

OPTIMIZING INDIVIDUALIZED DOSAGE REGIMENS OF POTENTIALLY TOXIC DRUGS.

Roger W. Jelliffe, Alan Schumitzky, Robert Leary, Andreas Botnen, Ashutosh Gandhi, Pascal Maire, Xavier Barbaut, Nathalie Bleyzac, and Irina Bondareva.

1. INTRODUCTION

The end product of drug development is the use of the drug in clinical therapy. When a drug has a narrow margin of therapeutic safety, we must steer its dosage between one that is too low, and likely to be ineffective on the one hand, or too high, and likely to be toxic, on the other. We must carefully plan and individualize the dosage for each patient, to achieve some desired target goal such as a serum concentration, or its profile over time. We must then observe the patient, and if needed, monitor serum concentrations at appropriate intervals. These intervals should be frequent enough so we can evaluate the patient when there are relatively small changes in the total amount of drug in the body between observations, so that if toxicity develops, we detect it in an early stage of its development so we can make the appropriate adjustment in dosage early, rather than later, after toxicity has become more severe and dangerous.

It is not useful to talk about dosage individualization without saying with respect to what. It is commonly said that one should individualize dosage to body weight and renal function, for example. But again, to what specific end, toward what specific goal? This is usually not explicitly stated. We usually do this, however, to control either the total amount of drug in the patient's body, or the serum concentration, for example, at a desired specific target value, usually within some general target "therapeutic range" of serum concentrations where most patients (but not all) do well, and where the incidence of toxicity is acceptably low.

However, this approach is appropriate only for the initial regimen, and it still ignores the opportunity to be gentle, moderate, or aggressive in the approach to the patient, according to each individual patient's need for the drug.

The expected incidence of toxicity should be no greater than that which is appropriate for the patient's need for the drug. In many cases, if the need for the drug is not great, or is not acute, the target goal should first be one that is associated with a low

Incidence (risk) of toxicity, leading to a gentle dosage regimen. Based on the patient's response, the target goal can then be revised upward and a higher dosage given to achieve it. This, for example, is behind the "start low, go up slow" dosage policy so well advocated by Cohen [1].

On the other hand, a firmer approach may be clinically indicated. If the patient has an acute and significant need for the drug, the "start low, go slow" approach is not warranted, and is not safe. Here a higher target goal must be selected, one which is more likely to be effective, and a greater risk of toxicity will have to be accepted in order to achieve such a higher goal.

The target goal must therefore be selected individually for each patient, according to that patient's individual need for the drug at that time. In this way, one can then develop a gentle approach to one patient, but a more firm or aggressive approach to another, as each patient's need dictates. Clearly, things are not the same at the bottom of the therapeutic range as at the top, or even further, if it is necessary to go to a still higher target if the patient's need dictates. Examples of this approach are the acceptance of a certain risk of toxicity with cancer chemotherapy, therapy for AIDS, and with the risk of toxicity with digoxin, aminoglycosides, vancomycin, and transplant chemotherapy, for example.

2. SET INDIVIDUALIZED TARGET GOALS FOR EACH PATIENT

The concept of a general "therapeutic range" of serum drug concentrations is therefore only a generalization. It is an overall range in which most patients, but certainly not all, do well. One must always check each individual patient to see if he or she is doing not only well, but optimally, on clinical grounds, regardless of whatever the serum concentration is actually found to be. This approach is quite different from much clinical teaching, but still is similar to what many clinicians give lip service to – "look at the patient, not just the serum concentration".

2.1. Problems with "Therapeutic Ranges"

Figure 1 shows the usual means by which therapeutic ranges appear to have been obtained. It is interesting that these ranges have never, to the authors' knowledge, been defined in a specific, quantitative, and explicitly described manner, but instead have simply been described as regions below which therapy is generally "ineffective", and above which "significant" toxicity has been observed. It is usually stated that first, there is a "significant" incidence of therapeutic effects with increasing serum drug concentrations. This defines the beginning of the therapeutic range. Later on the incidence of toxic effects also becomes "significant", and the "toxic range" has been entered.. For example, Evans [2] presents a definition of therapeutic range as "a range of drug concentrations within which the probability of the desired clinical response is relatively high and the probability of unacceptable toxicity is relatively low" However, the therapeutic range has never, to our knowledge, been defined quantitatively. The eye is drawn to the bends in each line in Figure 1, and the classification of the apparent "therapeutic range" has been developed, published, and accepted without apparent criticism or further thought. However, this procedure does not consider the need to

develop a gentle dosage regimen for a patient who needs only a gentle touch, or a more aggressive one for a patient who really needs the dosage “pushed”. Another problem has been that special populations of patients have not had their special needs recognized. For example, it is well known that patients with atrial fibrillation need higher serum concentrations, usually averaging 2.0 ng/ml, for full control of their ventricular rate, and yet this has not been followed up by setting a special therapeutic range for them. Special types of patients, the elderly, for example, may need special target serum concentrations selected for them.




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 classified in relation to these bends. This qualitative procedure of classification discards the important quantitative relationship of the incidence of toxic effects versus serum concentration.

In addition, the risks and benefits of using a therapeutic range or window are quite complex. One could, for example, develop a dosage regimen to maximize the probability of having the patient’s serum concentration be within some desirable therapeutic window. This sounds good at first, but the decision-making process rapidly becomes quite complex. One must also weigh the benefit and the probability of a desirable response against the risks, and their probabilities, of being outside that window, either below or above it.

Those risks associated with being below the window are usually associated with lack of therapeutic effect. Those associated with being above the window are usually associated with toxicity. Each of these outcomes - subtherapeutic, therapeutic, or toxic - thus has not only its associated probability but also its own positive or negative quantitative utility function of goodness or badness. Optimizing such a complex set of probabilities and utilities becomes a most complex process, and is poorly amenable to rigorous clinical decision analysis, especially at the bedside.

2.2 Setting specific target goals based on need.

A more intuitive and individualized approach is one in which the clinician evaluates the magnitude of each patient's individual clinical need for the drug in question, and selects an estimated risk of toxicity which is felt on clinical grounds to be justified by the patient's need. In this approach, there is no window of neutrality about a target, as in a "therapeutic range. Based on the relationship between serum concentration and incidence of toxicity shown in Figure 1, for example, one selects a specific target serum concentration goal to be achieved for a specific patient. One does not want the patient to run any greater risk of toxicity than is justified by the patient's clinical need for the drug. Within that constraint, however, one wants to give the patient as much drug as possible, to get the maximum benefit. This approach provides the rationale for selecting a specific target serum concentration goal, rather than some wider window, and then to attempt to achieve that target goal with the greatest possible precision, just as if one were shooting at any other target.

In this approach, the risks of being just slightly above the desired target goal are only minimally different from those associated with being just slightly below it, in the sense of an infinitesimal difference in calculus. Because of this, it appears easier, and also more intuitive for a clinician, to choose a desired target goal rather than some wider window, and then to attempt to achieve or hit that selected target goal with the greatest precision (least error) possible.

Without selecting such a specific target goal, there can be no truly individualized precise drug therapy. Individualized drug therapy therefore begins by setting a specific individualized target goal for each patient. The task of the clinician is to select, and then to hit, the desired target goal as precisely as possible. As soon as the initial regimen is given, the clinician's task is then to observe the patient's clinical response at appropriately frequent intervals, and to reevaluate whether the target goal was hit precisely enough or not, whether the target goal was correctly chosen or not, or if it should be changed and a new dosage regimen developed to hit the new target goal. This is the basis of the "target - oriented, model - based, approach to individualized drug dosage for each individual patient [3].

3. THE NEED FOR MODELS

Pharmacokinetic models, like other mathematical models, condense huge collections of experience into a form that can be easily grasped and understood. Newton's equations condensed wide ranges of experience and extensive data into his superb models of celestial mechanics. He did not make hypotheses. He simply discovered and then described with models these widely scattered relationships and events. These models are very scientific, in the deepest sense of the word.

In exactly the same way, pharmacokinetic models, like all models, can be extremely useful. No pharmacokinetic model is an exact description of reality, which is always more complex. Nevertheless, these models describe in useful, quantitative terms, the behavior of drugs when they are given to patients – their absorption, serum

concentrations, pathways and exchange rates to and from other compartments, and their various therapeutic and toxic effects. Often these processes and pathways may have important relationships to other clinical descriptors or covariates such as body weight, age, renal function, smoking status, gender, and genetic CYP450 makeup, for example.

Pharmacokinetic and pharmacodynamic (PK/PD) models also provide the tool to apply that recorded past experience to the care of new patients. Past experience with drug behavior is now usually stored in the form of a population PK/PD model which is then used to design the initial dosage regimen for the next patient who appears to belong to that particular population.

The dosage regimen to achieve the therapeutic target goal is computed and given. The patient is then monitored both clinically and by measuring serum concentrations. The serum concentrations are used not only to note if they are within some general “therapeutic range”, but most importantly to make a specific model of the behavior of the drug in that individual patient, based on information from the population PK/PD model, and using Bayes’ theorem to develop an individualized model that best describes the behavior of the drug in that individual patient, using both the patient’s individual data of serum concentrations, and balancing the patient’s individual data against the more general information from the population model.

In striking this balance, the relative credibility of both kinds of information must be weighed. Since individual data is often sparse, especially at the beginning of an experience with a patient, it never explains the entire picture, and it is usefully supplemented with the general population information. The more individual data is obtained, the more it dominates the picture. In this manner, using Bayes’ theorem, an individual model can be made, which has the property of being able to predict future serum concentrations at least somewhat better than models made without considering population information, using weighted nonlinear least squares, for example, which consider only the individual patient data.

One can then see what the patient’s probable serum concentrations were at all other times when they were not measured, even when he or she was not at all in a steady state, even during highly unstable clinical situations in very acutely ill patients with great changes in their clinical status and in their renal function over time. One can also reconstruct and see graphically the computed concentrations of drugs in a peripheral nonserum compartment or in various effect compartments. The patient’s individualized model permits one also to make dosage adjustments without having to wait for a steady state before sampling serum concentrations, to take into consideration practical clinical situations, such as handling data of different dosage regimens (with unequal doses and unequal dosing intervals), and totally arbitrary time intervals between drug doses and blood sampling.

These important physiological, pharmacological, pharmacokinetic and pharmacodynamic relationships cannot be seen or inferred at all without such models. In addition, the use of models gets around the need to wait for a steady state before monitoring serum concentrations. By comparing the clinical behavior of the patient with the behavior of the patient’s model, one can evaluate the patient’s clinical sensitivity to

the drug, and can adjust the target goal appropriately. For digoxin, for example, the inotropic effect of the drug correlates best with the computed concentrations of the drug in the peripheral compartment (ug/kg of body weight, for example) rather than with the serum concentrations. The excellent model made by Reuning and colleagues for digoxin [4] has been highly useful clinically [5].

4. CURRENT BAYESIAN INDIVIDUALIZATION OF DRUG DOSAGE REGIMENS

The Reverend Thomas Bayes, who died in 1761, was a mathematically inclined minister, and it is said that he was interested in “seeing through” the operations of chance, to better understand God’s design for the world. He described how we learn from life, by revising our expectations based upon our experience. His theorem describes in quantitative terms the important sequential relationship between:

1. the estimated probabilities of certain events (for example, in PK/PD terms, a patient’s apparent volume of distribution, and the clearance or rate constant for elimination of a drug), that are present in a patient before we have had any chance to know anything about the patient’s response to the drug (or the serum concentrations), because we have not given him/her the drug yet, and then

2. the measured serum concentrations that are found in that patient, and their precision, and then

3. the revised (Bayesian conditional posterior) probabilities of these PK/PD parameter values after the new information is obtained and evaluated (posterior to the new information).

In the beginning, before the drug is given, all we have is our past experience about the behavior of the drug in similar patients. This is why one of our most important tasks as clinicians is to store our experiences with patients in a form that can be used in the future to apply that experience optimally, usually using Bayes’ theorem above, to the care of the next similar patient. This is why it is important to make population PK/PD models of the behavior of the drug in the actual patients we treat, not just in research clinical trials, to obtain and store that important clinical past experience optimally.

All this decision making must be done before (prior to) giving the drug and before (prior to) being able to observe anything about the patient’s own individual clinical response and serum concentrations. Furthermore, since fitting a pharmacokinetic model only to the patient’s data is not optimal without supplementing it with information of general past experience as well, the Bayesian approach describes quantitatively the sequential relationship between prior probabilities (those present before obtaining any new information), the new information (serum concentrations, for example, and their precision), and the revised (posterior) probabilities after that information is taken into account.

Bayes' theorem can be used to describe this sequence of pharmacokinetic and pharmacodynamic events. It is widely used throughout the scientific and the military communities. The great majority of flight control, fire control, and missile guidance systems are Bayesian adaptive controllers. We use the same approach here to control the behavior of the patient as we treat him or her. We do this by selecting desired target therapeutic goals, by using the population model to compute the initial regimen to best hit the desired target, by monitoring the patient with serum concentrations, using that data to make an individual Bayesian posterior model, then by re-evaluating on clinical grounds whether or not the correct target goal was chosen (and changing it if needed), and finally by computing the new and adjusted dosage regimen to best hit the selected target once again. This cycle can then repeat as often as clinically indicated. This therapeutic process is called Bayesian adaptive individualization and control of drug dosage regimens.

