

Pharmacokinetically guided optimum adaptive dose selection in early phase clinical trials

M. Iftakhar Alam^{a,*}, Barbara Bogacka^b, D. Stephen Coad^b

^a*Institute of Statistical Research and Training, University of Dhaka, Dhaka-1000, Bangladesh*

^b*School of Mathematical Sciences, Queen Mary University of London, London E1 4NS, U.K.*

Abstract

A new statistical method is introduced for dose finding in phase IB/IIA trials, which, along with efficacy and toxicity as endpoints, also considers pharmacokinetic information in the dose-selection procedure. Following the assignment of a current best dose to a cohort of patients, the concentration of a drug in the blood is measured at the locally D -optimal time points. The dose-response outcomes are also observed for each patient. Based on the updated information, a new dose is selected for the next cohort so that the estimated probability of efficacy is maximum, subject to the condition that the estimated probability of toxicity is not more than a chosen constant. Another condition for the dose selection is related to the total exposure of the drug in the body, expressed by the area under the concentration curve over time, so that the curative purpose is likely to be achieved in the population without overdosing. Simultaneously to the maximisation of the estimated probability of efficacy, the mean area under the concentration curve for a chosen dose is not allowed to be more than a target value taking into account its inter-patient variability. The purpose is to investigate the gain in efficiency of using pharmacokinetic measures in the dose escalation. The proposed method is found to identify the optimal dose accurately without exposing many patients to toxic doses.

Keywords: Area under the concentration curve, Continuation ratio model, D -optimum sampling times, One-compartment pharmacokinetic model, Population pharmacokinetics.

1. Introduction

Interest has grown in recent years in the development of dose-finding methods incorporating both toxicity and efficacy as endpoints. The idea in these methods is to find a dose for further development which is both safe and efficacious. Thall and Russell (1998) developed a dose-finding method that satisfies both efficacy and safety requirements. The method treats a sufficient number of patients, like

*Corresponding author. Tel.: +8801716187606; fax: (880-2) 9667222.

Email address: iftakhar@isrt.ac.bd (M. Iftakhar Alam)

30 or 45, to estimate the rates of efficacy and toxicity at the selected dose with a given reliability, and also stops the trial if it is found that none of the doses are both safe and efficacious. The dose-response outcome in this case is trinomial, categorised as neutral, efficacious or toxic, as in Li et al. (1995). However, the method suffers from the limitation that, in the settings where all the doses have
10 acceptable toxicity with higher efficacy at the higher doses, it does not escalate to the more desirable doses with high probability. So it often fails to detect the best dose in the presence of a number of candidate doses. Thall and Cook (2004) proposed another method based on the trade-offs between treatment efficacy and toxicity, which provides a substantial improvement over the earlier version and also accommodates bivariate binary outcomes. The successive patients in a trial receive doses based
15 on a set of efficacy-toxicity trade-off contours that partition the two-dimensional outcome space. The method is known as EffTox.

Zhang et al. (2006) proposed another such method considering trinomial responses. The design selects a dose based on some optimal dose-selection criteria using the continual reassessment method
20 and is popularly known as TriCRM. Although the design follows a similar idea to that in Thall and Russell (1998), it uses the continuation ratio (CR) model with different dose-selection criteria and simplified stopping rules. Thall and Cook (2004) also extended the approach of Thall and Russell (1998) to the CR model, but it involves considerable effort to elicit priors. In that sense, TriCRM is a simple alternative to EffTox.

25 Dragalin and Fedorov (2006) suggested an adaptive procedure considering efficacy and toxicity as endpoints. The modelling of these bivariate binary endpoints is based on either Gumbel bivariate binary logistic regression or the Cox bivariate binary model. They express a dose-finding problem in terms of a penalised D -optimality criterion. The design maximises the information under the control
30 of a penalty function for treating patients at doses which are too low or too high. Thall et al. (2008) presented a dose-finding procedure based on bivariate outcomes that incorporates patients' covariates and dose-covariate interactions. This is an extension of the methodology in Thall and Cook (2004). Thall and Nguyen (2012) proposed a new approach for bivariate ordinal outcomes.

35 Phase I trials are small in size and as a consequence the dose-toxicity curves may not be well estimated. They often determine a dose which can be found either unacceptably toxic or ineffective in a later phase. Phase I/II trials are relatively larger and so the associated methods could potentially lead to more efficient dose selection. They are still not big enough for the selected dose to converge

to the best dose with probability one, even if the algorithm is shown to be convergent. The outcomes
40 in dose-finding studies are often dichotomised, and, as a result, we lose some information. In par-
ticular, an outcome which is not toxic may be just below the cut-off point or a toxic outcome may
be far above the cut-off point. Similarly, the efficacy is dichotomised. Usually these issues are not
considered in the dose-escalation methods. However, they can be considered implicitly by taking into
45 account continuous measures like the area under the concentration curve (AUC). This additional infor-
mation utilised during the process of dose selection may significantly improve the efficiency of the trial.

In this paper, we introduce a general approach for phase I/II trials, which, along with efficacy
and toxicity endpoints, also considers pharmacokinetic (PK) information in the dose selection, the
information often viewed as important for finding the best dose (Govindarajulu, 1988). Although
50 there is considerable work on modelling PK and pharmacodynamic (PD) data, not much effort has
been devoted to incorporating PK information into adaptive dose-finding studies. Piantadosi and Liu
(1998) demonstrated that the efficiency and accuracy of phase I clinical trials can be improved by
incorporating PK data into dose escalation. Along with dose, they incorporated the resulting AUC
as a covariate in the model. The implementation of the method requires accurate PK data, but the
55 theory of optimal design to efficiently collect blood samples has not been used. Nyberg et al. (2009)
considered simultaneous optimisation of dose and PK sampling times. A Bayesian approach for phase
I trials was presented in Whitehead et al. (2007), which is based on simultaneous monitoring of PK
and PD responses. The methodology in Zhou et al. (2008) is for phase I trials in healthy volunteers,
where, for each individual in the trial, the method monitors two continuous PK measures, AUC and
60 the maximum concentration, and a binary indicator variable for an undesirable event. The presented
methods that incorporate PK measurements in dose escalation are for phase I trials. Since these are
usually small, it is difficult to address the issue of population variability. In this paper, we propose a
dose-selection method for a seamless phase I/II clinical trial where it is possible to assess both toxic
and efficacious outcomes. In such trials there is particular interest in the PK parameters and we use
65 the optimal sampling times for their precise estimation.

Concentration of a drug in blood is generally modelled as a function of time for a given dose and the
PD effect is often modelled as a function of dose. In a combined approach, the PD response is modelled
as a function of concentration. Concentration is inherently more informative than dose because, unlike
70 dose, which is only a nominal mass administered to a patient, it gives biological information (Riviere,
2011). The PK/PD approach establishes the dose-concentration-effect relationship and is capable of

predicting the effect at any time after administering a dose. It also helps in estimating the dose and dosing interval to achieve the effect of a desired level. The PK/PD approach has been described in Meibohm and Derendorf (1997), Hooker and Vicini (2005), and Davidian (2010), among others. Since
75 characteristics such as age, body weight, ethnicity, physiological functions, genetics, disease status and sensitivity to treatment often vary from patient to patient, Jönsson (2004) developed methodologies for individualised dosing utilising population PK and PK/PD models. These did not employ optimal time points to collect the PK and PD responses, the use of which could save resources and provide more accurate parameter estimates.

80

More work is needed to increase the efficiency of clinical trials and here we present our attempt in this direction. In this paper, we do not combine the PK and PD models into one. Instead, we use a PD response to select the best dose with a restriction put on the drug exposure measured by the area under the concentration curve. We consider two models, one for the dose-response outcomes and the
85 other for the PK data. We begin with an up-and-down design for the first few cohorts, after which the model-based procedure starts. At that stage, we choose the dose which maximises the estimated probability of efficacy subject to the conditions that the estimated probability of toxicity at that dose is not more than an acceptable level and also the estimated mean AUC is not more than a target value. The latter constraint is defined so that it includes the inter-patient variability in the AUC.
90 The analytical expressions for the AUC and its variability are based on the assumed PK model. The use of the AUC based on a separate PK model is simpler than the PK/PD approach. It gives precise estimates of the PK parameters and leads to faster decision making regarding the dose. The aim of this paper is to develop an efficient dose-finding method that exposes not too many patients to either sub-therapeutic or toxic doses and recommends the best dose for further study in phase IIB or III.

95

The paper is organised as follows. Section 2 introduces the general algorithm for the new design, the PK and dose-response models, and the dose-selection criteria and stopping rules. Section 3 presents the simulation results comparing the new method with a special case of the one of Zhang et al. (2006). Finally, there is a discussion in Section 4. Details of the up-and-down initial stage, model approximation
100 and derivation of the approximate mean and variance of the AUC are given in appendices.

2. Methods

2.1. Basic idea

We assume that a set of dose levels is pre-specified from pre-clinical trials or very early clinical trials. The aim of the study is to select the dose level which can be recommended for further investigation in a larger group of patients, where it is compared with a placebo or another experimental treatment. We consider here an adaptive design where an interim analysis of the data is performed after treating each cohort of patients, and a new decision on the PK sampling times and on the dose allocation is made for the next cohort based on the updated knowledge of the responses. As mentioned in Section 1, we support the choice of dose levels by additional information on the concentration of the drug in the plasma, which is directly related to the dose response. Too high a concentration may cause toxicities while too low a concentration will not produce a desirable effect. The dose-optimisation criterion to be used in this paper is the maximisation of the estimated probability of efficacy subject to the constraints imposed on the estimated probability of toxicity at the selected dose and on the estimated mean AUC. The population D -optimal blood sampling times assure high accuracy in the estimation of the mean PK parameters and their variances. Our dose-response model is derived from a trinomial distribution for the outcomes: neutral, efficacious and toxic. A one-compartment mixed effects model with bolus input and first-order elimination is used for the PK data. In what follows, we present the detailed forms of the considered models, both for the concentration of the drug and for the drug effect. However, we would like to stress that the algorithm is more general, that is, it can be applied to other models as well, as it is indicated in Section 4.

2.2. Optimum adaptive clinical trial design

We assume that the patients enter a clinic sequentially and cohorts of the same size are treated with a dose level determined from the updated information. The choice of dose level for each cohort is model based and satisfies an optimisation criterion. There are various possible criteria, such as maximum tolerated dose, biologically optimum dose or simply D -optimum dose. The maximum tolerated dose is the dose level for which the probability of toxicity attains a maximum permissible value. This criterion is often used in oncology trials, as it is usually assumed that both the efficacy and toxicity probabilities increase with dose level. However, in cases where we can observe efficacy, it makes sense to consider a criterion which allows for the highest chance of efficacy, particularly if it does not increase with dose level. Alternatively, one can consider a criterion which in principle should lead to the best dose-response model prediction and so the best indication of the efficacious dose level after the trial.

