Biostatistics (2001), **2**, 1, *pp.* 47–61
Printed in Great Britain

Easy-to-implement Bayesian methods for dose-escalation studies in healthy volunteers

JOHN WHITEHEAD*, YINGHUI ZHOU
Medical and Pharmaceutical Statistics Research Unit, The University of Reading, UK
Email: J.R.Whitehead@reading.ac.uk

SCOTT PATTERSON
Glaxo SmithKline Pharmaceuticals, Philadelphia, PA, USA

DAWN WEBBER
Glaxo SmithKline Pharmaceuticals, Harlow, UK

STEPHEN FRANCIS
Astra Zeneca Pharmaceuticals, Macclesfield, UK

SUMMARY

In phase I clinical trials, experimental drugs are administered to healthy volunteers in order to establish their safety and to explore the relationship between the dose taken and the concentration found in plasma. Each volunteer receives a series of increasing single doses. In this paper a Bayesian decision procedure is developed for choosing the doses to give in the next round of the study, taking into account both prior information and the responses observed so far. The procedure seeks the optimal doses for learning about the dose–concentration relationship, subject to a constraint which reduces the risk of administering dangerously high doses.

Individual volunteers receive more than one dose, and the pharmacokinetic responses observed are, after logarithmic transformation, treated as approximately normally distributed. Thus data analysis can be achieved by fitting linear mixed models. By expressing prior information as 'pseudo-data', and by maximizing over posterior distributions rather than taking expectations, a procedure which can be implemented using standard mixed model software is derived. Comparisons are made with existing approaches to the conduct of these studies, and the new method is illustrated using real and simulated data.

Keywords: Bayesian decision procedure; Dose-escalation; First-into-man; Pharmacokinetics; Phase I trial.

1. DOSE-ESCALATION STUDIES

The safety and tolerability of all but the most toxic new drugs are investigated in phase I healthy volunteer studies, sometimes known as 'first-into-man studies', prior to experimental administration to patients. Details of the designs of such studies vary, and one commonly used procedure will be described in Section 2. Current practice uses simple statistical methods. Formal procedures of decision theory, Bayesian methodology or even mixed effects modelling have not yet received wide application. Volunteers

*To whom correspondence should be addressed.

are often treated in cohorts, each receiving several different pre-specified doses in ascending order. Based on the results observed at a given time, the investigator decides whether to continue dosing an individual subject, and whether to continue the trial. If all goes well, the trial proceeds until all planned doses have been investigated. References to some of the work which has sought to introduce novel statistical procedures to phase I studies include an application of Bayesian methodology (Racine-Poon and Dubois, 1989), and an account of frequentist design and analysis (Boon and Roes, 1999).

In a first-into-man study, many observations are taken. We focus on pharmacokinetic variables derived from the curve relating the concentration of the drug in plasma to the time since administration (Chow and Liu, 1999). Commonly used summaries such as the area under the curve (AUC) and the maximum concentration ($C_{\max}$) are often modelled by the normal distribution, after a logarithmic transformation (Westlake, 1988). It is desirable to avoid excessive drug concentrations, as reflected by large values of AUC or $C_{\max}$, although knowledge of suitable upper limits is often vague.

In this paper, the study is conducted according to a Bayesian sequential design, which is derived through the application of Bayesian decision theory (Berger, 1985; Lindley, 1971). We maximize the information about the dose–concentration curve subject to a safety constraint limiting the permissible range of doses to be used. As data accumulate, the permissible range will extend to higher doses. Criteria for the amount of information required can be used to determine stopping, although stopping rules are not explored here. The formal procedure described concerns only pharmacokinetic data, although in practice safety information would be an important part of the decision making process.

Prior information concerning the parameters of the dose–response relationship is needed before the procedure can begin. This can be formulated from pre-clinical data and expert opinion, and can be constrained so that the initial dosing of the first cohort proceeds as in conventional designs, and so that escalation does not proceed too quickly or too slowly. Two strategies are adopted to overcome computational difficulties, and thus to make the methods available for immediate implementation using standard software. Prior opinion is expressed, in part, as 'pseudo-data' and posterior modal estimates are used to determine parameter values.

The majority of recent statistical literature on dose-escalation studies has been written in the context of oncology. The drugs under investigation are highly toxic, and so the population studied comprises patients who have a potential to benefit from treatment. It is rare for patients to be treated with more than one dose, and (in the statistical accounts at least) responses are usually binary: toxicity or no toxicity. The best known formal procedure is the continual reassessment method (CRM) (O'Quigley *et al.*, 1990; Chevret, 1993). This can be expressed as a Bayesian decision procedure in which the difference between the dose to be administered and some ideal dose is minimized. O'Quigley and Shen (1996) explore a version of the CRM in which likelihood methods are used to estimate parameters: in the language of this paper, posterior modal estimates are used. Faries (1994) and Goodman *et al.* (1995) criticize the CRM on the grounds that it can recommend large doses too early. They introduce modifications, which play a similar role to the more formal constraint to the permissible range of doses suggested here. A similar constraint to that used here was introduced for trials yielding single binary responses by Babb *et al.* (1998). Bayesian decision procedures which generalize the CRM are described in Whitehead and Williamson (1998), and less formal procedures using posterior distributions for guidance have been described by Gatsonis and Greenhouse (1992). Although the details of dose-escalation in oncology differ from those underlying healthy volunteer studies, the principles of the papers cited above have motivated the approach described in this paper. An alternative approach, in the context of oncology, has been taken by Simon *et al.* (1997).

Several earlier authors have considered dose-finding procedures in which subject responses are quantitative, and assumed to follow a normal distribution. Usually, the objective has been to find the highest dose such that the probability of a response exceeding some pre-defined limit is controlled. Eichhorn and Zacks (1973, 1981) consider Bayesian schemes, although the variances of responses are assumed to be known. Robinson (1978) and Shih (1989) allow this variance to be unknown, but do not

Table 1. *The dosing schedule for the first cohort in a conventional dose-escalation study*

Period	Subject 1	Subject 2	Subject 3	Subject 4
1	d_1	d_1	d_1	Placebo
2	d_2	d_2	placebo	d_1
3	d_3	placebo	d_2	d_2
4	placebo	d_3	d_3	d_3

use a Bayesian approach. We are unaware of references to the application of formal decision theory to studies yielding *repeated* quantitative responses. The International Conference on Harmonization (of drug regulation) guideline on dose–response information (ICH, 1994) makes it clear that experience with designs estimating individual dose–response curves, and their analysis, is limited (section IB), and encourage openness towards novel Bayesian approaches (section IV 6).

