University of Southern California

School of Medicine

Laboratory of Applied Pharmacokinetics

Technical Report 92-2

Adaptive Control of Drug Dosage Regimens: Basic Foundations, Relevant Issues, and Clinical Examples

by

Roger W. Jelliffe, M.D.,Fred Sattler, M.D.,
Pilar Gomis, Pharm.D.,and Babak Tahani

Laboratory of Applied Pharmacokinetics,University of Southern California
School of Medicine,2250 Alcazar St., CSC 134-B,Los Angeles, CA 90033 USA

ADAPTIVE CONTROL OF DRUG DOSAGE REGIMENS: BASIC FOUNDATIONS, RELEVANT ISSUES, AND CLINICAL EXAMPLES.

Roger W. Jelliffe, M.D., Pascal Maire, Pharm.D., Fred Sattler, M.D., Pilar Gomis, Pharm.D., and Babak Tahani

Laboratory of Applied Pharmacokinetics, University of Southern California School of Medicine, 2250 Alcazar St, Los Angeles California 90033 USA

INTRODUCTION

In adaptive control (AC) of drug dosage regimens, we first must have a system or model of the process to be controlled, and then design a dosage regimen to control that model to achieve a desired goal. For individualizing drug therapy for patients, we begin by making a pharmacokinetic (PK) or pharmacodynamic (PD) model of the behavior of the drug and perhaps also of its effects such as possible toxicity on white blood cells or platelets, for example. Because of this, AC is more than therapeutic drug monitoring (TDM), although TDM is a key part of the overall process of Bayesian forecasting (BF) and AC. To begin with, let us examine the similarities and differences between AC, BF, and TDM. Some of the following discussion may be rather obvious, but it is often useful to consider the various basic issues involved.

THERAPEUTIC DRUG MONITORING AND THERAPEUTIC AND TOXIC RANGES

As currently done, TDM has consisted primarily of monitoring serum drug concentrations and classifying them into subtherapeutic, therapeutic, and toxic ranges. As shown in Figure 1, as the serum concentrations of a drug become greater, at first one sees an increasing incidence of clinical responses to the drug which we like to see in our patients. As the incidence of such behavior becomes appreciable, for example at the upward bend in the left hand curve in Figure 1, one can generally say the the concentrations no longer are subtherapeutic, but that the therapeutic range has now been entered. The therapeutic range extends from that rough point up until the incidence of toxic behavior becomes appreciable, usually as shown by the upward bend in the right hand curve of Figure 1. Beyond this lies the toxic range. The subtherapeutic range therefore is that in which the effects of the drug are small, infrequent, or both. The therapeutic range is that in which the therapeutic effects are easily seen and

frequent, but in which the toxic effects are still small and/or infrequent. The toxic range is that in which the toxic effects are easily seen, frequent, or both.
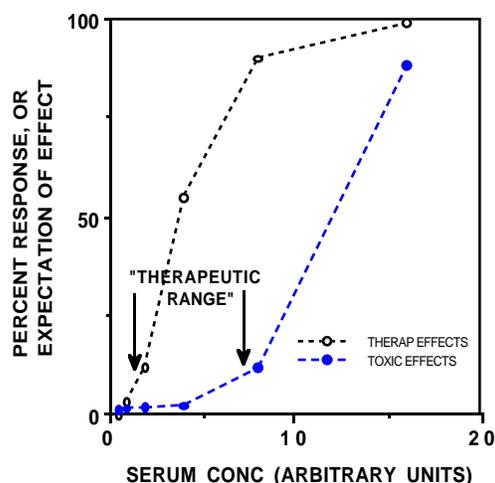

Figure 1. *Graph of the general relationship between incidence of therapeutic and toxic effects that leads to the development of a therapeutic range. The important quantitative relationship, however, is that between the serum level and the incidence (or risk) of toxicity.*

In the usual way TDM has been practiced, one has monitored serum drug concentrations, and has interpreted the meaning or significance of the results as subtherapeutic, therapeutic, or toxic. The clinician or clinical pharmacist then may adjust the dosage regimen in some manner, usually to better center the serum concentration within the therapeutic range.

A PROBLEM

All too often, however, dosage adjustment has been done with little or no reference to the actual clinical behavior of the patient. If the serum concentration is in the toxic range, that has often been enough to warrant lowering the dosage, regardless of whether the patient is behaving in a toxic manner or not. This problem of failing to look at the patient has become more common as TDM has

become more popular. It also does not consider the sensivity of the patient to the drug. Serum levels are often interpreted by laboratory personnel who do not see the patient. Nevertheless, they are deferred to by the clinician because "they are the ones who know what the therapeutic ranges are". This practice has inspired anxiety in the clinician, who fears he may be sued if his patient has a serum level interpreted as toxic, even if this is done by someone who has not seen the patient, and thus has no idea at all whether the patient is actually toxic or not.

A common example of this is the serum digoxin concentration. If the serum level is at all over 2.0 ng/ml, this frequently inspires fear in the clinician and clinical pharmacist, even if the patient should have atrial fibrillation for example, with a well-controlled ventricular rate and no evidence of toxicity, and even though it is well known that serum levels of 2.0 ng/ml and a dose of 0.5 mg/day are often required to achieve good rate control. The current practice of defining the therapeutic range of serum digoxin levels only for patients in sinus rhythm is a good example of the failure to think of the patient when therapeutic ranges have been developed. Instead of correlating the serum level with the behavior of each individual patient as reflecting that patient's therapeutic or toxic behavior, the clinician often becomes hesitant to continue with a relatively high serum level, as he also correctly knows that the margin of safety is less with higher levels, even though the patient is doing well. The serum level has now become a factor in the clinician's anxiety rather than the mirror to the patient's behavior.

Moreover, serum levels by themselves give us no dynamic information of the important relationships we need to know between the doses, the patient's body weight, his renal or cardiac function, the resulting serum levels, and the behavior of the drug in other parts (the nonserum compartment) of the body. The patient gets forgotten. He gets treated only as a member of a group of many patients which has a certain population therapeutic range of serum levels.

Further, traditional TDM gives us no information about what has gone on at all the other times when serum levels were **not** obtained. A steady state is the only clinical situation which is usually considered. Dynamic and changing states tend to be ignored. With some drugs, digoxin and digitoxin, for example, the important events appear to reflect not the serum levels, but rather the amounts or concentrations of drug in the peripheral, nonserum, compartment, and quite different clinical behavior can be observed at identical serum levels (see below).

In addition, many clinicians are afraid to give the higher drug doses such as digoxin doses of 0.5 mg/day or greater, for example, even though the patient under consideration may be having a clearcut problem of reduced bioavailability as evidenced by little clinical effect, low serum levels, and a large apparent central volume of distribution, which may reflect either poor bioavailability or poor compliance. Proper management of such a patient cannot be done without modeling the relationships described above, and either improving compliance or using the higher doses required to do the job, to approach and achieve the desired serum level and the desired clinical effect in the well-known series of steps which has been called "clinical titration". In this careful stepwise process, the use of AC brings a precision and a clinical confidence to this procedure which was not possible without AC, even when traditional TDM was used.

ADAPTIVE CONTROL, PHARMACOKINETIC AND DYNAMIC MODELS, AND THE ACHIEVEMENT OF INDIVIDUALIZED THERAPEUTIC GOALS.

In contrast, BF and AC use the data of the measured serum concentrations obtained by TDM first to make an individualized PK or PK/PD model of that drug, and then either to forecast or to control that model's behavior. This permits an intelligent reconstruction over time of the behavior of the drug in that individual patient. The clinical sensitivity of the patient to the drug can then be assessed by comparing his clinical behavior with the behavior of his drug model. In this way, specific therapeutic goals can then be selected for each patient, and a truly individualized regimen can be developed, usually using computer software, to achieve that goal as precisely as possible. This is where AC differs from BF, as BF forecasts the results of various drug regimens that one might intuitively consider, but does not explicitly set out to achieve a carefully chosen individualized therapeutic goal as precisely as possible, as does AC. The magnitude of this difference can be seen in the work of Bayard, Milman, and Schumitzky in this symposium.