In this paper, we consider the biologically optimum dose approach, as, apart from recommending an optimum dose for further studies, it generates dose levels during the trial considered as most effi-
135 cacious, and thus gives a good chance of the best treatment for the patients in the trial. We escalate the dose level in the trial to maximise the chance of efficacy, but not skipping too many dose levels, while controlling the toxic events. This is done in two ways: by putting a constant threshold on the probability of toxic outcomes and by restricting the total exposure of the drug in the body. The former restriction is based on a dichotomised response, yes or no for a toxic event, over a cohort, while the
140 latter one is based on a continuous PK measure, the population mean AUC. This additional information needs to be as accurate as possible and so we derive population D -optimum designs for blood sampling times for each cohort of patients. This gives us very precise estimates of the PK parameters at each stage of the adaptive trial.

145 Below we present the main steps of the adaptive design, where k represents the stage in a trial.

Step 1: Treat cohort k with the current best dose.

Step 2: Obtain the D -optimal sampling time points and observe the PK responses at these time points, when appropriate.

Step 3: Observe the dose-response outcomes.

150 **Step 4:** Estimate the model parameters and update the models.

Step 5: Select the best dose for the next cohort based on the chosen dose optimisation criterion and constraints.

Step 6: Stop if the stopping rule is met, otherwise set $k = k + 1$ and repeat Steps 1-5.

Step 7: Carry out a complete analysis of the data to recommend a dose for further study.

155

We need to estimate the PK parameters as well as the dose-response parameters at each stage of the adaptive trial. With a small cohort size, we need to collect data from a few cohorts to obtain some reliable parameter estimates. Hence, we start the trial with an up-and-down procedure, which is run for the first few cohorts before we start the fully adaptive parametric algorithm described above.
160 During the up-and-down procedure, we observe the PK responses and the dose-response outcomes, but there is no estimation of parameters. The up-and-down procedure is similar to that of Ivanova (2006), where a dose is increased, stays at the present dose or is decreased depending on the responses of the most recent cohort. Here, however, we take into account responses from all cohorts up to the most recent one. The basis of the method lies in the toxic outcomes, which can stop the trial if toxicity is

165 above the acceptable level. Details of this initial part of the trial are given in Appendix A.

In Section 3, we conduct a simulation study to illustrate the operating characteristics of the proposed design. The example used consists of the one-compartment PK model with random effects and the continuation ratio dose-response model, described in Sections 2.3 and 2.4, respectively.

170 2.3. Population PK model

We consider PK models for the concentration of a drug in the plasma. Often, the mechanistic part of such models is a solution of differential equations representing the distribution of the drug in the body's compartments. The drug is absorbed, distributed and eliminated, and these processes differ among patients. Hence, so-called population PK models need to be considered to account for this
175 variability.

We denote the population PK models by

$$y_{ij} = f(\boldsymbol{\theta}_i, t_{ij}) + \epsilon_{ij}, \quad i = 1, \dots, N, \quad j = 1, \dots, n_i, \quad (1)$$

where y_{ij} is a random variable representing the concentration of a drug in the blood for the i th individual at time t_{ij} , $\boldsymbol{\theta}_i$ is a p -dimensional vector of the model parameters for the i th individual, $f : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is a function which is non-linear with respect to the parameters and possibly also with respect to the
180 design variable t , ϵ_{ij} denotes the random error and n_i is the number of measurements taken on individual i . An additive random error model may lead to negative concentrations in the simulation study with large error observations. However, it works well in the presence of a small error variance. For a large error variance, it would be wise to consider a multiplicative error model $y_{ij} = f(\boldsymbol{\theta}_i, t_{ij}) \exp(\epsilon_{ij})$, so that $\log(y_{ij}) = \log(f(\boldsymbol{\theta}_i, t_{ij})) + \epsilon_{ij}$. The methodology presented in this paper can be applied to both
185 cases with some minor technical adjustments to the model approximation.

Furthermore, we assume that some or all of the model parameters are random and can be represented in an additive form as $\boldsymbol{\theta}_i = \boldsymbol{\beta} + \mathbf{b}_i$, where $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^T$ is a vector of the population mean parameter values and \mathbf{b}_i is a vector of random effects for subject i . We assume that \mathbf{b}_i has a multivariate normal distribution with zero mean vector and diagonal covariance matrix denoted by

$$E(\mathbf{b}_i) = \mathbf{0} \quad \text{and} \quad \text{Var}(\mathbf{b}_i) = \boldsymbol{\Omega} = \text{diag}(\omega_1, \dots, \omega_p).$$

We also assume that the random errors are independently normally distributed with zero mean and constant variance σ^2 and that \mathbf{b}_i and $\boldsymbol{\epsilon}_i$ are independent. Retout et al. (2001) also assumed an additive

model for the random parameters. Although the normality assumption may lead to negative values
of the parameters, this will be very rare unless the variability in the parameters is very large. This
190 possibility can be avoided if an exponential model $\theta_i = \beta \exp(\mathbf{b}_i)$ is assumed instead. Here, we consider
the additive random parameters to keep the PK model simple. However, a switch to the exponential
model would not alter the dose-finding methodology presented here.

We are interested in the efficient estimation of all the population parameters, denoted by $\Psi =$
 $(\beta^T, \lambda^T)^T$, where $\beta^T = (\beta_1, \dots, \beta_p)$ and $\lambda^T = (\omega_1, \dots, \omega_p, \sigma^2)$. The dimension of Ψ is $2p + 1$ if all
the parameters are random, otherwise it will be less. The Fisher information matrix (FIM) for the i th
individual and design $\xi_i = \{t_{i1}, \dots, t_{in_i}\}$ is given by

$$M_i(\Psi, \xi_i) = E \left\{ -\frac{\partial^2 l_i(\Psi | \mathbf{y}_i)}{\partial \Psi \partial \Psi^T} \right\}, \quad (2)$$

195 where $l_i(\Psi | \mathbf{y}_i)$ is the log-likelihood function for a given vector of observations $\mathbf{y}_i = (y_{i1}, \dots, y_{in_i})^T$
for individual i . Since our model is non-linear in the parameters, derivation of an analytic expression
for the log-likelihood function is not possible. Various approximations to the information matrix have
been proposed in the statistical literature. A comparison of different methods is presented in the PhD
thesis by Mielke (2012), who concludes that none of the methods is uniformly best.

200

As in Pinheiro and Bates (1995) or Retout et al. (2001), we approximate the log-likelihood function
using a first-order Taylor series expansion of the function $f(\theta_i, t_{ij})$. Here, however, we expand the
function about $\phi_i = (\beta^T, \mathbf{b}_i^T)^T$ at $\phi^0 = (\beta^{0T}, E(\mathbf{b}_i)^T)^T = (\beta^{0T}, \mathbf{0}^T)^T$, where β^0 is a vector of initial
values for the population means: see Appendix B. The population FIM for the design for all subjects
in the trial $\Xi = \{\xi_1, \dots, \xi_N\}$ is defined as the sum of the N individual Fisher information matrices in
(2), that is,

$$M(\Psi, \Xi) = \sum_{i=1}^N M_i(\Psi, \xi_i).$$

The FIM is an argument of many optimal design criteria, initially derived for fixed effects lin-
ear models, but more recently extended to non-linear models by Fedorov (1972); Fedorov and Hackl
(1997), and Atkinson et al. (2007), including random effects, as in Mentré et al. (1997). In this work,
we use the D -optimality criterion, that is, we maximise the determinant of the population FIM to
205 obtain sampling times for measuring the concentration of the drug in the plasma. The properties of
this criterion are well known and it is widely used in many applications. We find that our estimates
of the PK model parameters are very precise, as one would expect from data collected according to
a D -optimum design. Computations for the D -optimal time points are implemented in *PFIM 3.2*

(Bazzoli et al., 2010), an *R* package to evaluate and optimise designs in the context of population PK
210 experiments.

Since the FIM depends on the model parameters, to obtain the locally optimal time points we
need to assume some prior values for the parameters. The procedure starts with some informed guess
of the parameters and then replaces them at each stage with the current maximum likelihood esti-
215 mates. The covariance matrix of the maximum likelihood estimators approaches the inverse of the
FIM asymptotically. Therefore, by minimising the inverse of the determinant of FIM, the *D*-criterion,
we can minimise the asymptotic general variance of the estimated model parameters. To achieve this,
we obtain the maximum likelihood estimates using the *R* procedure *nlme* (Pinheiro and Bates, 2000).
It is also possible to use the posterior estimates of the PK parameters in the search for the optimal
220 time points, but that would involve more computational challenges and would require assumptions
regarding prior distributions of the parameters. Wakefield et al. (1994) developed a Bayesian approach
for the non-linear mixed-effects model. They used a three-stage model and a Markov Chain Monte
Carlo method to obtain the posterior density of the random effects, and hence the posterior estimates.

225 Apart from precise estimation of the model parameters and model prediction, we are interested
in using the PK information for guiding an adaptive dose-selection procedure. A patient's response
depends on the exposure to the drug. Hence, we use the AUC of the PK profile as a measure in the
guidance. Other measures could also be considered, such as the maximum concentration or the time
to maximum concentration. Although we focus here on the AUC, it is straightforward to adapt the
230 algorithm to another measure.

Let $h(x, \boldsymbol{\theta}_i)$ represent the AUC for individual i , where x is the dose received, $\boldsymbol{\theta}_i$ is the vector of
random PK parameters and h is a differentiable function of the parameters. Then the concentration
curve is a random variable. To derive approximations to its expectation and variance, we use a first-
235 order Taylor series expansion of $h(x, \boldsymbol{\theta}_i)$ about $\boldsymbol{\theta}_i$ at $\mathbf{E}(\boldsymbol{\theta}_i)$: see Appendix C. The expression derived
here for measuring the inter-patient variability in the area under the concentration curve is general.
The methodology can be adapted for the maximum concentration, another important PK parameter,
and, of course, for any underlying PK model.

Example

Consider the simple one-compartment PK model with bolus input and first-order elimination. Then model (1) becomes

$$y_{ij} = \frac{x}{V_i} \exp\left(-\frac{Cl_i}{V_i} t_{ij}\right) + \epsilon_{ij}, \quad i = 1, \dots, N, \quad j = 1, \dots, n_i,$$

240 where Cl_i and V_i denote the clearance and volume of distribution for individual i , so that $\boldsymbol{\theta}_i = (V_i, Cl_i)^T$ is the vector of parameters, with $\boldsymbol{\beta} = (V, Cl)^T$ and $\mathbf{b}_i = (b_{V_i}, b_{Cl_i})^T$. It is shown in Appendix C that the approximate mean and variance of the AUC over the range $[0, t_1]$ are

$$E\{h(x, \boldsymbol{\theta}_i)\} \cong h(x, \boldsymbol{\beta}) = \frac{x}{Cl} \left\{ 1 - \exp\left(-\frac{Cl}{V} t_1\right) \right\} \quad (3)$$

and

$$\text{Var}\{h(x, \boldsymbol{\theta}_i)\} \cong \left\{ \frac{x t_1}{V^2} \exp\left(-\frac{Cl}{V} t_1\right) \right\}^2 \omega_1 + \left\{ \frac{x}{Cl} \exp\left(-\frac{Cl}{V} t_1\right) \left(\frac{1}{Cl} + \frac{t_1}{V} \right) - \frac{x}{Cl^2} \right\}^2 \omega_2. \quad (4)$$

These approximations are used in the adaptive dose-selection procedure, which is described in
245 Section 2.5.

