Published in final edited form as:

Pharm Stat. 2017 March ; 16(2): 143–156. doi:10.1002/pst.1797.

A Simulation Study of Methods for Selecting Subgroup-Specific Doses in Phase I Trials

Satoshi Morita¹, Peter F. Thall², and Kentaro Takeda^{3,4}

¹Department of Biomedical Statistics and Bioinformatics, Kyoto University Graduate School of Medicine, Kyoto, Japan

²Department of Biostatistics, The University of Texas M. D. Anderson Cancer Center, Houston, TX, U.S.A

³Department of Biostatistics and Epidemiology, Yokohama City University, Yokohama, Japan

⁴Biostatistics Group, Data Science, Global Development, Astellas Pharma Inc., Tokyo, Japan

Summary

Patient heterogeneity may complicate dose-finding in phase I clinical trials if the dose-toxicity curves differ between subgroups. Conducting separate trials within subgroups may lead to infeasibly small sample sizes in subgroups having low prevalence. Alternatively, it is not obvious how to conduct a single trial while accounting for heterogeneity. To address this problem, we consider a generalization of the continual reassessment method (O'Quigley, et al., 1990) based on a hierarchical Bayesian dose-toxicity model that borrows strength between subgroups under the assumption that the subgroups are exchangeable. We evaluate a design using this model that includes subgroup-specific dose selection and safety rules. A simulation study is presented that includes comparison of this method to three alternative approaches, based on non-hierarchical models, that make different types of assumptions about within-subgroup dose-toxicity curves. The simulations show that the hierarchical model-based method is recommended in settings where the dose-toxicity curves are exchangeable between subgroups. We present practical guidelines for application, and provide computer programs for trial simulation and conduct.

Keywords

Bayesian study design; Conditionally independent hierarchical model; Continual reassessment method; Subgroup-specific dose-finding; Phase I clinical trial

1 Introduction

Patient heterogeneity may complicate phase I clinical trials in oncology. The goal may be either to determine a single optimal dose, or possibly different optimal doses within

*Address for correspondence: Satoshi Morita, PhD, Department of Biomedical Statistics and Bioinformatics, Kyoto University Graduate School of Medicine, Address: 54 Kawahara-cho, Shogoin, Sakyo-ku, Kyoto 606-8507, Japan, smorita@kuhp.kyoto-u.ac.jp, Phone: +81-75-751-4717, Fax: +81-75-751-4767.

Conflict of interest statement: There are no conflicts of interest in this study.

subgroups. For example, subgroups may be determined by disease subtypes, biomarkers targeted by the agent being studied, or known prognostic variables. Preclinical and clinical data often suggest that patient subgroups may have different dose-toxicity relationships, but the order in the tolerability of the subgroups is not known. Conventionally, even if such subgroups have been identified, most often a phase I trial ignoring subgroups is conducted. If the dose-toxicity curves differ between subgroups, however, a single dose chosen for all subgroups may be either sub-therapeutic or excessively toxic in some subgroups. Moreover, ignoring subgroup effects during the dose-finding process may lead to undesirable interim dose assignments and adaptive decisions. An alternative approach is to conduct a separate dose-finding study within each subgroup. This may not be feasible in subgroups for which the prevalence is too low to reliably identify an optimal dose within a reasonable time frame. A more refined approach is to conduct one trial with the goal to find optimal doses that may differ between subgroups. Ideally, the underlying model should borrow strength between subgroups so that the data obtained from each subgroup may help inform the decisions in the other subgroups.

We apply a hierarchical logistic regression model used in Morita, *et al.* [1] to illustrate the computation of a prior equivalent sample size in hierarchical models. The hierarchical model-based method generalizes the continual reassessment method (CRM) proposed by O'Quigley, *et al.* [2] by allowing different doses to be chosen within subgroups, while borrowing strength between subgroups. In this paper, we use the same model as Morita, *et al.* [1] and develop in detail its uses for subgroup-specific dose finding. We review the model and dose-finding method, which we call the HB-CRM, compare it to three alternative approaches, each based on a non-hierarchical model, give practical guidelines, and provide a computer program for simulation and trial conduct. All four methods considered here address the problem of determining an optimal dose, or optimal subgroup-specific doses, based on toxicity in settings where K subgroups have been identified. Each method uses a CRM-type criterion for optimality. For each subgroup, indexed by $k = 1, \dots, K$, denote the probability of toxicity with dose x by $\pi_k(x, \theta_k)$, where θ_k is the model's parameter vector. The HB-CRM assumes a hierarchical structure for $\theta_1, \dots, \theta_K$, which implies *a priori* that $\pi_1(x, \theta_1), \dots, \pi_K(x, \theta_K)$ are exchangeable for each x , and conducts a single trial including all subgroups. Each of the three alternative model-based comparators relies on a non-hierarchical model. The first alternative completely ignores subgroups and conducts a single trial using the same logistic dose-toxicity model for all subgroups. That is, it is the CRM based on a logistic dose-toxicity model. The second alternative assumes K different subgroup-specific models and conducts K separate trials. The third alternative conducts one trial, assuming a dose-toxicity model with K different parameters θ_k to account for the inter-subgroup variability, and allows different optimal doses to be chosen within subgroups.

To help motivate the problem, it is worthwhile to consider a simple example in which there are $K = 3$ subgroups, with true dose-toxicity probability curves $\pi_1(x)^{true}$, $\pi_2(x)^{true}$, $\pi_3(x)^{true}$, given by Figure 1. If the aim of a phase I trial is to find a dose having mean toxicity probability $\pi^* = .30$ then, as shown by Figure 1, the true optimal doses are different for the three subgroups. Any method that finds one optimal dose x^{opt} ignores this possibility, and

giving the same x^{opt} to all patients has the consequence of underdosing patients in subgroup 1, and overdosing patients in subgroup 3.

Several authors have addressed the problem of accounting for patient heterogeneity in phase I trials. O'Quigley, Shen and Gamst [3] and O'Quigley and Paoletti [4] proposed a parametric model-based two-sample CRM to find the optimal dose for each of two possibly ordered subpopulations of patients. Ivanova and Wang [5] proposed a non-parametric design with bivariate isotonic regression to address the same problem. Yuan and Chappell [6] compared three dose-finding methods, which respectively extended the up-and-down design [7], the CRM [2], and the isotonic design [8], to deal with multiple risk subgroups which can be ordered according to their risk of toxicity. Thall, Nguyen and Estey [9] proposed a phase I-II design including covariates accounting for patient subpopulations based on a trade-off between efficacy and toxicity. Liu *et al.* [10] proposed an extended CRM with multiple skeletons of toxicity probabilities to deal with dose-finding in different ethnic populations. All of these approaches assume that the probability of toxicity is monotonically ordered for subgroups, so that in particular the subgroups are not exchangeable and do not represent qualitatively different subgroups for which no ordering may be assumed. In particular, O'Quigley, Shen and Gamst [3] reported that the two-sample CRM in this setting was comparable to conducting two independent trials for each of the two subgroups separately using the one-sample CRM. This comparison is similar to that between our second and third alternatives. The main objective of our simulation study is to determine advantages and disadvantages of the hierarchical model-based method, in comparison with these alternative approaches, in a setting with more than two non-ordered subgroups. We also examine the effects of the subgroup proportions, or prevalences, on how each of the methods behaves.

In Section 2, we present probability models and prior specification for dose-finding methods to account for patient heterogeneity. We evaluate the operating characteristics of HB-CRM and each of the three alternative methods by simulation in Section 3. Section 4 gives guidelines for constructing study designs. We close with a brief discussion in Section 5.

