = Mean Weighted Error
MWSE = Mean Weighted Serum Error
%SLAP = percent of Serum Levels Accurately Predicted (within ± 20%).

Table 4: Percent of Serum Levels Accurately Predicted (within ± 20%) in the RCD and NCD Groups, Using the One-Compartment MAP Bayesian Fitted Model.

B1	RCD	NCD	p Value *
PSLAP	80%	22%	<0.01
TSLAP	30%	13%	N.S. **

* Comparing the Groups
** N.S. = Not Significant
PSLAP = Peak Levels Accurately Predicted
TSLAP = Trough Levels Accurately Predicted

Table 5: Percent of Serum Levels Accurately Predicted (within ± 20%) in the RCD and NCD Groups, Using the Two-Compartment MAP Bayesian Fitted Model.

B2	RCD	NCD	p Value *
PSLAP	90%	39%	<0.01*
TSLAP	80%	17%	<0.01*

* Comparing the groups
** N.S. = Not Significant
PSLAP = Peak Levels Accurately Predicted
TSLAP = Trough Levels Accurately Predicted

7.35 Implications for Therapy

It has been customary to emphasize the importance of accurately recording when doses were given, when infusions started and stopped, and when serum levels were drawn. Indeed, the importance of these nonpharmacokinetic factors in permitting or preventing precise achievement of clinically selected serum aminoglycoside therapeutic concentration goals was carefully evaluated in a rigidly controlled study in which Monte-Carlo simulations of a typical clinical scenario with tobramycin therapy were analyzed. In that study, it was shown that the ward care setting, in which the recording of data of dosage administration timing was well or poorly done, was the single most important environmental factor in permitting or preventing precise control of serum aminoglycoside concentrations to achieve specifically selected therapeutic goals, having a greater effect on predictions than the pharmacy (in preparing the doses precisely or not) or the laboratory assay errors [29].

The results of that simulation study were confirmed and extended by the results of the above clinical study, which show that when trained personnel are used to record the data of dosage administration times and serum sampling times, significantly greater therapeutic precision was obtained in the achievement of the desired therapeutic serum concentration goals associated with aminoglycoside therapy than when untrained nursing personnel recorded the data, and more of the serum levels which should have been obtained were actually drawn. Further, this was true even when only data of peak and trough serum concentrations was used, rather than the full data set [30]. These results emphasize the importance of careful attention to recording when doses are given and serum samples drawn. They stand in distinct contrast to the widespread policy of nursing personnel to

record a dose as having been given exactly when ordered if it is given within 30 minutes of the desired time.

8.0 OPTIMAL DESIGN OF DOSAGE REGIMENS

8.1 THE LIMITATIONS OF MAP BAYESIAN ADAPTIVE CONTROL.

MAP Bayesian adaptive controllers are based on single-point estimators of the PK parameter values. Dosage regimens based on such single point estimates will achieve the desired therapeutic goal exactly, for a patient having these exact parameter values. It never considers that the goal will not be exactly achieved. However, no patient's parameter values are ever known exactly. Because of this uncertainty, which is especially great with the initial dosage regimen, based solely on a population model, prior to obtaining any feedback from serum level data, the goal is never achieved exactly. Some error, great or small, is always present in the achievement of the desired goals. Until now, this fact has not been faced. Neither has the fact been faced that errors are present due to the environmental noise factors described earlier.

The diversity of patient behavior is captured in a discrete nonparametric probability distribution such as that developed by the NPEM2 programs. Each spike in Figure 17 above represents a candidate model of the patient, each with its parameter values, and each with its probability of representing the true patient. We thus have many possible versions or contending multiple models (MM) of the patient.

It is from this approach that the "Multiple Model" (MM) Adaptive Control Strategy was recently developed [41]. This strategy recognizes the fact that we can only give a patient a single dosage regimen at any given time, and that, because of this, it should be one which specifically achieves the desired goals as precisely as possible. In this case, the MM dosage designer described below develops a regimen specifically to minimize the expected value of the squared error in the achievement of the desired goals, when that regimen is given to all the multiple competing models of the patient. This approach is especially well suited for cancer chemotherapy, where one must often commit the patient to a single dose or course of therapy before any information becomes available from the serum levels or other measures of response.

8.2 "MULTIPLE MODEL" (MM) DOSAGE DESIGN

The scenario of initiating therapy with lidocaine, or with any other drug having a central and a peripheral compartment linear model, when the goal is to achieve and maintain a stable serum level with a series of continuous but stepwise constant intravenous infusions, has been examined with clinical simulations.