2.4. Dose-response model

We consider a trinomial response $\mathbf{Y} = (Y_0, Y_1, Y_2)^T$ for each patient, where Y_0 is a neutral response, Y_1 is an efficacious response and Y_2 is a toxic response. The probability of each of these outcomes depends on dose. It is commonly accepted that a drug's toxicity increases with dose. Efficacy also
250 increases in many cases, but it is possible for some drugs that the efficacy attains a plateau or increases and then decreases. For an experimental drug, we assume that the probability of a neutral response decreases monotonically with dose and that the probability of toxicity increases monotonically with dose. However, the probability of an efficacious outcome may be non-monotonic, increasing or decreasing. The corresponding probabilities are denoted by $\psi_0(x, \boldsymbol{\vartheta})$, $\psi_1(x, \boldsymbol{\vartheta})$ and $\psi_2(x, \boldsymbol{\vartheta})$, so that
255 $\psi_0(x, \boldsymbol{\vartheta}) + \psi_1(x, \boldsymbol{\vartheta}) + \psi_2(x, \boldsymbol{\vartheta}) = 1$, where $\boldsymbol{\vartheta}$ is the vector of parameters.

The proportional odds model is often used to establish the dose-response relationship in such a situation. However, it requires the assumption that dose has the same effect across the cumulative logits. Since this assumption is very difficult to satisfy, we plan to utilise the continuation ratio model
260 (Agresti (1990), Chapter 9) given by

$$\log \left\{ \frac{\psi_1(x, \boldsymbol{\vartheta})}{\psi_0(x, \boldsymbol{\vartheta})} \right\} = \vartheta_1 + \vartheta_2 x \quad \text{and} \quad \log \left\{ \frac{\psi_2(x, \boldsymbol{\vartheta})}{1 - \psi_2(x, \boldsymbol{\vartheta})} \right\} = \vartheta_3 + \vartheta_4 x.$$

The parameter ϑ_1 represents the baseline log-relative probability, ϑ_2 reflects the contribution of dose in the log-relative probability of having an efficacious outcome relative to a neutral one, ϑ_3 is the baseline log-odds and ϑ_4 is the contribution of dose in the log-odds of a toxic outcome relative to a neutral or efficacious one. Solving the above equations, we obtain

$$\psi_0(x, \boldsymbol{\vartheta}) = \frac{1}{(1 + e^{\vartheta_1 + \vartheta_2 x})(1 + e^{\vartheta_3 + \vartheta_4 x})},$$

$$\psi_1(x, \boldsymbol{\vartheta}) = \frac{e^{\vartheta_1 + \vartheta_2 x}}{(1 + e^{\vartheta_1 + \vartheta_2 x})(1 + e^{\vartheta_3 + \vartheta_4 x})}$$

265 and

$$\psi_2(x, \boldsymbol{\vartheta}) = \frac{e^{\vartheta_3 + \vartheta_4 x}}{1 + e^{\vartheta_3 + \vartheta_4 x}}.$$

To ensure that the above three non-linear functions exhibit the assumed behaviour of the responses, we restrict the parameter space to

$$\Theta = \{\boldsymbol{\vartheta} = (\vartheta_1, \vartheta_2, \vartheta_3, \vartheta_4)^T : \vartheta_1 \geq \vartheta_3, \vartheta_3 < 0 \text{ and } \vartheta_2, \vartheta_4 > 0\}.$$

Assume that we are at the k th stage in a trial. So k cohorts have been treated so far with dose levels from the set of ordered levels \mathcal{X} . Let \mathbf{x} be the $k \times 1$ dose vector with components x_l and let
 270 \mathbf{R} be the $k \times 3$ outcome matrix with $\mathbf{R}_l = (R_{l0}, R_{l1}, R_{l2})$ as the l th row, $l = 1, 2, \dots, k$. Note that $R_{l0} + R_{l1} + R_{l2} = c$, where c is the number of subjects in a cohort treated with dose x_l . The successive components of \mathbf{R}_l are the counts of neutral, efficacious and toxic responses for the l th cohort. Thus, the likelihood function is

$$L_k(\boldsymbol{\vartheta} \mid \mathbf{x}, \mathbf{R}) \propto \prod_{l=1}^k \{\psi_0(x_l, \boldsymbol{\vartheta})\}^{R_{l0}} \{\psi_1(x_l, \boldsymbol{\vartheta})\}^{R_{l1}} \{\psi_2(x_l, \boldsymbol{\vartheta})\}^{R_{l2}}.$$

The parameters $\boldsymbol{\vartheta}$ are estimated at each stage of the adaptive trial. Since maximum likelihood
 275 estimation is unsuitable because of small sample sizes at the early stages of the trial, we employ a Bayesian approach. The posterior estimate of $\boldsymbol{\vartheta}$ at the k th stage is

$$\hat{\boldsymbol{\vartheta}}_k = \frac{\int_{\Theta} \boldsymbol{\vartheta} p(\boldsymbol{\vartheta}) L_k(\boldsymbol{\vartheta} \mid \mathbf{x}, \mathbf{R}) d\boldsymbol{\vartheta}}{\int_{\Theta} p(\boldsymbol{\vartheta}) L_k(\boldsymbol{\vartheta} \mid \mathbf{x}, \mathbf{R}) d\boldsymbol{\vartheta}}, \quad (5)$$

where $p(\boldsymbol{\vartheta})$ is the joint prior distribution of the parameters. Let us assume that $0 < \vartheta_2 < u_1$, $0 < \vartheta_4 < u_2$, $w_1 < \vartheta_1 < w_2$ and $w_3 < \vartheta_3 < w_4$, and that the joint prior distribution is uniform. Then we obtain

$$p(\boldsymbol{\vartheta}) = \frac{2}{u_1 u_2 (w_2 - w_3)^2}, \quad \boldsymbol{\vartheta} \in \tilde{\Theta}, \quad (6)$$

280 where

$$\tilde{\Theta} = \{\vartheta : w_3 < \vartheta_3 \leq \vartheta_1 < w_2, \quad 0 < \vartheta_2 < u_1, \quad 0 < \vartheta_4 < u_2\}. \quad (7)$$

In Section 3, we choose a different $\tilde{\Theta}$ for each of the simulation scenarios.

2.5. Dose-selection criteria

At each stage in the trial, we select that dose for the next cohort for which the estimated probability of efficacy is maximum, subject to the condition that the estimated probability of toxicity is not more
 285 than a pre-specified acceptable level. Often, it is not advisable to skip the pre-specified dose levels when they are increased for application in the next cohort. In our method, we introduce an option of constraining the increase by any number of dose levels. However, we impose no such constraint on the levels when they are decreased. For example, the next best dose could be five levels higher than the previous one, but with a constraint of not skipping more than one dose level, we apply the dose
 290 two levels higher than the previous one. Another condition for the dose selection is related to the total exposure of the drug in the body so that the curative purpose is likely to be achieved without overdosing; this is expressed by a restriction on the area under the concentration curve over time.

Suppose that we are at the k th stage of the trial and based on the current data, we have the
 295 estimates $\hat{\Psi}_k$ and $\hat{\vartheta}_k$. Then, we select the dose x_{k+1} for the next cohort of patients so that

$$x_{k+1} = \arg \max_{x \in \mathcal{X}} \psi_1(x, \hat{\vartheta}_k),$$

subject to the conditions that

$$\psi_2(x_{k+1}, \hat{\vartheta}_k) \leq \gamma$$

and

$$\frac{\hat{E}\{h(x_{k+1}, \theta_i)\} - \text{AUC}^0}{\widehat{\text{SD}}\{h(x_k, \theta_i)\}} \leq \delta(x_k, \hat{\vartheta}_k), \quad (8)$$

where $\gamma \in (0, 1)$ is a given constant, usually allowing 20-30% chance of toxicity and $\hat{E}\{h(x, \theta_i)\} = h(x, \hat{\beta}_k)$ is the estimate of the approximate population mean AUC for dose x at stage k . The estimate
 300 of the approximate standard deviation of the AUC is denoted by $\widehat{\text{SD}}\{h(x_k, \theta_i)\}$. Also, AUC^0 is a value for the AUC that is considered to be desirable and $\delta(x_k, \hat{\vartheta}_k) = 1/\psi_1(x_k, \hat{\vartheta}_k)$. A desirable AUC is one for which the curative purpose is likely to be achieved, allowing some acceptable level of toxicity. The choice of such a value will require expert opinion. Previous studies of similar drugs or pre-clinical

studies can help in this context.

305

The left-hand side of (8) represents a relative difference between $h(x_{k+1}, \hat{\beta}_k)$ and AUC^0 . We constrain the choice of x_{k+1} so that, for large values of the estimated probability of efficacy $\psi_1(x_k, \hat{\vartheta}_k)$, this difference is small. This ‘forces’ convergence of the dose to the one giving the required exposure to the drug. On the other hand, when the estimated probability of efficacy is small, the constraint is weak, allowing for a wider choice for the next dose level. The PK constraint (8) is dynamic, that is, the value of $\delta(x_k, \hat{\vartheta}_k)$ changes during the trial according to the current estimate of the probability of efficacy, and the estimates of the expectation and the standard deviation of the AUC are updated after each cohort. This gives some flexibility to the algorithm. A fixed δ might lead to choosing a sub-optimal dose, and, in any case, it would be difficult to decide on its value.

315

It follows from the constraint that $\hat{E}\{h(x_{k+1}, \theta_i)\} \leq \text{AUC}^0 + \delta(x_k, \hat{\vartheta}_k) \widehat{\text{SD}}\{h(x_k, \theta_i)\}$. If $\psi_1(x_k, \hat{\vartheta}_k)$ attains the maximum possible value 1, then $\delta(x_k, \hat{\vartheta}_k)$ will have the value 1 and consequently we will choose a dose with a mean AUC within one standard deviation of the target, which accommodates the population variability. However, in the majority of cases, δ will have a larger value than 1. Therefore, we will usually be selecting a dose with a mean AUC within more than one standard deviation of the target value.

320

This constraint is introduced as an additional indicator of a curative effect, with a precaution against the allocation of too toxic doses. It works differently to the constraint on the probability of toxicity ψ_2 . The AUC, expressed as a function of the random PK parameters, not only takes into account the population variability, but also directly constrains the PK parameters responsible for the drug’s action, and so indirectly, the dose level.

325

2.6. Stopping rules

In a clinical trial, a variety of stopping rules are possible and the choice depends on the purpose of the trial. We employ a simple rule: stop the trial when the same dose is repeated for r cohorts or when the trial reaches the maximum number of m cohorts, whichever comes first. The idea is to terminate the trial early to save resources if it is found that the same dose is being selected repeatedly. Otherwise, it will run for the maximum number of cohorts available. For early stopped trials, the optimum dose (OD) is defined as the dose that has been repeated r times. However, for the trials that use the maximum number of cohorts m , we carry out a complete analysis of the data and define the OD as the dose that would be allocated to cohort $m + 1$ if that cohort were in the trial.