2 Dose-finding methods

2.1 Preliminaries

Denote the population proportions (prevalences) of the K subgroups by $\xi=(\xi_1, \dots, \xi_K)$, that is, a patient belongs to subgroup k with probability ξ_k . In the trial, each patient in each subgroup receives one of J doses, denoted by $d_1 < \dots < d_J$. We formulate the models using

standardized doses $x_j = \log(d_j) - J^{-1} \sum_{l=1}^J \log(d_l)$. For the i^{th} patient in subgroup k , denote the assigned dose by $x_{[k,i]}$, the indicator $Y_{k,i} = 1$ if the patient suffers toxicity, 0 if not, and the toxicity probability

$$\pi_k(x_{[k,i]}, \theta_k) = Pr(Y_{k,i}=1 | x_{[k,i]}, \theta_k), \quad k=1, \dots, K.$$

At any given point during the trial, let $\mathcal{D}_n$ denote the data for the first n patients and $\mathbf{n} = (n_1, n_2, \dots, n_K)$ the current sample sizes within the subgroups, so $n = n_1 + \dots + n_K$. Let $\mathbf{N} = (N_1, N_2, \dots, N_K)$ denote the final subgroup sample sizes at the end of a trial, so the final total

sample size is $N = N_1 + \dots + N_K$. Given a planned maximum total sample size, N_{max} , due to the use of early stopping rules it may be the case that $N < N_{max}$. An important point, which will play a central role in determining the properties of the designs, is that both $\mathbf{n}$ and $\mathbf{N}$ depend on $\boldsymbol{\xi} = (\xi_1, \dots, \xi_K)$, as well as the particular design being used. Given N , temporarily ignoring the effects of early stopping, the expected final subgroup sample sizes are $\xi_1 N, \dots, \xi_K N$. Due to both random variation and the use of adaptive rules, however, each achieved N_k may differ substantially from its mean.

2.2 Hierarchical Bayesian CRM (HB-CRM)

In the HB-CRM (Morita, *et al.*) [1], the parameter vector in subgroup k is $\boldsymbol{\theta}_k = (\alpha_k, \beta)$ and model's linear components are

$$\text{logit}\{\pi_k(x_{[k,i]}, \alpha_k, \beta)\} = \alpha_k + \beta x_{[k,i]} \quad (1)$$

for $k = 1, \dots, K$. For Level 1 priors, it is assumed that $\alpha_1, \dots, \alpha_K$ are i.i.d. $N(\tilde{\mu}_\alpha, \tilde{\sigma}_\alpha^2)$ and that β follows a $N(\tilde{\mu}_\beta, \tilde{\sigma}_\beta^2)$ prior. For Level 2 priors (hyperpriors), it is assumed that $\tilde{\mu}_\alpha$ follows a normal hyperprior and, following the recommendation of Gelman [11], $\tilde{\sigma}_\alpha$ follows a uniform prior on the interval .01 to U_ϕ denoted $\tilde{\sigma}_\alpha \sim U(0.01, U_\phi)$. In summary, the model assumptions are as follows :

Sampling model	$Y_{k,i} \sim \text{Bernoulli}(\pi_k(x_{[k,i]}, \alpha_k, \beta))$ indep. for all k
Priors	$\alpha_k \sim \text{i.i.d. } N(\tilde{\mu}_\alpha, \tilde{\sigma}_\alpha^2)$ for all k
	$\beta \sim N(\tilde{\mu}_\beta, \tilde{\sigma}_\beta^2)$
Hyperpriors	$\tilde{\mu}_\alpha \sim N(\mu_{\alpha,\phi}, \sigma_{\alpha,\phi}^2)$
	$\tilde{\sigma}_\alpha \sim U(0.01, U_\phi)$.

(2)

We do not impose the constraint that $\beta > 0$ with probability 1 to ensure that each $\pi_k(x, \alpha_k, \beta)$ increases in x or, alternatively, assume a lognormal prior for β . In practice, appropriate calibration of the hyperparameters $\tilde{\mu}_\beta$ and $\tilde{\sigma}_\beta^2$ in (2) ensures this monotonicity. Specifically, if $\tilde{\mu}_\beta$ is a large enough positive value and $\tilde{\sigma}_\beta^2$ is sufficiently small, then no constraint on β is needed since all computed posterior values of β will be positive.

While a hierarchical prior structure is assumed for $(\alpha_1, \dots, \alpha_K)$, the dose effect parameter β shared by all K marginal toxicity probabilities has a usual prior without an additional hyperprior on its hyperparameters $(\tilde{\mu}_\beta, \tilde{\sigma}_\beta^2)$, which are fixed. Collecting terms, the $K + 1$ sampling model parameters that characterize the marginal probabilities of toxicity in the K subgroups are $\boldsymbol{\theta} = (\alpha_1, \dots, \alpha_K, \beta)$, the four hyperparameters that characterize the priors on $\boldsymbol{\theta}$ are $\tilde{\boldsymbol{\theta}} = (\tilde{\mu}_\alpha, \tilde{\sigma}_\alpha, \tilde{\mu}_\beta, \tilde{\sigma}_\beta^2)$, and the fixed hyperparameters that characterize the hyperpriors on $(\tilde{\mu}_\alpha, \tilde{\sigma}_\alpha^2)$ are $\boldsymbol{\phi} = (\mu_{\alpha,\phi}, \sigma_{\alpha,\phi}, .01, U_\phi)$. Consequently, to complete the Bayesian model one must specify numerical values for a total of five parameters, the two hyperparameters $(\tilde{\mu}_\beta, \tilde{\sigma}_\beta^2)$ and for the three fixed hyperprior parameters $(\mu_{\alpha,\phi}, \sigma_{\alpha,\phi}, U_\phi)$. In expression (2) for the hierarchical model, “Priors” may be called “Level 1 Priors” since they are distributions

on the parameters θ of the sampling model, while “Hyperpriors” may be called “Level 2 Priors” since they are distributions on the two parameters $(\tilde{\mu}_\alpha, \tilde{\sigma}_\alpha^2)$ of the Level 1 priors.

Under this hierarchical model, *a priori*, the parameter vectors $(\theta_1, \dots, \theta_K)$ of the marginal toxicity probability models for the K subgroups are exchangeable. This is the property that, if the subgroup indices $(1, \dots, K)$ are replaced by any permutation $(j_1, \dots, j_K)$, then the joint prior distribution of the permuted vector $(\theta_{j_1}, \dots, \theta_{j_K})$ is the same as that of $(\theta_1, \dots, \theta_K)$. Exchangeability is a useful property for probability models of random quantities corresponding to qualitatively different objects, where the order in which the objects are indexed to identify them is arbitrary. A well known special case is a vector of random quantities that are independent and identically distributed (iid), which trivially must be exchangeable. The joint distribution of an exchangeable random vector is more general, however, since the random quantities need not be independent. The prior exchangeability of $(\theta_1, \dots, \theta_K)$ implies that, *a priori*, the toxicity probabilities $\{\pi_1(x, \theta_1), \dots, \pi_K(x, \theta_K)\}$ are exchangeable for each x .

The hierarchical model given above induces prior association among $\alpha_1, \dots, \alpha_K$. To see this, denoting the Level 1 priors by p_1 and Level 2 priors (hyperpriors) by p_2 , we obtain the unconditional prior of the sampling model parameters by averaging over the hyperprior $p_2(\tilde{\mu}_\alpha, \tilde{\sigma}_\alpha | \phi)$, as

$$\begin{aligned} & p(\beta, \alpha_1, \dots, \alpha_K | \tilde{\mu}_\beta, \tilde{\sigma}_\beta, \phi) \\ &= p_1(\beta | \tilde{\mu}_\beta, \tilde{\sigma}_\beta) \int \prod_{k=1}^K p_1(\alpha_k | \tilde{\mu}_\alpha, \tilde{\sigma}_\alpha) p_2(\tilde{\mu}_\alpha, \tilde{\sigma}_\alpha | \phi) d\tilde{\mu}_\alpha d\tilde{\sigma}_\alpha \\ &= p_1(\beta | \tilde{\mu}_\beta, \tilde{\sigma}_\beta) p_{1,2}(\alpha_1, \dots, \alpha_K | \phi), \end{aligned} \quad (3)$$

where $p_{1,2}$ denotes the marginal prior of $(\alpha_1, \dots, \alpha_K)$ obtained by averaging over the hyperprior of $(\tilde{\mu}_\alpha, \tilde{\sigma}_\alpha)$. The prior association among $\alpha_1, \dots, \alpha_K$ in $p_{1,2}$ induces prior association among the toxicity probabilities $\pi_1(x, \theta_1), \dots, \pi_K(x, \theta_K)$ for any x . In this regard, it also is important to note that the dose effect parameter β shared by these probabilities also induces positive association among them. These prior associations shrink the posteriors of the $\pi_k(x, \theta_k)$'s toward each other. Expression (3) also shows why numerical values of $(\tilde{\mu}_\beta, \tilde{\sigma}_\beta, \phi)$ must be specified to complete the model.