Clinically, one can begin AC by evaluating how much the patient actually **needs** the drug in question. If the need is small, then only a low or moderate serum level, one associated with a low or moderate risk of toxicity, is justified. If the need is greater, especially if previous regimens have not

3

resulted in the desired response, then a higher goal and a greater risk of toxicity may well be justified. In this way one can select a low, moderate, or aggressive therapeutic goal, carefully individualized to the needs of each individual patient. There can be no truly individualized drug therapy without first setting a thoughtfully individualized therapeutic goal for each patient.

Once the goal has been selected, there is no clinical zone of indifference concerning its achievement. One wants to achieve the goal with the greatest possible precision, as the patient should not run any greater risk of toxicity than is justified by the need, and should have, on the other hand, the greatest possible therapeutic effect within the constraint of the justifiable risk of toxicity.

Various issues are therefore involved in AC. They consist of the process of making structural PK or PK/PD models, choosing the best one to use to describe the behavior of the drug in question, using that model to find the population values of the various parameters in the model (population modeling), statistical issues such as the use of assay error patterns, process and measurement noise terms, parameter uncertainties and probability distributions, the evaluation of renal and cardiac function, methods (usually Bayesian methods) of revision of a population model to make it become an individual one for each patient, the selection of individualized therapeutic goals as described above, and how best to implement the ideal regimen in a given patient in various clinical and therapeutic environments, especially with various types of dosage preparations, solutions, and different types of available infusion apparatus.

MAKING PK/PD MODELS

In making models of the process to be controlled, one must begin with a knowledge of the relationships (or possible ones) between the inputs to the system (the doses), and the outputs (the serum concentrations, and other measures of response or effect). These form the basis for the structural model to be used, and the types of pathways that best describe the passage of drug (or other responding substance, such as the white blood cells, red cells, or platelets, for example). Further, one must specify the specific responses that will be measured (the outputs of the model).
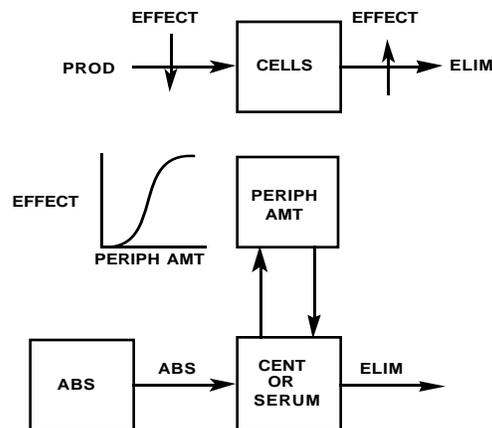

Figure 2. *Model of a drug with absorptive, central (serum), and peripheral (nonserum) compartments, with a Hill effect model related to the amount of drug in the peripheral compartment. That effect can then act either to reduce the production of a cell line or to accelerate its destruction.*

Relationships between model parameters and other descriptors or covariates must also be considered at this point, in deciding on the particular structural model (or various candidate models) to be employed. An example of such a model is shown in Figures 2 and 3.

Figure 2 shows a complex structural model which now can be easily made with the BOXES program [1] which saves much of the tedious labor of writing the differential and output equations by doing it automatically.

Figure 3 then shows the time course of that effect acting to impair white cell production, with the delayed onset, the sharp fall, and the delayed recovery after therapy is stopped.

Furthermore, if one can model such relationships for a single drug, one can also do this for whole classes of drugs that share common effects. This will lead to quantitative descriptions of drug interactions and effects, and improved ways to treat patients effectively, with coordinated regimens of multiple drugs, in the presence of these interactions.
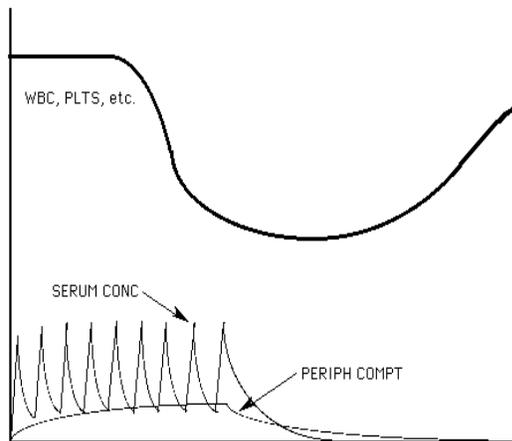
4

Figure 3. *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 that drug upon the production or destruction of a cell line. The relationship between serum levels and the effect on the cell line can be revised by adaptive control **during** the first course of therapy, not having to wait until after it is over to establish a correlation.*

POPULATION PK/PD MODELING

NAIVE POOLING

Probably the most primitive (and yet a very useful) population model is that which is made when all observations on all subjects in the population are pooled to make, as it were, a single subject. This approach is commonly used when there is only one observation per subject (destructive animal experiments with tissue observations, for example). In this case one gets an overall estimate of the population parameter values, but cannot estimate their variation within the population.

TWO-STAGE METHODS

The **Standard Two-Stage** (STS) method first employs a fitting procedure such as weighted nonlinear least squares to get the parameter estimates for each individual subject. Then the mean and standard deviation are found for each parameter. While this method is inferior to others to be discussed below, it is still a very good one, and provides the added benefit of giving the investigator a careful look at the data for each subject. The plots of the measured responses and the behavior of the fitted model during that drug regimen provide a good opportunity to check that subject's data for errors in the timing of doses or serum concentrations. Detection of such errors is a key to accurate population modeling, as they will not be found out by any method that looks only at population data without considering individual subjects.

One example of an **Iterative Two-Stage** (ITS) method is to begin with the STS method above. Then, using the population parameter values and their standard deviations, one uses the Maximum Aposteriori Probability (MAP) Bayesian method (see below) to find the Bayesian posterior parameter values for each individual subject. From these new parameter values a second population model is made from their means and standard deviations. Again, the individual Bayesian posterior parameter values are found for each subject. Another population model is made from these. This process is repeated until the population parameter means and standard deviations become stable by some criterion, when the process ends.

The **Global Two-Stage** (GTS) Method uses an extended least squares algorithm which also considers the covariance matrix of the parameters [2], while the STS and ITS methods do not.

THE HIERARCHICAL BAYESIAN METHOD

The **Hierarchical Bayesian** (HB) method, as implemented by Racine-Poon and Smith [3] is another method of parametric population modeling, yielding means and standard deviations for the population parameter values. The means and co-variances, however, are now regarded as random vari

ables, each with their own probability distributions.

CRITIQUE

In population modeling, any estimation procedure must satisfy two important criteria: 1) **compatibility with the data** - being able to handle routine clinical data, which may have as few data points as one per subject; and 2) **consistency** - as the number of subjects gets large, the population estimates should increasingly approach the true population values.

With regard to the above methods, the STS, ITS, GTS, and HB methods as described above all require enough data per subject not only to estimate the model parameter values, but also to estimate the errors in these estimates. None of these methods can handle the clinical condition of one data point per subject.

TRUE POPULATION MODELING METHODS

The methods below, in contrast, can handle clinical data of one point per subject, and therefore are generally more useful.

A PARAMETRIC METHOD: NONLINEAR MIXED EFFECT MODELING

This method, using first order linearization, developed by Beal and Sheiner, [4-6] was the first true population modeling program. It can analyze data of a single observation per subject. It gives means and standard deviations for the population parameter values, and their relationships to covariates or descriptors. However, the method at times has given results at variance with those of other methods when careful simulation studies were done [7]. The consistency of the first order linearization method has been examined by Spieler [8]. Improved versions of this algorithm are under development.

NONPARAMETRIC METHODS

The two methods described below are both consistent and can handle clinical data of only a single point per subject. Both methods compute the nonparametric maximum likelihood estimate of the unknown population density. They use different algorithms. The NPML method, using the algorithm of Mallet, has no theoretical limit to its applicability, and thus potentially can handle large models. However, it may be slower than the NPEM algorithm for small models. The NPEM algorithm runs comfortably on PC's.