335

3. Simulation results

3.1. Choice of design parameters

We assume that an experimental drug has 20 available doses from 0.5 to 10, increasing in steps of 0.5, that is, the set of doses is $\mathcal{X} = \{0.5, 1.0, \dots, 10.0\}$. Four hypothetical dose-response scenarios are considered: see Figure 1. Scenario 1 has a monotonically decreasing efficacy curve with dose 0.5 as

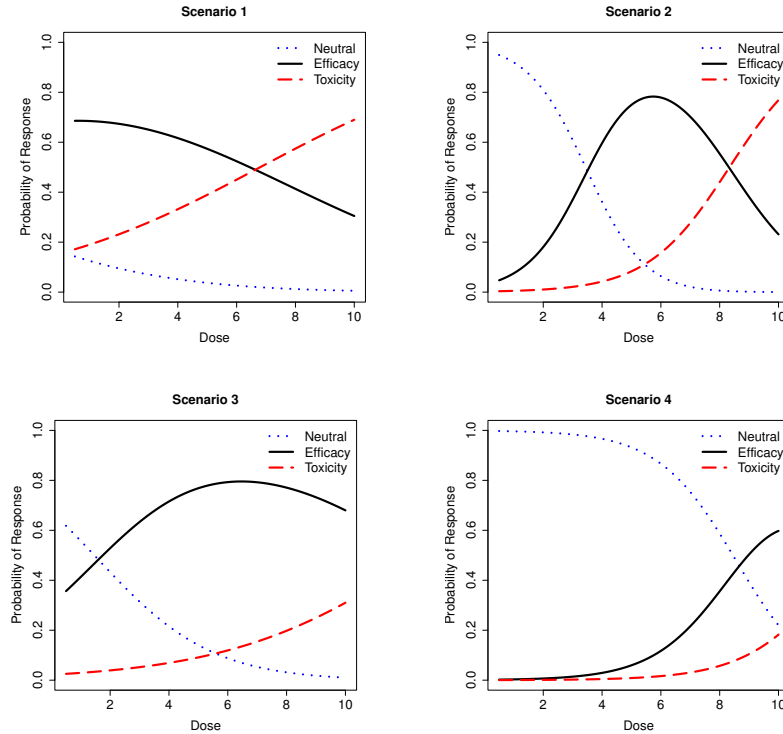


Figure 1: Dose-response scenarios for simulation study. The respective parameter values are: Scenario 1, $\boldsymbol{\vartheta} = (1.44, 0.26, -1.70, 0.25)^T$; Scenario 2, $\boldsymbol{\vartheta} = (-3.50, 1.00, -6.00, 0.72)^T$; Scenario 3, $\boldsymbol{\vartheta} = (-0.80, 0.50, -3.80, 0.30)^T$; and Scenario 4, $\boldsymbol{\vartheta} = (-6.50, 0.75, -8.00, 0.65)^T$.

the OD. Scenarios 2 and 3 depict non-monotonic efficacy curves with respective ODs of 5.5 and 6.5. A monotonically increasing efficacy curve is reflected in Scenario 4 with 10 as the OD. Two kinds of responses need to be generated for the simulation study: the concentration of the drug in the blood and the trinomial dose-response outcomes. The values of the PK parameters for the simulation study are $V = 0.5$ L, $Cl = 0.06$ L/hr, $\omega_1 = 0.004$, $\omega_2 = 0.00005$ and $\sigma^2 = 0.000225$. The parameter values are chosen such that the coefficient of variation is around 12%. The choice of a small error variance should be reasonable as we have found the observations to be small as well.

350 Each trial starts with the lowest dose of 0.5 mg/kg body weight. The acceptable toxicity level γ is taken to be 0.20. Doses to the first four cohorts in each trial are allocated according to the up-and-down design in Appendix A. The up-and-down stage need not have exactly four cohorts. There are five parameters to estimate in our PK model and, therefore, we have chosen four cohorts to provide initial values for the maximum likelihood estimates. One should choose this number depending on the complexity of the model. The value of AUC^0 is taken to be the AUC at the true OD in each scenario. 355 Although we consider the same γ for each scenario, we have different values of AUC^0 . The highest doses which meet the safety level are 1 mg/kg, 6 mg/kg, 8 mg/kg and 10 mg/kg for the respective scenarios. For each trial, we set the maximum number of cohorts to be $m = 20$, each cohort of size 3. To let the trials stop early when it is found that no further improvement in dose selection is possible, 360 we set $r = 6$.

3.2. Generation of PK and dose-response outcomes

A vector of random effects \mathbf{b}_i for individual i is generated from the bivariate normal distribution $N_2(\mathbf{0}, \mathbf{\Omega})$. The PK parameters for that individual are then obtained as $\boldsymbol{\theta}_i = \boldsymbol{\beta} + \mathbf{b}_i$. The next step is to find the individual concentrations at the D -optimal time points. The design region for the sampling 365 times is $\mathcal{T} = [0, t_1]$ hours, where $t_1 = 30$ in (3). To decide on the optimal number of sampling times, we use the relative efficiency defined in (9). Although we consider the same prior in each case, the designs have a different number of design points. It has been found that the efficiency of a 3-point design relative to a 2-point one is double. The efficiency of a 4-point design relative to a 3-point one falls to 1.15. As we increase the number of design points, such efficiencies become closer to 1: see Figure 2. This means that the gain is substantial if we consider 3 design points rather than 2.




Figure 2: Rationale for setting the number of design points in the one-compartment PK model with bolus input and first-order elimination. The locally D -optimum design points are obtained using the initial prior values Ψ^0 assuming that the lowest dose is given to a cohort of 3 patients.

370

Therefore, to avoid the difficulty of collecting many samples, three optimal time points are considered for the individuals in each cohort, that is, $n_i = 3$ for all i . The random errors are then generated from $N_3(\mathbf{0}, \sigma^2 \mathbf{I}_3)$ and added to the previously generated individual concentrations to produce the simulated PK responses for individual i . The same scheme is followed to simulate the responses for all individuals in each cohort. We chose the prior values $\Psi^0 = (0.1, 0.005, 0.0007, 0.0000006, 0.000004)^T$ to obtain the optimal points for the first four cohorts in the up-and-down stage of the trial. The prior values are quite far away from the true values of the parameters, which would normally be unknown and could be wrongly assumed at the beginning of the trial. For the fifth cohort onwards, the current maximum likelihood estimates of the PK parameters are used.




Figure 3: Simulated concentrations in a randomly chosen trial for Scenario 3.

Figure 3 gives the concentration data that are obtained in a randomly chosen trial for Scenario 3 from the simulation study of 1000 simulated trials. A total of 60 patients are recruited in this particular trial and the concentrations at the three D -optimal time points are shown for each patient and are denoted by circles. Though not presented, the other scenarios would have similar kind of concentration data. The progress of the same trial is summarised in Figure 4. Under each of the scenarios in Figure 1, we have specific probabilities at each dose to generate the trinomial dose-response outcomes. The




Figure 4: Summary of the progress in a random trial for Scenario 3.

outcomes obtained for each cohort of the presented trial are plotted against their numbers. Each dotted horizontal line that goes through the middle of the outcomes indicates the dose that the corresponding cohort received.

3.3. Model fitting

390 Once we have data on the concentration and the dose-response outcomes, we can update the model fits. We obtain the maximum likelihood estimates of the PK parameters using the *R* procedure *nlme* (Pinheiro and Bates, 2000). The posterior estimates of the dose-response parameters in (5), are obtained by numerical integration using *cubeature*, an *R* package (Johnson and Narasimhan, 2009). The package carries out the adaptive multidimensional integration over hypercubes. One needs to
 395 specify the tolerance limit and the maximum number of function evaluations desired. The smaller the tolerance limit or the larger the maximum number of function evaluations, the more accurate the estimate is. We set these to be 0.001 and 5,000, respectively, to keep the computational time down. We use a joint uniform prior distribution for ϑ , given in (6). The parameter space $\tilde{\Theta}$ is chosen for each scenario so that the true values of the parameters lie in the middle of the corresponding intervals. For
 400 instance, since Scenario 1 has the true parameters $\vartheta = (1.44, 0.26, -1.70, 0.25)^T$, $\tilde{\Theta}$ has $0 < \vartheta_2 < 0.52$, $0 < \vartheta_4 < 0.50$, $0 < \vartheta_1 < 2.88$ and $-3.40 < \vartheta_3 < 0$. More specifically, $u_1 = 0.52$, $u_2 = 0.50$, $w_1 = 0$, $w_2 = 2.88$, $w_3 = -3.40$ and $w_4 = 0$ in (7). The same approach is followed for the other scenarios. In evaluating the integrals in (5) with the uniform distribution specified in (6), the prior distributions in the numerator and denominator cancel out as they are constants. Once the posterior estimates are
 405 obtained, we update the dose-response functions.

3.4. Dose selection for the next cohort

Once we have the updated PK and dose-response parameter estimates, we can select the dose for the next cohort based on the criteria defined in Section 2.5. As an additional safety precaution, we

allow the design to skip only one level at a time when the dose is increased.

410 *3.5. Checking the stopping rules and the OD selection*

We continue the process of allocating doses to the cohorts until the stopping rules are satisfied. Once a trial reaches m cohorts, we carry out a complete analysis to find the OD. At the end of each trial, we also record the PK and dose-response parameter estimates, whether the trial stopped early and the OD in that case. Each of the four scenarios is investigated through 1,000 simulated trials.

415 All the computations are conducted in *R*. The simulations are implemented on a Dell PC with an Intel Core 2 Duo processor running at 3.00 GHz and RAM 4.00 GB. The processing time for 1,000 simulations is 8-10 hours.

3.6. Numerical results

We compare the operating characteristics of the new PK guided design, incorporating the AUC
420 constraint, with the one that does not take into account the PK information. The dose selection in the latter design is based on the probability of efficacy and the toxicity condition only, defined in Section 2.5. It is similar to the method presented by Zhang et al. (2006). The simulation results for these two designs are presented in Table 1 and Figures 5-11. In Figures 5-8, the first row shows the summaries when the additional PK constraint is employed. The summaries in the second row are based on the
425 toxicity constraint only.

Table 1 clearly shows the advantages of the additional constraint on the AUC profile. The gain in percentage of the doses correctly recommended for further studies depends on the scenario, but, in all cases considered, the PK guided designs are uniformly better in this respect.

430

Table 1: Percentage of best doses recommended as optimum for further studies (%BD), percentage of doses recommended as optimum for further studies, but carrying the probability of toxicity above the acceptable level (%TD), and percentage of cohorts treated at the best doses throughout the trials (%AD).

Scenario	Best Doses	%BD		%TD		%AD	
		PK	No PK	PK	No PK	PK	No PK
1	0.5	99.0	52.4	0.6	32.7	65.3	31.8
2	5.5 and 6.0	80.2	66.2	0.9	9.5	41.1	33.5
3	5.5-7.5	91.7	85.7	0.0	2.5	52.4	49.2
4	10.0	47.9	46.3	0.0	0.0	17.8	17.5

As seen in Table 1, as well as in the left panels of Figure 5, the largest benefit is shown in Scenario 1, where the best dose is the first one and small doses have a high probability of toxicity. The right panels of this figure show how the doses were allocated in the simulated trials. The new approach selects 0.5 as the OD in 99% of the trials, compared to 52.4% by the old approach. Only in 0.6% of the trials was a dose with a probability of toxicity above the limit γ chosen as the optimum in the PK guided trial, whereas it was 32.7% in the other case. Therefore, the PK guided design avoids doses with a high chance of toxicity, while the other design has not prevented this from happening. Furthermore, 65.3% of the cohorts were treated at the best dose in the PK guided trials, compared to only 31.8% in the trials without the PK guidance.