Medical settings where this hierarchical model is appropriate include trials in which there are K qualitatively different disease subtypes, different solid tumors types, or subgroups defined by biomarkers. The common feature is that one's prior uncertainty about $\theta_1, \dots, \theta_K$ would not be changed if the subgroups were re-indexed in a different order. In contrast, for example, the hierarchical model is not appropriate if the subgroups correspond to prognostic risk of toxicity from any agent, such as Good (low risk, $k = 1$), Intermediate ($k = 2$), and Poor (high risk, $k = 3$). In this case, $\pi_k(x, \theta_k)$ is stochastically increasing in k for any dose x , and the exchangeability assumption is not valid.

For each decision during trial conduct, the HB-CRM defines the optimal dose $x_{[k]}^{opt}$ in subgroup k to be that for which the posterior mean of $\pi_k(x_j, \theta_k)$ is closest to a given fixed target, π^* . Formally, given $\mathcal{D}_n$, the dose chosen for subgroup k is

$$x_{[k]}^{opt} = \arg \min_{j=1, \dots, J} |E\{\pi_k(x_j, \alpha_k, \beta) | \mathcal{D}_n\} - \pi^*|. \quad (4)$$

If desired, different target values for the subgroups may be used, although we will not explore that case here.

A safety rule imposed on the method is that, within each subgroup, the HB-CRM may not skip an untried doses when escalating. In addition, to control overdosing, HB-CRM does not escalate within subgroup k if

$$Pr(\pi_k(x_j, \theta_k) > \pi^{odc} | \mathcal{D}_n) > \psi^{odc}$$

where π^{odc} is a fixed upper limit and ψ^{odc} is a probability cutoff. This rule supersedes the criterion (4) for dose-escalation. The design parameter ψ^{odc} must be calibrated along with the prior parameters to obtain a design with desirable operating characteristics.

The HB-CRM in (2) relies on the prior assumption that the toxicity probabilities of the patient subgroups are exchangeable, since $\alpha_1, \dots, \alpha_K$ are conditionally i.i.d., given $\tilde{\mu}_\alpha$ and $\tilde{\sigma}_\alpha^2$. Since one must average over the hyperpriors of $\tilde{\mu}_\alpha$ and $\tilde{\sigma}_\alpha^2$ to compute posteriors, given the observed data $\mathcal{D}_n$ the intercept parameters $\alpha_1, \dots, \alpha_K$ are positively correlated. This, and the fact that the $\pi_k(x_j, \alpha_k, \beta)$'s share the common slope parameter β , induces positive correlation among all KJ toxicity probabilities. This induces association among the posterior means in (4), which in turn shrinks the chosen doses $x_{[1]}^{opt}, \dots, x_{[K]}^{opt}$ toward each other. In this way, conducting a single trial with this hierarchical model provides a basis for borrowing strength across patient subgroups.

The hyperpriors on $\tilde{\mu}_\alpha$ and $\tilde{\sigma}_\alpha^2$ play key roles in how the HB-CRM design behaves. Thus, their fixed parameters, $\mu_{\alpha, \phi}$, $\sigma_{\alpha, \phi}^2$ and U_ϕ must be calibrated carefully, along with the fixed Level 1 prior parameters $\tilde{\mu}_\beta, \tilde{\sigma}_\beta^2$. As our simulations will show, this may lead to more accurate within-subgroup dose selection compared to what is obtained by either conducting separate trials within subgroups or conducting one trial but ignoring subgroups to obtain a design with good operating characteristics.

2.3 Non-hierarchical models and designs

For the three alternative dose-finding methods, we assume non-hierarchical logistic models, with different parameterizations of their intercepts and slopes to account for subgroups. That is, for each of the following models, usual priors are assumed, and there are no hyperpriors. The first alternative method assumes complete patient homogeneity under the model $\pi(x_j, \alpha, \beta) = \Pr(Y_i = 1 | x_j, \alpha, \beta)$ with

$$\text{logit}\{\pi(x_j, \alpha, \beta)\} = \alpha + \beta x_j \quad (5)$$

for all subgroups, where α and β follow $N(\tilde{\mu}_\alpha, \tilde{\sigma}_\alpha^2)$ and $N(\tilde{\mu}_\beta, \tilde{\sigma}_\beta^2)$ priors, respectively. For this model, numerical values of the four prior hyperparameters $\tilde{\theta} = (\tilde{\mu}_\alpha, \tilde{\sigma}_\alpha, \tilde{\mu}_\beta, \tilde{\sigma}_\beta)$ must be specified. This method conducts a single trial and treats all newly enrolled patient at the same recommended dose, defined under the usual CRM criterion as the x_j having estimated posterior mean $E\{\pi(x_j, \theta) | \mathcal{D}_n\}$ closest to π^* , where $\theta = (\alpha, \beta)$. This is a usual CRM criterion, computed under a two-parameter logistic model. For comparability, and to ensure an ethical trial, the CRM as defined here also includes a “do-not-skip” rule and a rule for overdose control, but applied overall rather than within subgroups.

The second alternative method, which we denote by K -CRM-1-trial, uses the CRM in one trial accounting for K subgroups. This method accounts for patient heterogeneity by modeling the within-subgroup probability of toxicity using the same logistic form with linear term $\alpha_k + \beta x$ as given in Eq.(1), but without hierarchical borrowing of strength among $\alpha_1, \dots, \alpha_K$ through a common hyperparameter. For this model, the assumed priors are $\alpha_1, \dots, \alpha_K \sim \text{i.i.d. } N(\tilde{\mu}_\alpha, \tilde{\sigma}_\alpha^2)$ and $\beta \sim N(\tilde{\mu}_\beta, \tilde{\sigma}_\beta^2)$. Again, since there are only Level 1 priors and no Level 2 priors (hyperpriors) in this model, only numerical values of the four prior hyperparameters $\tilde{\theta} = (\tilde{\mu}_\alpha, \tilde{\sigma}_\alpha, \tilde{\mu}_\beta, \tilde{\sigma}_\beta)$ must be specified. This model also has prior exchangeability across patient subgroups, but it borrows strength between the subgroups only through the common slope parameter, β . Like the HB-CRM, the K -CRM-1-trial conducts a single trial and treats each newly enrolled patient in subgroup k at the subgroup-specific dose for which $E\{\pi_k(x_j, \alpha_k, \beta) | \mathcal{D}_n\}$ is closest to π^* . Thus, the chosen doses for the K subgroups are obtained by averaging the $\pi_k(x_j, \alpha_k, \beta)$'s with respect to the parameters $(\alpha_1, \dots, \alpha_K, \beta)$. This method applies the same subgroup-specific do-not-skip and overdose-control rules as those used in HB-CRM. That is, the K -CRM-1-trial uses precisely the same decision rules as the HB-CRM, but assumes a different, non-hierarchical model.

The third alternative method, 1-CRM- K -trials, conducts separate trials in the K subgroups using the ordinary one-sample CRM in each trial, and it does not assume exchangeability across patient subgroups or borrow strength in any way between subgroups. For subgroup k , this method assumes the model $\pi_k(x_j, \alpha_k, \beta_k) = \Pr(Y_i = 1 | x_j, \alpha_k, \beta_k)$ with

$$\text{logit}\{\pi_k(x_j, \alpha_k, \beta_k)\} = \alpha_k + \beta_k x_j, \quad (6)$$

with priors $\alpha_1, \dots, \alpha_K \sim \text{i.i.d. } N(\tilde{\mu}_\alpha, \tilde{\sigma}_\alpha^2)$ and $\beta_1, \dots, \beta_K \sim \text{i.i.d. } N(\tilde{\mu}_\beta, \tilde{\sigma}_\beta^2)$. For this third model, numerical values of the four prior hyperparameters $\tilde{\theta} = (\tilde{\mu}_\alpha, \tilde{\sigma}_\alpha, \tilde{\mu}_\beta, \tilde{\sigma}_\beta)$ must be specified. Like the HB-CRM and K -CRM-1-trial, the 1-CRM- K -trials design includes subgroup-specific do-not-skip and overdose-control rules.