6

A NONPARAMETRIC MAXIMUM LIKELIHOOD METHOD (NPML)

This algorithm was first described by Mallet [9]. It is very versatile, and can theoretically handle large complex population PK/PD models. It is consistent. It was the first method to compute the entire discrete joint probability density, without making any parametric assumptions such as means, standard deviations, and the like, for pharmacokinetic models. It has been well described previously [2].

A NONPARAMETRIC EM METHOD (NPEM)

This method, developed by Schumitzky for pharmacokinetic population modeling [10], uses an iterative EM algorithm having steps of both expectation and maximization. With each iteration, a joint probability density is computed which is more likely than than of the last one. The joint density at first is a continuous function, but as iterations go by, they become more narrow, and become discrete in the limit, approaching the Mallet solution. As currently implemented for a 1 - compartment (2 parameter) model with intravenous infusions of drug, it runs well on a PC and uses clinical data files available from routine care of patients, using commonly available software [11,12]. A program for larger population models is in development. An example of a discrete joint probability density function is shown in Figure 4 for patients receiving gentamicin [13].
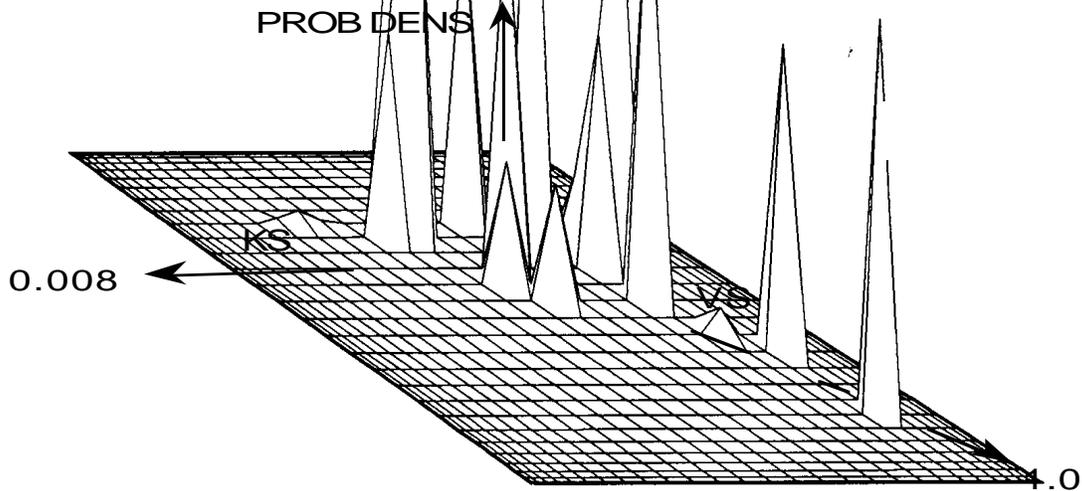

Figure 4. *3D plot of the population joint PDF found in adult patients on gentamicin. KS = slope of Kel with creatinine clearance. The mean and SD were $0.00292 \pm 0.001217\ hr^{-1}/(ml/min)$. The nonrenal Kintercept was $0.0069315\ hr^{-1}$. VS = slope of V to body weight. Mean and SD for VS was $0.36795 \pm 0.14004\ L/kg$.*

WHY ARE PARAMETER DISTRIBUTIONS DISCRETE?

The discreteness of these nonparametric algorithms was surprising at first. However, one might consider the case of a 2 - parameter model, to keep the logic simple, and might ask what form the joint parameter distribution would have if one actually knew exactly what the entire population parameter distribution was. In this case, each subject would have an exact known value for each parameter, such as a volume of distribution (V) and an elimination rate constant (K). If one then looked at the entire population distribution, he would simply see a scattergram of V and K, with one point, of unit height, for each patient. If 2 patients had identical values, then the height of that point would be 2 units, just as in a frequency distribution. Indeed, if the parameter values for every subject in a population were all exactly known , then the joint population density would in fact be a frequency distribution.

Since we do not have any way to know the exact parameter values as described above, we do studies on subjects in which we give them the drug and measure their serum concentrations, or other re-sponses, and estimate the joint parameter densities. The discreteness of the densities reminds us that all experience, after all, is finite, and that a joint population density, if known completely, is a frequency distribution of individual parameter values.

BAYES' THEOREM

The Reverend Thomas Bayes was extremely interested in how we learn from experience. He described the quantitative relationships between the probabilities of certain events before new information becomes available and the revised probabilities of these events after such new information is obtained and evaluated [14,15]. He wanted to be able to see through the operations of chance, to better understand God's design for the world. His work, now known as Bayes' Theorem, has become a standard method for revising the probabilities of events as new information becomes available.

Bayes' Theorem was introduced to the pharmacokinetic community by Sheiner and Beal [16]. In this form, the mean population parameter values and their standard deviations are considered, along with any observed serum concentrations, and **their** standard deviations (see below), that the subject might have. This variant of the common weighted least squares fit then yields the mode (the Maximum Aposteriori Probability {MAP} value) of the Bayesian posterior probability density function for each parameter. The standard deviations of the population parameter values and those of each measured serum concentration, provide the measures of credibility (the Fisher information [17]) which determine how far the fit goes toward the serum data or stays toward the population data. Many serum data points, especially when precisely measured, will dominate the model when present, while few serum data points, especially if measured with large standard deviations, will cause the fit to be dominated by the population parameter values, especially if these are obtained from a relatively uniform population which has small standard deviations about its parameter values.

The Maximum Aposteriori Probability (MAP) Bayesian posterior parameter values are found by minimizing the Bayesian objective function

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

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. The use of Bayesian methods to make individualized pharmacokinetic models of drug behavior in patients has led to improved prediction (and therefore AC) of future serum drug concentrations compared to traditional methods of linear regression [18,19]

PROCESS AND MEASUREMENT NOISE TERMS

Currently, as models are fitted to data, the weighting procedure used gives weights only to the noise or error with which the measurement of the serum concentrations or other responses are made. However, both the error of the assay used and of the errors in recording the times at which the serum levels are drawn in relation to the dose given play roles in determining the credibility of the response measures. Both factors properly belong in the measurement noise. It is for this reason that procedures such as extended least squares include the weights of the responses as extra parameters whose values need to be estimated.

However, other important factors are also operating, but their contribution more properly belongs as uncertainties in the dynamic process itself. They are the errors with which the various doses are prepared and the errors in recording when they were given. These factors constitute process noise. They convert deterministic models into stochastic ones, significantly increasing their complexity [20].

NONPHARMACOKINETIC CLINICAL FACTORS AND THEIR EFFECTS ON THERAPEUTIC PRECISION

The contributions of these various factors was evaluated in a Monte Carlo simulation study [21]. In this study of a typical scenario of ten days of tobramycin therapy in a representative hypothetical 70 kg patient having a stable creatinine clearance of 50 ml/min/1.73 M^2 body surface area, four clinical factors were evaluated - 1) the effects of good or poor **ward care** (the standard deviation of the discrepancy between when doses were given or infusions started and stopped versus when they were said to have happened; 2) the effects of a good or poor **pharmacy** - the standard deviation with which the doses were prepared; 3) the effects of a good or poor **phlebotomy service** - the standard deviation of the discrepancy between when serum specimens were obtained versus when they were said to have been drawn; and 4) the effects of a good or poor laboratory **assay**. In addition, the effect of a "smart" infusion pump, with a built in clock, was evaluated, as with this apparatus the times of infusions starting and stopping can be known quite

accurately. These various factors were defined as either good or bad as described previously [21].

One hundred identical simulated patients were treated in each group. The difference between the desired therapeutic goals of 7.0 and 1.5 ug/ml at the end of each infusion and at the trough respectively was noted by the computer with a precision which we can never achieve clinically, as our clinical perceptions of the true state of the patient are always corrupted by the four factors being evaluated here. The squared error of the difference was totaled for each simulated patient, for each course of therapy.