440

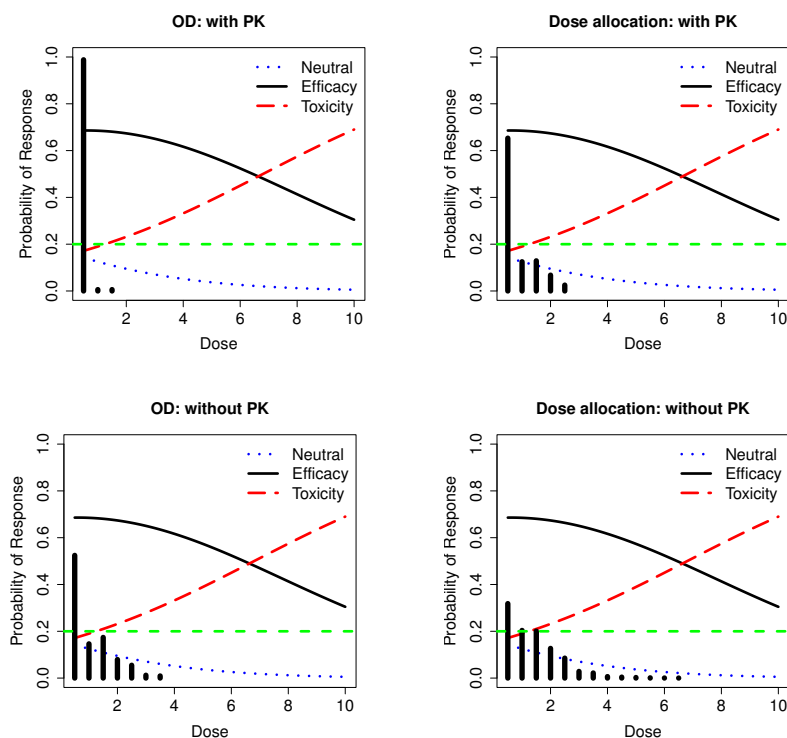


Figure 5: Scenario 1 with the OD at 0.5. The bars in the left panel represent the proportions of the doses selected as the OD in the simulations and those in the right panel represent the proportions of the cohorts treated at the allocated doses during the trials. The acceptable level for the probability of toxicity is indicated by the horizontal dashed line.

The new approach in Scenario 2 selects 5.5 as the OD in 39.7% of the trials. It selects dose 6.0 in 40.5% of the trials. This happens as the true probabilities of efficacy at these doses are quite close. These two figures together make 80.2% of what we call in Table 1 “best doses”. The corresponding figure for the old approach is 66.2%. Again, we observe that the PK guided design avoids recommending

445 doses with a high probability of toxicity, and, moreover, such doses are used much less in the simulated trials in this case: see Figure 6.

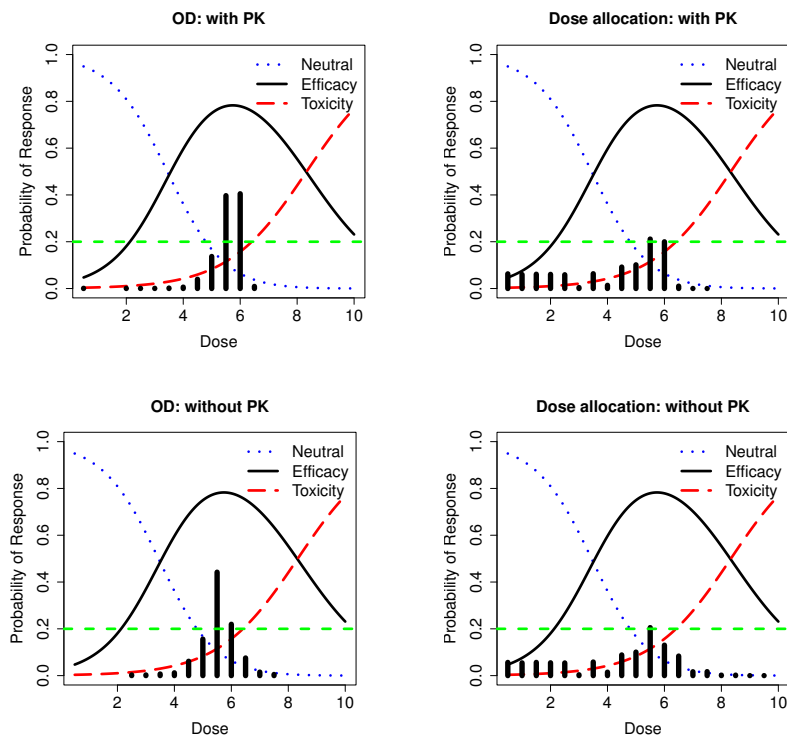


Figure 6: Scenario 2 with the OD at 5.5. The bars in the left panel represent the proportions of the doses selected as the OD in the simulations and those in the right panel represent the proportions of the cohorts treated at the allocated doses during the trials. The acceptable level for the probability of toxicity is indicated by the horizontal dashed line.

The new approach in Scenario 3 identifies exactly 6.5 as the OD in 17.8% of trials. Because of the flat shape of the efficacy curve, this scenario has a number of doses with probabilities of efficacy quite close to that for the OD. The doses are 5.5, 6.0, 6.5, 7.0 and 7.5, and these “best doses” are selected in 91.7% of trials. Although the old approach selects these doses in 85.7% of trials, it recommends doses above the toxicity probability threshold in 2.5% of cases. Furthermore, from Figure 7, the allocation of doses in the trials is again more ethical in the PK guided design.

455 There is little difference between the two designs in Scenario 4, as shown in Table 1 and also in Figure 8. The new approach selects 10.0 as the OD in 47.9% of the trials, while the percentage for the old approach is 46.3. This is the case where both the probability of efficacy and the probability of

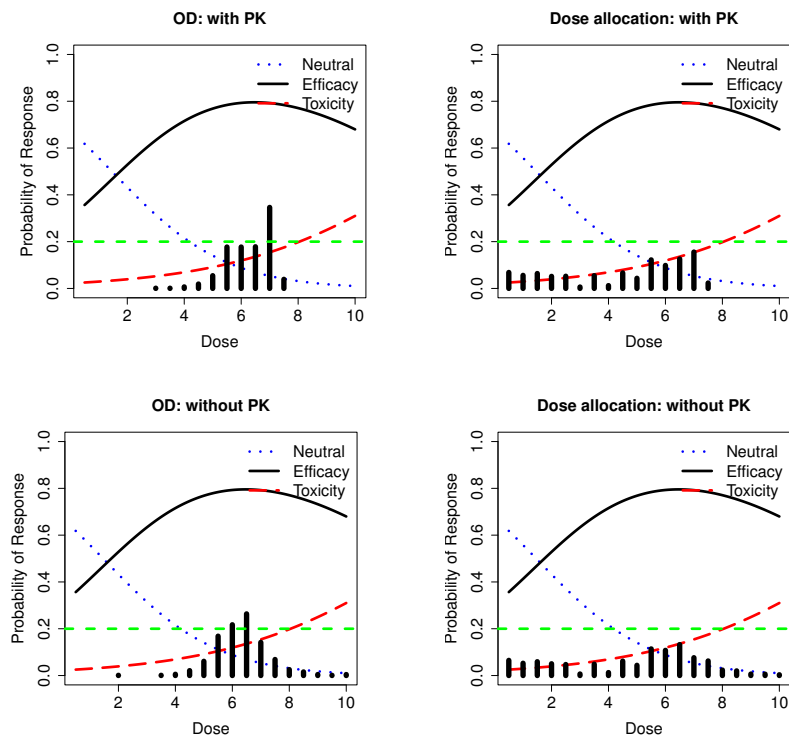


Figure 7: Scenario 3 with the OD at 6.5. The bars in the left panel represent the proportions of the doses selected as the OD in the simulations and those in the right panel represent the proportions of the cohorts treated at the allocated doses during the trials. The acceptable level for the probability of toxicity is indicated by the horizontal dashed line.

toxicity increase with dose, where only higher doses have a better chance of having an effect and all doses are below the toxicity threshold. This scenario illustrates a very cautiously chosen dose range. As
 460 a consequence, this leads to slow learning in the trial and requires the collection of a lot of information before a recommendation can be made.

We have found that, for this scenario, all available cohorts were used almost all the time. This is in contrast to Scenario 1, where, especially in the PK guided design, the learning process was fast
 465 and there were much smaller numbers of cohorts required in the trials. In fact, Scenario 1 is another extreme case, where the dose range is not well chosen. Smaller dose levels in this case could have been even better and the trial carries a high risk of toxic responses.

Figure 9 shows the distribution of the PK parameter estimates for all the scenarios. We have found
 470 the coefficient of variation for the mean PK parameter estimates to be around 2%, and those for the variance components and error variance to be 10-35%, except for Scenario 1 where it is high for the

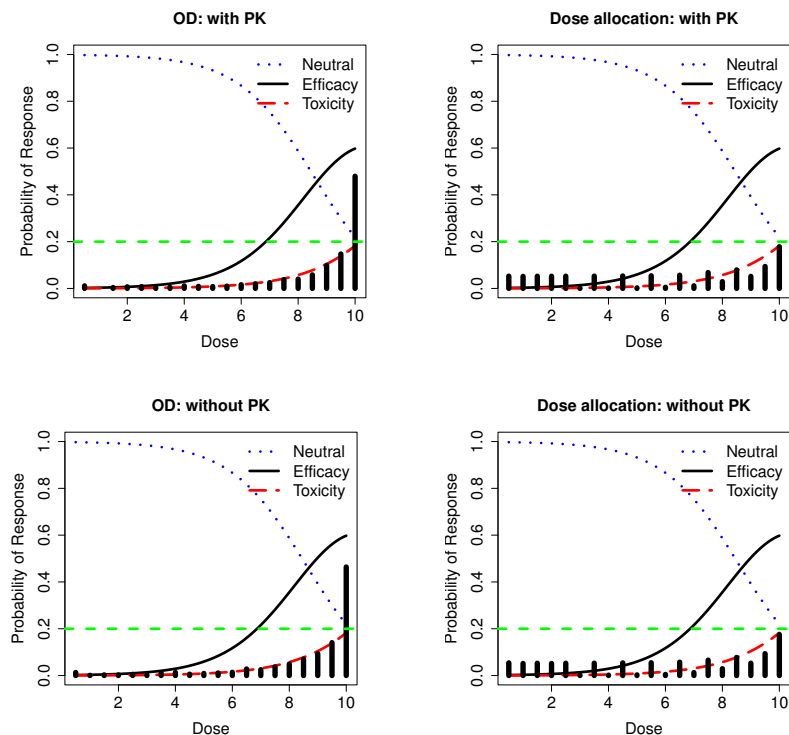


Figure 8: Scenario 4 with the OD at 10.0. The bars in the left panel represent the proportions of the doses selected as the OD in the simulations and those in the right panel represent the proportions of the cohorts treated at the allocated doses during the trials. The acceptable level for the probability of toxicity is indicated by the horizontal dashed line.

error variance because of the presence of some outliers in the estimates. Since the design employs the D -criterion for measuring PK responses, it is assuring high efficiency in the parameter estimation. From Figures 10 and 11, the dose-response parameter estimates obtained from the two approaches
 475 are similar. Obviously, they are not as good as the PK estimates. This is due to the fact that the information on the trinomial dose response is not gathered in a way that would be optimal for parameter estimation. Here, we focused on the criterion which would provide a good dose for further studies in an ethical trial, which is particularly important in classes of drugs where toxicity can be very serious.