Table 1 summarizes the four methods in terms of the linear terms of their logistic models, the structural assumptions for the parameters, and the rules for trial conduct, i.e. whether there is one trial or K separate trials that do not use each others' data to make decisions. While both the HB-CRM and K -CRM-1-trial are based on models that borrow strength between subgroups, the key difference is that the HB-CRM model has a hierarchical prior structure on $\alpha_1, \dots, \alpha_K$, while the model used by the K -CRM-1-trial does not.

2.4 Prior specification and numerical methods

Recall that, for the hierarchical model, to establish the prior the five fixed hyperparameters $(\tilde{\mu}_\beta, \tilde{\sigma}_\beta)$ and $(\mu_{\alpha,\phi}, \sigma_{\alpha,\phi}, U_\phi)$ must be specified. To establish the prior for each of the three non-hierarchical models, the four fixed hyperparameters $(\tilde{\mu}_\alpha, \tilde{\sigma}_\alpha, \tilde{\mu}_\beta, \tilde{\sigma}_\beta)$ must be specified. We recommend minimally informative priors to allow the data to dominate the decisions in general and, in the present context, to ensure a fair comparison among the different methods.

For the location parameters, either $(\tilde{\mu}_\alpha, \mu_{\alpha,\phi})$ for the hierarchical model or $(\tilde{\mu}_\alpha, \tilde{\mu}_\beta)$ for the non-hierarchical models, one can two elicited mean toxicity probabilities to solve for the fixed two fixed hyperprior means. Given the standardized doses, in each case the two fixed hyperprior means may be calculated by equating elicited values of toxicity probabilities at two different doses, say $x^{(1)}$ and $x^{(2)}$, with the corresponding formulas for $\pi(x^{(j)}, \theta)$, $j = 1, 2$, replacing θ by its mean, and solving the two equations for the two unknown hyperparameters. This is illustrated below. Given these fixed location parameters, variance parameters may be determined in terms of prior informativeness, quantified by prior effective sample size (ESS) [1],[13]. To speed up computation, one may use approximate ESS values [12] (computational details are given in the Appendix). One may set the values of the variance hyperparameters to control prior informativeness so that the per-subgroup ESS values are a small number, such as 1, 2, or 3.

CRM—For the CRM, first choose an overall ESS, and then divide it by K to obtain a common per-subgroup ESS value. As explained in Section 2.3, this method assumes

$\text{logit}\{\pi(x_j, \alpha, \beta)\} = \alpha + \beta x_j$ (5), with $\alpha \sim N(\tilde{\mu}_\alpha, \tilde{\sigma}_\alpha^2)$ and $\beta \sim N(\tilde{\mu}_\beta, \tilde{\sigma}_\beta^2)$. To obtain the means $\tilde{\mu}_\alpha$ and $\tilde{\mu}_\beta$, if, for example, one elicits the two prior mean toxicity probabilities $E\{\pi(x_2, \theta)\} = 0.10$ and $E\{\pi(x_5, \theta)\} = 0.50$, then the two resulting equations $\tilde{\mu}_\alpha + \tilde{\mu}_\beta x_2 = \text{logit}(0.10)$ and $\tilde{\mu}_\alpha + \tilde{\mu}_\beta x_5 = \text{logit}(0.50)$ yield $\tilde{\mu}_\alpha = -1.23$ and $\tilde{\mu}_\beta = 2.40$. Assuming that $\tilde{\sigma}_\alpha^2 = \tilde{\sigma}_\beta^2$, one then may compute the approximate overall ESS values for a suitable range of $\tilde{\sigma}_\alpha^2 (= \tilde{\sigma}_\beta^2)$, e.g., 0.01, 0.02, ..., 10. Finally, one may choose a value of $\tilde{\sigma}_\alpha^2 (= \tilde{\sigma}_\beta^2)$ so that the overall ESS value is closest to 4 ($= 1 \times 4$), that is, the per-subgroup ESS value nearly equals 1, resulting in $\tilde{\sigma}_\alpha^2 = \tilde{\sigma}_\beta^2 = 1.25$.

1-CRM- K -trials—This method conducts K separate trials and assumes $\text{logit}\{\pi_k(x_j, \alpha_k, \beta_k)\} = \alpha_k + \beta_k x_j$ (6) with $\alpha_1, \dots, \alpha_K \sim \text{i.i.d. } N(\tilde{\mu}_\alpha, \tilde{\sigma}_\alpha^2)$ and $\beta_1, \dots, \beta_K \sim \text{i.i.d. } N(\tilde{\mu}_\beta, \tilde{\sigma}_\beta^2)$. Thus, one may perform the same calculations as done above for the CRM, but within each subgroup. Given $\tilde{\mu}_\alpha = -1.23$ and $\tilde{\mu}_\beta = 2.40$, one may choose $\tilde{\sigma}_\alpha^2 (= \tilde{\sigma}_\beta^2)$ so that the per-subgroup ESS takes a value close to 1, which in this case gives $\tilde{\sigma}_\alpha^2 = \tilde{\sigma}_\beta^2 = 5.92$.

K-CRM-1-trial—This method conducts a single trial and assumes $\text{logit}\{\pi_k(x_j, \alpha_k, \beta)\} = \alpha_k + \beta x_j$ (6) with priors $\alpha_1, \dots, \alpha_K \sim \text{i.i.d. } N(\tilde{\mu}_\alpha, \tilde{\sigma}_\alpha^2)$ and $\beta \sim N(\tilde{\mu}_\beta, \tilde{\sigma}_\beta^2)$. Since it is assumed that, *a priori*, $\alpha_1, \dots, \alpha_K$ have the same mean $\tilde{\mu}_\alpha$, it is reasonable to derive a single value of $\tilde{\mu}_\alpha$. Thus, as above, $\tilde{\mu}_\alpha = -1.23$ and $\tilde{\mu}_\beta = 2.40$ would be derived. In the evaluation of approximate ESS values with fixed $\tilde{\mu}_\alpha$ and $\tilde{\mu}_\beta$, one may compute the within-subgroup probability of toxicity $E\{\pi_k(x_j, \alpha_k, \beta) | \tilde{\mu}_\alpha, \tilde{\mu}_\beta, \tilde{\sigma}_\alpha^2, \tilde{\sigma}_\beta^2\}$ for a suitable range of $\tilde{\sigma}_\alpha^2 (= \tilde{\sigma}_\beta^2)$. Then, choose a value of $\tilde{\sigma}_\alpha^2 (= \tilde{\sigma}_\beta^2)$ so that the per-subgroup ESS takes a value being closest to 1, resulting in $\tilde{\sigma}_\alpha^2 = \tilde{\sigma}_\beta^2 = 5.92$. Although ESS is not necessarily additive over subgroups because β is the common slope parameter for patient subgroups in this method, we avoid more complicated ESS computations to facilitate practical application.