The Maximum A posteriori Probability (MAP) Bayesian approach to individualization of drug dosage regimens was introduced to the pharmacokinetic community by Sheiner et al. [6]. In this approach, parametric population models are used as the Bayesian priors. In these models, the parameters in the structural model (the apparent volume of distribution, clearances, rate constants, and their variances, are described by other parameters of their means, standard deviations (SD's), and the correlations between them. The credibility of these population models (their parameter SD's) is then evaluated in relationship to the SD's of the measured serum concentrations as they are obtained. The contribution of these two types of data and their SD's to the MAP Bayesian posterior individualized patient model is shown in the MAP Bayesian objective function below,

$$\sum \frac{(C_{\text{obs}} - C_{\text{mod}})^2}{SD^2_{(C_{\text{obs}})}} + \sum \frac{(P_{\text{pop}} - P_{\text{mod}})^2}{SD^2_{(P_{\text{pop}})}} \quad (1)$$

where C_{obs} is the collection of observed serum concentrations, $SD_{(C_{\text{obs}})}$ is the collection of their respective SD's, and C_{mod} is the model estimate of each serum concentration at the time it was obtained. Similarly, P_{pop} is the collection of the various population model parameter values, $SD^2_{(P_{\text{pop}})}$ is the collection of their respective SD's, and P_{mod} is the collection of the Bayesian posterior model parameter values. Each data point can be given a weight according to its Fisher information, the reciprocal of its variance (the square of the SD), and so is each population parameter value, according to its variance. Population models in which there is greater diversity, and therefore greater variance, contribute less to the individualized model than do population models having smaller variances. Similarly, a precise assay will draw the fitting procedure more closely to the observed concentrations, and a less precise assay will do the opposite. The more serum data are obtained, the more that information dominates the determination of the MAP Bayesian posterior parameter values (P_{mod}) in the patient's individualized pharmacokinetic model.

In contrast with older methods of fitting of PK/PD data, the MAP Bayesian method can fit using only a single serum concentration data point if needed. This is because the MAP Bayesian procedure already has one data point for each parameter. Those data points are the collection of the population parameter values themselves. Because of this, the MAP Bayesian procedure can start to fit with only a single serum concentration. This feature of the MAP Bayesian method allows one to handle the often very poor and sparse data usually present in clinical strategies of therapeutic drug monitoring and dosage adjustment.

Having made the patient's individualized model, one then uses it to reconstruct the past behavior of the drug in the patient during his therapy to date. One can examine a graphical plot of the behavior of this model over the duration of the past therapy. One can thus evaluate the clinical sensitivity of the patient to the drug, by looking at the patient clinically and comparing the patient's clinical behavior with that of the patient's individualized pharmacokinetic model. In that way, one can evaluate whether the initial target goal was well chosen or not. One can choose a different goal if needed, and once again one can compute the dosage regimen to achieve it. In this way, the model can be individualized and dosage can continue to be adjusted to the patient's body weight, renal function, and available serum concentrations, for example, to achieve the desired target goal, usually with increasing precision during the course of the patient's therapy.

5. COMPARISON WITH OTHER METHODS OF FITTING DATA

The MAP Bayesian fitting procedure has been shown to be generally better in predicting future serum concentrations than the method of weighted nonlinear least squares. MAP Bayesian fitting is also significantly better than the earlier traditional but now obsolete method of linear regression on the logarithms of the concentrations (see below).

5.1 Weighted Nonlinear least Squares Regression

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

$$\sum \frac{(C_{\text{obs}} - C_{\text{mod}})^2}{SD^2(C_{\text{obs}})} \quad (2)$$

Because of this, only the patient's serum data are considered in the fitting procedure, and this information is not supplemented by the additional information from the population parameter values. Because of this, fitted models made using weighted nonlinear least squares have been shown to predict future serum concentrations slightly less well than those made using MAP Bayesian fitting [7].

Like the MAP Bayesian procedure, this method can fit the model to data of doses and serum concentrations acquired over many dose intervals, usually the patient's entire dosage history. There is no longer any reason to do the traditional "single dose" pharmacokinetic study. Further, there is no need for the patient to be in a steady state or for the serum data to be only post-distributional. Studies and population pharmacokinetic / pharmacodynamic modeling can be done on the actual patients being treated, as they are receiving their therapy. This is a second, and very important, function of therapeutic drug monitoring. The algorithm of Nelder and Mead [8] is a good one for fitting the data in both the least squares and the MAP Bayesian fitting procedures. A very useful nonmathematical description of this method has been given in BYTE magazine [9].

Secondly, like the MAP Bayesian method, weighted nonlinear least squares can provide correct weighting of serum concentration data according to its credibility or Fisher information [10]. It thus has the potential for obtaining good estimates of the pharmacokinetic parameter values.

However, this method cannot take into account population information that is generally known about how that drug usually behaves in patients like the individual patient under consideration. As the procedure moves from the starting population parameter values to others which fit the data better, it discards all the general information used to begin the fitting procedure, instead of supplementing it with the individual patient's data. Since no fitting procedure ever explains the entire relationship between doses given and concentrations found, discarding the general population information is a suboptimal feature. It may well be because of this feature that the nonlinear least-squares method, while "fitting" serum concentration data "best", has been shown to be a slightly poorer predictor of subsequent serum concentrations than the MAP Bayesian method [7]. In contrast to the MAP Bayesian procedure, this method, like linear least squares regression (see below), requires at least one serum concentration for each parameter to be fitted, or at least two serum concentrations in the models considered here, as will be discussed further below. The MAP Bayesian method, in contrast, can begin to fit using only a single serum concentration data point if needed. This is because the MAP Bayesian procedure already has one data point for each parameter. Those data points are the collection of the population parameter values themselves. Because of this, the MAP Bayesian procedure can start to fit with only a single additional data point, the very first serum concentration. This is an important feature that is extremely helpful in the "practical" aspects of therapeutic drug monitoring for patient care.

5.2 Linear Least Squares Regression

Another method used to fit serum concentrations to make individual patient models has been the old traditional but now obsolete method of linear regression on the logarithms of the serum concentrations (see below). This method was the traditional one in which a pharmacokinetic model (restricted to only a single compartment) was fitted to data obtained during only a single dose interval, specifically to the logarithms of the serum concentrations. No weighting of the serum data was used. The method was simple, and it has been widely implemented on hand calculators. It was generally the community standard for monitoring serum gentamicin concentrations ever since Sawchuk and Zaske showed its utility to individualize aminoglycoside dosage regimens [11].

The method requires at least 2 serum concentrations. It cannot handle anything more than a 1-compartment pharmacokinetic model. Distribution of the drug after a dose must be complete before a “meaningful” serum sample can be obtained, where the ratio between the serum concentration and that in any other compartment is constant. The method takes advantage of the fact that one can linearize the solution of a first-order linear differential equation for such a model if one transforms the serum concentration values to their logarithms. However, the method has three important weaknesses.

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




Figure 2. Error pattern assumed using fitting by linear regression on logarithms of serum concentrations. Note the much greater weighting given to the lower concentrations.

Second, linear regression contains the assumption that the assay error is a constant percent of the measured concentrations. The lower the concentration, the more accurately it is assumed to be known. Because of this, if the assay has any other error pattern over its working range (and it almost always does), this method greatly

overestimates the credibility of low serum concentrations over high ones. This can be seen if one considers two serum samples, one of 8.0 ug/ml for example, and one of 1.5 ug/ml, as shown in Figure 2. One usually wishes to attach approximately equal credibility (weight) to these data points. One might thus assume that their laboratory error is approximately equal. Since the Fisher information (an index of credibility) of a data point having a normally distributed error is proportional to the reciprocal of the variance of that data point, the relative weights given by linear regression to serum concentrations of 8.0 and 1.5 ug/ml would be proportional to the reciprocal of their squares [7,10]. Because of this, the method of linear least squares, which assumes that the error bars are equal on the logarithmic scale, arbitrarily gives the value of 1.5 ug/ml a weight of $8^2/1.5^2 = 64/2.25 = 28.4$ times the weight of the concentration of 8.0 ug/ml. A concentration of 0.1 has 100 times the weight of a concentration of 1.0, and 10,000 times the weight of a concentration of 10.0 units! Because of this assumption, this method often obtains model parameter values that are significantly different from those obtained by other methods [7]. Third, this method ignores all population data, and therefore all past general experience, concerning the behavior of the drug.

5.3 Conclusions

The MAP Bayesian method [6] appears to be the best of these three [7]. As with nonlinear least squares, it can provide correct weighting of serum concentration data according to the known laboratory assay error, and it can analyze such data over many dose intervals. In addition, it supplements population data (general knowledge) with specific information about each patient, instead of discarding it. Because of this, the method has been a slightly better predictor of future serum concentrations [7]. Lastly, the method requires only a single serum concentration to begin the analysis, no matter how many parameters are present in the population pharmacokinetic model. As more serum concentrations are obtained, the fitted model gradually becomes less of a population model and more of a patient-specific model. Both general and patient-specific data are combined intelligently in the M.A.P. Bayesian procedure to provide the most probable single-point estimates of the parameter values given both types of data and their respective standard deviations.

Finally, one other fitting procedure, now coming on the scene, holds promise of doing still better than the MAP Bayesian method. This is the "Multiple Model" method of fitting data and designing drug dosage regimens [12]. It is a stochastic rather than a deterministic method, and is based on nonparametric population models [13,14] and their individualized Bayesian posterior pharmacokinetic models. This method of dosage design will be discussed more fully later on in this chapter.

6. EXAMPLES OF MAP BAYESIAN TARGET-ORIENTED, MODEL – BASED, APPROACHES TO PATIENT CARE

6.1 Gentamicin Therapy

With a 1-compartment pharmacokinetic model in which the elimination rate constant (K_{el}) was composed of a nonrenal component (K_{nr}) and a renal component

having a slope (K_{slope}) relationship to creatinine clearance (CCr) so that $K_{el} = K_{nr} + K_{slope} \times CCr$, the MAP Bayesian procedure resulted in significantly better prediction of future serum concentrations (see Figure 3) than predictions made using linear regression (see Figure 4). 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 [7].

Because the software used in that study [7,15] was specifically designed to operate in the presence of significant changes in renal function from dose to dose, it has also been useful in the analysis and management of aminoglycoside therapy for patients who must undergo periodic hemodialysis.

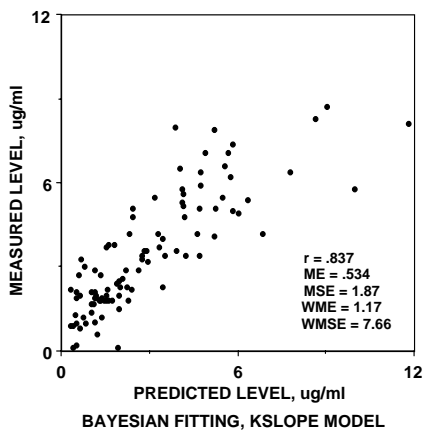


Figure 3. - Predicted versus measured serum Gentamicin concentrations found with M.A.P. Bayesian fitting and the K_{slope} model. r = correlation coefficient, ME = mean error, MSE = mean squared error. WME = mean weighted error. WMSE = weighted mean squared error. See text for discussion.




Figure 4. - Predicted versus measured serum concentrations found with linear regression on the logarithms of the serum concentrations. Other symbols as in Figure 3.

Managing aminoglycoside therapy can also be done quite well in patients who must be on periodic hemodialysis. The key is to get serum samples of both the drug and of serum creatinine concentrations both before and after dialysis. For current dialysis equipment, and for the aminoglycoside antibiotics, the apparent increase in creatinine clearance during dialysis is about 50 ml/min above the patient's baseline value.

The baseline value of CCr can be estimated by examining the most recent rising pair of serum creatinine values, after one dialysis for example, and just before the next one. This is why the pair of serum creatinine samples pre- and post- dialysis is useful. In addition, when a patient goes on dialysis, one can record this as giving a dose of the drug, where the amount of the dose is 0.0 mg. With this dose of zero, the infusion time can be stated as very short, 0.1 hr, for example, and the creatinine clearance can be directly entered as being 50 ml/min above that of the patient's baseline. Finally, when the patient goes off the dialysis, another dose of zero is entered at that time, and the patient's creatinine clearance is set back to the baseline value [18].

6.2 Timing the Aminoglycoside Dose and the Dialysis

Another corollary for dialysis patients is that while most of them are given their dose of drug soon after the dialysis, this can cause a problem. The serum aminoglycoside concentrations in such patients have extremely long half-times, and these patients are often the ones who have the greatest incidence of renal toxicity and ototoxicity, because their serum concentrations stay so high for so long after each dose, even though the doses themselves are adjusted to keep the total area under the serum concentration curve at an appropriate value constrained by the desirable target peak and trough goals.

Instead of this, it may be more prudent and useful to give the dose before dialysis, about 2 or 3 hours before dialysis. In this case, one gets the desired peak value. Then the dialysis helps to mimic the renal function of a patient with more normal (or less abnormal) renal function, reducing the serum concentration more rapidly, and helps to achieve a serum concentration profile somewhat more like that of a patient with better renal function.

7. CLINICAL STUDIES OF OUTCOME AND COST

7.1. Gentamicin Therapy

Probably the best examination to date of the utility of the MAP Bayesian approach to individualize drug dosage regimens for patients has been the work of van Lent-Evers et al [3]. They compared the model-based, target goal approach to aminoglycoside therapy with a more conventional therapeutic drug monitoring strategy. The mean peak and trough concentrations in the study group were 10.6 ± 2.9 ug/ml and 0.7 ± 0.6 ug/ml respectively versus 7.6 ± 2.2 and 1.4 ± 1.3 ug/ml respectively, both significant differences. The peaks were significantly higher and the troughs significantly lower in the study group. Overall mortality was 9 of 105 (9%) in the study group versus 18 of 127 (14%) in the control group, not a significant difference ($P = 0.26$). However, in

those patients who had obvious infections present on admission, mortality was only 1 of 48 in the study group versus 9 of 62 in the control group, a significant difference ($p = 0.023$). In addition, nephrotoxicity was only 2.9% in the study group versus 13.4% in the control group.