In the next section, conventional designs for dose-escalation studies will be described, and straightforward methods of analysis outlined. A Bayesian approach to such an analysis is introduced in Section 3, and in Section 4 it is used in the formulation of a decision-theoretic procedure for choosing doses. Section 5 contains illustrations of the new procedure, and Section 6 is a discussion of the use of the new procedure in practice together with possible extensions. Further illustrations of how the procedure would be implemented in practice and of the computational details are included in Patterson *et al.* (1999).

2. Conventional approaches to the design and analysis of dose-escalation studies

The clinical objectives for dose-escalation studies of a new drug in healthy volunteers are usually the assessment of safety and tolerability following rising single oral doses and the characterization of the pharmacokinetic profile. Dose escalation has to proceed in a way that protects the safety of the volunteers. A design commonly used within SmithKline Beecham for a first-into-man study is a placebo-controlled, dose-rising, four-period crossover study of two to six cohorts of four healthy volunteers. In the first cohort, each volunteer receives three active doses in an ascending order and a placebo dose, inserted in a random position in the sequence, as illustrated in Table 1. The doses are denoted by $d_1, d_2, \ldots$, where $d_1 < d_2 < \cdots$. Administration of the doses is separated by a wash-out period. The placebo is administered to each subject so that any adverse effects following administration of active drug can be put into context, and to introduce an element of blindness.

The pharmacokinetic data for each volunteer are expressed in terms of the concentrations of active drug substance in the plasma against time after administration. The study may be terminated or the dosing regimen altered if volunteers exceed some pre-defined exposure level or if an unacceptable adverse event profile is seen. The maximum exposure level is defined prior to the start of the study based on the toxicity profile for the compound observed in the most sensitive animal species.

If it is deemed safe to continue with the dose escalation, then the next cohort of volunteers will receive three active doses (d_3, d_4, d_5) and placebo arranged in a similar pattern to the first cohort. The lowest dose used in the second cohort is the highest dose used in the previous one. This procedure continues until all planned doses have been administered, or the study is terminated due to an unacceptable adverse event profile or to values of AUC or $C_{\max}$ in excess of the set limits. The starting and top doses in first-into-man studies are usually fixed based on pre-clinical and toxicology data for that drug (Boxembaum and DiLea, 1995).

These studies can take a long time to complete. Each cohort usually takes a month to complete the

Table 2. *Results from a dose-escalation study showing actual doses administered and AUC values recorded*

Cohort	Period	Actual dose, AUC			
		Subject 1	Subject 2	Subject 3	Subject 4
1	1	2, *	2, *	2, *	placebo
	2	5, 4.25	5, 3.93	placebo	2, *
	3	10, 10.52	placebo	5, 6.38	5, 3.98
	4	placebo	10, 8.18	10, 13.20	10, 8.22
		Subject 5	Subject 6	Subject 7	Subject 8
2	1	10, 10.04	10, 8.38	10, 8.29	placebo
	2	25, 24.64	25, 22.87	placebo	10, 12.17
	3	40, 45.35	placebo	25, 20.30	25, 19.07
	4	placebo	40, 35.01	40, 31.57	40, 42.33
		Subject 9	Subject 10	Subject 11	Subject 12
3	1	40, 44.67	40, 46.27	40, 68.20	placebo
	2	60, 56.41	60, 66.57	placebo	40, 49.62
	3	80, 74.96	placebo	50, 78.68	50, 54.45
	4	placebo	80, 73.56	60, 67.72	80, 61.16

dosing schedule, and there may be up to six cohorts of volunteers. Furthermore, information is not always gathered at the most appropriate doses. Often, the majority of the information is gathered at very low doses, with few volunteers receiving the higher doses likely to be used when developing the drug further. There are no formal guidelines for determining dose escalation, which proceeds by informal evaluation of the data observed so far. The results are generally summarized descriptively and graphically and not analysed formally.

Although formal analysis of dose-escalation data is not usually performed, straightforward and appropriate methods exist. These will now be described, and will then be built upon in the construction of Bayesian decision procedures. We consider the natural logarithm of either AUC or $C_{\max}$, which will be denoted by y. The linear mixed model (Laird and Ware, 1982) for the response y_{ij}, for the jth observation on the ith subject, is

$$y_{ij} = \boldsymbol{\theta}_1 + \boldsymbol{\theta}_2 \ell_{ij} + s_i + \varepsilon_{ij}. \tag{2.1}$$

Placebo administrations are ignored in this analysis, as there will be no drug detected in plasma, thus ℓ_{ij} is the logarithm of the jth *active* dose received by the ith subject, and j runs from 1 to 3. The term s_i is the random effect relating to the ith subject. The s_i and ε_{ij} are mutually independent, normally distributed random variables with mean zero, and variances τ^2 and σ^2 respectively. The linear relationship described in (2.1) corresponds to a power relationship on the original scale. While a random slope term might make scientific sense, the data set will generally be too small to warrant its inclusion.

A SmithKline Beecham data set is presented in Table 2 and displayed in Figure 1. Logarithmic transformation of both AUC and dose reveals a linear relationship with roughly constant scatter. Results are shown for the eight doses 2, 5, 10, 25, 40, 50, 60 and 80 mg. For safety reasons, subjects were not to be dosed again if their AUC exceeded 100 μg h mL^{-1}. The trial illustrates the flexible nature of many such exploratory studies. When administration of the 40 mg dose to subjects in cohort 3 gave AUC values sufficiently high to cause concern to the investigators, the 50 mg dose was added to the original dose schedule in order to be more cautious.