HB-CRM—As explained in Section 2.2, this method conducts a single trial assuming a hierarchical model with $\text{logit}\{\pi_k(x_j, \alpha_k, \beta)\} = \alpha_k + \beta x_j$ (Eq.(1)) for subgroup $k = 1, \dots, K$. For the Level 1 priors, it is assumed that $\alpha_1, \dots, \alpha_K$ are i.i.d. $N(\tilde{\mu}_\alpha, \tilde{\sigma}_\alpha^2)$ and that $\beta \sim N(\tilde{\mu}_\beta, \tilde{\sigma}_\beta^2)$. For Level 2 priors, it is assumed that $\tilde{\mu}_\alpha \sim N(\mu_{\alpha,\phi}, \sigma_{\alpha,\phi}^2)$ and $\tilde{\sigma}_\alpha \sim U(0.01, U_\phi)$. The location parameters in this method are $\tilde{\mu}_\beta$ in Level 1 and $\mu_{\alpha,\phi}$ in Level 2. Since $\mu_{\alpha,\phi}$ represents the overall mean of $\alpha_1, \dots, \alpha_K$, it may be acceptable not to differentiate the prior levels of $\tilde{\mu}_\alpha$ and $\mu_{\alpha,\phi}$ in order to simplify the derivation of their estimates. Thus, similarly to the non-hierarchical methods, $\mu_{\alpha,\phi} = -1.23$ and $\tilde{\mu}_\beta = 2.40$ are obtained using the two elicited prior means $E\{\pi(d_2, \theta)\} = 0.10$ and $E\{\pi(d_5, \theta)\} = 0.50$. Next, three parameters $(\tilde{\sigma}_\beta^2, \sigma_{\alpha,\phi}^2, U_\phi)$ are specified using the ESS computation process. To simplify computation, we use a simplified algorithm to obtain these three parameters sequentially. First, use the same value of $\tilde{\sigma}_\beta^2$ specified in K-CRM-1-trial, which has a study design similar to that of HB-CRM, regardless of whether a hierarchical structure is assumed or not. Second, for U_ϕ referring to Morita, *et al.*, 2012) [1], evaluate ESS at two values, $U_\phi = 2$ and $U_\phi = 5$. Then, for a suitable range of $\sigma_{\alpha,\phi}^2$, e.g., 0.01, 0.02, ..., 10, compute the within-subgroup prior mean probability of toxicity $E\{\pi_k(x_j, \alpha_k, \beta) | \mu_{\alpha,\phi}, \tilde{\mu}_\beta, \tilde{\sigma}_\beta^2, \sigma_{\alpha,\phi}^2, U_\phi\}$ to obtain approximate ESS values. Finally, choose a pair of values of $\sigma_{\alpha,\phi}^2$ and U_ϕ so that the per-subgroup ESS takes a value close to 1. If both values of U_ϕ in combination with some value of $\sigma_{\alpha,\phi}^2$ yield the per-subgroup ESS value 1, choose the smaller value, 2 for U_ϕ , thereby obtaining an appropriately informative hyperprior of $\tilde{\sigma}_\alpha$. It is expected that HB-CRM with a suitably informative hyperprior for $\tilde{\sigma}_\alpha$ that controls the between-subgroup variability in the intercepts $(\alpha_1, \dots, \alpha_K)$ will do a better job of dealing with differences in toxicity probabilities by borrowing strength between subgroups.

Because these specifications of fixed prior parameters involve some arbitrary choices, one should evaluate the operating characteristics of the design via simulation, and if necessary adjust the numerical hyper-parameter values on that basis. To compute the posteriors, we use Markov chain Monte Carlo (MCMC) [14], because the joint posterior distribution of the intercept and slope parameters is not readily available in closed form.

3 Simulations and case-by-case examples

3.1 Simulation study design

We compared the four methods in several cases chosen to evaluate the advantages and disadvantages of HB-CRM and the three other designs in terms of how accurately each estimates optimal doses (ODs).

To evaluate the performance of the four methods fairly, we used the same basic setup with respect to the dose levels $J=6$ with $(d_1, \dots, d_6) = (100, 200, 300, 400, 500, 600)$, starting dose (d_1), target toxicity level $\pi^* = .33$, and the number of subgroups $K=4$. For the subgroups, we assumed two different distributions of population proportions: either $\xi = (.25, .25, .25, .25)$ or $\xi = (.40, .30, .20, .10)$, named “equal” and “different” prevalence patterns, respectively. In addition, we evaluated each design’s operating characteristics using four maximum sample sizes ($N_{max} = 48, 72, 96, 120$). We chose the minimum and maximum values of N_{max} , 48 and 120, taking into account that their corresponding expected per-subgroup sample-sizes under the equal population proportions, 12 and 30, often may be used in an ordinary phase I trial for each subgroup.

As explained in Section 2.4, we set-up the priors of the four methods to ensure reasonably fair comparisons. The details are given in the Supplementary Materials. First, we used prior estimates of $E\{\pi(d_2, \theta)\}$ and $E\{\pi(d_5, \theta)\}$ to solve for the location parameters, $\tilde{\mu}_\alpha$ and $\tilde{\mu}_\beta$. With $E\{\pi(d_2, \theta)\} = .10$ and $E\{\pi(d_5, \theta)\} = .50$, the location parameters were specified as $\tilde{\mu}_\alpha = -1.23$ and $\tilde{\mu}_\beta = 2.40$ for all four methods (except for HB-CRM, which does not use a fixed $\tilde{\mu}_\alpha$). Given the location parameters, the scale parameters were specified as

$\tilde{\sigma}_\alpha^2 = \tilde{\sigma}_\beta^2 = 1.25$ for the CRM so that the overall ESS value was close to 4 ($= 1 \times 4$), that is, the per-subgroup ESS values nearly equaled 1. For 1-CRM- K -trials and K -CRM-1-trial, the scale parameters were specified as $\tilde{\sigma}_\alpha^2 = \tilde{\sigma}_\beta^2 = 5.92$ to obtain per-subgroup ESS value close to 1. The priors and hyperpriors of HB-CRM were specified with $\tilde{\mu}_\beta = 2.40$, $\tilde{\sigma}_\beta^2 = 5.92$, $\mu_{\alpha, \phi} = -1.23$, $\sigma_{\alpha, \phi}^2 = 4.85$, and $U_\phi = 2$. The location parameters were specified based on the prior estimates $E\{\pi(d_2, \theta)\} = .10$ and $E\{\pi(d_5, \theta)\} = .50$ and the dispersion parameters were specified according to prior ESS so that the per-subgroup ESS value was close to 1. For the design parameters of the overdose control rule, based on preliminary investigation, we determined that the combinations of $\pi^{odc} = .50$ and $\psi^{odc} = .25$ gave designs with good operating characteristics.

We constructed four different dose-toxicity scenarios by specifying values of the true toxicity probabilities in each subgroup, shown in Figure 1 and Supplementary Table 1 (STable 1). These scenarios were chosen to illustrate how the methods behave in a variety of settings in terms of inter-subgroup difference of the dose-toxicity relationship. The scenarios are not based on any of the models. Each scenario is characterized by the true probabilities of toxicity, $\pi_{j,k}^{true}$, for the six dose levels $j = 1, \dots, 6$ and four subgroups, $k = 1, \dots, 4$. Thus, the assumed true dose-toxicity curve within subgroup k is characterized by the 6-dimensional vector $\pi_k^{true} = (\pi_{1,k}^{true}, \dots, \pi_{6,k}^{true})$. Under Scenario 1, given the target probability $\pi^* = .33$, d_4 is the OD in all the four subgroups. In contrast, Scenario 2 represents a case

where the ODs differ between subgroups, with respective ODs d_4 , d_6 , d_1 , and d_4 , in the four subgroups, respectively. Scenario 3 also has different ODs for the subgroups, but they are closer to each other than in Scenario 2. Scenario 4 is a difficult case where the four subgroup-specific dose-toxicity curves π_1^{true} , π_2^{true} , π_3^{true} , and π_4^{true} differ substantially between subgroups, and moreover each subgroup's curve has two dose levels with true toxicity probabilities equal to or close to the target $\pi^* = .33$. In this case, two doses both are good choices for each subgroup, but these two doses differ between subgroups.

To evaluate and compare the four designs under each of the dose-toxicity scenarios in the simulations, we use the following weighted average of the dose selection probabilities.