The study involved eight groups of 100 identical simulated patients. In group 1, all factors were defined as "good". In group 2, all were good except for the phlebotomy service. In group 3, the laboratory was poor, in Group 4 the pharmacy, and in Group 5 the ward care. Putting these together, in Group 6 all the various factors were poor. In Group 7, all were poor, but the smart pump, with its built in clock and good records of infusion times, was used. Finally, in Group 8, all factors were good once again, but in addition the smart pump was also used.

The results showed that when all factors were good, the mean total squared error (MTSE) in the achievement of the therapeutic goals was 39.4 ± 26.1 units. Since there were 10 days of therapy, doses were given every 12 hours, and both peak and trough goals existed for each dose, there were 40 goals in total. Dividing the MTSE by 40 and taking the square root yields an average deviation (AD) from the goal of 0.99 ug/ml. When the phlebotomy service alone was poor the MTSE and AD increased slightly. When the laboratory alone was poor, the MTSE and AD increased somewhat more, but not significantly so.

When the Pharmacy alone was poor, however, the MTSE and AD were significantly increased, to 80.8 ± 43.5 units and to 1.42 ug.ml respectively. Finally, when the ward care alone was poor, its effect was the greatest. The MTSE was significantly increased to 159.4 ± 84.8 units, and the AD to 2.00 ug/ml. When all poor factors were combined, the MTSE and AD were significantly increased to 233.8 ± 94.0 units and to 2.42 ug/ml respectively.

In contrast, when all factors were poor but the smart pump was used, the errors were significantly reduced to essentially the same values found in Group 1 when all factors were good. The MTSE was 37.8 ± 21.9 units, and the AD was 0.97 ug/ml. Further, when all factors were good and the smart pump was used, the MTSE was significantly reduced

below that of Group 1, to only 7.5 ± 3.6 units, with an AD of 0.43 ug/ml.

These results show that while much attention has been paid to the laboratory assay and the phlebotomy service, other factors which are less easily studied in the clinical setting such as the pharmacy and the ward care and recording of dosage times actually appear to play more major roles in permitting or preventing precise drug therapy [21].

The above results are also supported by the recent work of Charpiat et al., showing that clinical data of doses, dose times, and serum sampling times, when collected by trained pharmacy residents specifically assigned to this task, yielded MAP Bayesian fitted pharmacokinetic models which gave more precise prediction of future serum concentrations after fitting to an initial set of serum levels than did similar data which was recorded by nurses as part of their usual daily tasks [22].

CLINICAL ISSUES IN ADAPTIVE CONTROL

WHEN TO GET SERUM SAMPLES: D-OPTIMAL MONITORING AND EXPERIMENTAL DESIGNS

The best times to get samples for serum level monitoring are those times which contain the most information about the patient's model and its pharmacokinetic parameter values. These times appear to be those when serum concentrations are most sensitive to, and therefore changed by, changes in the pharmacokinetic parameter values. These times are those (1 time for each parameter) which maximize the determinant of the Fisher information matrix, which is a function of the partial derivatives of the predicted serum levels with respect to the model parameter values. They are known as *D-optimal* times because they maximize the above *determinant*.
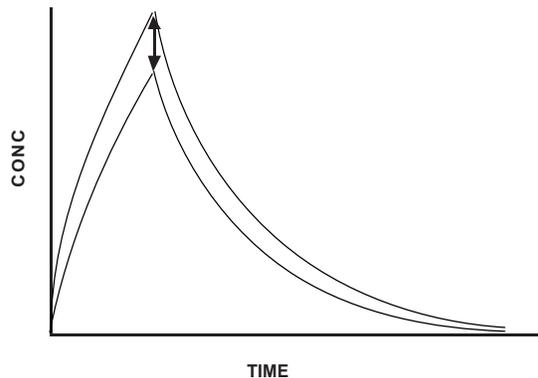
Optimal times to monitor serum drug levels. A change in the volume of distribution causes the greatest change in the serum level when the level is the highest (at the true peak).

While D-optimal concepts are totally general, for any model [23], they can be seen and understood graphically for a 1-compartment model, as shown in Figures 5 and 6. In Figure 5, when intermittent IV therapy is given, as for aminoglycosides, a decrease in the volume of distribution by a certain percent, for example, results in a corresponding percent increase in all the serum levels. Because of this, the greatest absolute increase is seen with the true peak level. This is why a sample drawn at that time contains the most information about the volume of distribution of the drug.
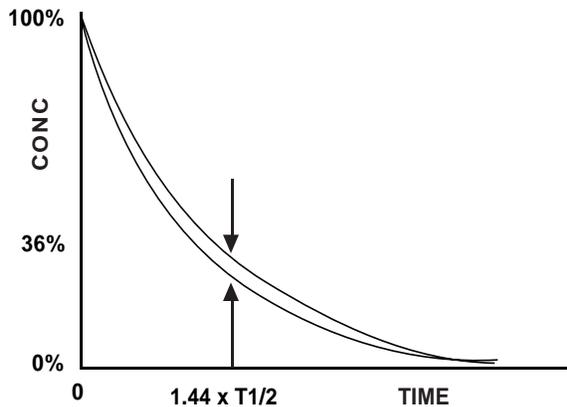

Figure 6. *Optimal times to monitor serum drug levels. A change in the elimination rate constant causes the greatest change in the serum level at 1.44 drug half-times after the end of the infusion, when the level has fallen to 36 % of the original true peak.*

Similarly, a change in the elimination rate constant causes the greatest change in the serum levels when the levels have fallen to 36 percent of their original peak. This can easily be seen by taking a hand

calculator, selecting any value for a rate constant, and calculating the value e^{-kt} for a variety of times. This column of values can then be compared with those computed from a slightly different value for the rate constant. The difference between the values in the two columns of e^{-kt} is greatest when the value of e^{-kt} is 0.36. This is shown graphically in Figure 6.

These pairs of serum levels are the ones that appear to tell us the most about what is going on with our patient with respect to his drug model. There are thus two reasons to monitor serum levels. The first is the traditional reason, simply to see what they are. The second is the real reason, to help us control the patient's model to get the levels we want the patient to have. D-optimal design helps us know the patient's model better, to most closely achieve the therapeutic goals which have been selected for him.

Suppose instead that we wanted to draw 4, or 10, not just 2 serum levels. When should we get the others? An interesting corollary of the D-optimal design is that once we have **decided to use a certain model** to describe the drug, there is no point in getting data at any other times than the optimal ones. This is the principle of replication, and it tells us that if we want more data, simply keep getting it at the optimal times, running replicates of the serum levels until we have obtained the desired number of data points. These D-optimal monitoring concepts were evaluated in a study of patients receiving gentamicin therapy [18]. Figure 7 is re-drawn from that study. Predictions of future serum concentrations made with D-optimally chosen pairs of levels were less biased and more precise than conventional peaks and troughs, and more cost-effective.
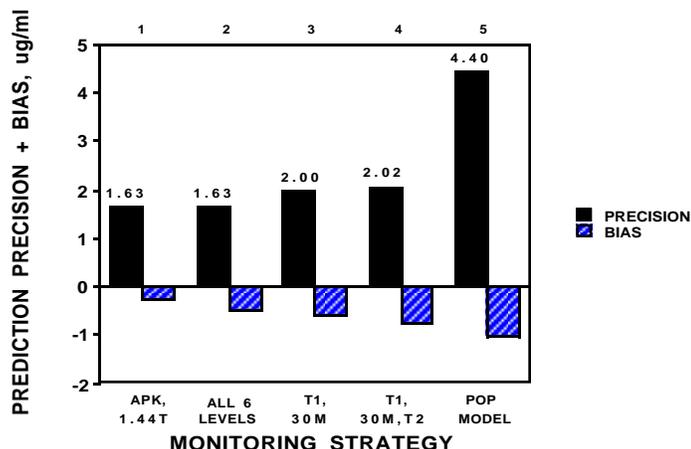
Figure 7. *Optimal times to monitor serum drug levels. The D-optimal pair, extreme left, has less bias and greater precision in predicting future levels than any other combination of data points except the full cluster of 6 points. APK, 1.44T = the D-optimal pair, drawn at APK, the absolute peak, and at 1.44 half-times after the end of the IV infusion. T1 = trough before the infusion. 30M = delayed peak sample drawn 30 min after the end of the infusion. T2 = next trough. Pop model = population Kslope model. All levels = T1, APK, 30M, 1.44T, and T2 together, plus another at 23 hrs into the regimen [19].*

DETERMINING LABORATORY ASSAY ERRORS

In the MAP Bayesian objective function shown earlier, 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 also true for the data of the measured serum drug concentrations. However, the usual practice of most clinical laboratories has been simply 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 assay 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 an important and practical matter in the fitting of serum drug concentration data for patients.