While the clinical outcome was significantly improved (more effective, less toxic) with the use of this model-based, target-oriented approach to monitoring and dosage individualization, it was interesting to see that the hospital stay was also significantly reduced, from 26.3 ± 2.9 days overall in the control group to 20.0 ± 1.4 days in the study group ($p = 0.045$). For patients with infections present on admission, the stay was similarly reduced, from 18.0 ± 1.4 days in the control group to 12.6 ± 0.8 days in the study group. Thus in both patient groups, those with and also without clearcut infections on admission, hospital stay was reduced by about 6 days with the use of this approach to serum concentration monitoring and model-based dosage individualization.

Further, despite the added effort and cost to implement this therapeutic approach, the overall cost per patient was reduced from $16,882 \pm 17,721$ Dutch florins in the control group to $13,125 \pm 9,267$, a significant difference ($p < 0.05$). In the patients with infections on admission, the cost was reduced from $11,743 \pm 7,437$ Dutch florins to $8,883 \pm 3,778$ florins, an even more significant difference ($p < 0.001$). Thus in a sizeable group of patents, the model based, target oriented method of monitoring and individualizing aminoglycoside dosage regimens not only resulted in better outcomes, but also in shorter hospital stays, at a net cost savings of about \$1000 per patient [3].

7.2. Amikacin Therapy

MAP Bayesian target-oriented, model-based adaptive control has been used to manage amikacin therapy in geriatric patients, often for extended periods, by Maire et al [16]. In their patients, whose renal function was often quite reduced but who were generally clinically stable, visibly better prediction (and therefore control) of serum concentrations was seen with MAP Bayesian analysis than with their unfitted population model].




Figure 5 – Left: Predicted versus measured serum Amikacin concentrations found with M.A.P. Bayesian fitting, 1 compartment Kslope model (B1). Right: Predicted versus measured serum Amikacin concentrations found with A Priori population 1 compartment Kslope model (AP1).

The results of Maire et al. [16] in these clinically more stable patients are shown in Figure 5, left. They are better than those found in the gentamicin patients with unstable renal function [7] shown in Figure 3 above. Further, Figure 5, right, shows the much poorer predictions based simply on the population model for Amikacin, without any fitting to the serum data.

7.3. Vancomycin Therapy

Vancomycin therapy was evaluated by Hurst et al [17] using a two - compartment (central plus peripheral compartment) model. Using traditional linear regression, extremely poor prediction was found, as shown in Figure 6, left. In contrast, the 2 compartment model, coupled with MAP Bayesian fitting, led to significantly better prediction of future serum concentrations than did the linear regression method, as shown in Figure 6, right.

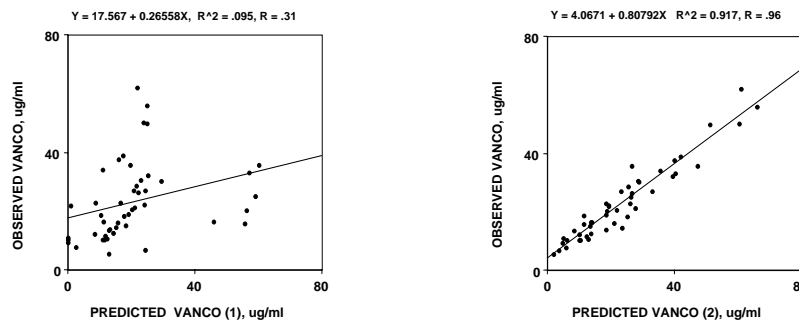


Figure 6.- Left: Predicted versus measured serum Vancomycin concentrations found with Linear regression. Right: - Predicted versus measured serum Vancomycin concentrations found with a 2 compartment Kslope model and MAP Bayesian fitting).

7.4. Digoxin Therapy

The digoxin population model used in the USC*PACK MAP Bayesian software [18] is based on that described by Reuning, Sams, and Notari [4]. That two - compartment model uses both a central (serum) and a peripheral (nonserum) compartment. Computed concentrations of drug in the peripheral compartment correlate much better with inotropic effect than do serum concentrations [4,19]. The USC*PACK digoxin software [18] not only uses this model, but also develops dosage regimens to achieve desired target goals in either the central (serum concentration) compartment or in the peripheral (tissue or effect) compartment.

The following example is illustrative. A 58 year old man developed rapid atrial fibrillation at another center, after missing his usual daily dose of 0.25 mg. He was clinically titrated with several intravenous doses of digoxin, and converted to sinus rhythm. The problem then was to select a successful dosage regimen for the patient. He was placed back on his original oral maintenance dosage. After a day, atrial fibrillation

recurred, showing that his digoxin requirements had changed. He again was titrated with several doses of intravenous digoxin and again converted to sinus rhythm. Again, the problem was to select a successful dosage regimen for the patient. Once again, he was placed on his original oral maintenance dosage, and once again, after about two days, atrial fibrillation recurred. For a third time he was titrated with several intravenous doses of digoxin, and for a third time he converted to sinus rhythm. A week of hospital time had been consumed during this phase of his care. The same question remained – now that sinus rhythm had been restored, what digoxin dosage regimen should this patient receive?




Figure 7. 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 10-13 ug/kg (right hand scale, and not mg/kg as labeled). Selection of a therapeutic goal of 11.5 ug/kg in the peripheral compartment led to a dosage regimen of 0.5 and 0.625 mg/day. On that regimen, the patient could be discharged home in sinus rhythm and was still in sinus rhythm when seen in clinic 2 weeks later.

At this point the USC*PACK MAP Bayesian digoxin software [18] was used, in a telephone consultation, to analyze the patient's situation. The raw data of three serum concentrations, all taken during the post-distributional phase after a dose, showed almost no correlation with the patient's clinical behavior. As shown in Figure 7, he was back in atrial fibrillation when the first serum concentration of 1.0 ng/ml was obtained (the first black rectangle) and had converted again to sinus rhythm when the second and third serum concentrations of 1.0 and 1.2 ng/ml were obtained (the 2nd and 3rd black rectangles) with a lapse back into atrial fibrillation in between them. An important question is: how can it be that the patient was in atrial fibrillation at one time with a serum concentration of 1.0 ng/ml, and in sinus rhythm at another time, with exactly the same serum concentration? It is experiences of this type that have made many cardiologists feel that serum monitoring of digoxin is not useful.

Two questions need to be asked here. 1) Was the patient in a steady state at the time the serum samples were drawn? The answer clearly is no – the patient was not in a steady state at all when either of the two samples was drawn. 2) Were the serum samples obtained at the same time after the dose? Again, the answer is no. Because of this, there is no justification for using the conventional interpretation of serum concentrations which is based on these premises. Without the use of a model, the raw data of the serum concentrations is of little use.

However, when the 2 - compartment digoxin population model was fitted to the data of his various doses and these serum concentrations, the resulting fitted model was very informative, and quite good correlation was seen between the computed peripheral tissue concentrations and the patient's clinical status.

As shown in Figure 7, sinus rhythm was present in this patient whenever his peripheral compartment tissue concentrations were between 10.0 to 13.0 ug/kg. Based on this correlation, a therapeutic goal of 11.5 ug/kg was chosen for the desired peripheral compartment peak body concentration. The resulting regimen to achieve the target goal was 0.25 mg for the first day, and then averaged 0.57 mg/day.

Following this consultation, he was given 0.25 mg the first day, and then placed on a maintenance regimen of 0.5 and 0.625 mg on alternating days. On this regimen he remained in sinus rhythm. 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.

When one considers that an entire week had elapsed before the above consultation, with three successful conversions to sinus rhythm but two failures to arrive at a useful regimen to maintain that conversion,, all involving the patients, the physicians', and the ward personnel's time, effort, and money, one can see that the above pharmacokinetic consultation and recommendation was extremely cost effective.

7.5. Lidocaine Therapy

A target-oriented, model-based approach was used to manage lidocaine therapy [20]. Patients with myocardial infarcts who had arrhythmias requiring Lidocaine were retrospectively evaluated. Seventy eight patients received conventional nonpharmacokinetically oriented therapy, and an equal number of patients received pharmacokinetically designed, target-oriented, model-based infusion regimens. Of the conventional therapy patients, eight developed ventricular fibrillation, one became toxic, and 33 required additional lidocaine to control their arrhythmias. In contrast, only two patients receiving the target and model based regimens developed ventricular fibrillation, a suggestive though not significant difference ($P=0.12$), one patient became toxic, and the two who developed ventricular fibrillation were the only patients who required more lidocaine to control their arrhythmias ($p<0.001$). The pharmacokinetically designed regimens not only achieved more effective serum concentrations [19] but also suggestively reduced the incidence of ventricular fibrillation, and significantly reduced the incidence of breakthrough arrhythmias [20].

7.6. Busulfan Therapy

In a comparison with an equal number of historical control subjects by Bleyzac et al. [21], children who underwent bone marrow transplantation for various hematologic problems and malignancies had their Busulfan therapy given as a standard regimen during the process of transplantation versus having it given using MAP Bayesian adaptive control. Twenty nine patients composed each group. The patients having Busulfan given by MAP Bayesian adaptive control had an incidence of veno-occlusive disease reduced from 24.1 to 3.4 percent ($p < 0.05$). In addition, graft failure was reduced from 12.0 percent to zero percent. Furthermore, survival was increased from 65.5 to 82.8 percent. Because of this, MAP Bayesian adaptive control of the Busulfan dosage regimens in this study made visible and significant improvements in the care of these patients [21].

8. WHY WE REALLY MONITOR SERUM CONCENTRATIONS: FOR MODEL-BASED, TARGET-ORIENTED INDIVIDUALIZED DRUG THERAPY

Traditional approaches to therapeutic drug monitoring were originally designed for use only in steady state situations, and usually employed only 1 - compartment models. They developed dosage regimens only for such steady state situations, and were oriented to keeping serum concentrations within a general therapeutic range rather than to achieving a specific target goal for a specific individual patient. Such approaches made it impossible to deal with patients in some of their most important clinical moments, as, for example, during changing renal function or dialysis, or when certain "golden clinical moments" must be captured, and a dosage regimen developed to achieve and maintain a desired target goal immediately, without waiting for a steady state, as in the case of the above patient receiving digoxin.

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

That patient's case also emphasizes the fact that one does not use serum concentrations 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 concentrations is rather to find out how each patient actually handles the drug, how the drug (and its model) really behaves in each individual patient, especially in non-steady-state situations, and to correlate the behavior of each patient's fitted model with his/her own clinical behavior. Only then can one optimally evaluate each patient's clinical sensitivity to, and specific need for, a drug. MAP Bayesian adaptive control, in the context of model based, target-oriented individualized drug therapy, brings a precision and capability to drug dosage which is

not possible with older obsolete approaches based on linear regression or simply on raw data of the serum concentrations alone.

8.1 Optimal monitoring strategies

The issue of what are the best times to obtain serum concentrations is also important. Often samples are obtained at the trough, just before the next dose, after distribution is complete, and the errors in recording the time at which the previous dose was given and the time at which the sample was drawn make the least difference in the value of the measured serum concentration. It is not generally realized that because of this, one has deliberately selected the time containing the least information concerning the actual behavior of the drug. There is minimal information about the processes of absorption, distribution, elimination, and their relationship to the actual time course and profile of the serum concentrations, and to the time course of the drug effects.

It is often much better to obtain serum samples when they contain the most information about the various processes described above. One can use a model, and can make small variations in the model parameter values, and note their effect upon the profile of the serum concentrations. At what time do changes in the model parameter values cause the greatest changes in the serum concentration profile? These are the times when the serum concentrations are maximally sensitive to changes in the various parameter values. These are the times when getting the serum samples lets one best "see through" the many clinical uncertainties, and best understand the behavior of the drug, by permitting the most precise parameter estimates to be made for the chosen model. These times can be calculated using the well-known D-optimal sampling strategies, based on the work of D'Argenio, for example [37].

These strategies can easily be employed in routine clinical care. For the aminoglycosides, for example, one can start by getting a peak sample, out of the opposite arm at the end of the intravenous infusion. It is then useful to wait, whatever the dose interval is, until about 21 hours into the regimen, when the patient's creatinine clearance at least 40 ml.min/1.73 M². Because of this, it is easy to center the patient's aminoglycoside doses about three hours after routine morning blood drawing time. In that way it is easy to make the routine blood sample at that time be quite close to a D-optimal sample. In general, considerations of D-optimal sampling strategies also suggest that it is useful to obtain at least one sample for each parameter to be fitted in the patient's model when doing therapeutic drug monitoring.

8.2 More general comments

Further, we need to monitor drug therapy better in general. It is distressing to see patients with multidrug resistant TB, for example, dosed without such monitoring. Since many patients with multidrug resistant TB absorb the drugs poorly, it wastes the patient's lives by treating them with an unmonitored regimen and waiting to see if their sputum smears and cultures eventually become negative. It is much more useful to know early in the course of therapy whether the serum concentrations achieved on a given regimen are likely to be effective.