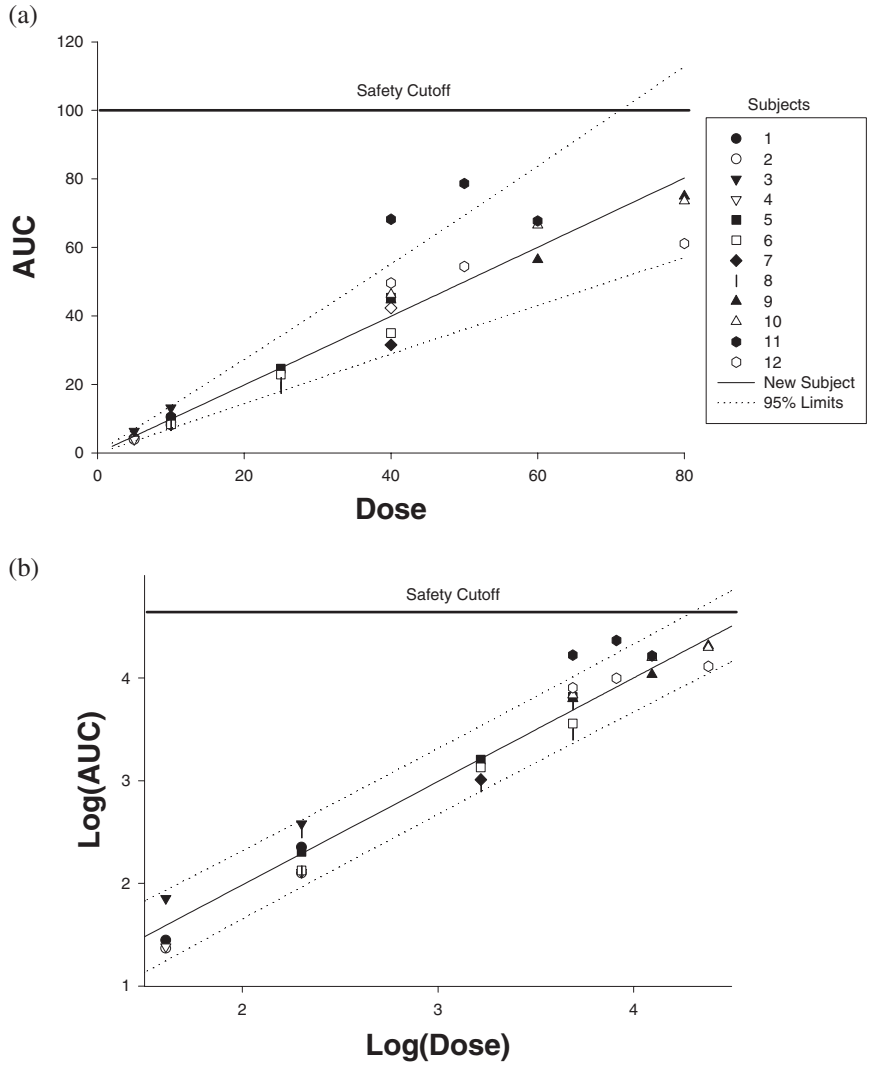

Fig. 1. The relationship between AUC and dose in a dose-escalation study conducted by SmithKline Beecham: (a) on untransformed scales (b) after logarithmic transformation, and showing the fitted model. The dotted lines in (a) and (b) represent 95% confidence limits for a new subject.

In order to illustrate the use of model (2.1), it has been fitted to these data. An immediate problem concerns the approach to non-quantifiable concentrations corresponding to low doses. Here, the results from an initial cohort involving the doses 0.05, 0.5 and 2 mg which yielded no measurable plasma concentrations at all have all been ignored, and further non-quantifiable concentrations corresponding to use of the 2 mg dose in cohort 1 (denoted by asterisks in Table 2) are treated as missing. Model (2.1) is an approximate representation of the dose–response relationship, and the data provide no reassurance that the linearity implied extends to the doses 0.05 and 0.5: our approach is to fit a model that claims validity only

for doses between 2 and 80 mg. If the non-observable concentrations at dose 2 are replaced by fixed values such as 1 μg h mL^{-1}, or by partial observations of the AUC which are available, then these imputed values prove to be very influential and lead to very small estimates of the between-subject variance τ^2 which are inconsistent with other SmithKline Beecham datasets that we have analysed. Multiple imputation methods may provide an alternative form of analysis, but we have not investigated such an approach. From a SAS PROC MIXED analysis (Littell *et al.*, 1996), maximum likelihood estimates of the parameters in model (2.1) are $\hat{\boldsymbol{\theta}}_1 = -0.031$, $\hat{\boldsymbol{\theta}}_2 = 1.008$, $\hat{\sigma}^2 = 0.022$ and $\hat{\tau}^2 = 0.021$. The fitted relationship is shown in Figure 1.

3. A BAYESIAN ANALYSIS OF DOSE-ESCALATION DATA

The objective of this paper is to present designs for first-into-man studies derived as Bayesian decision procedures. As a prerequisite, this section is devoted to the Bayesian analysis of the resulting data, which could be performed at any stage during the study.

Model (2.1) of Section 2 will be adopted, and the parameters $\boldsymbol{\theta}$ and ν will be treated as random, where $\boldsymbol{\theta}$ is the vector made up of the slope and intercept; $\boldsymbol{\theta}' = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$, and ν is the within subject precision, $\nu = \sigma^{-2}$. In order to avoid complication, the correlation ρ between two responses on the same subject will be treated as fixed and known. In Bayesian terms, this amounts to the imposition of a prior with zero variance. A value of $\rho = 0.6$ was found to be consistent with most of the data sets studied, and has been imposed throughout this work. The implications of this choice, and of the principle of fixing ρ, are discussed in Section 6. As $\rho = \tau^2/(\sigma^2 + \tau^2)$, this strategy avoids the need to seek distributions to describe the variation of τ^2. A normal-gamma prior distribution is used to express prior information on $\boldsymbol{\theta}$ and ν, as this is easy to fit and is also conjugate (Bernardo and Smith, 1994, section 5.2).

Let $\mathbf{y}$ denote the vector of responses with elements $y_{11}, y_{12}, y_{13}, y_{21}, y_{22}, y_{23}, \ldots$ and $\mathbf{s}$ denote the vector of subject effects with elements $s_1, s_2, \ldots$. The design matrix $\mathbf{X}$ has rows $(1, \ell_{11}), (1, \ell_{12}), (1, \ell_{13}), (1, \ell_{21}), (1, \ell_{22}), (1, \ell_{23}), \ldots$, and a matrix $\mathbf{U}$ is defined as having a 1 in the ith column of rows $(3i - 2)$, $(3i - 1)$ and $3i$, $i = 1, 2, \ldots$, and zeros elsewhere. The identity matrix is denoted by $\mathbf{I}$. In matrix form, the Bayesian model has the following hierarchical structure:

$$\begin{aligned} \mathbf{y}|\mathbf{s}, \boldsymbol{\theta}, \nu &\sim \mathrm{N}(\mathbf{X}\boldsymbol{\theta} + \mathbf{U}\mathbf{s}, \nu^{-1}\mathbf{I}), \\ \mathbf{s}|\boldsymbol{\theta}, \nu &\sim \mathrm{N}\left(\mathbf{0}, \frac{\rho}{\nu(1-\rho)}\mathbf{I}\right), \\ \boldsymbol{\theta}|\nu &\sim \mathrm{N}(\mathbf{m}, (\nu\mathbf{Q})^{-1}), \\ \nu &\sim Ga(\alpha, \beta), \end{aligned} \tag{3.1}$$

where N denotes a normal distribution, Ga a gamma distribution, and the values of $\mathbf{m}$, $\mathbf{Q}$, α and β are chosen to represent prior knowledge.

Expert opinion is used to derive an informative prior distribution. A technical reason for doing this rather than relying on non-informative priors is to obtain a proper Bayesian predictive distribution, before the study has commenced. Less formally, in the absence of any prior information, it would be impossible to identify a safe starting dose. Usually clinical pharmacologists will have relevant data from animal experiments and other investigations to draw on in formulating a prior distribution.