Motivated by the idea that smaller values of $|\pi_{j,k}^{true} - \pi^*|$ correspond to d_j being more desirable in subgroup k , equivalently larger values of $1 - |\pi_{j,k}^{true} - \pi^*|$ are more desirable in that subgroup, we define the weights

$$w_{j,k} = \frac{1 - |\pi_{j,k}^{true} - \pi^*| - \min_{r=1, \dots, J} \{1 - |\pi_{r,k}^{true} - \pi^*|\}}{\max_{r=1, \dots, J} \{1 - |\pi_{j,k}^{true} - \pi^*|\} - \min_{r=1, \dots, J} \{1 - |\pi_{r,k}^{true} - \pi^*|\}} \quad (7)$$

Then, we define the *subgroup-specific weighted probability of selection*

$$WPS_k = \sum_{j=1}^J w_{j,k} \cdot Pr(x_{j,k} \text{ is selected as the OD in subgroup } k), \quad (8)$$

for each $k = 1, \dots, K$. We subtract the smallest value of $1 - |\pi_{r,k}^{true} - \pi^*|$ among the j doses in the numerator and denominator of (7) so that $w_{j,k} = 0$ for the least desirable dose in that subgroup in order to give greater relative weights to the doses having $\pi_{j,k}^{true}$ closer to π^* in that subgroup. Particularly, Eq(7) gives weight 1 to the dose having $\pi_{j,k}^{true}$ closest to π^* . We do not take the alternative approach of using $|\pi_{j,k}^{true} - \pi^*|^{-1}$ as a basis for constructing weights because this takes on the value ∞ if $\pi_{j,k}^{true} = \pi^*$. We also evaluate the statistic PCS_k , the probability of correctly selecting the dose that minimizes $|\pi_{j,k}^{true} - \pi^*|$ in subgroup k , which gives weight 1 to the dose having $\pi_{j,k}^{true}$ closest to the target and weight 0 to all other doses.

Under each scenario, we simulated the trial 1,000 times using each method. The SAS program to implement HB-CRM is provided in the Supplementary Materials (SAS for Windows release 9.3; SAS Institute Inc., Cary, NC, USA).

3.2 Simulation results

The operating characteristics for the four methods are summarized by toxicity scenarios and the results are shown in terms of the WPS and PCS only for $N_{max} = 96$ and except for Scenario 1. Figures 2 and 3 show the results under the assumptions of equal and different

subgroup proportions (.25, .25, .25, .25) and (.40, .30, .20, .10), respectively. The selection probabilities of x_j as the OD in subgroup k in Eq.(8) were computed as the percentage of times that each of the methods selected x_j as the OD in each subgroup. More complete results are shown in the supplementary materials (S-Table 2).

Overall, the simulation study reconfirmed that ignoring subgroups resulted in undesirably low probabilities of selecting ODs, especially when the dose-toxicity relationships were largely different between subgroups. It also was reconfirmed that K -CRM-1-trial and 1-CRM- K -trials behaved about the same in a setting with more than two non-ordered subgroups. HB-CRM showed worse performance than K -CRM-1-trial and 1-CRM- K -trials in several cases. However, when the subgroup proportions were different, HB-CRM gave much better results in the subgroups with small proportions $\xi = 0.1$ or 0.2 .

As shown in Figure 2, under the assumption of equal subgroup proportions, under Scenario 2, K -CRM-1-trial and 1-CRM- K -trials both performed best in subgroup 3, but in the other subgroups the performance of HB-CRM was very similar to these two methods. As in Scenario 3, the performance of HB-CRM was quite favorable when the ODs were close to each other between patient subgroups. It appears that, in such a case, HB-CRM effectively borrows strength between subgroups through its hierarchical structure. Under Scenario 4, the WPS and PCS values of HB-CRM were higher in subgroups 2 and 3, for which two doses were good choices, compared to those of K -CRM-1-trial and 1-CRM- K -trials. Under the different subgroup proportions (Figure 3), HB-CRM yielded much higher WPS values in subgroup 4 under each of Scenarios 2, and 3 compared to K -CRM-1-trial and 1-CRM- K -trials. That is, the desirable effect of borrowing strength between subgroups in HB-CRM appeared to be more pronounced in subgroups with smaller numbers of patients. In other aspects, overall, the results were similar to those obtained in the case of equal subgroup proportions.

Figures 4a and 4b shows the WPS values for the four sample sizes by toxicity scenario under the equal and different subgroup proportions, respectively. The four columns correspond respectively to HB-CRM, K -CRM-1-trial, 1-CRM- K -trials, and CRM, starting from the left. Overall, the performance of HB-CRM improves as N_{max} increases much more than those of K -CRM-1-trial and 1-CRM- K -trials. Particularly, under Scenario 1, HB-CRM yielded high values of WPS even with $N_{max} = 48$, and the WPS values of HB-CRM with $N_{max} = 48$ were comparable to those of K -CRM-1-trial and 1-CRM- K -trials for $N_{max} = 120$. This may be due to the way that HB-CRM borrows strength between subgroups. Under this scenario, it is not surprising that CRM provides the highest values of WPS, because the patients come from one population rather than multiple subpopulations.

Since similar conclusions were obtained from the simulation results in WPS and PCS, it might be acceptable to use the PCS to evaluate the operating characteristics of a dose-finding study design. However, because the toxicity probabilities for one or more other doses may be close to that of the optimal dose in one or more subgroups, the WPS may be a more suitable index to compare the performances between methods, especially in a difficult case like Scenario 4.

4 Guidelines for Constructing Designs

To construct a study design using the HB-CRM method, the following steps may be taken.

1. Determine the definition of toxicity, the target toxicity probability π^* , and the dose levels ($d_1, \dots, d_J$) to be tested.
2. Specify the patient subgroups ($1, \dots, K$), anticipated subgroup prevalences $\xi = (\xi_1, \dots, \xi_K)$, and patient accrual rates.
3. Set-up the priors of the dose-finding model to be minimally informative in terms of the prior ESS, following the approach described in Section 2.4.
4. Determine the total maximum sample size, N_{max} , by running the computer program (provided in the supplementary materials) for a range of feasible values of N_{max} so that the study design has a sufficiently good performance in terms of the WPS of the subgroups.

As a guide, in Step 4 above, the values of N_{max} may range from $12 \times K$ to $30 \times K$ when the number of dose levels J is 4 to 6. If more dose levels are examined, one may consider increasing the per-subgroup sample-size, as in an ordinary dose-finding trial. In some cases, some subgroups may be very small, that is, the corresponding values of $(\xi_1, \dots, \xi_K)$ take very small values, e.g., 0.05 or smaller. This is likely to occur if the number of subgroups is large, e.g., $K = 10$ or 20 . In such a case, due to the limited number of patients in a phase I trial, we strongly recommend reducing the number of subgroups, K , to a number that allows the proposed methodology to be applied in a practical way. In addition, one should take the patient accrual rates in subgroups into account when combining subgroups, in order to improve trial feasibility and simplify trial conduct, that is, to complete the trial within a realistic time frame. If different patient subgroups have different toxicity targets, $\pi_1^*, \dots, \pi_K^*$, one can design the trial in the same way as a trial with a common target π^* , although one should investigate the operating characteristics of the study design carefully.

5 Illustration

For illustration, based on two real phase I trials, we show how HB-CRM and K-CRM-1-trial may work in practice, via simulations. To carry out these simulations, we assumed true toxicity probabilities based on the empirical data observed in the two clinical trials. The first example (Example 1) is a case where exchangeable toxicity probabilities may reasonably be assumed between patient subgroups, while exchangeability clearly should not be assumed in the second example (Example 2).

Example 1 is based on Minami *et al.* [15]. They examined three doses (400, 600, 800 mg once daily) of sonidegib (LDE225), a selective protein inhibitor, in $N = 45$ Asian patients with advanced solid tumors, including two racial subgroups. Subgroup 1 consisted of $N_1 = 21$ Japanese and subgroup 2 consisted $N_2 = 24$ of Hong Kong / Taiwanese, following the health authority's request. The MTDs were reported to be 400 mg in both subgroups. For these patients, we consider it reasonable to assume, *a priori*, that the toxicity probabilities are exchangeable between the two racial subgroups. In this trial, dose-limiting toxicities (DLTs) were evaluated during the first treatment cycle. Table 2 shows the observed DLT data.