We spend a great deal of money of expensive treatments for patients with cancer and AIDS, and we follow the viral load, and the measures of hematological toxicity. But we are not yet optimizing this process, and we should. We treat cancer patients with methotrexate to a desired area under the serum curve (AUC), but we usually do not ask if that AUC is really optimal for each individual patient. Monitoring serum concentrations and determining their relationship to the hematocrit, leukocyte count, and platelet count, for example, would permit therapy to be optimized within the constraints of tolerable measures of toxicity. There is a great deal to be done in this area!

9. SPECIAL CASES: ENTERING INITIAL CONDITIONS - CHANGING POPULATION MODELS DURING THE FITTING PROCEDURE.

Most pharmacokinetic analyses deal with patients (and their pharmacokinetic models), who have had stable values for their various parameters such as volume of distribution, rate constants, clearances, etc.. However, this is not always so, even though one can express a rate constant as an intercept plus a slope times a descriptor of elimination such as creatinine clearance or cardiac index [34], so that renal function or cardiac index can change from dose to dose during therapy, and the patient's drug model can keep up with these changes as they take place.

Probably the most serious problem in analyzing pharmacokinetic data in patients is caused by sudden significant changes in a patient's volume of distribution (Vd) of the central (serum concentration) compartment, without any change in any currently known clinical descriptor.. It is generally known, for example, that patients in an ICU setting have larger values for the Vd of gentamicin and other aminoglycosides than do general medical patients. Indeed, young very healthy people who suddenly require an aminoglycoside for a perforated or gangrenous appendix often have even smaller values for Vd [18].

9.1. An Aminoglycoside Patient with a sudden Change in Clinical Status and Volume of Distribution

An interesting 54 year old woman in Christchurch, New Zealand, was seen through the courtesy of Dr. Evan Begg in the fall of 1991. She was 69 in tall, weighed 80 kg, and her serum creatinine on admission was 0.7 mg/dL. She had a pyelonephritis, and was receiving tobramycin 80 mg approximately every 8 hours. She had a measured peak serum concentration of 4.6 and a trough of 0.4 ug/ml respectively, and had been felt by all to be having a satisfactory clinical response. During this time, her Vd was 0.18 l/kg, based on those two serum samples. However, on about the 6th day, she suddenly and most unexpectedly relapsed and went into clear-cut septic shock.

Following her surprising relapse on therapy, she was aggressively treated with much larger doses, 300 mg every 12 hours during this time. Her serum concentrations rose to peaks of 10.1 ug/ml. During this period of sudden septic shock, her serum creatinine also rose, from 0.7 to 3.7, and her estimated CCr fell to 18 ml/min/1.73m². After about another 10 days she improved. At that time, her serum tobramycin

concentrations rose to a peak of 16, and it was necessary to sharply reduce the dose to 140 mg about every 12 to 24 hours. Her serum creatinine fell to 1.1 to 1.3 mg/dL, and her CCr rose to 57 ml/min/1.73m².

It was simply not possible to get a good MAP Bayesian fit to all the serum data over the entire time period. Most samples were obtained during her second, sickest phase, and they dominated the fit. The ones at the beginning, prior to the sepsis, and at the end, after her improvement, were not at all well fitted.

Because of this, the data was divided into three parts - an initial one before her relapse into sepsis, a second one when she was septic, and a third one following improvement, but before it was felt safe to discontinue therapy. Each data set was fitted separately, using the USC*PACK programs [18].

During the first data set, the first 6 days, when her clinical behavior was that of a general medical patient, not gravely ill, her Vd was 0.18 L/kg as described above. The problem then was to pass on the ending values of the serum and peripheral compartment concentrations as initial conditions for the fitting process for the second data set. This was done, using that feature of the USC*PACK clinical software [18], which was developed specifically for this purpose.

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

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

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

At the Cleveland Clinic, Drs Marcus Haug and Peter Slugg [22] have spoken of "Vd collapse", when the Vd would drop from a larger to a smaller value. They showed that this change was a clinical indicator of incipient recovery of the patient. The present patient not only demonstrated such Vd collapse later on, as she got better, but also its opposite, Vd expansion, as she made the earlier transition from being a general medical

patient with a pyelonephritis to a seriously ill ICU patient with life-threatening septic shock.

We see, therefore, that not only do different populations of aminoglycoside patients have different values of V_d , but that each individual patient goes through these transitions, as demonstrated by this patient. The analysis of this patient's data was greatly facilitated, and indeed was only possible, using the MAP Bayesian approach, by breaking the dosage history up into several parts. Each part was then analyzed, and the ending concentrations from one part were passed on to the next data set as initial conditions or concentrations of drug present prior to the first dose given in the next data set, with the appropriate population model, if needed, as well.

9.2. A Patient on Digoxin when Quinidine was Added.

Another example of the utility of using initial conditions, provided through the courtesy of Dr. Marcus Haug, is that of a 72 year old woman, 4 ft 10 in tall, who weighed only 75 pounds. She was admitted to another hospital with congestive heart failure and atrial fibrillation. Her estimated creatinine clearance on admission was 38 ml/min/1.73m², falling to 23 after admission. She had been receiving 0.25 mg of digoxin daily. This was continued after admission to the hospital.

A serum digoxin concentration was 1.8 ng/ml on admission. Following this, her serum creatinine rose to 1.8 mg/dL, and her digoxin concentration after 5 days rose to 2.5 ng/ml. Her digoxin was stopped, though she had no clinical manifestations of toxicity. The next day her serum concentration had fallen to 2.0, and the next day it was down to 1.4 ng/ml, as shown in Figure 8.

At this point her ventricular rate with her atrial fibrillation had become rapid again, and she was restarted on her digoxin, again at 0.25 mg/day, to control the rate. However, quinidine was also added to her regimen at the same time. Her creatinine clearance at that time was 22 ml/min/1.73m².

Five days later her serum digoxin concentration was measured and found to be 7.6 ng/ml for a trough, and 10.0 ng/ml two hours after the next dose was given.

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




Figure 8. Plot of serum and peripheral compartment digoxin concentrations of patient admitted receiving digoxin. She was receiving 0.25 mg of digoxin daily, and weighed only 75 lb. Solid rectangles - measured serum concentrations. Solid line and left hand scale - digoxin serum concentrations. Dashed line and right hand scale - digoxin peripheral (nonserum) compartment concentrations. Using the MAP Bayesian approach, the population model for digoxin was fitted to the patient's data of doses and serum concentrations. Serum concentrations rose as her renal function worsened. Digoxin was stopped after the serum concentration of 2.5 ng/ml was obtained, after which her serum concentrations fell to 1.4, and, in the fitted model, finally to 1.19 ng/ml at the end of this plot, when digoxin was begun again, but along with quinidine.

The clinical problem was analyzed as follows. First, her original dosage history on digoxin alone was fitted to her serum concentrations, using the 2-compartment population model for digoxin made from the work of Reuning, Sams, and Notari [4,18]. This included the three measured serum concentrations. At the end of that part of her history, just before her first dose of quinidine was added, her fitted and predicted central compartment (serum) concentration was 1.19 ng/ml, and her peripheral (nonserum) compartment concentration was 7.58 ug/kg.

These two ending values from this first phase of her analysis were passed on as initial concentrations of drug already present in those compartments of her pharmacokinetic model at the time her digoxin was restarted, but now with quinidine as well. A population model for digoxin with quinidine [18] was now used. This model was not fitted to her subsequent serum concentrations, but merely used to supply predictions of those high measured concentrations. If the prediction was good, the interpretation

would be that the interaction would quantitatively account for the measured concentrations found. If not, then another explanation would have to be considered.




Figure 9. Plot of serum and peripheral compartment digoxin concentrations of patient admitted receiving digoxin. In this plot, digoxin was restarted at 0.25 mg/day, but along with quinidine. Solid rectangles - measured serum concentrations. Solid line and left hand scale - digoxin serum concentrations predicted using the population model for digoxin with quinidine [18]. Dashed line and right hand scale - predicted digoxin peripheral (nonserum) compartment concentrations. This plot begins with initial conditions equal to the final concentrations found at the end of the plot in Figure 10.

As shown in Figure 9, the predicted concentration of 7.2 ng/ml closely corresponded to the measured one of 7.6 ng/ml. In addition, the measured concentration of 10.0 ng/ml was predicted as 9.9 ng/ml. Because of these good predictions, it was felt that the digoxin-quinidine interaction adequately explained the measured concentrations found, and that no other alternative explanation was needed. This is a good example of how pharmacokinetic analyses can be used to evaluate clinical experiences with drugs, and can provide strong evidence for or against a particular clinical question or issue, much more than a clinical opinion made only judgmentally, unsupported by the quantitative evidence which can be obtained from such a model. The use of initial conditions was the key to being able to change from one population model to another in the middle of this patient's clinical history. In the same way, one can make the transition from regular theophylline to a long-acting preparation, or from conventional carbamazepine and valproate to their sustained-

release formulations, for example. With the use of initial conditions, one can thus follow the patient as he or she goes from one clinical situation to another, passing on the information from one data set to another.

10. LINKED PHARMACODYNAMIC MODELS: DIFFUSION OF DRUGS INTO ENDOCARDIAL VEGETATIONS, AND MODELS OF POSTANTIBIOTIC EFFECT AND BACTERIAL GROWTH AND KILL

In this section we will describe the linkage of nonlinear pharmacodynamic models to the basic linear pharmacokinetic model, and show some applications in clinical software describing drug diffusion into endocardial vegetations, the simulation of a postantibiotic effect, and the modeling of bacterial growth in the absence of a drug and its kill by an antibiotic.

10.1 MODELING DRUG DIFFUSION

A problem in the treatment of patients with infectious endocarditis is that it is difficult to estimate whether or not the drug is able to kill the organisms all the way into the center of a vegetation. Because of this, a diffusion model was made of this process.

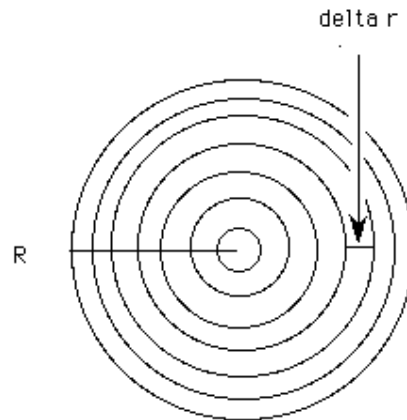


Figure 10. Diagram of the concentric layers of the spherical model of the endocardial vegetation.

A spherical shape was assumed for the vegetation, as shown in Figure 10, and it was modeled having several concentric layers, with diffusion taking place from layer to layer, δr , as shown in Figure 10. The sphere was assumed to be homogeneous, with equal diffusion in all directions, and as having a constant coefficient of diffusion throughout. The diffusion was assumed to be dependent on the concentration of drug in the surrounding medium, such as the serum concentration, and its time course. The diameter of endocardial vegetations can be measured by transesophageal echocardiography.

The following equation was used to model the diffusion:

$$\frac{\partial C}{\partial t} = \frac{1}{r^2} \times \frac{\partial}{\partial r} [D \times r^2 \times \frac{\partial C}{\partial r}] \quad (3)$$

where C represents the concentration in the sphere at time t, at a distance r from the center of the sphere, and D represents the coefficient of diffusion in the sphere, and \times indicates multiplication.

When D is assumed constant, the equation becomes

$$\frac{\partial C}{\partial t} = D \times \left[\frac{\partial^2 C}{\partial r^2} + \frac{2}{r} \times \frac{\partial C}{\partial r} \right] \quad (4)$$

The vegetation is assumed to be continuously immersed in the surrounding medium, and the drug concentration in that medium is assumed to attain a value which results in equilibrium with the very outer layer of the sphere. The medium then undergoes the changes in concentration with time that constitute the serum level time course. This time course is thus presented as the input to the spherical model [23].

The diffusion coefficient found by Bayer, Crowell, et al. for aminoglycosides in experimental endocarditis [24,25] was used. The model has become part of the USC*PACK clinical programs for individualizing drug dosage regimens [18]. The model can also be used to simulate behavior inside an abscess, and, by appropriate choice of sphere diameter and diffusion coefficient, to simulate the post-antibiotic effect of a certain desired duration.

10.1.1 Examples: Simulated Endocardial Vegetations of Various Diameters

Suppose one were to develop an amikacin dosage regimen for a hypothetical 65 year old man, 70 in tall, weighing 70 kg, with a serum creatinine of 1.0 mg/dL. Let us assume that he has a vegetation seen by echocardiography on his aortic valve that might be either 0.5, 1.0, or 2.0 cm in diameter. We wish to examine the ability of an amikacin regimen designed to achieve serum peaks of 45 ug/ml and troughs of approximately 5.0 ug/ml to reach effective concentrations within the vegetation in these three cases. Let us apply the findings of Bayer et al [24,25] to compute the time course of probable amikacin concentrations in the center of these three vegetations of different diameters, to examine their possible ability to kill an organism having an estimated minimum inhibitory concentration (MIC) of 8.0 ug/ml, for example.

Using the Amikacin program in the USC*PACK collection [18], let us estimate, from the patient's age, gender, height, weight, and serum creatinine concentration, that his creatinine clearance (CCr) is about 69 ml/min/1.73M². This method of estimating CCr is described elsewhere [15]. We enter the target goal for the peak serum concentration of 45 ug/ml and an initial trough concentration of about 5.0 ug/ml. The ideal dose interval to

achieve that peak and trough exactly, adjusted for the patient's renal function, employing a planned duration of the IV infusion of 0.5 hr, turns out to be 10.231 hrs. Let us approximate this in a practical manner by choosing a dose interval of 12 hrs. The dosage regimen to achieve the peak goal with such a dose interval is, when revised to practical amounts, 850 mg for the first dose, followed by 750 mg every 12 hrs thereafter.