THE PRACTICAL DETERMINATION OF ASSAY ERROR PATTERNS

What is needed is a practical means to determine the standard deviation of each serum drug concentration as it is routinely measured by the clinical laboratory. A convenient and practical way to achieve the goal of computing the probable SD with which a single determination of a serum drug concentration is measured is to do replicate measurements of some representative samples of such concentrations (at least in quadruplicate) and to determine the mean and SD of each sample. This can be done, for example, on a blank sample, a low sample, an intermediate one, a high one, and a very high one, so that the entire assay range, subtherapeutic, therapeutic, and toxic, is covered. One can then fit this data of the relationship between measured concentration and SD 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 [24] the error pattern of the EMIT gentamicin assay in use at the Los Angeles County - USC Medical Center was determined, and its polynomial equation was found to be

$$SD \text{ (ug/ml)} = 0.56708 - 0.10563C + 0.016801 C^2$$

As shown in Figure 8, this assay has an SD of 0.567 ug/ml at 0.0 ug/ml (the blank), yielding a variance of 0.321 and a weight (1/variance) of 3.11. The SD falls to 0.403 ug/ml at 3.0 ug/ml, yielding a variance of 0.162 and a weight of 6.16, double that of the blank. The SD then rises to 0.797ug/ml at 8.0 ug/ml and to 1.71 ug/ml at a concentration of 12.0 ug/ml, when the variance is 2.96 and the weight has fallen to 0.338. Note that the weights range from a high of 6.16 to a low of 0.338, a factor of 18 in the credibility given to the serum concentration data points within this range.
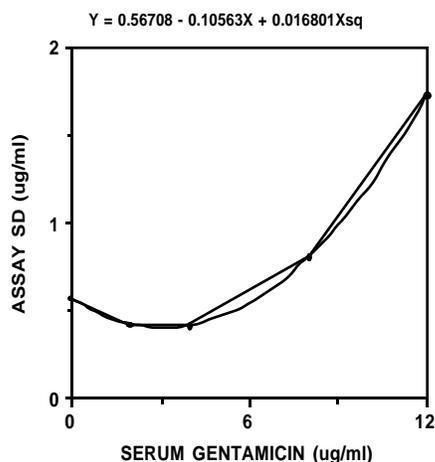
Figure 8. *Error pattern of an EMIT assay for Gentamicin and its associated polynomial equation.*

The coefficients of the polynomial equation are then stored with the Gentamicin population model in the USC*PACK clinical program for adaptive control of Gentamicin dosage regimens [11,12] so that correct weighting of each measured serum concentration can be implemented during the Bayesian fitting procedure [25].

Similarly, the error pattern of the Abbott TDx assay for lidocaine used by our Laboratory of Applied Pharmacokinetics was found to be

$$SD (ug/ml) = 0.040415 + 0.008604C + 0.004786C^2$$

This assay is more precise than that of the above EMIT assay of gentamicin. Its SD is 0.04 ug/ml at the blank, with a variance of 0.0016 and a weight of 612.2. The SD then rises and the weight progressively falls to 0.077 ug/ml and 169.7 respectively at a concentration of 2.0 ug/ml, to 0.203 and 24.3 at 5.0 ug/ml, and to 0.416 and 5.79 respectively at a concentration of 8.0 ug/ml. With this assay the weights thus vary by a factor of 612.2/5.79 = 111.5 over the range from 0.0 to 8.0 ug/ml.

It is easy and inexpensive for a clinical laboratory to determine its own assay SD over its entire working range in the above manner, thus arriving at a practical mixture of within - run and between - run errors which are useful in the fitting of pharmacokinetic models [25].

ANALYSIS 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 reports 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 (BT and RJ) examined the 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, Theophylline, and Vancomycin. We took the means and SD's of concentrations found for the various specimens and fitted them with a polynomial, usually of second order, occasionally of first order, to provide a library of error patterns for the above assays. These can be used to estimate the SD of individual serum samples for Bayesian fitting of pharmacokinetic models until laboratories can determine their own error patterns.

RESULTS

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. Results for the Abbott TDx, The Dupont ACA, and the Syva Emit assays were evaluated. 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 " = $0.46475 + 0.028131C + 0.0021305C^2$, $R^2 = 0.939$
Syva Emit SD " = $0.23237 + 0.047015C + 0.0016876C^2$, $R^2 = 0.965$
All Methods SD " = $0.32272 + 0.018365C + 0.0012051C^2$, $R^2 = 0.983$

The Abbott TDx assay was the most precise. The Dupont ACA and Syva Emit assays were less so. As was also the case with all subsequent results, 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 Abbott TDx, Dupont ACA, and the Syva Emit assays were evaluated. The Baxter Stratus assay was not evaluated as the number and range of data points

was much less. 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   "   = 0.25719 - 0.016215C + 0.0081998C^2, R^2 = 0.982
Syva Emit SD   "   = 0.14078 - 0.002263C + 0.0184060C^2, R^2 = 0.991
All Methods SD   "   = 0.09114 - 0.043524C + 0.0045964C^2, R^2 = 0.992

It was noteworthy here that two samples were found by all laboratories to have concentrations significantly less than those labeled by the College. Except for this, all other samples had general good agreement between the target values stated by the College and the means found by the various laboratories. Again, the Abbott TDx assay was the most precise.

Theophylline

Seventeen sample means ranging from 3.0 to 30.0 ug/ml, and their SD's, obtained from 3682 to 4696 reporting laboratories, provided data. The Abbott TDx, Baxter Stratus, Dupont ACA, HPLC, and Syva Emit assays were evaluated. The following polynomial equations were obtained.

Abbott TDx SD (ug/ml) = 0.22605 + 0.023955C + 0.00056926C^2, R^2 = 0.978
Baxter Stratus SD "   = 0.0.0789 + 0.083394C, R^2 = 0.985
Dupont ACA SD   "   = 0.29967 + 0.010201C + 0.1379800C^2, R^2 = 0.963
HPLC assay SD   "   = 1.04060 - 0.120450C + 0.0093092C^2, R^2 = 0.707
Syva Emit SD   "   = 0.21770 + 0.057018C + 0.07131800C^2, R^2 = 0.972
All Methods SD   "   = 0.25463 + 0.039573C + 0.00088179C^2, R^2 = 0.976

The Abbott TDx assay was the most precise, while the HPLC assay was the least. The Dupont ACA assay was next most precise, and the Baxter ACA and Syva Emit assays were intermediate and of about equal precision. The HPLC polynomial had the lowest value of R^2.

Vancomycin

Fifteen sample means ranging from 4.9 to 40.0 ug/ml, and their SD's, obtained from 645 to 862 reporting laboratories, provided data. The Abbott

TDx and Syva Emit assays were evaluated. The following polynomial equations were obtained.

Abbott TDx SD (ug/ml) = 0.57694 + 0.012816C + 0.00058286C^2, R^2 = 0.971
Syva Emit SD   "   = 0.93214 + 0.023689C + 0.00177690C^2, R^2 = 0.971
All Methods SD   "   = 0.59421 + 0.012291C + 0.00071299C^2, R^2 = 0.979

Both error patterns were well characterized by their equations, with R^2 values over 0.97. The Abbott TDx assay was the more precise.