A prior distribution for $\boldsymbol{\theta}$, conditional on ν, will be formed by combining a non-informative prior with some imaginary responses $\mathbf{y}_0$ of volunteers at dose levels corresponding to a design matrix $\mathbf{X}_0$. The matrix $\mathbf{U}_0$ contains the pseudo subject random effects. Combining the first two elements of model (3.1) gives

$$\mathbf{y}_0|\boldsymbol{\theta}, \nu \sim \mathrm{N}\left(\mathbf{X}_0\boldsymbol{\theta}, \nu^{-1}\left(\mathbf{I} + \frac{\rho}{1-\rho}\mathbf{U}_0\mathbf{U}_0^T\right)\right), \tag{3.2}$$

as $\tau^2 = \nu^{-1}\rho/(1-\rho)$. The posterior distribution formed by combining these pseudo-data with a non-informative prior for $\boldsymbol{\theta}$ given ν is used as a prior. This prior has density proportional to the density of $\mathbf{y}_0$ given $\boldsymbol{\theta}$ and ν expressed in (3.2). It follows that

$$\boldsymbol{\theta}|\mathbf{y}_0, \nu \sim \mathrm{N}((\mathbf{X}_0^T\mathbf{P}_0\mathbf{X}_0)^{-1}\mathbf{X}_0^T\mathbf{P}_0\mathbf{y}_0, (\nu\mathbf{X}_0^T\mathbf{P}_0\mathbf{X}_0)^{-1}), \tag{3.3}$$

where $\mathbf{P}_0 = [\mathbf{I} + \{\rho/(1-\rho)\}\mathbf{U}_0\mathbf{U}_0^T]^{-1}$. This provides values to be used for $\mathbf{m}$ and $\mathbf{Q}$ in the third part of model (3.1).

Note that the mean of this distribution is equal to the usual maximum likelihood estimates based on $\mathbf{y}_0$ for intercept and slope; the variance is just the variance of these estimates. In the simplest case, assumed here, imagine that the pseudo-data consist of just two representative responses on a single subject: y_a corresponding to log-dose a, and y_b corresponding to log-dose b. Then

$$\mathbf{P}_0 = \frac{1}{1+\rho}\begin{pmatrix} 1 & -\rho \\ -\rho & 1 \end{pmatrix}$$

and it follows that

$$\boldsymbol{\theta}|\mathbf{y}_0, \nu \sim \mathrm{N}\left(\frac{1}{(b-a)}\begin{pmatrix} by_a - ay_b \\ -(y_a - y_b)\end{pmatrix}, \frac{1}{\nu(b-a)^2}\begin{pmatrix} \dfrac{a^2+b^2-2\rho ab}{1-\rho} & -(a+b) \\ -(a+b) & 2 \end{pmatrix}\right). \tag{3.4}$$

The parameters α and β of the gamma prior for ν are more difficult to choose. However, as will be seen in Section 4, constraints may be imposed on the prior distributions which enable suitable values for α and β to be determined from considerations that are familiar to clinical investigators.

An advantage of the use of pseudo-data to define prior distributions is that posterior distributions of a relatively simple form are obtained. First consider how study data on the ith subject change our belief about that person's subject effect s_i (conditional on $\boldsymbol{\theta}$ and ν). Suppose that n_i observations have already been made on this subject. Let $\bar{y}_i$ be the mean response observed and $\bar{\ell}_i$ the mean log-dose administered. Then it can be shown that the mean of the distribution of s_i is shifted from zero and its variance is decreased:

$$s_i|\boldsymbol{\theta}, \nu, \bar{y}_i \sim \mathrm{N}\left(w_i\,(\bar{y}_i - \boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\bar{\ell}_i), \frac{w_i}{\nu n_i}\right) \tag{3.5}$$

where

$$w_i = \frac{n_i\rho}{1+(n_i-1)\rho}.$$

Given real data on several subjects, each of the subject effects has, independently, a distribution of the form given in (3.5). The posterior distribution for $\boldsymbol{\theta}$ conditional on ν is given by (3.3), with $\mathbf{y}_0$, $\mathbf{X}_0$, $\mathbf{U}_0$ and $\mathbf{P}_0$ replaced by $\mathbf{y}_*$, $\mathbf{X}_*$, $\mathbf{U}_*$, and $\mathbf{P}_*$ representing a combination of the pseudo-data with the real data.

The posterior for ν is well known to be ν is well known to be

$$\nu \sim Ga\left(\alpha + \frac{n}{2}, \beta + \frac{\mathbf{y}_*^T\mathbf{P}_*\mathbf{y}_* - \mathbf{y}_*^T\mathbf{P}_*\mathbf{X}_*(\mathbf{X}_*^T\mathbf{P}_*\mathbf{X}_*)^{-1}\mathbf{X}_*^T\mathbf{P}_*\mathbf{y}_*}{2}\right), \tag{3.6}$$

where n is the number of *real* observations. The second term in the expression for the second parameter of the gamma distribution in (3.6) reduces to zero when based only on the prior pseudo-data.

The posterior modal estimates $\hat{\boldsymbol{\theta}}$ and $\hat{\sigma}^2$ for $\boldsymbol{\theta}$ and σ are given by

$$\hat{\boldsymbol{\theta}} = (\mathbf{X}_*^T\mathbf{P}_*\mathbf{X}_*)^{-1}\mathbf{X}_*^T\mathbf{P}_*\mathbf{y}_*$$

and

$$\hat{\sigma}^2 = \frac{2\beta + \mathbf{y}_*^T \mathbf{P}_* \mathbf{y}_* - \mathbf{y}_*^T P_* \mathbf{X}_* (\mathbf{X}_*^T P_* \mathbf{X}_*)^{-1} \mathbf{X}_*^T \mathbf{P}_* \mathbf{y}_*}{2\alpha + n}.$$

These estimates are related to maximum likelihood estimates $\boldsymbol{\theta}_*$ and σ_*^2 obtained from an analysis of pseudo- and real data combined according to $\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}_*$ and $\hat{\sigma}^2 = \{2\beta + (n+2)\sigma_*^2\}/(2\alpha + n)$, allowing their calculation from any standard mixed model software in which $\rho = \tau^2/(\sigma^2 + \tau^2)$ can be fixed.

4. CONSTRAINED OPTIMAL DOSE-ESCALATION

In this section, Bayesian decision theory is applied to derive a dose-escalation scheme which is optimal, in the sense of maximizing some gain function within certain safety constraints. As in the conventional design, it is assumed that cohorts of four volunteers are observed, each volunteer receiving three active doses and a placebo. The scheduling of the placebo dose is predetermined, with the $(4c - j + 1)$th subject, who is in the cth cohort, receiving placebo in the jth period $c = 1, 2, \ldots$; $j = 1, \ldots, 4$ (as in Table 1 for $c = 1$). Doses are chosen from amongst a predetermined set $d_1 < \cdots < d_k$.