In this clinical simulation, the MM control strategy described above was compared with that of the current MAP control in current clinical use. A model of lidocaine having 4 parameters, with 3 values for each parameter, each with its own probability, was developed, resulting in 81 possible combinations of parameters (multiple models {MM} of the patient). The mean values for each parameter were also found. When the dosage

regimen developed by the current MAP control strategy, designed to achieve and maintain a serum level goal of 3 ug/ml (resulting in exact control of the "mean" patient) was given to the 81 possible versions of the patient in the MM model, the trajectories of serum levels actually ranged from approximately 0.9 up to as high as 9.7 ug/ml, showing great diversity in patient response.

In contrast, when the MM dosage strategy was used, the resulting variation in serum levels was visibly reduced, with a range only from 0.7 to 6.0 ug/ml. The MM control strategy, specifically designed to minimize the expected value of the squared error about the achievement of the stated goal, did so unequivocally, specifically minimizing the diversity in patient response around a selected goal [41].

In an extension of that work, a further simulated Lidocaine scenario was examined. However, it now incorporated state and parameter updating using Bayesian feedback. As above, the 81 point multiple model discrete prior was used, with two-compartment dynamics, and the same type of piecewise constant continuous IV input was considered.

In this simulation study, the feedback cycle (event horizon) was set up so that the intravenous infusion rates were changed at times t=0, 5, 20, 50, 110, 170, and 230 minutes. Serum levels were obtained at t=5, 20, 50, 110, and 170 minutes, and made available prior to 230 minutes. At 230 minutes a new regimen of similar format was developed and given, and a new similar event horizon began. This process was repeated for a total of 4 such event horizons. The variability in patient response was plotted, and could be seen to converge as the therapy continued. Specifically, to 99% probability, the patient's serum concentration was within ± 1.0 ug/ml of the desired 3.0 ug/ml goal by the second horizon, and within ± 0.2 ug/ml by the third horizon.

Still another simulation study has now been done by our laboratory in which Vancomycin regimens were developed to achieve and maintain a stable serum level of 15 ug/ml. In contrast to the lidocaine simulation above, this one was based on an actual population model of Vancomycin made using the NPEM2 program [42]. In this study, three 2 hour Vancomycin infusion steps were followed by three 6 hour steps, to complete a 1 day feedback cycle of therapy (event horizon). Simulated serum levels were obtained at 2, 4, and 8 hours into the regimen. At the end of the cycle, a new regimen based on feedback from the serum levels was implemented, and the cycle repeated. This continued for 4 such simulated days of therapy. This scenario was also compared with MAP Bayesian adaptive control during the same 4 day period.

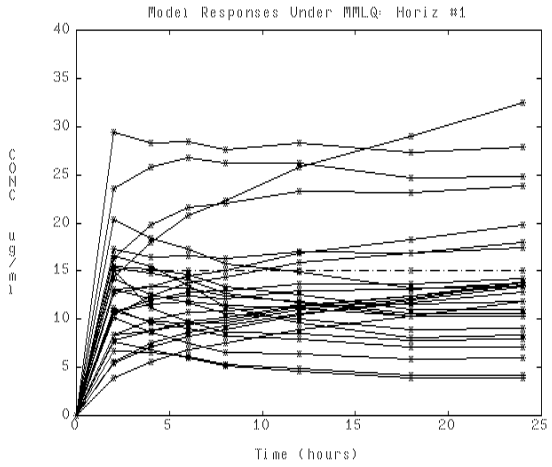
Model Responses Under MMLQ: Horiz #1
CONC ug/ml
Time (hours)
0
5
10
15
20
25
30
35
40

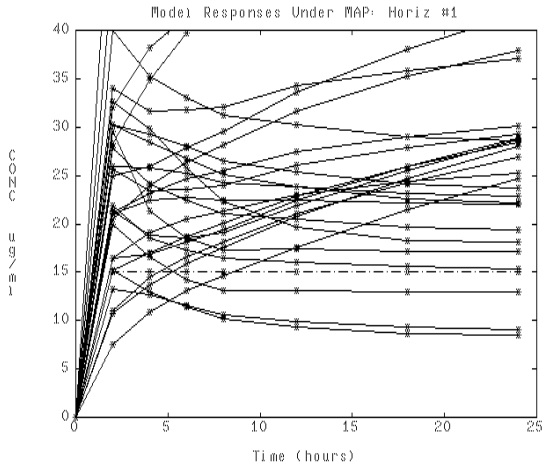
Model Responses Under MAP: Horiz #1
CONC ug/ml
Time (hours)
0
5
10
15
20
25
30
35
40