On this regimen, predicted serum concentrations are 43 ug/ml for the peak and 3.2 ug/ml for the trough. The peak is 542 % of the stated MIC, and serum concentrations are predicted to be at least the MIC for 66 % of each dose interval. The AUC/MIC ratio for the first 24 hours is 48.8. The plot of these predicted serum concentrations is shown in Figure 11.




Figure 11. Predicted time course (the first 6 days) of serum Amikacin concentrations for the patient described. Upper horizontal dotted line - initial stated target peak serum concentration of 45 ug/ml. Lower horizontal dashed line - the estimated organism MIC of 8.0 ug/ml.

The important question now is whether or not this predicted serum concentration profile will result in adequate penetration of the vegetation in each of the three cases, and whether or not the regimen will kill effectively there, as well as in the central (serum level) compartment.

Figure 12 now shows the predicted amikacin concentrations in the center of the simulated vegetation having a diameter of 0.5 cm. As shown, concentrations rise rapidly above the MIC and stay there, suggesting that the above regimen should probably be able

to kill organisms having an MIC of about 8.0 ug/ml fairly promptly in the center of the vegetation. The time lag of concentrations in the center of the sphere is modest, about 3-4 hrs, behind the serum concentrations.




Figure 12. Predicted time course (the first 6 days) of Amikacin concentrations (dashed line) in the center of a simulated endocardial vegetation of 0.5 cm. Solid line - Predicted serum concentrations, and other lines and symbols as in Figure 11. The predicted endocardial concentrations rise promptly, and are consistently above the estimated MIC of 8.0 ug/ml.

On the other hand, if the vegetation were 1.0 cm in diameter instead, the drug would take about 12 hours to diffuse to the center and reach the MIC, and the rise and fall of drug concentrations would be much more damped, as shown in Figure 13.




Figure 13. Predicted time course (the first 6 days) of Amikacin concentrations (dashed line) in the center of a simulated endocardial vegetation of 1.0 cm. Solid line - Predicted serum concentrations, and other lines and symbols as in Figure 11. Predicted endocardial concentrations rise more slowly, are more damped, with smaller oscillations from peak to trough, but once the estimated MIC is reached, are consistently above 8.0 ug/ml.

Further, if the diameter of the vegetation were 2.0 cm, all this would take still longer, and the time course of the computed concentrations in the center would be as shown in Figure 14. The drug would take considerably longer, about 48 hours, to reach the MIC in the center of the vegetation, and significant growth of organisms might well take place before that. For every doubling of the diameter of the sphere, the equations show that it will take 4 times as long (the square of the ratio of the diameters) to reach an equal concentration in the center of the sphere.




Figure 14. Predicted time course (the first 6 days) of Amikacin concentrations (dashed line) in the center of a simulated endocardial vegetation of 2.0 cm. Solid line - Predicted serum concentrations, and other lines and symbols as in Figure 11. The predicted endocardial concentrations rise much more slowly and are much more damped, with essentially no oscillations from peak to trough. Once the estimated MIC is reached, the concentrations are consistently above 8.0 ug/ml, but two full days are required before the MIC is reached.

10.1.2 Another Example: Simulating an Abscess or a Post-Antibiotic Effect.

Another use for such a spherical model might, of course, be an abscess. If we could know the diffusion coefficient into abscesses of different sizes, we could similarly begin to model and compute the concentrations of drug diffusing into the abscess. There might well be different diffusion coefficients through the wall, into the bulk of the abscess, and into the center. All this is theoretically capable of being modeled. Effects of oxygen tension and pH upon bacterial growth and response to drugs can also be determined by careful needle aspiration done at carefully documented times just prior to incision and drainage of them, with careful cultures and determination of pH, pO₂, numbers of viable organisms, and rates of growth and kill from different parts of the abscesses. In this way, useful models of events taking place within an abscess can be made. One can see visually, in Figure 14, for example, why abscesses much over 1 cm in

diameter usually need to be incised and drained, as the diffusion into abscesses is very likely to be poorer than that seen here for endocardial vegetations.

Figure 15 shows computed drug concentrations in the center of a small sphere simulating a microorganism having a diameter of 0.1 micron, 3 simulated layers of diffusion, and a diffusion coefficient of 1.5×10^{-14} . This particular sphere has the property that in its center, the concentrations of drug fall below the MIC about 6 hrs after the serum levels do, thus simulating (without making any suggestions or conclusions about mechanism of action) a post-antibiotic effect of about 6.0 hrs, as the organisms will not begin to grow again for about 6 hrs after the serum concentrations fall below the MIC.




Figure 15. Plot of computed amikacin concentrations (the first 6 days) in the center of a simulated microorganism. Diffusion coefficients in the very small sphere diffusion model are adjusted so that concentrations in the center of the organism lag behind the serum concentrations and, if they fall below the MIC, would do so approximately 6 hrs after the serum concentrations do, thus simulating a post-antibiotic effect of about 6 hrs.

The effects of these computed concentrations in the center of these spheres will be discussed below in the section on modeling bacterial growth and kill. We see here that the process of diffusion into and out of spherical porous objects such as endocardial vegetations and small microorganisms can be described with reasonably simple models. The equations describing this process are the same as those for release of drug from a sustained-release preparation formulation.

11. MODELING BACTERIAL GROWTH AND KILL: CLINICAL APPLICATIONS

11.1 General Considerations

Let us assume that an organism is in its logarithmic phase of growth in the absence of any antibiotic. It will have a rate constant for this growth, and a doubling time. The killing effect of the antibiotic can be modeled as a Michaelis-Menten or Hill model. The model generates a rate constant for this effect. The rate of growth or kill of an organism depends upon the difference between these two rate constants. The killing effect will be determined by the E_{max} , representing the maximum possible rate constant for kill, the EC_{50} , the concentration at which the effect is half maximal, and the time course of the serum concentrations achieved with the dosage regimen the patient is given. Both the growth rate constant and the E_{max} can be found from available data in the literature for various organisms. The general growth versus kill equation is

$$\frac{dB}{dt} = (K_g - K_k) \times B \quad (5)$$

and

$$K_k = \frac{(E_{max} \times C_t^n)}{(EC_{50}^n + C_t^n)} \quad (6)$$

where B is the number of organisms (set to 1 relative unit at the start of therapy), K_g is the rate constant for growth, K_k is the rate constant for killing, E_{max} is the maximum possible effect (rate of killing), EC_{50} is the concentration at which the killing rate is half maximal, n is the Hill or sigmoidicity coefficient, and C_t is the concentration at the site of the effect (serum, peripheral compartment, effect compartment, or in the center of a spherical model of diffusion), at any time t , and x indicates multiplication..

The EC_{50} can be found from the measured (or clinically estimated) minimum inhibitory concentration (MIC) of the organism. This relationship was developed by Zhi et al. [26], and also independently by Schumitzky [27]. The MIC is modeled as a rate of kill that is equal to but opposite in direction to the rate constant for growth. The MIC thus offsets growth, and at the MIC there is neither net growth nor decrease in the number of organisms. At the MIC,

$$\frac{dB}{dt} = 0, \quad \text{and } K_k = - K_g \quad (7)$$

and

$$MIC = \left(\frac{K_g \times EC_{50}^n}{E_{max} - K_g} \right)^{1/n} \quad (8)$$

In this way, the EC_{50} can be found from the MIC, and vice versa.

The input to this effect model can be from either the central or the peripheral compartment concentrations of a pharmacokinetic model, or from the center (or any other layer) of one of the spherical models of diffusion. The sphere may represent an endocardial vegetation, an abscess, or even a small microorganism. In the latter case, one can adjust the sphere diameter and the diffusion coefficient so that the concentrations in the center of the small sphere lag behind the serum concentrations and cross below the MIC about 6 hours after the serum concentrations do, to simulate a post-antibiotic effect of about 6 hours, for example. The effect relationship was modeled by Bouvier D'Ivoire and Maire [28], from data obtained from Craig and Ebert [29], for pseudomonas and an aminoglycoside.

Let us first examine these effect models with relationship to the dosage regimen of amikacin developed in the previous section on analyzing concentrations in spherical diffusion models. In that section we had considered a hypothetical 65 year old man, 70 in tall, weighing 70 kg, having a serum creatinine of 1.0 mg/dL. We also assumed that he had a vegetation on his aortic valve, seen by echocardiography, that might be either 0.5, 1.0, or 2.0 cm in diameter, and we wanted to examine the ability of an amikacin regimen designed to achieve serum peaks of 45 and troughs of about 5 ug/ml to reach effective concentrations within the vegetation in these three cases. We applied the findings of Bayer et al [24,25] to predict the time course of amikacin concentrations in the center of the above three different vegetations. Let us now examine the results of these analyses.

The patient's dosage regimen consisted of an initial dose of 850 mg of amikacin followed by 750 mg every 12 hours thereafter. On that regimen, predicted serum concentrations were 43 ug/ml for the peak and 3.2 for the trough, possibly a bit low, as the MIC of the organism was stated to be 8.0 ug/ml. Figure 16 is a plot not only of the predicted time course (the first six days) of serum amikacin concentrations for the patient described here, but also of its ability to kill microorganisms using the model made by Bouvier D'Ivoire and Maire [28], based on the data of Craig and Ebert [29]. In Figure 16 there is no assumption of any post-antibiotic effect. The serum concentration profile alone is presented as the input to the bactericidal effect model.

The model always assumes an initial inoculum of one relative unit of organisms. The scale of the relative number of organisms is shown on the right side of Figure 16, while the scale of the serum concentrations is on the left. As shown in the figure, the serum concentration profile resulting from that regimen appears to be able to kill such an organism well in this particular patient. As the serum concentrations fall below the MIC with the first dose, however, the organisms begin to grow again, but the second dose kills them again, with slight regrowth once again toward the end of that dose interval. The third dose reduces the number of organisms essentially to zero. Use of this effect model

suggests that such a serum concentration profile should be effective in killing an organism having an MIC of 8.0 ug/ml, even though the serum concentrations are below the MIC about one third of the time, as the high peaks are effective in the killing.

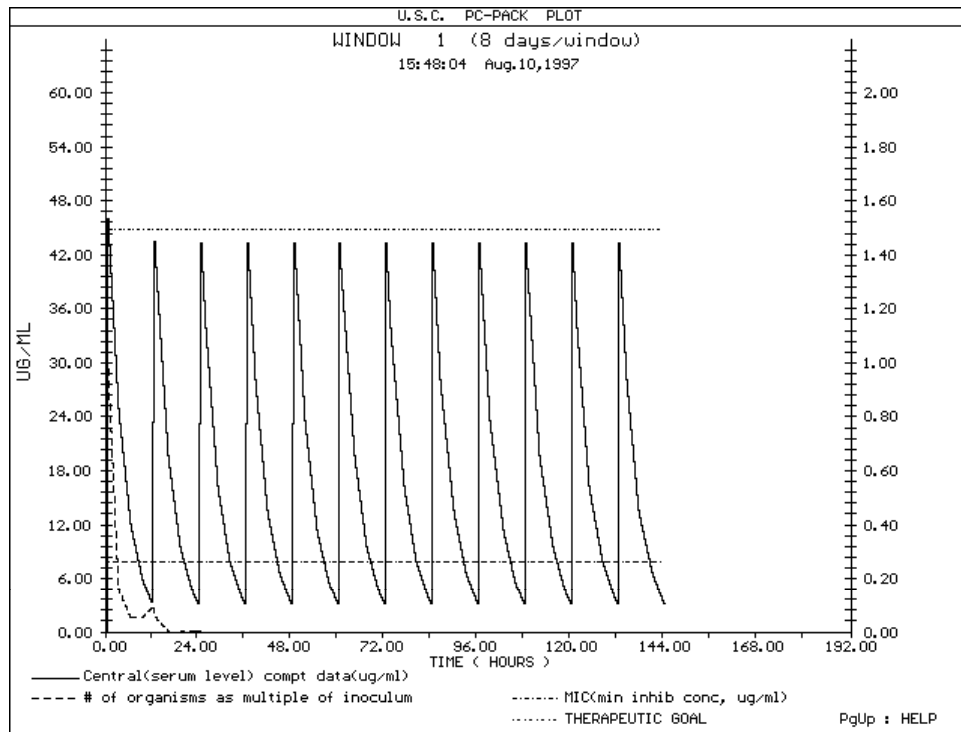


Figure 16. Predicted Killing effect of the regimen. Input from the central (serum) compartment profile of serum concentrations. The regimen is likely to kill well for a bloodstream infection (sepsis). Solid line and left hand scale - serum concentrations. Dashed line and right hand scale - relative numbers of organisms, with 1.0 relative unit present at the start of therapy. Upper horizontal dotted and dashed line - original peak serum goal of therapy. Lower horizontal dashed line: the patient's MIC of 8.0 ug/ml.

Figure 17 now shows the computed results in the center of the simulated endocardial vegetation having a diameter of 0.5 cm. The solid line shows the computed time course of amikacin which has diffused into the center of the vegetation. The dotted line again represents the computed relative number of organisms. Here the organisms grow almost 4-fold, to almost four relative units, before the concentrations in the center of the vegetation reach the MIC, after which killing begins. The organisms are reduced essentially to zero by 24 hours, suggesting that such a regimen would probably kill well in the center of an endocardial vegetation of approximately 0.5 cm diameter.