It is assumed that the doses given in period j for the cth cohort are administered simultaneously. The decision problem consists of using the prior information, and the real data from previous cohorts and from previous periods within the current cohort, to determine the three doses to be administered during the next period. The general principles of such Bayesian decision procedures have been discussed by Whitehead (1997).

In this paper, two optimality criteria will be explored, both being operated within the same safety constraint. The first option, which will be referred to as *maxsafe*, simply treats each subject at the highest dose permitted by the safety constraint. This will have the effect of gathering information efficiently about both the maximum dose which can be safely administered and about the response to that dose, and it is similar to the strategy used in Eichhorn and Zacks (1973). However, this criterion has not formally been expressed as a gain function. The second criterion, which will be referred to as *optsafe*, is D-optimality and corresponds to an objective of estimating the parameters $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ as precisely as possible. The optsafe criterion will now be studied in some detail, and then the safety constraint will be explored. The practical choice of an optimality criterion is discussed further in Section 6.

For D-optimality, the optimal choice of doses is that which minimises the determinant of the variance–covariance matrix of their joint posterior distribution based on pseudo- and real data, or equivalently maximizes the determinant of $\nu\ \mathbf{X}_*^T\ \mathbf{P}_* \mathbf{X}_*$. Hence the gain function can be identified with det $(\nu\ \mathbf{X}_*^T\ \mathbf{P}_* \mathbf{X}_*)$. Now

$$\begin{aligned}
\mathbf{X}_*^T \mathbf{P}_* \mathbf{X}_* &= \mathbf{X}_*^T \mathbf{X}_* - \sum_i \begin{pmatrix} w_i n_i & w_i n_i \bar{\ell}_i \\ w_i n_i \bar{\ell}_i & w_i n_i \bar{\ell}_i^2 \end{pmatrix} \\
&= \begin{pmatrix} \dfrac{1-\rho}{\rho} \displaystyle\sum_i w_i & \dfrac{1-\rho}{\rho} \displaystyle\sum_i w_i \bar{\ell}_i \\ \dfrac{1-\rho}{\rho} \displaystyle\sum_i w_i \bar{\ell}_i & \displaystyle\sum_i \sum_j (\ell_{ij} - \bar{\ell}_i)^2 + \dfrac{1-\rho}{\rho} \displaystyle\sum_i w_i \bar{\ell}_i^2 \end{pmatrix}
\end{aligned}$$

and

$$\det(\mathbf{X}_*^T\mathbf{P}_*\mathbf{X}_*)$$
$$= \left(\frac{1-\rho}{\rho}\right)^2 \left(\sum_i w_i\right) \sum_i w_i \left(\bar{\ell}_i - \frac{\sum_i w_i \bar{\ell}_i}{\sum_i w_i}\right)^2 + \frac{1-\rho}{\rho}\left(\sum_i w_i\right) \sum_i \sum_j (\ell_{ij} - \bar{\ell}_i)^2, \qquad (4.1)$$

where all sums are over pseudo- and real observations.

This determinant can be seen to depend on two components: a weighted between-subject sum of squares of log-doses, and a within-subject sum of squares of log-doses. Thus, it is desirable to have both contrasting patterns of doses between subjects and a wide spread of doses within each individual subject. If D-optimality is used without a safety constraint, then one or more subjects will typically be allocated very high doses in order to achieve the latter contrast, even at relatively early stages of a trial. This is unlikely to be acceptable in practice, where dose escalation has usually proceeded in a gradual manner. Caution is needed to prevent subjects from being exposed to excessive dose concentrations. When less information is available, greater caution is necessary. The strategy adopted here is to operate the sequential D-optimal design within the confines of a safety constraint.

Suppose a limiting level L of response (log AUC or log $C_{\max}$) has been specified: larger values than this are considered undesirable and should be avoided. In terms of this 'safety cutoff', a reasonable criterion for using a candidate dose d_f for the ith subject can be based on a Bayesian predictive probability:

$$\Pr(y_{if} > L) \leqslant c_0 \qquad (4.2)$$

where y_{if} is the future response corresponding to dose d_f, and the probability accounts for the uncertainty of the unknown parameters inherent in the current posterior distributions. If the dose d_f satisfies (4.2), then it can be administered. The dose which gives equality in (4.2) will be denoted by d_f*, and is considered to be the maximum safe dose. Now the distribution of y_{if}, given that the subject effect s_i, $\boldsymbol{\theta}$ and ν are known, is

$$y_{if}|s_i, \boldsymbol{\theta}, \nu \sim \mathrm{N}(\boldsymbol{\theta}_1 + \boldsymbol{\theta}_2\ell_f + s_i, \nu^{-1})$$

where $\ell_f = \log d_f$.

Successively integrating out the posterior densities of s_i given $\boldsymbol{\theta}$ and ν, and $\boldsymbol{\theta}$ given ν, we obtain

$$y_{if}|\nu \sim \mathrm{N}(\hat{\boldsymbol{\theta}}_1 + \hat{\boldsymbol{\theta}}_2\ell_f + \hat{s}_i, \nu^{-1}(R_i + V_{if}))$$

where

$$\begin{aligned} R_i &= 1 + w_i/n_i, \\ V_{if} &= (1 - w_i, \ell_f - w_i\bar{\ell}_i)(X_*^T P_* X_*)^{-1}(1 - w_i, \ell_f - w_i\bar{\ell}_i)^T \\ &= \nu \operatorname{var}(\boldsymbol{\theta}_1 + \boldsymbol{\theta}_2\ell_f + w_i(\bar{y}_i - \boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\bar{\ell}_i)|\nu), \end{aligned}$$

$\bar{y}_i$ and ℓ_i are the means of the responses and log-doses already observed on the ith subject, and $\hat{s}_i = w_i(\bar{y}_i - \hat{\boldsymbol{\theta}}_1 - \hat{\boldsymbol{\theta}}_2\bar{\ell}_i)$ is the predicted value of the random subject effect. Further integrating out the gamma distribution for ν leads to the predictive distribution for forecasting a future observation. This is a t-distribution (parametrized in terms of location, precision and degrees of freedom):

$$y_{if} \sim t(\hat{\boldsymbol{\theta}}_1 + \hat{\boldsymbol{\theta}}_2\ell_f + \hat{s}_i, \{(R_i + V_{if})\hat{\sigma}^2\}^{-1}, n + 2\alpha).$$

The predictive distribution provides the safety criterion (4.2) which a dose d_f must satisfy before it can be considered suitable to use next. Then, for those combinations of doses for the next period that are allowable, the one which maximizes $\det(\mathbf{X}_*^T\mathbf{P}_*\mathbf{X}_*)$ will be the optsafe selection. Often, but not always, this will be a combination of the highest doses that are considered safe.