DISCUSSION

Sources of Error

The errors reported by the College survey are a mixture of within - run and between - run laboratory errors, as well as within - laboratory and between - laboratory errors. They are useful for purposes of therapeutic drug monitoring and Bayesian pharmacokinetic modeling until such time as a clinical laboratory determines its own assay error patterns for its own drugs and develops its own set of polynomial equations. Fortunately, this is easily done.

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.

The Importance of Reporting Low Concentrations Below "Detectable Limits"

While the above practice is useful in toxicological analysis in making a firm decision as to whether a substance is present in the body or not, it is a distinct obstacle to optimal therapeutic drug monitoring. In therapeutic drug monitoring there is no question that the drug has been given. One clearly knows this from the history, the orders, and the nurses' notes. Indeed, many clinical laboratories will not measure a serum drug concentration unless the time since the last dose is 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 its concen-

tration. Low trough aminoglycoside concentrations for example, below those clearly detectable, are not only useful but necessary for therapeutic drug monitoring and for proper Bayesian pharmacokinetic modeling and AC. To withhold such results renders that measurement useless for AC, and should not be charged to the patient's bill when done for therapeutic drug monitoring rather than for toxicology. A vital data point is absent.

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 the secure detectable limits of 0.5 ug/ml", for example. This procedure will answer both the needs of the toxicologists and the pharmacokineticists, and is therefore more generally useful.

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.

On the other hand, it is equally important to know know the errors of concentrations found well into the toxic range. Because of this, when 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.

The Importance of Improving Assay Precision at the High End

When doing Bayesian modeling, 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.

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 very heteroschedastic. 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 approach the high concentrations as closely as one might perhaps wish.

This is also true for the polynomial equations described above. The difference here is that the polynomial equations are derived from carefully measured SD's over the 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. 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 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 resulting change in the error pattern toward homoschedasticity.

This analysis of the results of the survey by the College of American Pathologists now provides a library of polynomial equations which characterize the error patterns of several assays over their working range. These equations can be used by clinical pharmacokinetic laboratories to improve their precision of Bayesian fitting of pharmacokinetic models until that laboratory can determine its own error patterns as described above. The procedure is easy and inexpensive to do, and can be repeated as indicated from time to time.

EVALUATION OF RENAL FUNCTION

A MODEL OF CREATININE KINETICS

As previously described [26] 128 measurements of serum creatinine were reviewed in 15 patients (9 men and 6 women) who had undergone renal

14

transplantation at the Los Angeles County/USC Medical Center. Using stepwise multiple regression, it was found that the reciprocal of the serum creatinine concentration was by far the most strongly associated with the carefully measured creatinine clearance (CCr), and that sex, the daily change in serum creatinine concentration, and age were also important variables. Body weight was less important. Using this information, a dynamic model of creatinine production and excretion was made. It used the general view that the daily change in the total amount of creatinine in a patient's body is the difference between the daily creatinine production (P) and the daily excretion (E). This may be described by

$$V(C2 - C1) = P - E$$

where V is the apparent volume of distribution of serum creatinine (in hundreds of ml), C1 and C2 are the first and second serum creatinine values taken 1 day apart (in mg/dL), and P and E are in mg. Since V is less than total body water, it was approximated as 40% of the patient's body weight (in hundreds of grams).

CALCULATION OF CREATININE PRODUCTION

Adjustment for Age.

The data of Siersbaek-Nielsen et al. [27] of the effect of age upon the carefully measured urinary creatinine excretion was shown to be described by

$$E = 29.305 - 0.203A$$

where E is the measured urinary creatinine excretion (in mg/kg/day) and A is the age (in years). That data was obtained in hospitalized patients who were clinically free of any renal disease. Since the patients were all quite stable,

$$E = P.$$

In this way, one can use this carefully measured data of excretion to make a reasonable estimate of daily creatinine production. This can be further refined as described below. It should also be noted that in these patients, the average serum creatinine their patients was 1.1 mg/dL. This will be useful below.

Adjustment for Degree of Uremia.

It has also been shown by Goldman [28] that uremic patients also have a decreased excretion (and therefore production) of creatinine. Using data from

that report, creatinine production (CP, in mg/kg/day) is related to serum creatinine (C, in mg/dL) by

$$CP = 1344.4 - 43.76C$$

One can thus adjust the first estimate of creatinine production for age to the average value (Cavg) of C1 and C2 by the ratio R, where

$$P1 = 1344.4 - 43.76 \times Cavg,$$

where Cavg = the patient's average serum creatinine,

$$P2 = 1344.4 - 43.76 \times 1.1,$$

where 1.1 = the average serum creatinine of Siersbaek-Nielsen's patients, as described above. Then,

$$R = P1 / P2, \text{ and}$$

the adjusted $P = P \times R$

In this way creatinine production can be estimated for men, based on a great deal of careful measurements of 24 hour urinary creatinine excretion, and adjusted to the patient's age, weight, and degree of uremia. In further adjustments, 90% of this value was then taken if the patient was female, and 85% of that for either men or women if the patient was a dialysis patient. Further, if a patient's body mass is clearly below normal, as may be the case with cirrhotic patients, one can simply make a rough clinical estimate of the patient's body (muscle) mass as a percent of normal, if desired, to make a further final adjustment of P. This last adjustment for body mass was not done in the study described here [26].

CALCULATION OF CREATININE EXCRETION

In the usual calculation of creatinine clearance,

$$\mathbf{C = UV/P,}$$

where $\mathbf{U}$ is urinary creatinine concentration, $\mathbf{V}$ is the urine volume, and $\mathbf{P}$ is the plasma or serum creatinine concentration, and $\mathbf{C}$ is creatinine clearance. This can be rearranged to show that what comes out of the body is equal to what was cleared from the body. Thus

$$\mathbf{CP = UV.}$$

Because they are numerically equal, $\mathbf{PC}$ can therefore be substituted for $\mathbf{UV}$, the measured 24 hour excretion. Thus

E = **PC**, and

E = **PC** = Cavg x CCr/100 x 1440,

where E is expressed in mg/day, Cavg is in mg/dL, CCr is in ml/min, and 1440 represents the minutes in 1 day.

The final relationship to calculate creatinine clearance from unstable serum creatinine values, and without requiring a urine specimen may now be written as

.4W(C2 - C1) = P - Cavg x CCr/100 x 1440

After this, the raw creatinine clearance (CCr) above can be corrected to that of an average patient having a body surface area of 1.73 square meters. The above equation thus represents a dynamic model of creatinine kinetics, and permits estimation of CCr from routine clinical data of age, sex, height, weight, and either a pair of changing serum creatinine levels or a single stable serum creatinine, all without having to collect a urine specimen, which is an extremely unreliable procedure in all but research situations.

COMPARISON OF THE ESTIMATES WITH MEASURED CREATININE CLEARANCES

In a first set of 128 observations on 15 patients [26], the algorithm was shown to have an accuracy essentially equal to that of Jadrny [29]. In an additional set of 250 observations on a group of 14 patients who had just undergone renal transplantation, the standard error of the estimate was again slightly better that the equations of Jadrny, with an overall scatter of about ± 25% between the estimated and the measured values.

It is also interesting to consider the errors present in the classical measurement of CCr. If one can measure a serum creatinine level with a coefficient of variation of 5%, as is the case with common autoanalyzer methods, and if one measures urinary creatinine concentrations with a coefficient of variation of 8%, as is also common, then if one can collect a 24 hour urine specimen with a coefficient of variation of 5%, the resulting value of measured creatinine clearance will have a coefficient of variation of 11%, and a 95% confidence limit of ± 22%. This closely corresponds to the scatter found between the estimated and the measured values. Because of this, it is likely that this method of estimating CCr is of about equal precision with the classical measurement of it. It addition, it is practical in clinical situations. It is also probably better at sensing changes in renal function in response to sudden changes in serum creatinine than are the more simple formulas of Jelliffe [30], or of Cockcroft and Gault [31], both of which are designed only for use when serum creatinine is stable.