Figure 17. Killing effect as predicted in the center of the 0.5 cm diameter vegetation. Good and fairly prompt killing is seen. Solid line and left hand scale - drug concentrations in the center of the vegetation. Dashed line and right hand scale - relative numbers of organisms, with 1.0 relative unit present at the start of therapy. Upper horizontal dotted and dashed line - original peak serum goal of therapy. Lower horizontal dashed line: the patient's MIC of 8.0 ug/ml.

Figure 18 is a similar plot, but for the simulated vegetation having a diameter of 1.0 cm. Things here are not quite so good. There is a significant lag time of about 3 to 4 hours before any visible concentrations are reached in the center of the vegetation. The MIC is not reached until approximately 10 hours. During that time, the number of organisms has increased from one to approximately 150 relative units. However, after the MIC is reached, killing begins, although not quite so rapidly as with the smaller vegetation, due to the slower rate of rise of drug concentration in the center of the larger vegetation. However, killing appears to be essentially complete after approximately 40 hours. This suggests that the above regimen may be adequate to kill in the center of a 1.0 cm vegetation, but probably does so with less confidence of success than with the vegetation of 0.5 cm diameter. The doubt about this is suggested by the slower rate of killing and by the longer time required to reduce the number of organisms essentially to zero.




Figure 18. Killing effect computed for the center of a simulated vegetation with 1.0 cm diameter. The effect is delayed due to the slower diffusion into the center, but finally became adequate. Solid line and left hand scale - drug concentrations in the center of the vegetation. Dashed line and right hand scale - relative numbers of organisms, with 1.0 relative unit present at the start of therapy. Upper horizontal dotted and dashed line - original peak serum goal of therapy. Lower horizontal dashed line: the patient's MIC of 8.0 ug/ml.

As shown in Figure 19, things are much worse for the simulated endocardial vegetation of 2.0 cm. Diffusion into the center is a great deal (4 times) slower. Visible concentrations are not achieved until after approximately 12 hours, and about 48 hours are required before they reach the MIC. During this time, the number of organisms has increased astronomically, from 1 relative unit to over 1 million such units. However, after about five days, due to the continued presence of drug concentrations in the center of the vegetation approaching 12 to 15 ug/ml, killing in fact does seem to take place, and after about six days the number of organisms appears to be close to zero. However the behavior of this model strongly suggests that such a dosage regimen might very likely be inadequate in the center of a 2.0 cm simulated vegetation and might, at a minimum, require much more aggressive therapy with higher doses and serum concentrations, surgery, much more prolonged therapy, or all of these.

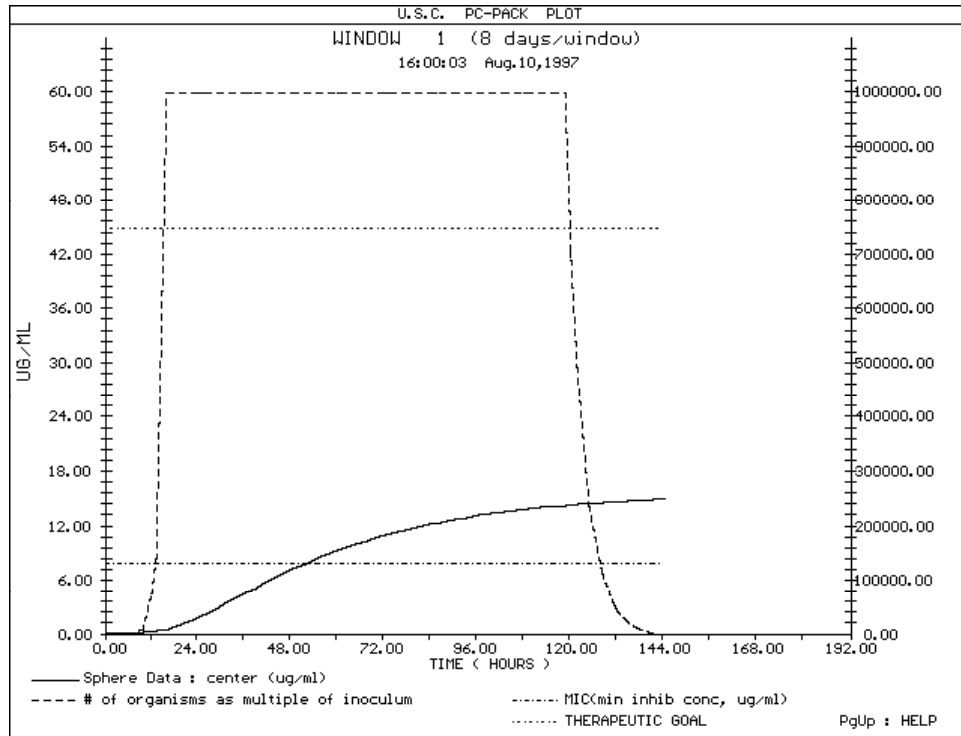


Figure 19. Killing effect computed for the center of a 2.0 cm simulated vegetation. Diffusion to the center is much prolonged while bacterial growth continues. Killing is delayed very significantly. Solid line and left hand scale - drug concentrations in the center of the vegetation. Dashed line and right hand scale - relative numbers of organisms, with 1.0 relative unit present at the start of therapy. Upper horizontal dotted and dashed line - original peak serum goal of therapy. Lower horizontal dashed line: the patient's MIC of 8.0 ug/ml.

Figure 20 shows the computed concentrations in the small hypothetical microorganism used in the previous paper to simulate the time course of the post-antibiotic effect (PAE). There is a lag of about 6 hours between the fall of the serum concentrations and that of the concentrations in the center of this hypothetical microorganism. Because of this, if the dosage interval were to be greater so that the concentrations in the hypothetical microorganism would fall below the MIC, one would see that they would do so approximately six hours after the serum concentrations fall below the MIC, thus simulating a post-antibiotic effect of approximately six hours.




Figure 20. Computed amikacin concentrations in the center of a hypothetical microorganism in which concentrations fall below the MIC about 6 hr after the serum concentrations do, thus simulating (regardless of mechanism) a post-antibiotic effect of about 6 hrs. Solid line and left hand scale - serum drug concentrations. Dashed line and right hand scale - computed concentrations in the center of the microorganism simulating the post-antibiotic effect. Upper horizontal dotted and dashed line - original peak serum goal of therapy. Lower horizontal dashed line: the patient's MIC of 8.0 ug/ml.

What would be the contribution (if any) of such a PAE to overall therapy? As shown in Figure 21, the outcome is not very different from that shown in Figure 16. In both cases, killing is rapid and prompt. One can see that due to the diffusion model, there may be a delay of approximately six hours before the concentrations in the hypothetical microorganism reach the MIC. During that time, the number of organisms has grown from 1 to about 5.4 relative units. However, after that time, the concentrations are always above the MIC, and killing at a significant rate begins and continues, with the organisms being reduced essentially to zero by about 36 hours.




Figure 21. Killing effect predicted for the simulated post-antibiotic effect of 6 hrs, using the computed concentrations in the center of the simulated microorganism as input to the effect model. Solid line and left hand scale - drug concentrations in the center of the microorganism simulating the post-antibiotic effect. Dashed line and right hand scale - relative numbers of organisms, with 1.0 relative unit present at the start of therapy. Upper horizontal dotted and dashed line - original peak serum goal of therapy. Lower horizontal dashed line: the patient's MIC of 8.0 ug/ml.

11.2 An Interesting Case

These diffusion and effect models were of interest when they were used to analyze retrospectively the data obtained much earlier, back in 1991, from the patient described earlier, in section 9.1, from Christchurch, New Zealand, seen through the courtesy of Dr. Evan Begg. She had a pyelonephritis and received tobramycin, 80 mg approximately every 8-12 hours. She was having a satisfactory clinical response to therapy when, on about the 6th day of therapy, she suddenly and unexpectedly relapsed and went into septic shock. She then received much more aggressive tobramycin, and eventually recovered. The MIC of her organism was 2.0 ug/ml. The analysis of bacterial growth and kill described below was not done until several years later, after the models of diffusion and growth and kill had been developed.

Figure 22 shows the computed concentrations of drug in the center of the 0.1 micron sphere representing a hypothetical organism having a PAE of 6 hours, in the patient's first phase, when she appeared to be a general medical patient (not an ICU

patient) having a satisfactory clinical response to her tobramycin therapy. However, at the end of this time (from 120-148 hours into therapy) she unexpectedly relapsed on about day 6 of therapy, went into septic shock, and clearly became an ICU patient.




Figure 22. Patient receiving Tobramycin. This figure shows the measured serum concentrations (small rectangles), and the Bayesian fitted model. Small solid rectangles - measured serum concentrations. Solid line and left hand scale - fitted serum drug concentrations. Dashed line and right hand scale - concentrations in the small organism simulating the post-antibiotic effect. Horizontal dashed line: the patient's MIC of 2.0 ug/ml.

Note the damped response in the center of the small sphere to the sharp peaks and troughs of the serum concentrations. While one may have many different views as to what the mechanism of the PAE is, this diffusion model appears to do a reasonable job of describing the effect itself. As data accumulate about diffusion into endocardial vegetations and abscesses, this diffusion model will permit modeling of these events during a patient's clinical care in a way that is now becoming possible, as illustrated in the present case.

Now, consider the computed rate of growth and kill of her organisms during her treatment. We have a hint in that her peak serum concentrations were low, only about 5.0 ug/ml. Her measured serum peak was 4.5 ug/ml, and her trough was 0.4 ug/ml. Figure 23 describes the growth and kill of the organisms in response to events in her serum concentration compartment, while Figure 24 shows the same events as viewed with the sphere model simulating the PAE.




Figure 23. Patient receiving Tobramycin. This figure shows the measured serum concentrations (small solid rectangles), and her individualized Bayesian fitted model. Solid line and left hand scale: fitted serum drug concentrations. Dashed line and right hand scale: relative numbers of organisms. The plot always begins with 1.0 relative units of organism. Horizontal dashed line: the patient's MIC of 2.0 ug/ml.




Figure 24. Graph of effects found with the model simulating the PAE of about 6 hours. Solid line and left hand scale: drug concentrations in the microorganism simulating the post-antibiotic effect. Dashed line and right hand scale: relative numbers of organisms. The plot always begins with 1.0 relative units of organism, as shown on the right hand scale. Horizontal dashed line: the patient's MIC of 2.0 ug/ml.

Note in both figures that there appear to be few organisms present at the outset of therapy. Growth becomes visible in Figure 23 after the first hiatus between doses, and then becomes exponential during day 6 of therapy, after the last dose on that plot, which

ends just before the next dose which was given during his next period, that of septic shock. Figure 24 extends this examination to the possible contribution of the PAE, where the concentrations in the center of the small sphere are evaluated with respect to their ability to kill the organisms. Note that the regimen also appeared to be effective at first, but that the organisms grew out exponentially when the concentrations fell below the MIC for a significant time, but slightly later than in Figure 23, showing events without any PAE. The exponential growth of organisms escaping from control as shown in Figures 23 and 24 correlated well with the patient's clinical relapse at just that time, with development of septic shock.

Figure 25 shows the subsequent course of this patient in her next phase of acute urosepsis. There was essentially no carry-over of drug from the last dose shown in Figures 22 to 24 to the patient's next dose, which was given at time zero in Figure 25. Figure 25 shows the many serum concentrations measured during this second phase of her hospital course. It also shows the results of Bayesian fitting based on the population model of tobramycin for ICU patients [18], with its much larger central volume of distribution (now that she had become a seriously ill ICU patient) in both her central (serum) and peripheral (nonserum) compartments. During this time, the patient's serum creatinine also rose from 0.7 to 3.7 mg/dL. One can see that it took about two days, as new serum concentrations were obtained, for ward personnel to react to her suddenly much increased volume of distribution (from 0.18 to 0.51 L/kg), and to her much decreased renal function, and to give her the much larger doses required to achieve effective peak serum concentrations, despite her rising serum creatinine. Note also that her trough concentrations rose from about 0.3 up to 2.0 ug/ml during this time, so that the time that serum concentrations were below the MIC was greatly reduced.

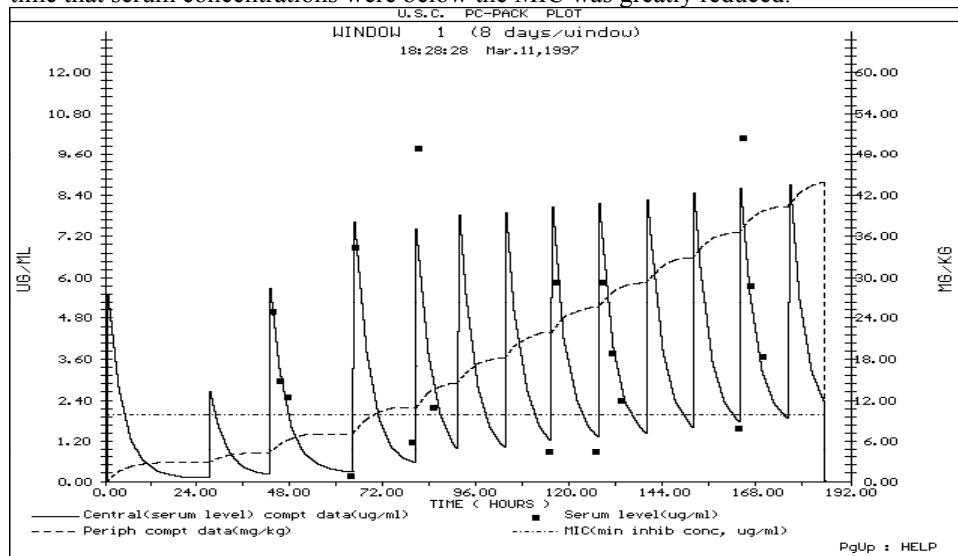


Figure 25. Plot of serum and peripheral compartment concentrations during the time of the patient's sepsis and his recovery. Small solid rectangles - measured serum concentrations. Solid line and left hand scale - fitted serum concentrations. Dashed line and right hand scale - peripheral compartment concentrations, also fitted from the serum data. Horizontal dashed line: the patient's MIC of 2.0 ug/ml.