As mentioned in Section 3, the safety constraint can be used to facilitate the choice of α and β. Suppose that it is only just safe to start dosing (on new, previously untested subjects) at dose d_1, and that it is only just unsafe to use dose d_2:

$$\Pr(y_{11} > L) = c_1 \qquad \text{and} \qquad \Pr(y_{12} > L) = c_2, \tag{4.3}$$

where $c_1 \leqslant c_0 < c_2$.

These two requirements may be solved to obtain values for α and β. In examples, we have set c_0 and c_1 equal to 0.05, and found that the value 0.067 for c_2 leads to sensible escalation schemes. The choice of these constraints is discussed further in Section 6.

The study may be conducted until a fixed number of cohorts has been observed, in which case its main advantage will be the concentration of experimentation at higher, more interesting, doses. Alternatively, the study might proceed until the highest dose d_k is reached, and this might happen sooner using the proposed scheme. Stopping rules could be devised based on a required value of $\det(\mathbf{X}_*^T\mathbf{P}_*\mathbf{X}_*)$ or by tracking d_f^* until its rate of increase slows down.

5. ILLUSTRATIONS OF THE PROCEDURE

This section presents two illustrations based on the results from the SmithKline Beecham study presented in Table 2 and discussed in Section 2. These are based on the maxsafe and optsafe criteria respectively: in both, a complete run of the new procedure is simulated in order to demonstrate its advantages.

In the simulations, the doses 0.05 and 0.5 which led to non-quantifiable plasma concentrations in the actual study have not been included, but the dose of 50 mg which was introduced during the study has. For these illustrations, a prior distribution has been retrospectively selected, based on values of AUC of 5 μg h mL^{-1} and 60 μg h mL^{-1} for the two doses, 5 mg and 60 mg respectively. The limiting AUC value of 100 μg h mL^{-1}, consistent with the actual stopping criteria described above, was used so that $L = \log 100$, and the value of c_0 in constraint (4.2) was set at 0.05. It was assumed that the risk of the response to the lowest dose (2 mg) exceeding the limiting AUC value should also be 0.05, and that the risk for second lowest dose (5 mg) should be 0.067, giving $c_1 = 0.05$ and $c_2 = 0.067$ in (4.3).

Tables 3 and 4 show the progress of simulated studies, based on maxsafe and optsafe respectively, and Figure 2 shows some of the prior and posterior densities associated with Table 3. Responses were simulated according to the model which was fitted to the actual study (the maximum likelihood estimates were presented at the end of Section 2 and are given in the first row of Table 5). In the model used to simulate the data, the value of ρ was chosen to be 0.488, consistently with the fitted model, although the escalation method adopts the universally imposed value of 0.6 as in practice ρ would be unknown to the investigators. Three responses were simulated for each subject at each of the available doses and then, as the run progressed, the relevant values were input as responses for the purposes of the dose-escalation procedure. The same full potential data set was used in the runs leading to Tables 3 and 4 so that when a subject receives the same dose in the same period both tables record the same response: this approach was adopted in order to enhance comparability.

The results illustrate how, after giving the first subjects in the first period the bottom dose, some doses are then skipped, thereby achieving the desired acceleration of dose-escalation. The top dose of 80 mg is never used, but 50 and 60 mg are demanded frequently. The optsafe criterion dictates that the bottom

Table 3. *A simulated dose-escalation study based on the maxsafe criterion*

Cohort	Period	Dose, AUC				d_f^*
		Subject 1	Subject 2	Subject 3	Subject 4	
1	1	2, 1.57	2, 1.92	2, 2.07	placebo	25.52
	2	25, 20.95	25, 23.53	placebo	25, 22.09	42.48
	3	50, 39.87	placebo	40, 51.59	50, 63.52	44.44
	4	placebo	50, 47.97	40, 51.68	50, 48.54	49.11
		Subject 5	Subject 6	Subject 7	Subject 8	
2	1	40, 38.45	40, 48.68	40, 53.34	placebo	50.36
	2	60, 62.56	50, 56.58	placebo	50, 58.78	52.66
	3	60, 67.30	placebo	50, 57.11	50, 59.46	54.57
	4	placebo	50, 50.44	50, 55.43	50, 51.38	56.70
		Subject 9	Subject 10	Subject 11	Subject 12	
3	1	50, 58.40	50, 54.64	50, 46.74	placebo	58.38
	2	60, 74.39	60, 55.05	placebo	50, 59.27	59.21
	3	60, 80.96	placebo	60, 71.54	60, 69.07	59.26
	4	placebo	60, 65.45	60, 61.31	60, 57.91	60.25

Table 4. *A simulated dose-escalation study based on the optsafe criterion*

Cohort	Period	Dose, AUC				d_f^*
		Subject 1	Subject 2	Subject 3	Subject 4	
1	1	2, 1.57	2, 1.92	2, 2.07	placebo	25.52
	2	25, 20.95	25, 23.53	placebo	25, 22.09	42.48
	3	50, 39.87	placebo	40, 51.59	2, 2.48	46.08
	4	placebo	50, 47.97	50, 64.71	50, 48.54	50.40
		Subject 5	Subject 6	Subject 7	Subject 8	
2	1	50, 48.14	50, 60.95	50, 66.79	placebo	51.17
	2	2, 2.03	2, 2.21	placebo	2, 2.29	53.57
	3	60, 67.30	placebo	2, 2.23	50, 59.46	55.00
	4	placebo	2, 1.97	2, 2.16	2, 2.01	56.59
		Subject 9	Subject 10	Subject 11	Subject 12	
3	1	50, 58.40	50, 54.64	50, 46.74	placebo	58.36
	2	2, 2.42	2, 1.79	placebo	2, 2.31	59.00
	3	60, 80.96	placebo	2, 2.32	50, 57.47	59.31
	4	placebo	60, 65.45	60, 61.31	60, 57.91	60.58

dose will regularly be required (even when most of the available doses are believed safe), thus breaking the current convention of non-decreasing dosing.

Now, according to the model being used for simulation, a future response on a *new* subject i has distribution $y_{if} \sim \mathrm{N}(-0.031 + 1.008\ell_f, 0.043)$. It is easy to deduce from (4.2) that the true 'maximally tolerated' dose d_f^* is 71. Thus the avoidance in the simulated runs of the 80 mg dose is appropriate. The estimate of d_f^* based on the simulated data rises to a value of about 60 in each of the tables, and would be expected to continue to rise slowly towards 71 if new cohorts were to be treated.