For Example 2, we use the data reported by Bendell *et al.* [16]. They conducted a phase I trial to test six dose levels, 12.5, 25, 50, 80, 100, 150 *mg* once daily of BKM120, a pyrimidine-derived pan-PI3K inhibitor with specific and potent activity against class I PI3Ks. 100 *mg* was estimated to be the MTD in this trial. As a clinical background characteristic of the $N = 35$ patients enrolled in this trial, Bendell *et al.* [16] reported the number of prior therapies for two subgroups, with > 3 in subgroup 1, and ≤ 3 in subgroup 2. The sizes of the two subgroups were nearly the same, with $N_1 = 18$ and $N_2 = 17$. Because, in general, patients who are more heavily pre-treated are more likely to experience toxicity, it clearly is not appropriate to assume that these two patient subgroups have exchangeable prior toxicity probabilities. Table 2 shows the DLTs observed during the first treatment cycle in this trial.

For each of the two examples, we simulated toxicity data for each of four hypothetical total sample sizes, $N = 35, 45, 70$, and 100 , with subgroup proportions $\xi = 0.467$ in Example 1 and $\xi = 0.514$ in Example 2, to mimic the proportions in the reported data. Since the implicit target range for the DLT probability was 0.16 to 0.33 to determine the MTD in both trials, we defined 0.25 as the target DLT probability for both examples. Table 2 shows the true DLT probabilities derived from the empirical data of the two clinical trials. In Example 1, where subgroups were determined by race, we assumed true DLT probabilities such that dose level 1 (400 *mg*) was the MTD in subgroup 1, while dose level 2 (600 *mg*) was the MTD in subgroup 2. For Example 2, where subgroups were determined by number of prior therapies, toxicity data within the patient subgroups were not given by Bendell *et al.* [16]. Thus, we assumed true DLT probabilities so that dose levels 4 (80 *mg*) and 6 (150 *mg*) were the respective MTDs in subgroups 1 (> 3 prior therapies) and 2 (≤ 3 prior therapies).

The simulation results for Examples 1 and 2 are summarized in Table 2 for two maximum sample sizes, $N = 45$ and 100 and $N = 35$ and 100 , respectively. The table gives the percentages of times that each method selected each dose as the MTD in each subgroup. Correct selection percentages are given in boldface. Corresponding results for other maximum N values are summarized in the supplementary material.

In Example 1 (exchangeable case), HB-CRM performed better than K-CRM-1 overall in the two subgroups. In this example, Table 2 shows within-subgroup correct selection percentages of 71.8% and 80.5% for HB-CRM even with $N = 45$, although these high values are due in part to the fact that there were only three the number of dose levels. Still, the example illustrates the ability of HB-CRM to reliably choose different optimal doses within subgroups. In Example 2 (non-exchangeable case), for $N = 35$, Table 2 shows that K-CRM-1 performed better than HB-CRM, especially in subgroup 2. With larger N , the performances of both methods improved, and the difference in performance between the two methods diminished. A key point is that, for phase I trials with patient heterogeneity and six or more doses levels, N should be larger than conventional values to obtain reliable subgroup-specific dose selections.

6 Discussion

Our simulation studies suggest that HB-CRM works well in situations where the dose-toxicity curves are expected to be similar or not largely different between multiple patient subgroups, and the exchangeability assumption is valid. This arises commonly in settings whether qualitatively different disease subgroups are included, and there is no prior knowledge about the comparative risks of toxicity in the subgroups. In terms of overall performance quantified by the WPS (weighted probability of selection) or PCS (probability of correct selection), Figure 2 shows that the HB-CRM method (solid line) does well across all subgroups. In contrast, the non-hierarchical model-based methods K -CRM-1 and 1-CRM- K that account for subgroups may perform well for some subgroups but not as well for others. For example, in Scenario 2, the K -CRM-1 method outperforms the HB-CRM method slightly in terms of WPS and markedly for subgroup 3 in terms of PCS, but K -CRM-1 has the same or inferior performance compared to HB-CRM in Scenarios 3 and 4. Not surprisingly, the CRM that ignores subgroups has greatly inferior performance for several subgroups in each of Scenarios 2, 3 and 4. Figures 4a and 4b illustrate the extremely poor performance of the CRM for many subgroups in the presence of heterogeneity.

Since, in practice, one cannot know the true toxicity curves, there are two main messages. First, one certainly should account for known patient heterogeneity, since failure to do so is very likely to produce a selected dose that is far below optimal in some subgroups. Second, when the underlying assumptions are appropriate, the HB-CRM performs well consistently across a broad range of different dose-toxicity-subgroup scenarios, and it may be preferable to non-hierarchical model based methods that choose subgroup-specific doses.

An important caveat is that the HB-CRM based method is not appropriate when it is known that the subgroups are not exchangeable, with an important case being that where the risk of toxicity is known to be ordered by subgroup. If preclinical or clinical data identify multiple patient subgroups that are likely to have substantially different dose-toxicity relationships, but the HB-CRM is not appropriate, then one should use a non-hierarchical dose-finding method that chooses subgroup-specific doses.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors thank Dr. Peter Müller for his helpful advice on an earlier draft of this manuscript. We also thank an Editor, Associate Editor and two referees for their constructive comments and suggestions. Satoshi Morita's research was partially supported by a Grant-in-Aid for Scientific Research C-24500345 from the Ministry of Health. Peter Thall's research was supported by NIH/NCI grant RO1 CA 83932.

Appendix

The priors for the models underlying each of the four methods CRM, 1-CRM- K -trials, K -CRM-1-trial, HB-CRM may be constructed in the following two steps.

Step 1: Location hyper-parameters

Determine the location hyper-parameters of the priors, $(\tilde{\mu}_\alpha, \mu_{\alpha,\phi})$ for the HB-CRM and $(\tilde{\mu}_\alpha, \tilde{\mu}_\beta)$ for the three non-hierarchical methods, by first obtaining numerical values of the mean probability of toxicity at each of two doses. A convenient choice consists of the second lowest and second highest dose, denoted by $\pi(d_2, \boldsymbol{\theta})$ and $\pi(d_{J-1}, \boldsymbol{\theta})$, although other dose pairs may be used. These prior mean probabilities may be obtained by elicitation from the physicians, or based on historical data. For the HB-CRM model, the prior elicitation process does not distinguish between patient subgroups, because the hierarchical model prior assumes that the toxicity probabilities of the patient subgroups are exchangeable. For each model, given the two elicited prior mean probabilities, the two equations are solved for the two location parameters.

Step 2: Dispersion hyper-parameters

Given the prior hyper-means, determine numerical values of the dispersion parameters controlling the informativeness by using prior ESS, as described in Section 2.4. These dispersion parameters are $(\tilde{\sigma}_\beta, \sigma_{\alpha,\phi}, U_\phi)$ for the HB-CRM model and $(\tilde{\sigma}_\alpha, \tilde{\sigma}_\beta)$ for each of the three other models. To speed up computation, compute approximate ESS values based on the fact that a $\text{beta}(a, b)$ distribution has $\text{ESS} = a + b$, by approximating the prior of any probability $\pi(\boldsymbol{\theta})$ by a $\text{beta}(a, b)$ and matching the means and variances. One then solves the two equations

$$E\{\pi(\boldsymbol{\theta}|\tilde{\boldsymbol{\theta}})\} = a/(a+b)$$

and

$$\text{var}\{\pi(\boldsymbol{\theta}|\tilde{\boldsymbol{\theta}})\} = ab/\{(a+b)^2(a+b+1)\}$$

for $\text{ESS} \doteq a + b$.

Step 2a

To compute an approximate ESS within each subgroup, use the above approach to compute the beta-approximated value ESS_j of the prior $\pi(\boldsymbol{\theta}|\tilde{\boldsymbol{\theta}})$ for each dose $d_j, j = 1, \dots, J$, and use the mean of these J values as a per-subgroup approximate ESS.

Step 2b

Multiply the value of the per-subgroup ESS by the number of subgroups, K , to obtain an overall ESS.

References

1. Morita S, Thall PF, Mueller P. Prior effective sample size in conditionally independent hierarchical models. *Bayesian Analysis*. 2012; 7:591–614.