EVALUATION OF CARDIAC FUNCTION

Brandfonbrener, Landowne, and Shock [32] used dye dilution methods to measure the cardiac output and index of hospitalized patients who were clinically free of cardiovascular disease, afebrile, and ambulatory. They found an essentially linear fall in cardiac index with age. A regression relationship was developed relating these by Rodman el al. [33] and used to adjust the elimination rate constant of lidocaine to a patient's normal cardiac index. It then appeared that patients with failure and rales, for example, had behavior of lidocaine consistent with a cardiac index of about 75% of normal for age, and that those in frank shock still had cardiac indices about 65% of normal for their age.

Based on these preliminary relationships, Rodman then developed a more explicit protocol to estimate cardiac function as a percent of normal, based on clinical factors which are easy to evaluate at the bedside [34]. When cardiac index was estimated for age from the relationship

CI = 4.18575 - 0.02328 x Age in years, when age is at least 20,

and then further adjusted by the algorithm of Rodman, reasonable predictions of serum lidocaine concentrations became possible from a pharmacokinetic model based on this relationship, and pharmacokinetically designed lidocaine infusion regimens were designed which significantly improved the ability to achieve and maintain effective serum concentrations early in therapy, with a significant reduction in breakthrough arrhythmias and a suggestive reduction in the incidence of ventricular fibrillation [33].

Thus, while thermodilution and similar methods are obviously the best for measuring cardiac index, use of the age - adjusted relationship based on the data of Brandfonbrener and colleagues, coupled with the clinical evaluation of the patient with failure having an index of about 75% of normal, a patient in shock having one of about 65% of normal, provides a practical way to make reasonable clinical estimates of cardiac index at the bedside, without having to do thermodilution determinations. The basic relationship between age and cardiac index is shown below.

AGE	CARDIAC INDEX (L/min/M^2)
20	3.72
30	3.49
40	3.25
50	3.01
60	2.78
70	2.54
80	2.31
90	2.07
100	1.84

MANAGING INTRAPATIENT VARIABILITY: SETTING INITIAL CONDITIONS

Pharmacokinetic behavior is described by differential equations. These require initial conditions to be set before they embark on their trajectories over time as doses are given and serum levels are drawn. Usually the initial conditions are set to zero. That is, before the first dose is given in a data set, it is assumed that there was no drug present in any pharmacokinetic compartment.

It is useful, however, to be able to set nonzero initial conditions. A patient may change formulations of a drug at some time. It is useful to end the data set of the first drug, and to take the end - of - interval (usually the trough) values of serum concentrations, unabsorbed amounts in the gut, and body concentrations of drug in the peripheral, nonserum, compartment and to enter them as the initial conditions present at the start of the next data set. This is a useful feature in clinical software for managing AC of drug dosage regimens.

Haug and Slugg [35] have coined the term "VD collapse" to describe the rather sudden reduction in the apparent volume of distribution of gentamicin that they have often seen as a patient gets well. They have used this "VD collapse" as an indicator of impending recovery of the patient. One of us (RJ) has seen similar behavior in a patient who was managed by the Clinical Pharmacology Division at Christchurch Hospital, Christchurch, New Zealand, under the direction of Evan Begg, M.D, who generously allowed us to analyze the data. That patient began as a general medical patient with a urinary tract infection. He received tobramycin and appeared to make a good recovery. During that first data set, his Vd was 0.185 L/kg with the 1 compartment model and 0.165 L/kg (central volume (Vc) only) with the 2 compartment model. However, the patient relapsed after therapy was ended, and developed clearcut septic shock. During that second data set his 1 compartment Vd was much higher, 0.618 L/kg, and his 2 compartment Vc was 0.578 L/kg. His nonzero ending values present in that data set were passed on as initial conditions to the third data set, which also represented therapy given while he was quite ill. During that time his 1 compartment Vd was 0.559 L/kg and his 2 compartment Vc was 0.557 L/kg. Again the last values in the various compartments were passed on to the fourth data set, which was obtained during his obvious clinical recovery. During that time, his 1 compartment VD decreased to 0.157 L/kg, and his 2 compartment VC to 0.147 L/kg.

This experience with the patient of Dr. Begg, prompted by the earlier experience of Haug and Slugg, documents in a single patient the significant changes in distribution volume that may both come and go, with "Vd expansion and collapse", at least with aminoglycoside therapy, as a patient's clinical status may change. This can be detected as the failure to obtain a good fit of all the data over time, and one can usually find a reasonable place to break the data into stepwise sets to analyze as described above.

The use of software that permits easy setting of initial conditions passed over from a previous data set also permits the use of different population models from one time to another, including population models of different formulations of the same drug (regular theophylline to sustained release form, for example).

In addition, if one has a patient with an unreliable dosage history, as is often the case with prior oral theophylline therapy, for example, in a patient with an acute exacerbation of asthma, one must begin by treating his acute problem, including a regimen of theophylline if indicated. One can start by estimating what percent of the dosage regimen the patient is most likely to have taken, and develop an initial regimen based on that. However, one can also get a baseline serum level just prior to starting therapy, and at least two others, to check the progress of therapy and to make an individualized model for the patient. One then can fit, not only the measured serum levels, but one can also enter, as initial conditions, reasonable estimates of the unabsorbed drug in the gut at time zero (some fraction of the supposed last dose, for example). The baseline serum level at time zero can also be entered as the initial condition for the serum compartment. Since these estimates are uncertain, especially in the gut, one can put very wide standard deviations upon them in the Bayesian fitting procedure. This method provides the clinician with extra flexibility in making individualized pharmacokinetic models for patient care.

17

AC OF GENTAMICIN AND OTHER AMINOGLYCOSIDE THERAPY

The use of MAP Bayesian strategies for AC of Gentamicin therapy has been shown to be superior to that of linear regression [18]. Using this approach, Haug and colleagues [36] employed aggressive strategies to achieve peak goals of 10-14 ug/ml, and appear to have improved quality of care. Similarly, AC of Amikacin therapy has shown the ability to obtain quite precise achievement of predicted serum levels, and thus to achieve prolonged therapy of patients with osteomyelitis, for example, without evidence of toxicity [37,38]. Furthermore, these approaches and software have been useful in reconstructing the entire therapeutic history of patients, even when their clinical status and central apparent volume of distribution have been changing significantly as their clinical status changed, taking the last values of serum and peripheral body concentrations of drug and passing them on as initial conditions for the next data set, as described earlier.

AC OF VANCOMYCIN THERAPY

The MAP Bayesian data analysis yielded superior predictions (and therefore probably control) of

Vancomycin serum levels, while linear regression was simply not able to cope with a drug having such marked 2 compartment behavior [19].

AC OF TRIMETHOPRIM THERAPY

Pneumocystis carinii pneumonia (PCP) occurs in many patients with the acquired immunodeficiency syndrome (AIDS), and has been a significant cause of death. In an attempt to optimize dosage regimens of trimethoprim (TMP) in these patients, we analyzed a previous study [39] in which patients with AIDS who had PCP had their dosage regimens intuitively altered during therapy to maintain serum TMP levels within a targeted range of 5 to 8 ug/ml to minimize serious hematologic toxicity. In that previous study, 36 patients were initially treated with TMP - SMX at 15 to 20 mg/kg/day of the TMP component. The steady state dosage was then adjusted based on serum levels found to maintain TMP levels in the desired range.

With dosage adjustment, the final steady state dosage of TMP was 12 ± 2.7 mg/kg/day, well below the standard dosage above. Nevertheless, only 12% of patients died, despite an expected mortality of over 40% based on their initial blood gas values. Moreover, all of the 33 surviving patients were able to complete a 3 week course of therapy without developing absolute neutrophil counts below $500/mm^3$ or other treatment - terminating toxicity.