Table 5 shows, in addition to the model used for the simulation, maximum likelihood estimates

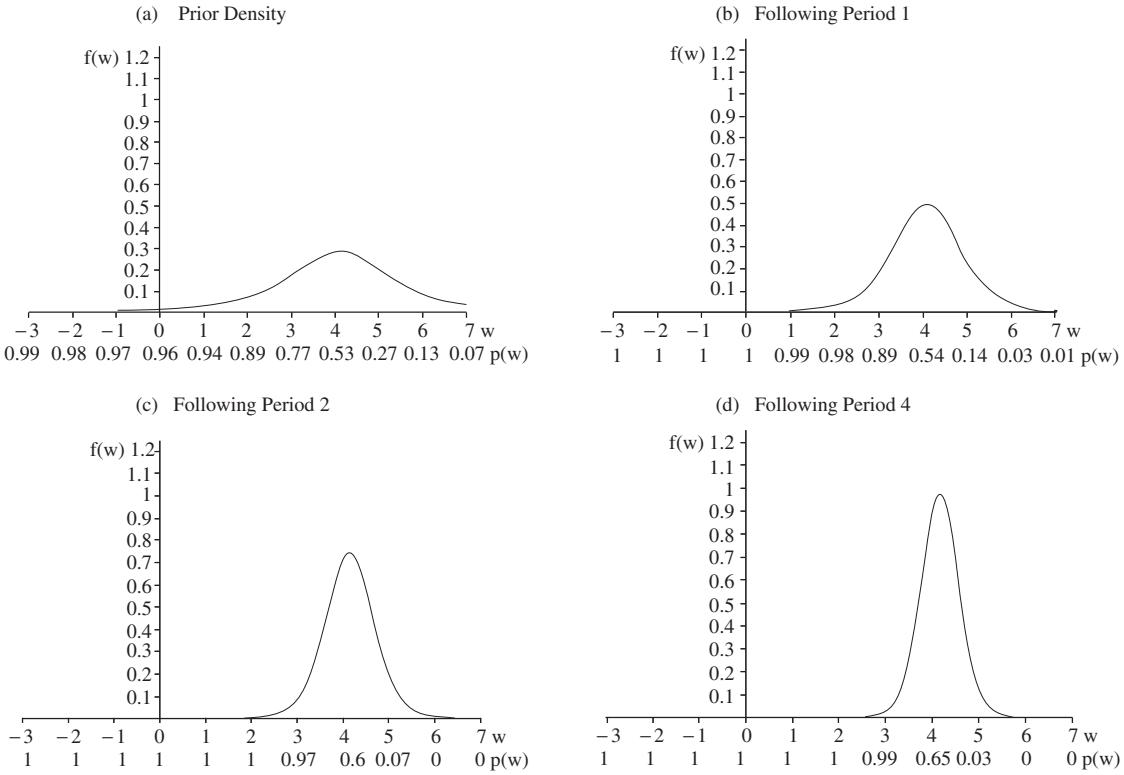

Fig. 2. Prior and posterior densities for log(AUC) for the simulated example of Table 3. The figures relate to the log(AUC) of a new, untested subject, administered a dose of 60 mg. The log(AUC) is denoted by w, and $p(w)$ denotes the subjective probability that log(AUC) exceeds w. Normal approximations to the true densities are shown for the situations (a) prior to observation of data, (b) after period 1 of cohort 1, (c) after period 2 of cohort 1 and (d) after period 4 of cohort 1.

Table 5. *Maximum likelihood estimates (standard errors)*

	θ_1		θ_2		σ^2	τ^2	d_f^*
Model fitted to data in Table 2 (and used for simulation)	−0.031	(0.139)	1.008	(0.043)	0.022	0.021	71
Model fitted to data in Table 3 (maxsafe)	−0.093	(0.082)	1.044	(0.021)	0.010	0.007	73
Model fitted to data in Table 4 (optsafe)	0.033	(0.043)	1.008	(0.012)	0.010	0.009	75

resulting from fitting model (2.1) to the datasets shown in Tables 3 and 4. The optsafe criterion has led to smaller standard errors for the estimates of θ_1 and θ_2, in accordance with its objective. Both σ^2 and τ^2 happen to be underestimated in the fitted models: recall that these are analyses of just two small and closely related datasets, so that they should be expected to be similar and not to be particularly precise.

6. Discussion

A computer program which performs the calculations described here is available from the corresponding author (JW). Within SmithKline Beecham this program is being used to allow investigators who conduct first-into-man studies to become familiar with the new method, and to influence its further development.

The purpose of the procedure described in this paper is to recommend the next round of doses to be given to the current cohort of volunteers. In practice, the decision as to whether to accept these recommendations will lie with the investigating clinician. Recommendations might be overruled because of safety data, or other outcomes not modelled by the procedure. Alternatively, the clinician's experience may indicate that different doses are appropriate. An advantage of the procedure is that it can learn from the results of administering any doses. As users become familiar with the procedure it is expected that it will be overruled less frequently, and that persistent discrepancies between recommendations and practice will lead to a fine tuning of the options used in the procedure.

The procedure described escalates doses most quickly early on, and then slows down. When the constraints c_1 and c_2 in (4.3) are set close to one another, early escalation is indeed fast; when they are further apart it becomes very slow in later cohorts. The choice of c_1 and c_2 can be aided by performing simulated runs in order to observe the likely rate of escalation and its deceleration. This use of frequentist properties as an aid to the choice of a Bayesian prior is in the spirit of the 'stylized Bayesian' approach described by Fisher (1996). Relative to the conventional procedure, simulated runs (not reported here) have indicated that appropriate settings of c_1 and c_2 lead to faster escalation early on, and slower escalation later. There do not appear to be settings which reproduce the conventional dosing patterns completely. This would seem to indicate that the conventional procedure begins cautiously but takes too much reassurance from early low AUC values, becoming rather cavalier in later choices of dose. Furthermore, the conventional procedure never returns to low doses in order to improve estimation. It might be possible to reproduce conventional dose-escalation patterns via a progressive raising of the permitted limit L, and thus to identify the conditions under which it is optimal.

The simulation results of Section 5, and further simulations not presented here, indicate that the optsafe procedure, by increasing the frequency of observation at the lowest dose, improves the precision of estimation of both slope and intercept. It must be stressed that while this overall result is certainly true when the underlying model is valid, it is unlikely to be robust to departures from linearity. In practice, the calls from the optsafe criterion for administration of a low dose might be tempered by realism concerning the assumption of linearity, leading to choice of a compromise dose which is low but in the region of therapeutic interest rather than being the lowest available. It is apparent that the estimation d_f^* itself is achieved with similar accuracy using the two criteria.