480 It has already been mentioned that we stop a trial early if the same dose is repeated for six cohorts and call the associated dose the OD. We have found that, as the location of the OD moves from left to right in the dose region of a scenario, more cohorts are needed to stop early. Most of the early stopped trials identify the OD accurately. It has been found that the PK guided approach utilises fewer cohorts than the other approach.

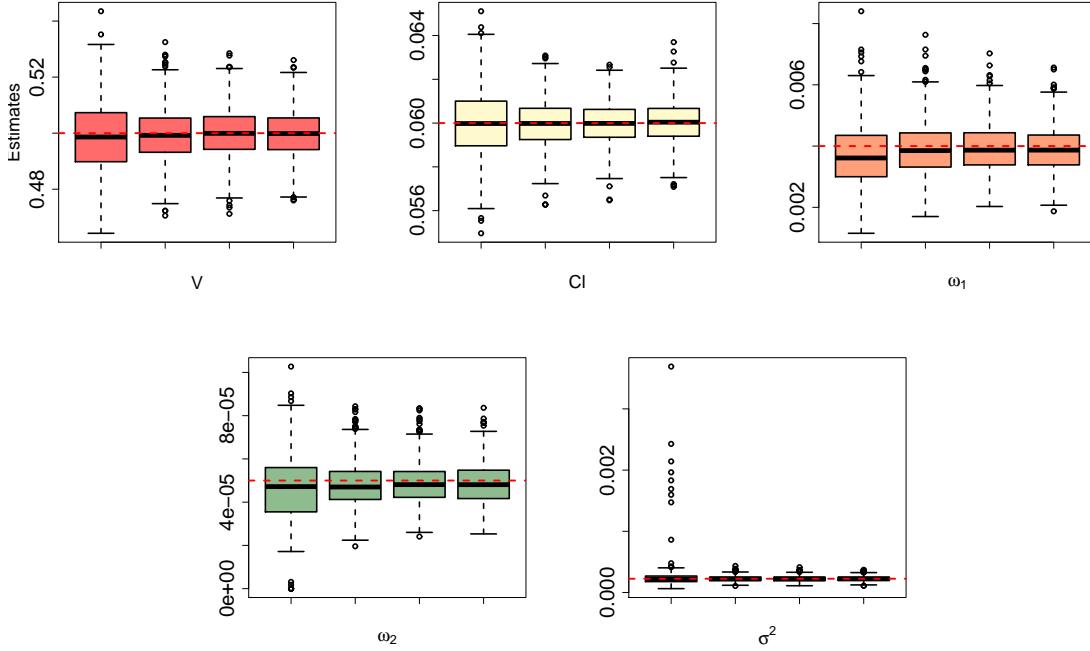


Figure 9: Boxplots of the PK parameter estimates obtained from the simulations. The horizontal dashed lines indicate the true parameter values. For each parameter, the successive boxes are for Scenarios 1, 2, 3 and 4, respectively.

To investigate the sensitivity of the PK sampling times to the assumed prior values, we define the relative D -efficiency of a design ξ_k^* to ξ_{true}^* as

$$\text{Relative Efficiency} = \left(\frac{|M(\hat{\Psi}_k, \xi_k^*)|}{|M(\Psi_{\text{true}}, \xi_{\text{true}}^*)|} \right)^{\frac{1}{p}}, \quad (9)$$

where ξ_k^* is the optimum design obtained at the k th stage of a trial using the current estimates of the parameters Ψ , ξ_{true}^* is the optimum design obtained for the true values of the parameters and p is the number of parameters in the model. In both cases, the dose remains fixed and it is the one administered to the cohort at the k th stage. Since both designs depend on the parameter values, it is possible, with completely different parameter values to the true ones, to have a design for which the relative efficiency is very high. A larger determinant of the information matrix means a smaller generalised variance of the estimators. But, in our case, that will mean that the variance is underestimated. Hence, we want the numerator to be close to the denominator in (9). This, in turn, means that we want to have an optimum design which is obtained for values around the true values of the

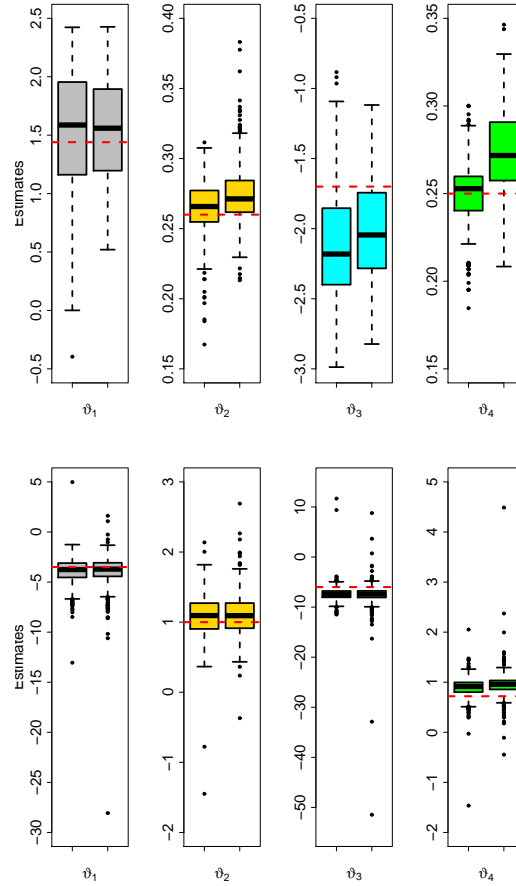


Figure 10: Boxplots of the dose-response parameter estimates obtained from the simulations, the top panel for Scenario 1 and the bottom panel for Scenario 2. The horizontal dashed lines indicate the true parameter values. For each parameter, the left box corresponds to the design which takes into account the AUC and the right box to the one which ignores it.

parameters. This can be achieved in a trial as we update the parameter estimates at each stage. After a sufficient number of stages, the estimates will be stable and so will be the design points: see Figure 12.

500 Figure 13 shows the relative D -efficiency of the designs computed at each of the stages using (9) for Scenario 2. Though not presented, we have found underestimated variances for the initial four stages. Recall that we use the up-and-down design for the first four cohorts and that blood samples for these cohorts are collected at the optimal time points which are based on an initial guess about the parameter values. From the fifth cohort onwards, we use the current estimates obtained from the
 505 trial data. Here, we observe a decreasing trend in the relative efficiency. Since the estimates stabilise as the trial proceeds, there is not much change at the later stages and also the efficiency approaches one.

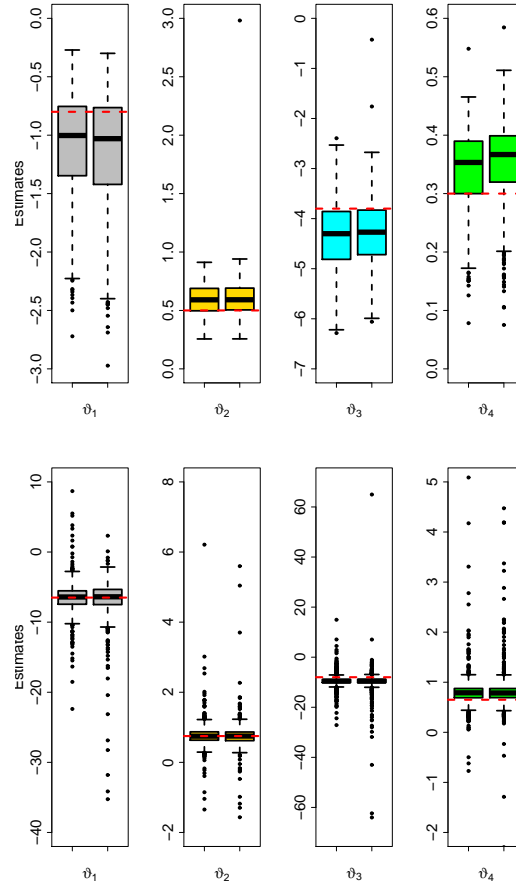


Figure 11: Boxplots of the dose-response parameter estimates obtained from the simulations, the top panel for Scenario 3 and the bottom panel for Scenario 4. The horizontal dashed lines indicate the true parameter values. For each parameter, the left box corresponds to the design which takes into account the AUC and the right box to the one which ignores it.

The PK guided design depends on the target value of the AUC. In our simulations so far, we have considered it as the one at the true OD. To assess the sensitivity of the design to the target value, we set it at doses other than the true OD. All of the scenarios are studied for this purpose. Table 2 gives a summary of the results. The notation in these tables is defined as follows: dose used for calculating AUC^0 (Dose), percentage of best doses recommended for further study (%BD), percentage of doses recommended as optimum, but carrying the probability of toxicity above the maximum allowed threshold (%TD), and percentage of cohorts treated at the best doses throughout the trials (%AD).

515

The figures in the tables indicate that the design is sensitive to the choice of target value for the




Figure 12: Optimal PK sampling time points in a trial. The left plot shows the points for the true values of the PK parameters and the one on the right gives the points using the current estimates.




Figure 13: Relative D -efficiency in a randomly selected trial for Scenario 2.

AUC. Scenarios 1 and 2 are more sensitive than Scenarios 3 and 4. As the target moves further away from that at the true OD, the %BD is affected. As reflected in Scenarios 2 and 3, if we choose a target below that at the true OD, the design will avoid recommending a toxic dose as the OD. It will also not allocate toxic doses to the cohorts often. But that may have a negative impact on the correct identification of the optimum dose, as in Scenario 2. Similarly, if the target is above that at the true OD, the design may not refrain from recommending a toxic dose as the OD. Scenario 3 is not that affected by the target AUC, as the dose-response curve is more flat here. Since this scenario has many doses with similar probabilities of efficacy, the target AUC chosen at any dose in the neighbourhood of the true OD will not affect the design much. In any of these four scenarios, a small misspecification of the target AUC should not significantly affect the outcome of the trials.

Table 2: Sensitivity of the design to the assumed target for the AUC.