We retrospectively analyzed 179 serum TMP levels determined in 35 patients from the above study, and used the NPEM software to make a population pharmacokinetic model of intravenous TMP in these patients. In the Kslope and Vslope parameterization of this model, the nonrenal component of elimination was set at 0.01155 hr^{-1}, equivalent to a half time of 60 hours in an anephric patient. The mean, median, and standard deviation for Kslope were 0.00136, 0.001, and 0.00136 hr-1 respectively, and for Vslope were 1.17, 0.8016, and 0.642 L/kg respectively. This population model, using the median values and the standard deviations described above, has been added to the USC*PACK PC collection of software for modeling and AC of drug dosage regimens [11,12]. The current population therapeutic goals appear to be peak concentrations of 10 and troughs of 6 ug/ml of intravenous TMP at this time. It is interesting that the current debate whether 15 or 20 mg/kg/day is better turns out, with this population model, to be the difference between patients having a creatinine clearance of 100 ml/minute (requiring about 20 mg/kd/day to achieve the above therapeutic goals), and those having the slightly lower one of 75 ml/minute (requiring about 15 mg/kg/day).

ANTIBIOTIC CONCENTRATIONS IN ENDOCARDIAL VEGETATIONS

A general model describing the diffusion of substances into spherical objects such as endocardial vegetations has been developed by Maire et al. [40]. It uses as inputs the computed time course of the serum concentrations provided either by a population pharmacokinetic model or by the patient's MAP Bayesian fitted model. The diffusion model then computes the time course of antibiotic concentrations at various points (just under the surface, halfway in, and at the center) within the sphere. Initial computations have been made using the diffusion coefficient of o.7 x 10^{-6} for Amikacin obtained by Crowell and colleagues [41]. These preliminary computations suggest that Amikacin penetrates fairly well into the center of a 1.0 cm simulated spherical vegetation when serum concentrations are reasonable, but that much poorer penetration is achieved when the sphere is over 1.5 cm diameter. This plot is now an option in the USC*PACK clinical collection for use during the analysis of past drug therapy.

It is also interesting to speculate beyond this. For example, one can compute possible concentrations within a simulated abscess of stated size. It is likely, however, that a diffusion coefficient for Amikacin into a vegetation may be much greater than that into an abscess, with its fibrotic wall, and may well represent a "best case scenario" for diffusion into abscesses.

These calculations nevertheless suggest that while it may be possible to achieve reasonable trough concentrations in the canter of a 1 cm. spherical object, when the diameter of the sphere becomes 5 cm, as with a simulated abscess, almost no drug gets into the center. It is a graphic demonstration of why abscesses of any significant size need to be surgically opened and drained This new program, called SPHERE [40], should stimulate investigators to determine diffusion coefficients with greater accuracy for several drugs into vegetations of various sizes.

AC OF LIDOCAINE THERAPY

Pharmacokinetically designed dosage regimens of lidocaine have been shown to increase the serum levels early in therapy from the lower part to about the middle of the therapeutic range, to significantly reduce breakthrough arrhythmias, and to reduce the incidence of ventricular fibrillation in CCU patients [33].

AC OF DIGITALIS THERAPY

An illustrative clinical example of AC as opposed to BF or TDM is shown below, in which a patient on oral digoxin 0.25 mg/day suddenly developed rapid atrial fibrillation. He was given serial IV digoxin doses, in the classical clinical titration, and converted to sinus rhythm. He was placed on his previous maintenance dose again, but slipped back into atrial fibrillation. A serum level at this time was 1.0 ng/ml. He was titrated with IV doses a second time, and for a second time he converted to sinus rhythm. A serum level at this time was also 1.0 ng/ml, showing that serum levels by themselves did not correlate with the patient's different clinical behavior at this time. When he was again placed on his old maintenance dose, he again went back into atrial fibrillation. For a third time, after a third titration with IV doses, he again converted to sinus rhythm. A serum level at this time was 1.2 ng/ml. At this point the MB program in the USC*PACK collection was run, and the 2 compartment digoxin population model was fitted to the data of his doses and his three digoxin serum concentrations. A plot of that fitted model is shown in Figure 9.

Note that while the serum levels did not correlate with the patient's clinical behavior, the concentrations in the peripheral (nonserum) compartment did, and provided the key guide to the selection of the therapeutic goal in this patient. This type of model, first described and documented by Reuning and colleagues, has a good correlation between the behavior of the nonserum compartment and ejection time indices, which reflect the inotropic effect of the drug [42]. Here, the patient's fitted model similarly shows good correlation between the behavior of the drug in the nonserum compartment and the drug's clinical effect in him. Inspection of the patient's fitted model shows that the patient consistently was in sinus rhythm when peripheral (nonserum) body concentrations of digoxin were 10 to 12 ug/kg, and that he was in atrial fibrillation when they were less. Because of this, a therapeutic peak goal of 11.5 ug/kg in the nonserum compartment was selected for this patient, and a dosage regimen was developed by the software to achieve it.

A hint as to what the correct regimen might be is also given by his fitted model. On his previous dose of 0.25 mg/day, peak and trough serum levels were stable at about 1.0 and 0.5 ng/ml. More importantly, though, the peak nonserum body concentrations were only about 5.0 ug/kg. Since the usual peak goal in the peripheral compartment for most patients in sinus rhythm is about 7.0 ug/kg [43], the patient may well have been somewhat underdigitalized prior to the onset of his atrial fibrillation. There was also some physiologic change, as it then became

necessary repeatedly to achieve peripheral compartment concentrations over 10.0 ug/kg in order to achieve conversion to sinus rhythm.

The selected goal of 11.5 ug/kg represents 11.5/5 of the original dose of 0.25 mg/day, suggesting that the proper maintenance dose was about 0.25 x 11.5/5, or 0.575 mg/day. The program's regimen was 0.327 mg for the first day of the new regimen, 0.599 and 0.574 mg for days 2 and 3, and 0.577 mg/day thereafter. This was given clinically as 0.375 mg on the first day, and then 0.5 and 0.625 mg on alternate days thereafter. The patient left the hospital in sinus rhythm, and was still in sinus rhythm when seen in the clinic two weeks later.
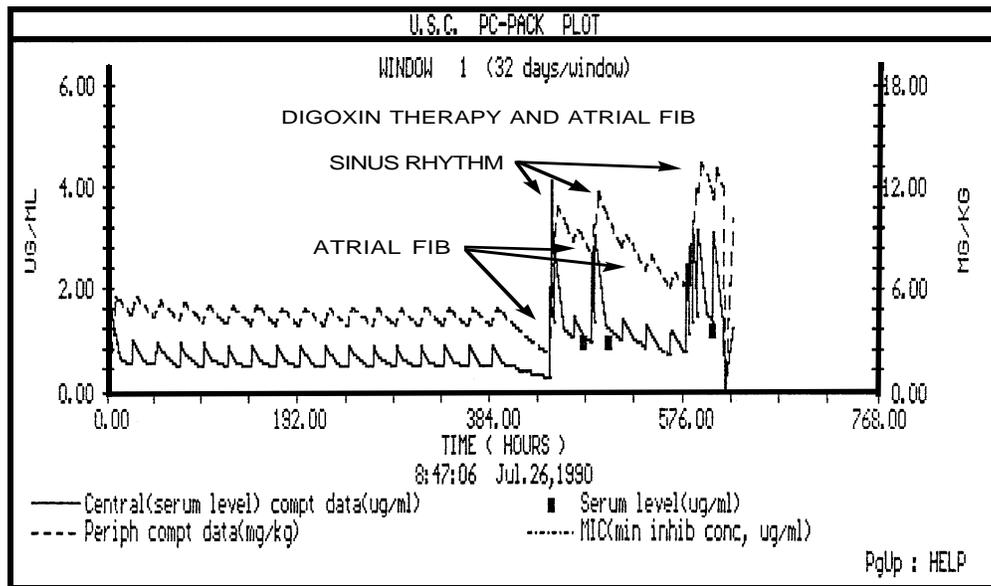

Figure 9. *Screen plot of patient with atrial fibrillation who was successfully converted to sinus rhythm with IV digoxin three separate times, but who relapsed into atrial fibrillation twice when put back on his previous maintenance dose. Sinus rhythm was consistently present when peripheral body glycoside concentrations were 10-12 ug/kg (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.*