Throughout this paper, the simplifying assumption that $\rho = 0.6$ has been made. Of course other fixed values could be imposed. Ideally, ρ should be treated as unknown in the same way as all of the other model parameters. To do so would introduce major complications, both in eliciting prior opinion concerning its value and in fitting the ensuing models. In principle, it would also be ideal to allow for uncertainty in the assumptions of linearity and of normality and for the often arbitrary selection of a limiting level L of log AUC: proceeding too far down this road leads to an expression of such uncertainty that escalation can hardly take place!

The principles underlying the procedure are open to flexible interpretation, and so various components can easily be varied. The modelling of the responses y_{ij} described in equation (2.1) could be amended in various ways. Instead of the subject effect being expressed as a random intercept, it could be included as a random slope: indeed both random terms could be included. However, we note that in our own fitting of models to pharmacokinetic data, random slopes have not been found to improve the model significantly. Covariates relating to the volunteers can be easily introduced, as could a factor relating to period effects.

Complex procedures, perhaps based on multiple imputation or on exact treatment of censored values, could be used to overcome the problem of non-observable pharmacokinetic responses. However, the desire to allow for the many potential influences on outcome must be tempered by the realization that the data available are very few, especially in the early cohorts. It will seldom be desirable to use models which are much more sophisticated than (2.1).

Other simple modifications include the use of criteria other than maxsafe or optsafe. In particular, the optimal estimation of the upper part of the dose–response relationship might be an improvement on the more general optsafe criterion. The procedure as a whole could be applied with other pharmacokinetic endpoints (e.g. maximal concentration, half-life) and/or pharmacodynamic responses, such as blood pressure or pulse, serving as the y_{ij}. For pharmacodynamic data, the placebo doses also have corresponding responses to be taken into account. Generalization to multivariate responses would be more difficult.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the input, advice, support, and review of Mick Ireson, Névine Zariffa, Darryl Downing, Kate Howland, Kerry Pocock and Muhammad Jalaluddin of Biostatistics and Data Sciences, SmithKline Beecham Pharmaceuticals.

REFERENCES

BABB, J., ROGATKO, A. AND ZACKS, S. (1998). Cancer phase I clinical trials: efficient dose escalation with overdose control. *Statistics in Medicine* **17**, 1103–1120.

BERGER, J. O. (1985). *Statistical Decision Theory and Bayesian Analysis*. New York: Springer.

BERNARDO, J. M. AND SMITH, A. F. M. (1994). *Bayesian Theory*. Chichester: Wiley.

BOON, P. C. AND ROES, K. C. B. (1999). Design and analysis issues for crossover designs in phase I clinical studies. *Journal of Biopharmaceutical Statistics* **9**, 109–128.

BOXEMBAUM, H. AND DILEA, C. (1995). First time in human dose selection: allometric thoughts and perspectives. *Journal of Clinical Pharmacology* **35**, 957–966.

CHEVRET, S. (1993). The continual reassessment method in cancer phase I clinical trials: a simulation study. *Statistics in Medicine* **12**, 1093–1108.

CHOW, S. C. AND LIU, J. P. (1999). *Design and Analysis of Bioavailability and Bioequivalence Studies*. Amsterdam: Dekker.

EICHHORN, B. H. AND ZACKS, S. (1973). Sequential search of an optimal dosage, 1. *J. Am. Statist. Assoc.* **68**, 594–598.

EICHHORN, B. H. AND ZACKS, S. (1981). Bayes sequential search of optimal dosage: linear regression with parameters unknown. *Comm. Statist. - Theory and Methods* A**10**, 931–953.

FARIES, D. (1994). Practical modifications of the continual reassessment method for phase I cancer clinical trials. *Journal of Biopharmaceutical Statistics* **4**, 147–164.

FISHER, L. D. (1996). Comments on Bayesian and frequentist analysis and interpretation of clinical trials. *Controlled Clinical Trials* **17**, 423–434.

GATSONIS, C. AND GREENHOUSE, J. B. (1992). Bayesian methods for phase I clinical trials. *Statistics in Medicine* **11**, 1377–1389.

GOODMAN, S. N., ZAHURAK, M. L. AND PIANTADOSI, S. (1995). Some practical improvements in the continual reassessment method for phase I studies. *Statistics in Medicine* **14**, 1149–1161.

ICH (1994). Dose-response information to support drug registration. *Topic E4* ICH Technical Coordination. EMEA, London.

LAIRD, M. N. AND WARE, J. H. (1982). Random effects models for longitudinal data. *Biometrics* **38**, 963–974.

LINDLEY, D. V. (1971). *Making Decisions*. London: Wiley.

LITTELL, R. C., MILLIKEN, G. A., STROUP, W. W. AND WOLFINGER, R. D. (1996). *SAS System for Mixed Models*. Cary, NC: SAS Institute Inc., pp. 253–266.

O'QUIGLEY, J., PEPE, M. AND FISHER, L. (1990). Continual reassessment method: a practical design for phase I clinical trials in cancer. *Biometrics* **46**, 33–48.

O'QUIGLEY, J. AND SHEN, L. Z. (1996). Continual reassessment method: a likelihood approach. *Biometrics* **52**, 673–684.

PATTERSON, S., FRANCIS, S., IRESON, M., WEBBER, D. AND WHITEHEAD, J. (1999). A novel Bayesian decision procedure for early phase dose finding studies. *Journal of Biopharmaceutical Statistics* **9**, 583–597.

RACINE-POON, A. AND DUBOIS, J. P. (1989). Predicting the range of plasma carbemazepine concentrations in patients with epilepsy. *Statistics in Medicine* **8**, 1327–1337.

ROBINSON, J. A. (1978). Sequential choice of an optimal dose: a prediction intervals approach. *Biometrika* **68**, 75–78.

SHIH, W. J. (1989). Prediction approaches to sequentially searching for an optimal dose. *Biometrics* **45**, 623–628.

SIMON, R., FREIDLIN, B., RUBINSTEIN, L., ARBUCK, S. G., COLLINS, J. AND CHRISTIAN, M. C. (1997). Accelerated titration designs for phase I clinical trials in oncology. *Journal of the National Cancer Institute* **89**, 1138–1147.

WESTLAKE, W. J. (1988). Bioavailability and bioequivalence of pharmaceutical formulations. In Peace, K. E. (ed.), *Biopharmaceutical Statistics for Drug Development*, New York: Dekker, pp. 329–352.

WHITEHEAD, J. (1997). Bayesian decision procedures with application to dose finding studies. *Int. J. Pharm. Med.* **11**, 201–208.

WHITEHEAD, J. AND WILLIAMSON, D. (1998). An evaluation of Bayesian decision procedures for dose-finding studies. *Journal of Biopharmaceutical Statistics* **8**, 445–467.

[*Received August 16, 2000*; *first revision January 31, 2000*; *second revision August 31, 2000; accepted for publication September 1, 2000*]