Scenario	Dose	%BD	%TD	%AD
1	0.5	99.0	0.6	65.3
	1.0	44.7	0.0	40.5
	1.5	46.4	44.4	36.2
	2.0	50.5	37.6	35.6
	2.5	53.0	36.1	33.3
2	4.5	0.5	0.0	2.0
	5.0	72.7	0.1	33.6
	5.5	80.2	0.9	41.1
	6.0	72.0	3.5	34.2
	6.5	62.9	8.5	34.0
3	5.5	91.9	0.0	45.1
	6.0	91.1	0.0	49.8
	6.5	91.7	0.0	52.4
	7.0	86.4	0.0	51.6
	7.5	82.4	1.5	48.5
4	8.0	9.0	0.0	6.3
	8.5	29.5	0.0	13.2
	9.0	49.7	0.0	18.8
	9.5	48.7	0.0	19.3
	10.0	47.9	0.0	17.8

4. Discussion

The presented design is conceptually similar to that of Zhang et al. (2006), but their design does not incorporate PK information. They find a dose that maximises the difference between the estimated probability of efficacy and λ times the estimated probability of toxicity, given that the estimated probability of toxicity is smaller than a pre-specified level, where $0 \leq \lambda \leq 1$. The value of λ can be varied to weigh the importance of toxicity in the criterion, but no recommendation is made about its value. However, in many real scenarios, such a difference with a non-zero λ may lead to doses which are not optimum. Therefore, we decide to use $\lambda = 0$, which also helps to avoid double dependence on the probability of toxicity.

In this paper, along with dose-response outcomes, we have considered an important PK measure, AUC, and its inter-patient variability in the dose escalation. The main purpose of this study was to investigate the role of PK measures in dose finding and, by means of detailed comparisons, we showed
540 that utilising the PK information can be very beneficial. The original version of Zhang et al. (2006) is a fully model-based procedure. To make it comparable with our design, we have employed the up-and-down procedure in that design too. The up-and-down stage is required in our method to facilitate the maximum likelihood estimation of the PK model parameters. The fully model-based version of our design should be possible if we apply the Bayesian approach to PK parameter estimation as well.

545

The PK information is commonly collected in early clinical trials and it is often analysed for the purpose of dose selection. Here, we propose a systematic method with a two-fold aim: first, to obtain the best dose level for further studies in phase IIB or III trials with the minimum chance of toxic responses during the trial and, second, to obtain the most efficient estimates of the population PK
550 parameters. The second goal is achieved by using the population D -optimum design for the sampling times. A PK guided trial will make sense only if the PK information is accurate and we assure this by the choice of design.

The simulation results from four different dose-response scenarios indicate that the incorporation
555 of such measures can improve the accuracy of dose finding studies. It is also shown that the method is capable of limiting overdosing by a considerable amount depending on the location of the OD. The proposed PK guided approach can therefore be used in situations where more careful escalation is essential to avoid toxicity.

As indicated in the up-and-down stage, a trial is terminated if an unacceptable level of toxicity is
560 found at the very beginning. Once the model-based procedure applies, we stop early only if the same dose is repeated for r cohorts. At any stage, when the constraints are not satisfied, we allocate the lowest dose to the next cohort. In our design, it is possible to terminate a trial early for lack of efficacy by using an additional constraint that would ensure that the selected dose has at least a prespecified
565 level of efficacy. Such a constraint could be used after a reasonable number of stages, as using it from the beginning may force the design to stop before learning much about the dose-response relationship.

The priors that have been chosen for Bayesian estimation of the dose-response parameters are

vague enough to produce any possible extreme scenario. This reflects lack of knowledge of the shape
570 of the response curves. If it is possible to choose more informative priors, the results will improve
even further. The information on point prior values of the PK parameters is often available, and so
we apply maximum likelihood estimation for the PK model. However, these parameters are treated as
random to allow for the variability in the population.

575 We have assumed a single PK profile for each of the four dose-response scenarios. Though not
presented here, we have found Scenarios 1 and 4 to be insensitive to the assumed coefficient of vari-
ation for the PK parameters. The other two scenarios can be affected if we increase it greatly. As
seen in (4), the variance of the AUC depends on both the PK parameters and the dose. If the dose
is small, the increase in the variance components will have a slight effect on the variance of the AUC
580 compared to that if the dose is high. Since in Scenario 1, the true OD lies towards the beginning
of the dose region, the impact of a change in the variance components is negligible. As there is no
dose above 10 in Scenario 4, the impact of the PK profile is negligible for this scenario too. The
consideration of a 12% coefficient of variation is not unrealistic in that here both mean parameters
have the same coefficient of variation. The coefficient of variation in one parameter is very likely to be
585 different from the other in real situations and any increase will have less impact than what we have now.

Since the proposed design is for seamless phase I/II trials, assuming a maximum sample size of 60
patients is not unreasonable. Also, an early stopping rule is in use to stop a trial if convergence to
a dose occurs. A smaller sample size can be used if a good prior distribution for the dose-response
590 parameters exists. Also, we have assumed 20 available doses of the experimental drug to investigate
the properties of the method. In practice, a smaller number of dose levels may be available.

Implementation of the method requires a reliable value for the target AUC. Since we did not have
one for our example in Section 3, for simulation purposes, we set the target AUC as the one at the true
595 OD. In real trials for new drugs, it may be elicited from the experiences of the clinicians. Previous
studies of similar drugs or extrapolation from pre-clinical studies of the same drug could be some
possible options as well.

Further work is needed, so that other available information can also be used to benefit clinical
600 trials. Inclusion of covariates into the models or considering additional PD responses, such as changes
in some biomarkers, can potentially help in capturing the variability in the observed drug action and

so further improving the dose-selection trials. Other optimality criteria may be considered too. The approach presented in this paper would be the same, although the level of difficulty of implementing it might increase.

605 Acknowledgements

This work was carried out whilst the first author was in receipt of a Doctoral Training Account Studentship from Queen Mary University of London. The authors wish to thank two referees for their comments, which have led to a much improved paper.

Appendix A. Up-and-down design

Assume that we are at the k th stage of the up-and-down procedure of a trial and that the successive k cohorts have received doses from a pre-specified sequence of doses \mathcal{X} . Let us denote the proportion of toxic responses up to cohort k by \hat{p}_k , that is,

$$\hat{p}_k = \frac{1}{kc} \sum_{i=1}^k R_{i2},$$

610 where c is the cohort size and R_{i2} is the number of toxic responses for the i th cohort after receiving a dose. The algorithm starts with the lowest dose from a pre-specified sequence of ordered doses. Then for the given thresholds p_L , p_M and p_U , we increase, stay at the same dose level, decrease or stop the trial depending on the value of \hat{p}_k . In the simulation study, we set $p_L = \gamma/3$, $p_M = 2\gamma/3$ and $p_U = \gamma$, where γ is the maximum acceptable level for the probability of toxicity. More specifically, the
 615 algorithm has the following structure:

$$\hat{p}_k \begin{cases} \leq p_L & \text{increase the dose to the next level if not at the highest} \\ & \text{level, otherwise stay at the highest level,} \\ \in (p_L, p_M) & \text{stay at the current dose,} \\ \in [p_M, p_U) & \text{decrease the dose by one level if not at the lowest} \\ & \text{level, otherwise stay at the lowest level,} \\ \geq p_U & \text{stop the trial.} \end{cases}$$

Although we have chosen the thresholds so that the width is $\gamma/3$, other choices can be made as recommended by a clinician.

Appendix B. Fisher information matrix for the population PK model

The Taylor series expansion gives the approximate function as

$$\begin{aligned}
 f(\boldsymbol{\theta}_i, t_{ij}) &\cong f(\boldsymbol{\theta}_i, t_{ij}) \Big|_{\boldsymbol{\phi}^0} + \left(\frac{\partial f(\boldsymbol{\theta}_i, t_{ij})}{\partial \boldsymbol{\phi}_i} \right)^{\text{T}} \Big|_{\boldsymbol{\phi}^0} (\boldsymbol{\phi}_i - \boldsymbol{\phi}^0) \\
 &= f(\boldsymbol{\theta}_i, t_{ij}) \Big|_{\boldsymbol{\phi}^0} + \left(\frac{\partial f(\boldsymbol{\theta}_i, t_{ij})}{\partial \boldsymbol{\beta}} \right)^{\text{T}} \Big|_{\boldsymbol{\phi}^0} (\boldsymbol{\beta} - \boldsymbol{\beta}^0) + \left(\frac{\partial f(\boldsymbol{\theta}_i, t_{ij})}{\partial \mathbf{b}_i} \right)^{\text{T}} \Big|_{\boldsymbol{\phi}^0} (\mathbf{b}_i - \mathbf{0}) \\
 &= \mu_{ij} + \left(\frac{\partial f(\boldsymbol{\theta}_i, t_{ij})}{\partial \boldsymbol{\beta}} \right)^{\text{T}} \Big|_{\boldsymbol{\phi}^0} \boldsymbol{\beta} + \left(\frac{\partial f(\boldsymbol{\theta}_i, t_{ij})}{\partial \mathbf{b}_i} \right)^{\text{T}} \Big|_{\boldsymbol{\phi}^0} \mathbf{b}_i,
 \end{aligned}$$

where μ_{ij} is a constant. Writing this in matrix notation for n_i observations, we have a linear mixed effects model for individual i given by

$$\mathbf{y}_i \cong \boldsymbol{\mu}_i + \mathbf{H}_i \boldsymbol{\beta} + \mathbf{L}_i \mathbf{b}_i + \boldsymbol{\epsilon}_i,$$

where $\mathbf{H}_i = (\partial f(\boldsymbol{\theta}_i, t_{ij}) / \partial \boldsymbol{\beta})^{\text{T}} \Big|_{\boldsymbol{\phi}^0}$ and $\mathbf{L}_i = (\partial f(\boldsymbol{\theta}_i, t_{ij}) / \partial \mathbf{b}_i)^{\text{T}} \Big|_{\boldsymbol{\phi}^0}$ are $(n_i \times p)$ -dimensional matrices, $\boldsymbol{\mu}_i$ is the vector of constants μ_{ij} and $\boldsymbol{\epsilon}_i$ is the vector of random errors ϵ_{ij} . This gives

$$\mathbf{E}_i = \mathbb{E}(\mathbf{y}_i) \cong \boldsymbol{\mu}_i + \mathbf{H}_i \boldsymbol{\beta},$$

and, by the assumption of independence of \mathbf{b}_i and $\boldsymbol{\epsilon}_i$, we have

$$\mathbf{V}_i = \text{Var}(\mathbf{y}_i) \cong \mathbf{L}_i \boldsymbol{\Omega} \mathbf{L}_i^{\text{T}} + \boldsymbol{\Sigma}_i,$$

620 where $\boldsymbol{\Sigma}_i$ is an $n_i \times n_i$ diagonal matrix with all the diagonal elements equal to σ^2 .