2. O'Quigley J, Pepe M, Fisher L. Continual reassessment method: a practical design for phase I clinical trials in cancer. *Biometrics*. 1990; 46:33–48. [PubMed: 2350571]
3. O'Quigley J, Shen LZ, Gamst A. Two-sample continual reassessment method. *J Biopharm Stat*. 1999; 9:17–44. [PubMed: 10091908]
4. O'Quigley J, Paoletti X. Continual reassessment method for ordered groups. *Biometrics*. 2003; 59:430–440. [PubMed: 12926728]
5. Ivanova A, Wang K. Bivariate isotonic design for dose-finding with ordered groups. *Biometrics*. 2006; 25:2018–2026.
6. Yuan Z, Chappell R. Isotonic designs for phase I cancer clinical trials with multiple risk groups. *Clinical Trials*. 2004; 1:499–508. [PubMed: 16279290]
7. Storer BE. Design and analysis of phase I clinical trials. *Biometrics*. 1989; 45:925–937. [PubMed: 2790129]
8. Leung DHY, Wang YG. Isotonic designs for phase I trials. *Control Clin Trials*. 2001; 22:126–138. [PubMed: 11306151]
9. Thall PF, Nguyen H, Estey EH. Patient-specific dose-finding based on bivariate outcomes and covariates. *Biometrics*. 2008; 64:1126–1136. [PubMed: 18355387]
10. Liu S, Pan H, Xia J, Huang Qin, Yuan Y. Bridging continual reassessment method for phase I clinical trials in different ethnic populations. *Stat in Med*. 2015; 34:1681–1694. [PubMed: 25626429]
11. Gelman A. Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis*. 2006; 1:515–533.
12. Thall PF, Herrick RC, Nguyen H, Venier JJ, Norris JC. Effective sample size for computing prior hyperparameters in Bayesian phase I-II dose finding. *Clinical Trials*. 2014; 11:657–666. [PubMed: 25179541]
13. Morita S, Thall PF, Mueller P. Determining the effective sample size of a parametric prior. *Biometrics*. 2008; 64:595–602. [PubMed: 17764481]
14. Gilks, W., Richardson, S., Spiegelhalter, D. *Markov Chain Monte Carlo in Practice*. Chapman & Hall; London: 1996.
15. Minami H, Ando Y, Ma BB, Lee JH, Momota H, Fujiwara Y, Li L, Fukino K, Ito K, Tajima T, Mori A, Lin CC. A phase I, multicenter, open-label, dose-escalation study of sonidegib in Asian patients with advanced solid tumors. *Cancer Sci*. 2016; [Epub ahead of print]. doi: 10.1111/cas.13022
16. Bendell JC, Rodon J, Burris HA, de Jonge M, Verweij J, Birle D, Demanse D, De Buck SS, Ru QC, Peters M, Goldbrunner M, Baselga J. Phase I, dose-escalation study of BKM120, an oral pan-Class I PI3K inhibitor, in patients with advanced solid tumors. *J Clin Oncol*. 2012; 30:282–290. [PubMed: 22162589]

Figure 1.

An example of three patient subgroups with different true dose-toxicity curves (y-axis: toxicity probability, x-axis: dose). Given the fixed target toxicity probability, $\pi^*=0.30$, the three subgroups have different true optimal doses.

Figure 2.

Subgroup-specific dose-toxicity curves assumed in the simulations, presented in terms of the true dose-toxicity probabilities $\pi_{1,k}^{true}, \dots, \pi_{6,k}^{true}$ for each subgroup (Sg): Sg 1, diamond and solid; Sg 2, square and dashed; Sg 3, triangle and dashed-dotted; Sg 4, star and dotted. Figures 1a, 1b, 1c, and 1d correspond to scenarios 1, 2, 3, and 4, respectively. Optimal doses are indicated by open circles.

Figure 3.

Subgroup-specific weighted probability of selection (WPS) of optimal dose and probability of correctly selecting (PCS) the optimal dose for HB-CRM (diamond and solid), K-CRM-1-trial (square and dashed), and 1-CRM-K-trials (triangle and dashed-dotted), and CRM (star and dotted) when the total sample size $N_{max} = 96$ with assuming equal subgroup proportions $\xi = (\xi_1, \dots, \xi_4) = (.25, .25, .25, .25)$.

Figure 4.

Subgroup-specific weighted probability of selection (WPS) of optimal dose and probability of correctly selecting (PCS) the optimal dose for K -CRM-1-trial (square and dashed), and 1-CRM- K -trials (triangle and dashed-dotted), and CRM (star and dotted) when the total sample size $N_{max} = 96$ under the dose-toxicity scenarios 2 to 4 with assuming different subgroup proportions $\xi = (.40, .30, .20, .10)$.

Table 1

Summary of study designs in terms of hierarchical versus non-hierarchical model and the linear term of the logistic model for the probability of toxicity as a function of dose x and prognostic subgroup $k = 1, \dots, K$. In the linear terms, α and β denote the intercept and slope parameters, respectively.

Bayesian model structure	Single trial		K separate trials
	$\alpha + \beta x$	$\alpha_k + \beta x$	$\alpha_k + \beta_k x$
Non-hierarchical	CRM ¹	K -CRM-1-trial ²	1-CRM- K -trials ⁴
Hierarchical	–	HB-CRM ³	–

¹ Ordinary CRM ignoring subgroups, conduct one trial.

² K -subgroup CRM in one trial, assuming different intercepts $\alpha_1, \dots, \alpha_K$ without a hierarchical structure, conduct one trial.

³ Hierarchical model based CRM assuming different intercepts $\alpha_1, \dots, \alpha_K$ with a hierarchical structure, conduct one trial.

⁴ Ordinary CRM conducted in each of K separate trials, assuming independent subgroup-specific parameters (α_k, β_k) .

Illustration of HB-CRM and K-CRM-1-trial for two real phase I trials, the sonidegib trial [15] and the BKM120 trial [16]. For each trial, the observed toxicity data and assumed true DLT probabilities used in the simulation first are summarized by subgroup, followed by simulation results for each method for each of two maximum sample sizes. Percentages of correct selection of the MTD within each subgroup are given in boldface.

Table 2

Sonidegib trial		Dose (mg)					
Subgroup		400	600	800			
1 (Japanese)	No. of DLTs	2	5	0			
	No. of patients	12	9	0			
	Empirical Pr(DLT)	0.17	0.56	–			
	Assumed true Pr(DLT)	0.15	0.55	0.65			
2 (Taiwanese)	No. of DLTs	2	1	2			
	No. of patients	12	8	4			
	Empirical Pr(DLT)	0.17	0.13	0.50			
	Assumed true Pr(DLT)	0.15	0.20	0.50			
N	Method	Subgroup					
45	HB-CRM	1	80.5	19.5	0.0		
		2	20.1	71.8	8.1		
	K-CRM-1	1	87.2	12.8	0.0		
		2	30.8	60.5	8.7		
100	HB-CRM	1	91.6	8.4	0.0		
		2	16.2	80.6	3.2		
	K-CRM-1	1	95.5	4.5	0.0		
		2	27.1	69.6	3.3		
BKM120 trial		Dose (mg)					
		12.5	25	50	80	100	150
	No. of DLTs	0	0	0	1	4	1
	No. of patients	1	1	3	6	16	3
	Empirical Pr(DLT)	0.00	0.00	0.00	0.17	0.25	0.33
Subgroup							

BKM120 trial			Dose (mg)					
			12.5	25	50	80	100	150
1 (>3 prior trts)	Assumed true Pr(DLT)		0.05	0.07	0.10	0.25	0.35	0.55
2 (3 prior trts)	Assumed true Pr(DLT)		0.01	0.02	0.05	0.10	0.15	0.25
N	Method	Subgroup						
35	HB-CRM	1	0.9	8.1	39.0	30.2	18.0	3.8
		2	0.3	2.1	18.3	13.1	33.7	32.5
	K-CRM-1	1	1.2	12.6	33.1	31.6	17.9	3.6
100		2	0.3	6.0	12.1	9.1	35.6	36.9
		1	0.2	3.8	25.1	55.7	14.2	1.0
	HB-CRM	2	0.0	0.4	6.9	5.3	38.5	48.9
		1	0.3	7.4	25.0	53.7	12.7	0.9
	K-CRM-1	2	0.1	2.6	5.4	3.1	38.5	50.3