Figure 26 shows the plot of the computed bacterial growth and kill based on the input from the serum concentration profile during this time, without any aid from the simulated PAE. The organisms grow out of control in the first two days (again correlating with the patient's relapse into sepsis the day before, and the time required for ward personnel to perceive the problem and to adjust her dosage sharply upward to achieve serum peaks in the range of 7 to 9 ug/ml). As these higher and more effective concentrations were achieved, however, bacterial killing could finally be seen at about the sixth day in this figure, and appeared to be effective after that, thus correlating with the patient's subsequent clinical recovery.

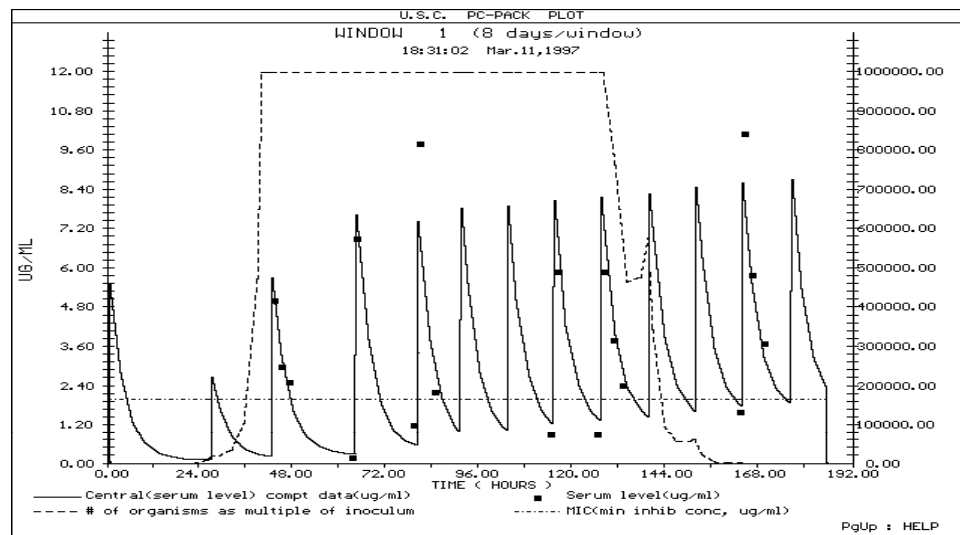


Figure 26. Plot of the effect on growth and kill using input from the serum concentration profile. The organisms grow out of control when serum concentrations are lower, but kill again when they are higher. These events correlated well with the patient's relapse at the beginning of the plot, and his recovery about one week later. Small solid rectangles - measured serum concentrations. Solid line and left hand scale - serum concentrations. Dashed line and right hand scale - relative numbers of organisms, with 1.0 relative unit present at the start of therapy. Horizontal dashed line: the patient's MIC of 2.0 ug/ml.

Figure 27 below shows the same events, but now using the diffusion model of the small microorganism and its simulated PAE. Using this model, the concentrations in the center of the small simulated microorganism do not exceed the MIC until almost 72 hours. Significant killing can be seen to begin slightly earlier in this figure, at about 110 hours in the figure, compared to about 130 hours (about 1 day later) in Figure 26, and appeared to be effective after that.

In general, models of bacterial growth and kill permit one to incorporate known *in vitro* data of the logarithmic growth rate of the organism and the maximum kill rate achieved with the antibiotic, to integrate it with data of the MIC of each individual patient's organism, and to model the growth and decline of the relative numbers of organisms. These Zhi models have correlated well, in this patient, with her unexpected

relapse from having an apparently satisfactory response to therapy to becoming a seriously ill patient with septic shock, and with her subsequent recovery later on, as effective serum concentrations were achieved and maintained.




Figure 27. Plot of the effect on growth and kill using input from the center of the organism, using its computed concentrations as input to the effect model. The post-antibiotic effect helps somewhat to delay the relapse and to augment the kill. Solid line and left hand scale - computed concentrations in center of microorganism simulating the post-antibiotic effect. Dashed line and right hand scale - relative numbers of organisms, with 1.0 relative unit present at the start of therapy. Horizontal dashed line: the patient's MIC of 2.0 ug/ml.

The Zhi model does not describe the decline of bacterial growth rate seen over time, reaching a maximum number of organisms, as found by Mouton, Vinks, and Punt [30]. The organisms are always assumed to be in their logarithmic growth phase, the maximum possible. In addition, the model does not account for the increase in bacterial resistance found over time, and the emergence of resistant organisms. However, one can use the maximum possible MIC which the emerging resistant organism is expected to reach, and examine the behavior of the model. In this case, the Zhi model becomes a useful example of the "worst case" scenario, with the resistant organisms being so from the very beginning of therapy, and with the logarithmic growth rate always being in effect, never slackening. If a given dosage regimen, generating a certain serum concentration profile, can kill well using the Zhi model, one might expect it probably to do at least as well in clinical circumstances, where the growth rate may (or may not) slacken with time and may (but may not) reach a maximum number of organisms, and the resistant organisms emerge more slowly with time.

Clearly, further work in this area is needed, but models of this type are beginning to provide a useful new way to perceive, analyze, and evaluate the efficacy of

antibiotic therapy. Similar approaches may also be useful in analyzing therapy of patients with AIDS, using the PCR assays, and with cancer.

12. LIMITATIONS OF CURRENT MAP BAYESIAN ADAPTIVE CONTROL

The maximum a posteriori probability (MAP) Bayesian approach to adaptive control and dosage individualization is straightforward and robust. However, it does not represent an optimal approach to dosage individualization. It has two significant limitations.

The first one is that the pharmacokinetic model parameter values used to describe the behavior of the drug are assumed to be either normally or log-normally distributed. This is often not so. Many drugs, for example, have clusters of both rapid and slow metabolizers within the population, and therefore may well have multimodal population parameter distributions for the elimination rate constant. Furthermore, the volume of distribution for drugs such as the aminoglycosides is affected by the patient's clinical state as a general medical patient or a patient in an intensive care unit, for example. Because of this, parameter distributions are often asymmetrical, neither normally or lognormally distributed, and are therefore not optimally described by mean, median, or mode values. This point reflects the significant problems associated with making parametric population models, and with using mean or median parameter values to develop dosage regimens. The problem is largely overcome by making nonparametric population models which describe the entire joint parameter distribution within the population, with up to one support point (set of parameter values, and its estimated probability) for each subject studied in the population [13,14].

The second limitation is that there is no tool in the MAP Bayesian strategy to estimate and predict the precision with which a desired dosage regimen developed to hit a desired target goal actually is likely to do so. The method lacks a vital performance criterion.

The separation, or heuristic certainty equivalence, principle is well known among the stochastic control community, but less so among the pharmacokinetic community. It states [31] that when the task of controlling the behavior of a system is separated into the steps of:

1. Obtaining the best single point parameter values in the model describing the behavior of the system, and then,
2. Using these single point values to design the inputs to control the system,

that, the task of hitting the target goal is usually performed suboptimally. Yet this is exactly what the MAP Bayesian, and all methods which estimate single values for each model parameter, do.

There is no performance criterion to optimize in the MAP Bayesian dosage strategy (such as the estimated precision with which the desired target will be hit, for

example) as there is only one set of parameter values, and the target is simply assumed to be hit exactly.

13. OVERCOMING THE LIMITATIONS: "MULTIPLE MODEL" DESIGN OF MAXIMALLY PRECISE DRUG DOSAGE REGIMENS

The two limitations above are overcome by the combination of nonparametric population models [13,14] and the "multiple model" design of dosage regimens [32]. Nonparametric population models have been discussed in another chapter. Their strength is that they are consistent, statistically efficient, and have good properties of statistical convergence [35]. They are not limited by the assumption that the parameter distributions must be Gaussian or lognormal, as in parametric methods. Instead of simply estimating parameter means, variances, and correlations between them, as point estimates of a distribution, the nonparametric methods estimate the entire parameter distributions themselves. These distributions are discrete, not continuous. They consist of discrete sets of parameter estimates, along with an estimate of the probability of each set [13,14]. Up to one set (support point of the distribution) is obtained for each subject studied in the population. This closely approaches the ideal population model (which can never be attained), which would consist of the correct structural model of the drug system, along with the exact value of each parameter in each subject if it would somehow be possible to know those values.

When a parametric population model is used as the Bayesian prior to design an initial dosage regimen for the next patient one encounters, one usually has only a single estimated value for each parameter. Because of this, only one prediction of future concentrations can be made. The action taken is therefore based only on the estimates of the central




Figure 28. Using lidocaine population mean parameter values, an infusion regimen designed to achieve and maintain a target goal of 3 ug/ml does so exactly when the patient, as here, has exactly the mean population parameter values.

tendencies of the parameter distributions, and not on the entire parameter distributions, which may well be multimodal, due to genetic polymorphism in the distribution and metabolism of drugs. The dosage regimen is simply assumed to achieve the target goal exactly, as shown in Figure 28. Figure 28 shows the results of an infusion regimen of lidocaine, based on the mean population parameter values for that drug, which was designed to achieve and maintain a target serum concentration of 3 ug/ml. As shown, this regimen, based on the single mean population parameter values, hits the target exactly, but only when the patient has parameter values which are exactly the population mean values. However, as shown in Figure 29, when the regimen used in Figure 28 was given to the combination of the actual 81 diverse nonparametric population support points from which these mean parameter values were obtained, an extremely wide distribution of predicted serum concentrations was seen, due to the diversity in the nonparametric population support points from which the mean parameter values were obtained, representing the diversity in the parameter values in the patient population. The predicted serum concentrations actually covered much more than the usual therapeutic range of 2 to 6 ug/ml.

In contrast, if one has a nonparametric population model [13,14], with its multiple sets of model parameter values (81 in this case), one can make multiple predictions, instead of only one, forward into the future from any candidate dosage regimen which is "given" to all the models in the population discrete joint density. The richer and more likely population parameter joint density reflects much better the actual diversity among the subjects studied in the past population. Based on these multiple models in the population (the discrete joint density), one can compute




Figure 29. Result when the above lidocaine infusion based on population mean parameter values is given to the 81 diverse support points from which the population mean values were obtained. Great diversity in the predicted responses is seen.

the weighted squared error with which any candidate regimen is predicted to fail to achieve the desired target goal at a target time. Other regimens can then be considered, and the optimal regimen can be found which is specifically designed to achieve the desired target goal with the least weighted squared error [32].

This approach, using the multiple models of the patient provided by the nonparametric population model, avoids much of the limitations of the separation principle. This is the real strength of the combination of nonparametric population models coupled with "multiple model" dosage design [32].




Figure 30. Predicted response of the 81 support points (models) when the regimen obtained by multiple model dosage design is given. The target is achieved with visibly greater, and optimal, precision.

As shown in Figure 30, the multiple model (MM) dosage regimen, based on the same nonparametric population model with its 81 support points, obtained a much more precise achievement of the target goal, because it was specially designed to do so. The error in the achievement of the therapeutic target goal is much less, and the dispersion of predicted serum concentrations about the target goal is much less.

13.1 Obtaining "Multiple Model" Bayesian Posterior Joint Parameter Distributions

With the MAP Bayesian approach to posterior parameter values, the single most likely value for each parameter is obtained when they altogether minimize the objective function shown in equation (1). In contrast, the MM Bayesian approach, using the nonparametric joint densities, preserves the multiple sets of population parameter values, but specifically recomputes their Bayesian posterior probability, based upon the serum concentrations obtained. Those combinations of parameter values that predict the measured concentrations well become more probable. Those that predict them less well

become less so. In this way, the probabilities of all the nonparametric population model support points become revised, using Bayes' theorem [33]. A smaller number of significant points, or perhaps even only one, is usually obtained. When the regimen for the next cycle is developed, these revised models, containing their revised MM Bayesian posterior probabilities, are used to develop it. The regimen is again specifically designed to achieve the desired target goal with maximum precision (minimum weighted squared error).

13.2 Other Bayesian Approaches

Three sequential Bayesian approaches have also been used by us to incorporate feedback from measured serum concentration data. The first is the sequential MAP Bayesian approach, in which the MAP posterior parameter values are sequentially updated after each serum concentration data point is obtained. This procedure improves the tracking of the behavior of the drug through each data set. However, at the end of each full feedback cycle, (after each new full cluster of data points), at the time the next regimen is to be developed, this method has learned no more with respect to developing the next new dosage regimen, than if it had fitted all the data together at once, even though it estimates changing MAP Bayesian parameter values sequentially. This is because the method only estimates a single set of parameter values that fit the objective function best, over all the data points.

The second approach is the sequential MM Bayesian one [33]. Here the MM Bayesian posterior joint density is also sequentially updated after each data point. Still, at the end of each feedback cycle, this procedure similarly has learned no more with respect to developing the next dosage regimen than if all the data in that cluster were fitted simultaneously. The procedure is still looking for a hypothetical single model (support point, set of parameter values) which best describes all the data. When this fails to be the case, combinations of support points are found which fit best. Still, the procedure estimates a fixed and unchanging single model, or combination of models (support points), which best fit the data, even though the posteriors are fitted sequentially.