Since \mathbf{b}_i and $\boldsymbol{\epsilon}_i$ are assumed to be normal, the log-likelihood function is approximated by

$$\ell_i(\boldsymbol{\Psi} \mid \mathbf{y}_i) \cong \log \left[(2\pi)^{-\frac{n_i}{2}} |\mathbf{V}_i|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} (\mathbf{y}_i - \mathbf{E}_i)^{\text{T}} \mathbf{V}_i^{-1} (\mathbf{y}_i - \mathbf{E}_i) \right\} \right].$$

Then the FIM for individual i can be approximated by the block diagonal matrix

$$\mathbf{M}_i(\boldsymbol{\Psi}, \boldsymbol{\xi}_i) \cong \begin{bmatrix} \mathbf{A}_i & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_i \end{bmatrix},$$

where elements (m, n) of \mathbf{A}_i and \mathbf{B}_i have the forms

$$(\mathbf{A}_i)_{mn} = \left(\frac{\partial \mathbf{E}_i}{\partial \beta_m} \right)^{\text{T}} \mathbf{V}_i^{-1} \frac{\partial \mathbf{E}_i}{\partial \beta_n} \quad \text{for } m, n = 1, \dots, p$$

and

$$(\mathbf{B}_i)_{mn} = \frac{1}{2} \text{tr} \left(\frac{\partial \mathbf{V}_i}{\partial \lambda_m} \mathbf{V}_i^{-1} \frac{\partial \mathbf{V}_i}{\partial \lambda_n} \mathbf{V}_i^{-1} \right) \quad \text{for } m, n = 1, \dots, p+1.$$

Appendix C. Calculation of the approximate mean and variance of the AUC

625 The first-order Taylor series expansion of $h(x, \boldsymbol{\theta}_i)$ about $\boldsymbol{\theta}_i$ at $E(\boldsymbol{\theta}_i)$ gives

$$\begin{aligned} h(x, \boldsymbol{\theta}_i) &\cong h(x, \boldsymbol{\theta}_i) \Big|_{E(\boldsymbol{\theta}_i)} + \left(\frac{\partial h(x, \boldsymbol{\theta}_i)}{\partial \boldsymbol{\theta}_i} \right)^\top \Big|_{E(\boldsymbol{\theta}_i)} (\boldsymbol{\theta}_i - E(\boldsymbol{\theta}_i)) \\ &= h(x, \boldsymbol{\beta}) + \left(\frac{\partial h(x, \boldsymbol{\theta}_i)}{\partial \boldsymbol{\theta}_i} \right)^\top \Big|_{E(\boldsymbol{\theta}_i)} \mathbf{b}_i, \end{aligned}$$

where $E(\boldsymbol{\theta}_i) = \boldsymbol{\beta}$. Therefore,

$$E\{h(x, \boldsymbol{\theta}_i)\} \cong h(x, \boldsymbol{\beta})$$

and

$$\text{Var}\{h(x, \boldsymbol{\theta}_i)\} \cong \left(\frac{\partial h(x, \boldsymbol{\theta}_i)}{\partial \boldsymbol{\theta}_i} \right)^\top \Big|_{E(\boldsymbol{\theta}_i)} \boldsymbol{\Omega} \left(\frac{\partial h(x, \boldsymbol{\theta}_i)}{\partial \boldsymbol{\theta}_i} \right) \Big|_{E(\boldsymbol{\theta}_i)}.$$

For our PK model, the AUC for an individual i over the range $[0, t_1]$ is defined as

$$\begin{aligned} h(x, \boldsymbol{\theta}_i) &= \int_0^{t_1} f(\boldsymbol{\theta}_i, t) dt \\ &= \int_0^{t_1} \frac{x}{V_i} \exp\left(-\frac{Cl_i}{V_i} t\right) dt \\ &= \frac{x}{Cl_i} \left\{ 1 - \exp\left(-\frac{Cl_i}{V_i} t_1\right) \right\}. \end{aligned}$$

Assuming that $E(V_i) = V$ and $E(Cl_i) = Cl$, we obtain

$$E\{h(x, \boldsymbol{\theta}_i)\} \cong h(x, \boldsymbol{\beta}) = \frac{x}{Cl} \left\{ 1 - \exp\left(-\frac{Cl}{V} t_1\right) \right\}$$

and

$$\begin{aligned} \text{Var}\{h(x, \boldsymbol{\theta}_i)\} &\cong \left(\frac{\partial h(x, \boldsymbol{\theta}_i)}{\partial \boldsymbol{\theta}_i} \right)^\top \Big|_{E(\boldsymbol{\theta}_i)} \boldsymbol{\Omega} \left(\frac{\partial h(x, \boldsymbol{\theta}_i)}{\partial \boldsymbol{\theta}_i} \right) \Big|_{E(\boldsymbol{\theta}_i)} \\ &= \left(\begin{array}{cc} \frac{\partial h(x, \boldsymbol{\theta}_i)}{\partial V_i} & \frac{\partial h(x, \boldsymbol{\theta}_i)}{\partial Cl_i} \end{array} \right)_{E(\boldsymbol{\theta}_i)} \begin{pmatrix} \omega_1 & 0 \\ 0 & \omega_2 \end{pmatrix} \begin{pmatrix} \frac{\partial h(x, \boldsymbol{\theta}_i)}{\partial V_i} \\ \frac{\partial h(x, \boldsymbol{\theta}_i)}{\partial Cl_i} \end{pmatrix} \Big|_{E(\boldsymbol{\theta}_i)} \\ &= \left\{ \frac{\partial h(x, \boldsymbol{\theta}_i)}{\partial V_i} \Big|_{E(\boldsymbol{\theta}_i)} \right\}^2 \omega_1 + \left\{ \frac{\partial h(x, \boldsymbol{\theta}_i)}{\partial Cl_i} \Big|_{E(\boldsymbol{\theta}_i)} \right\}^2 \omega_2. \end{aligned}$$

Also, we have

$$\frac{\partial h(x, \boldsymbol{\theta}_i)}{\partial V_i} \Big|_{E(\boldsymbol{\theta}_i)} = -\frac{x t_1}{V^2} \exp\left(-\frac{Cl}{V} t_1\right)$$

630 and

$$\left. \frac{\partial h(x, \boldsymbol{\theta}_i)}{\partial Cl_i} \right|_{\mathbb{E}(\boldsymbol{\theta}_i)} = \frac{x}{Cl} \exp\left(-\frac{Cl}{V}t_1\right) \left(\frac{1}{Cl} + \frac{t_1}{V}\right) - \frac{x}{Cl^2}.$$

It follows that

$$\text{Var}\{h(x, \boldsymbol{\theta}_i)\} \cong \left\{-\frac{xt_1}{V^2} \exp\left(-\frac{Cl}{V}t_1\right)\right\}^2 \omega_1 + \left\{\frac{x}{Cl} \exp\left(-\frac{Cl}{V}t_1\right) \left(\frac{1}{Cl} + \frac{t_1}{V}\right) - \frac{x}{Cl^2}\right\}^2 \omega_2.$$

References

- Agresti A. *Categorical Data Analysis*. Wiley, New York, 1990.
- Atkinson A, Donev A, Tobias R. *Optimum Experimental Designs, with SAS*. Oxford University Press,
635 Oxford, 2007.
- Bazzoli C, Nguyen T, Dubois A, Retout S, Comets E, Mentré F. *PFIM 3.2 User Guide*; 2010. Université Paris Diderot and INSERM.
- Davidian M. *An Introduction to Nonlinear Mixed Effects Models and PK/PD Analysis*; 2010. American Statistical Association Biopharmaceutical Section webinar. <http://www4.stat.ncsu.edu/~davidian/webinar.pdf>.
640
- Dragalin V, Fedorov V. Adaptive designs for dose-finding based on efficacy–toxicity response. *Journal of Statistical Planning and Inference* 2006;136(6):1800–23.
- Fedorov V. *Theory of Optimal Experiments*. Academic Press, New York, 1972.
- Fedorov V, Hackl P. *Model-Oriented Design of Experiments*. Lecture Notes in Statistics. Springer,
645 New York, 1997.
- Govindarajulu Z. *Statistical Techniques in Bioassay*. Karger, Basel, 1988.
- Hooker A, Vicini P. Simultaneous population optimal design for pharmacokinetic-pharmacodynamic experiments. *The AAPS Journal* 2005;7:E759–85.
- Ivanova A. Dose-finding in oncology - nonparametric methods. In: Ting N, editor. *Dose Finding in Drug Development*. Springer, New York; 2006. p. 49–58.
650
- Johnson S, Narasimhan B. *Cubature: Adaptive Multivariate Integration Over Hypercubes*. 2009.
- Jönsson S. Estimation of dosing strategies for individualisation. Ph.D. thesis; Acta Universitatis Upsaliensis; 2004.

- Li Z, Durham S, Flournoy N. An adaptive design for maximization of a contingent binary response. In:
655 Flournoy N, Rosenberger W, editors. Adaptive Designs. Hayward, CA: Institute of Mathematical
Statistics; 1995. p. 179–96.
- Meibohm B, Derendorf H. Basic concepts of pharmacokinetic/pharmacodynamic (PK/PD) modelling.
International Journal of Clinical Pharmacology and Therapeutics 1997;35(10):401–13.
- Mentré F, Mallet A, Baccar D. Optimal design in random-effects regression models. *Biometrika*
660 1997;84(2):429–42.
- Mielke T. Approximation of the Fisher Information and Design in Nonlinear Mixed Effects Models.
Ph.D. thesis; Otto-von-Guericke University Magdeburg; 2012.
- Nyberg J, Karlsson M, Hooker A. Simultaneous optimal experimental design on dose and sample
times. *Journal of Pharmacokinetics and Pharmacodynamics* 2009;36:125–45.
- 665 Piantadosi S, Liu G. Improved designs for dose escalation studies using pharmacokinetic measurements.
Statistics in Medicine 1998;15(15):1605–18.
- Pinheiro J, Bates D. Approximations to the log-likelihood function in the nonlinear mixed-effects
model. *Journal of Computational and Graphical Statistics* 1995;4(1):12–35.
- Pinheiro J, Bates D. *Mixed-Effects Models in S and S-PLUS*. Springer, New York, 2000.
- 670 Retout S, Duffull S, Mentré F. Development and implementation of the population fisher information
matrix for the evaluation of population pharmacokinetic designs. *Computer Methods and Programs
in Biomedicine* 2001;65(2):141–51.
- Riviere JE. *Comparative Pharmacokinetics: Principles, Techniques, and Applications*. Wiley-
Blackwell, 2011.
- 675 Thall P, Cook J. Dose-finding based on efficacy–toxicity trade-offs. *Biometrics* 2004;60(3):684–93.
- Thall P, Nguyen H. Adaptive randomization to improve utility-based dose-finding with bivariate
ordinal outcomes. *Journal of Biopharmaceutical Statistics* 2012;22(4):785–801.
- Thall P, Nguyen H, Estey E. Patient-specific dose finding based on bivariate outcomes and covariates.
Biometrics 2008;64(4):1126–36.
- 680 Thall P, Russell K. A strategy for dose-finding and safety monitoring based on efficacy and adverse
outcomes in phase I/II clinical trials. *Biometrics* 1998;54(1):251–64.

Wakefield J, Smith A, Racine-Poon A, Gelfand A. Bayesian analysis of linear and non-linear population models by using the gibbs sampler. *Applied Statistics* 1994;;201–21.

Whitehead J, Zhou Y, Hampson L, Ledent E, Pereira A. A Bayesian approach for dose-escalation
685 in a phase I clinical trial incorporating pharmacodynamic endpoints. *Journal of Biopharmaceutical
Statistics* 2007;17(6):1117–29.

Zhang W, Sargent D, Mandrekar S. An adaptive dose-finding design incorporating both toxicity and efficacy. *Statistics in Medicine* 2006;25(14).

Zhou Y, Whitehead J, Korhonen P, Mustonen M. Implementation of a Bayesian design in a dose-
690 escalation study of an experimental agent in healthy volunteers. *Biometrics* 2008;64(1):299–308.