A third approach is the interacting multiple model (IMM) approach [36]. This method permits the true patient being sought for actually to jump from one model or support point to another during the sequential Bayesian analysis. Because of this, the IMM method, originally designed to track missiles and aircraft taking evasive action, permits detection of changing pharmacokinetic parameter densities during the sequential analysis procedure. It thus provides an improved method to track the changing parameter densities and the behavior of a patient during the evolution of his/her clinical therapy. For example, it permits an improved ability to detect and to quantify changes in the volume of distribution of aminoglycoside drugs during changes in a patient's clinical status which are not captured by the use of conventional clinical descriptors.

Using carefully simulated models in which the true parameter values changed during the data collection, the integrated total error in tracking a simulated patient was very similar with the sequential MAP and sequential MM Bayesian procedures. However, the integrated total error of the sequential IMM procedure was only about one half that of the other two [36].

13.3 Clinical Applications

Nonparametric population parameter joint densities, MM dosage design and IMM Bayesian posterior joint densities appear to offer significant improvements in the ability to track the behavior of drugs in patients throughout their care, especially when the patients are unstable and have changing parameter values. These approaches also develop dosage regimens which are specifically designed to achieve target goals with maximum precision. These methods make optimal use of all information contained in the past population data, coupled with whatever current data of feedback may be available up to that point, to develop that patient's most precise dosage regimen.

13.4 Analyzing the Tobramycin Patient with MM and IMM Sequential Bayesian Methods. Implementation into Clinical Software

The above tools have now been implemented in clinical software, the MM-USCPACK package [33]. Figure 31 shows the plot of the fit to the data from the patient described in sections 9.1 and 11.2, and in Figure 22, when analyzed using the sequential MM Bayesian approach. The data was very poorly fitted, due to the changing parameter values in this highly unstable patient, as she changed from being a general medical patient with a pyelonephritis before 150 hours to an acutely ill and highly unstable patient with severe septic shock after that time.




Figure 31: Fit to data of patient described in section 9.1, 11.2, and Figure 22, analyzed with the MM Bayesian approach. Note the very poor fit to the data, due to the patient's changing parameter values as her clinical status changed significantly, going from someone with a pyelonephritis before 150 hours, to someone with clearcut septic shock afterward, becoming an acutely and severely ill intensive care patient.

On the other hand, Figure 32 shows the result when the IMM algorithm was used to analyze the patient's data. The fit was greatly improved, and the IMM algorithm was able to track the changing behavior of tobramycin in this acutely and severely ill, highly unstable patient.




Figure 32: Fit to data of patient described in section 9.1, 11.2, and Figure 22, analyzed with the IMM Bayesian approach. Note the very much improved fit to the data, as the IMM approach tracks the changing parameter values taking place in this acutely ill and highly unstable patient.

14. THE FUTURE OF INDIVIDUALIZED DRUG THERAPY

The clinical software above incorporates all the strengths of the nonparametric population models and the multiple model dosage design, and the MM and IMM Bayesian analysis of individual patient data. Work is now under way to develop similar MM dosage designers for the large and often nonlinear models of drugs having that behavior, such as Phenytoin, Carbamazepine, and many drugs used in the treatment of patients with AIDS, transplants, and cancer, for example. Nonparametric population models can be made now of large, nonlinear, interacting and multiple drug systems such as those found in the combination chemotherapy of many of the above patients, where the concentrations of one drug may either compete with a metabolite, or may increase or decrease the rate of metabolism of another drug. As MM dosage design is developed for such large multidrug systems, it will be possible, for the first time, to develop truly coordinated, maximally precise combination chemotherapy for patients with such problems, maximizing effectiveness while constraining toxicity within specifically selected quantifiable limits.

ACKNOWLEDGMENTS

Supported by US Government grants LM 05401, RR 01629, RR 11526, and GM65619.

References

1. Cohen J: *Make Your Medicine Safe*. ISBN: 0-380-79075-0. Avon Books Inc., New York, 1998.
2. Evans W: General Principles of Applied Pharmacokinetics, Chapter 1, in *Applied Pharmacokinetics: Principles of Therapeutic Drug Monitoring*, ed. By Evans W, Schentag J, and Jusko W. Applied Therapeutics Inc., Vancouver B.C., 1992.
3. van Lent-Evers N, Mathot R, Geus W, van Hout B, and Vinks A: Impact of Goal-Oriented and Model-Based Clinical Pharmacokinetic Dosing of Aminoglycosides on Clinical Outcome: A Cost-Effectiveness Analysis. *Therap. Drug Monit.* 221: 63-73, 1999.
4. Reuning R, Sams R, and Notari R: Role of Pharmacokinetics in Drug Dosage Adjustment. 1. Pharmacologic Effects, Kinetics, and Apparent Volume of Distribution of Digoxin. *J. Clin. Pharmacol.* 13: 127-141, 1973.
5. Jelliffe R, Schumitzky A, Van Guilder M, Liu M, Hu L, Maire P, Gomis P, Barbaut X, and Tahani B: Individualizing Drug Dosage Regimens: Roles of Population Pharmacokinetic and Dynamic Models, Bayesian Fitting, and Adaptive Control. *Therap. Drug Monit.*, 15: 380-393, 1993.
6. Sheiner LB, Beal S, Rosenberg B, and Marathe V: et al.: Forecasting Individual Pharmacokinetics. *Clin. Pharmacol. Ther.*, 31: 294-305, 1979.
7. Jelliffe R, Iglesias T, Hurst A, Foo K, and Rodriguez J: Individualizing Gentamicin Dosage Regimens: A Comparative Review of Selected Models, Data Fitting Methods, and Monitoring Strategies. *Clin. Pharmacokinet.* 21: 461-478, 1991.
8. Nelder JA and Mead R: A Simplex Method for Function Minimization. *Computer Journal* 7: 308-313, 1965.
9. Caceci MS and Cacheris WP: Fitting Curves to Data: The Simplex Algorithm is the Answer. *BYTE Magazine*, May 1984, pp. 340-362.
10. De Groot MH: *Probability and Statistics*, Second Edition, Addison-Wesley Publishing Co., Reading, MA, 1989, pp. 422-423.

11. Sawchuk R and Zaske D: Pharmacokinetics of Dosing Regimens which Utilize Multiple Intravenous Infusions: Gentamicin in Burn Patients. *J. Pharmacokin. Biopharm* 4: 183-195, 1976.
12. Jelliffe R, Schumitzky A, Bayard D, Milman M, Van Guilder M, Wang X, Jiang F, Barbaut X, and Maire P: Model-Based, Goal-Oriented, Individualised Drug Therapy: Linkage of Population Modelling, New "Multiple Model" Dosage Design, Bayesian Feedback, and Individualized Target Goals. *Clin. Pharmacokinet.* 34: 57-77, 1998.
13. Mallet A: A Maximum Likelihood Estimation Method for Random Coefficient Regression Models. *Biometrika.* 73: 645-656, 1986.
14. Schumitzky A: Nonparametric EM Algorithms for Estimating Prior Distributions. *App. Math. and Computation.* 45: 143-157, 1991.
15. Jelliffe R : Estimation of Creatinine Clearance in Patients with Unstable Renal Function, without a Urine Specimen. *Am. J. Nephrol.* 22: 320-324, 2002.
16. Maire P, Jelliffe R, Dumarest C, Roux D, Breant V, Charpiat B, Vermeulen E, Brazier J, and Courpron P: Controle Adaptatif Optimal des Posologies: Experience des Aminosides en Geriatrie. in *Information et Medicaments. Comptes Rendus du Colloque AIM-IF et IRT, Paris, December 1989*, ed. by Venot A and Degoulet P, Volume 2 of *Informatique et Sante*, directed by Degoulet P, Springer Verlag, Paris, 154-169, 1989.
17. Hurst A, Yoshinaga M, Mitani G, Foo K, Jelliffe R, and Harrison E.: Application of a Bayesian Method to Monitor and Adjust Vancomycin Dosage Regimens. *Antimicrob. Agents Chemother.*, 34: 1165-1171, 1990.
18. Jelliffe R, Schumitzky A, Van Guilder M, and Jiang F: User Manual for Version 10.7 of the USC*PACK Collection of PC Programs. December 1, 1995. Laboratory of Applied Pharmacokinetics, University of Southern California School of Medicine, Los Angeles, CA.
19. Jelliffe R: Clinical Applications of Pharmacokinetics and Control Theory: Planning, Monitoring, and Adjusting Dosage Regimens of Aminoglycosides, Lidocaine, Digitoxin, and Digoxin. In Maronde R, ed: *Topics in Clinical Pharmacology and Therapeutics*, Sprnnger-Verlag, New York, 1986, pp. 26-82.
20. Rodman J, Jelliffe R, Kolb E, Tuey D, de Guzman M, Wagers P, and Haywood L: Clinical Studies with Computer-Assisted Initial Lidocaine Therapy. *Arch. Int. Med.* 144: 703-709, 1984.
21. Bleyzac N, Souillet G, Magron P, Janoly A, Martin P, Bertrand Y, Galambrun C, Dai Q, Maire P, Jelliffe R, and Aulagner G: Improved Clinical Outcome of Paediatric Bone Marrow Recipients using a Test Dose and Bayesian Pharmacokinetic Individualization of Busulfan Dosage Regimens. *Bone Marrow*

- Transplantation 28: 743-751, 2001.
22. Marcus Haug, Pharm.D., and Peter Slugg, M.D., Personal communication.
 23. Maire P, Barbaut X, Vergnaud JM, El Brouzi M, Confesson M, Pivot C, Chuzeville M, Ivanoff N, Brazier J, and Jelliffe R: Computation of Drug Concentrations in Endocardial Vegetations in Patients during Antibiotic Therapy. *Int. J. Bio-Med. Comput.*, 36: 77-85, 1994.
 24. Bayer A, Crowell D, Yih J, Bradley D, and Norman D: Comparative Pharmacokinetics and Pharmacodynamics of Amikacin and Ceftazidime in Tricuspid and Aortic Vegetations in Experimental *Pseudomonas* Endocarditis. *J. Infect. Dis.*, 158: 355-359, 1988.
 25. Bayer A, Crowell D, Nast C, Norman D, and Borelli R: Intravegetation Antimicrobial Distribution in Aortic Endocarditis Analyzed by Computer-Generated Model: Implications for Treatment. *Chest*, 97: 611-617, 1990.
 26. Zhi J, Nightingale CH, and Quintiliani R: Microbial Pharmacodynamics of Piperacillin in Neutropenic Mice of Systemic Infection due to *Pseudomonas Aeruginosa*. *J Pharmacokin. Biopharm.* 4: 355-375, 1988.
 27. Schumitzky A: personal communication.
 28. Bouvier D'Ivoire M, and Maire P: Dosage Regimens of Antibacterials: Implications of a Pharmacokinetic/Pharmacodynamic Model. *Drug Invest.* 11: 229-239, 1996.
 29. Craig W, and Ebert S: Killing and Regrowth of Bacteria in Vitro: a Review. *Scand J Infect Dis Suppl* 74: 63-70, 1991.
 30. Mouton J, Vinks AATMM, and Punt N: Pharmacokinetic-Pharmacodynamic Modeling of Ceftazidime during Continuous and Intermittent Infusion. Chapter 6, pp 95-110, in the Ph.D. Thesis of Vinks AATMM: Strategies for Pharmacokinetic Optimization of Continuous Infusion Therapy of Ceftazidime and Aztreonam in Patients with Cystic Fibrosis, November, 1996.
 31. Bertsekas D: Dynamic Programming: deterministic and stochastic models. Englewood Cliffs (NJ): Prentice-Hall, pp. 144-146, 1987.
 32. Jelliffe R, Bayard D, Milman M, Van Guilder M, and Schumitzky A: Achieving Target Goals most Precisely using Nonparametric Compartmental Models and "Multiple Model" Design of Dosage Regimens. *Therap. Drug Monit.* 22: 346-353, 2000.
 33. Jelliffe R, Bayard D, Schumitzky A, Milman M, Jiang F, Leonov S, Gandhi A, and Botnen A: A New Clinical Software Package for Multiple Model (MM) Design of Drug Dosage Regimens for Planning, Monitoring, and Adjusting

Optimally Individualized Drug Therapy for Patients. Presented at the 4th International Meeting on Mathematical Modeling, Technical University of Vienna, Vienna, Austria, February 6, 2003.

34. Jelliffe R: A Mathematical Analysis of Digitalis Kinetics in Patients with Normal and Reduced Renal Function. *Math. Biosci.* 1: 305-325, 1967.
35. Leary R, Jelliffe R, Schumitzky A, and Van Guilder M: A Unified Parametric/Nonparametric Approach to Population PK/PD Modeling. Presented at the Annual Meeting of the Population Approach Group in Europe, Paris, France, June 6-7, 2002.
36. Bayard D and Jelliffe R: Bayesian Estimation of Posterior Densities for Pharmacokinetic Models having Changing Parameter Values. Presented at the Annual Meetings of the Society for Computer Simulation, San Diego CA, January 23-27, 2000. Published in the Proceedings, Health Sciences Simulation, pp. 75-83.
37. D'Argenio D: Optimal Sampling Times for Pharmacokinetic Experiments. *J. Pharmacokin. Biopharmacol.*, 9: 739-756, 1981.