



*Appl. Statist.* (2017)

# Parametric dose standardization for optimizing two-agent combinations in a phase I–II trial with ordinal outcomes

Peter F. Thall, Hoang Q. Nguyen and Ralph G. Zinner

*University of Texas MD Anderson Cancer Center, Houston, USA*

[Received August 2015. Revised April 2016]

**Summary.** A Bayesian model and design are described for a phase I–II trial to optimize jointly the doses of a targeted agent and a chemotherapy agent for solid tumours. A challenge in designing the trial was that both the efficacy and the toxicity outcomes were defined as four-level ordinal variables. To reflect possibly complex joint effects of the two doses on each of the two outcomes, for each marginal distribution a generalized continuation ratio model was assumed, with each agent's dose parametrically standardized in the linear term. A copula was assumed to obtain a bivariate distribution. Elicited outcome probabilities were used to construct a prior, with variances calibrated to obtain small prior effective sample size. Elicited numerical utilities of the 16 elementary outcomes were used to compute posterior mean utilities as criteria for selecting dose pairs, with adaptive randomization to reduce the risk of becoming stuck at a suboptimal pair. A simulation study showed that parametric dose standardization with additive dose effects provides a robust reliable model for dose pair optimization in this setting, and it compares favourably with designs based on alternative models that include dose–dose interaction terms. The model and method proposed are applicable generally to other clinical trial settings with similar dose and outcome structures.

**Keywords:** Adaptive design; Bayesian design; Combination trial; Ordinal variables; Phase I–II clinical trial; Utility

## 1. Introduction

This paper was motivated by the problem of designing an early phase clinical trial of a three-agent combination for treatment of cancer patients with advanced solid tumours. The first agent is a novel molecule M designed to inhibit the protein kinase complexes mTORC1 and mTORC2, and thus to interfere with cancer cell proliferation and survival, among other cancer properties. M also has antiangiogenic properties, through which it deprives the cancer of essential blood vessels that invest the tumours. The other two treatment components are the widely used chemotherapeutic agents carboplatin and paclitaxel. Paclitaxel, when given weekly, has been shown to act as an angiogenesis inhibitor as well. The property of antiangiogenesis that is shared by M and weekly paclitaxel motivates this combination regimen, through which a more powerful antiangiogenic, and therefore anticancer, effect is hypothesized. All three drugs also are expected to target the cancer cells directly through additional different mechanisms, thereby complementing each other.

For the three-agent regimen in this trial, carboplatin is administered at a fixed dose based on the patient's age, weight and kidney function. The doses of the two agents that are varied

*Address for correspondence:* Peter F. Thall, Department of Biostatistics, University of Texas MD Anderson Cancer Center, 1515 Holcombe Boulevard, Houston, TX 77030, USA.  
E-mail: rex@mdanderson.org

**Table 1.** Definitions of overall toxicity severity levels by grades of individual toxicities<sup>†</sup>

Illness	Grades for the following overall toxicity severity levels $Y_T$ :			
	Mild	Moderate	High	Severe
Fatigue	Grade 1	Grade 2	Grade 3	Grade 4
Nausea	Grade 1	Grade 2	Grade 3	—
Neuropathy	—	Grade 1	Grade 2	Grade $\geq 3$
Hyperglycaemia	Grade 2	Grade 3	Grade 4	—
Rash	Grade 1	Grade 2	Grade 3	Grade 4
Diarrhoea	Grade 1	Grade 2	Grade 3	Grade 4
Stomatitis	Grade 1	Grade 2	Grade 3	Grade 4
Pneumonitis	Grade 1	Grade 2	Grade 3	Grade 4
Febrile neutropaenia	—	—	Grade 3	Grade 4
Other non-haematologic illness	Grade 1	Grade 2	Grade 3	Grade 4
Hyperlipidaemia	Grade 1	Grade 2	Grade 3	Grade 4
Anaemia	Grade 3	Grade 4	—	—
Thrombocytopenia	Grade 2	Grade 3	—	—
Neutropaenia	Grade 3	Grade 4	—	—
Liver toxicity AST/ALT <sup>‡</sup>	Grade 2	Grade 3	Grade 4	—
Blindness	—	—	—	Grade 4
Myocardial infarction	—	—	—	Grade 4
Stroke	—	—	—	Grade 4
Regimen-related death	—	—	—	Grade 5

<sup>†</sup>Overall toxicity is scored as the maximum individual severity level. Grades are defined by using National Cancer Institute criteria.

<sup>‡</sup>AST/ALT is the ratio between the concentrations of the enzymes aspartate transaminase, AST, and alanine transaminase, ALT, in the blood, used as an index of liver toxicity.

are  $d_M = 4, 5, 6$  mg of M given orally each day, and  $d_P = 40, 60, 80$  mg  $m^{-2}$  of paclitaxel given intravenously twice weekly. A total of nine (M, paclitaxel) dose pairs  $\mathbf{d} = (d_M, d_P)$  are studied, with the goal to find the optimal  $\mathbf{d}$ . Our proposed method will define ‘optimal’  $\mathbf{d}$  by assigning joint utilities to toxicity and efficacy, assuming a Bayesian model, and identifying the  $\mathbf{d}$  having largest posterior mean utility. Toxicity is defined as a four-level ordinal variable  $Y_T$ , with possible levels  $y_T \in \{\text{mild, moderate, high, severe}\}$ . As shown in Table 1,  $Y_T$  is defined in terms of the severity grades of many qualitatively different toxicities, with the level of  $Y_T$  determined by the highest level of any individual toxicity experienced by the patient. Reducing the many toxicities in Table 1 to the four-level ordinal outcome  $Y_T$  required many subjective decisions by the clinical oncologist planning the trial (the third author of this paper, RGZ). Efficacy is a four-level ordinal variable  $Y_E$ , with possible values  $y_E \in \{\text{PD, SD1, SD2, PR/CR}\}$ , where PD = [progressive disease] = [ $> 20\%$  increase in tumour size], SD1 = [stable disease level 1] = [0–20% increase in tumour size], SD2 = [stable disease level 2] = [0–30% reduction in tumour size] and PR/CR = [partial or complete response] = [ $> 30\%$  reduction in tumour size]. This is a refinement of the commonly used three-category definition where SD1 and SD2 are combined as SD  $\equiv$  stable disease, sometimes with the 30% replaced by 20% so that SD is a change of 20% or less in tumour size in either direction. In the trial, both  $Y_E$  and  $Y_T$  are scored within 42 days from the start of treatment. Thus, a criterion for determining an optimal dose pair must be defined in terms of the joint effect of  $\mathbf{d}$  on  $\mathbf{Y} = (Y_E, Y_T)$ , which has 16 possible values.

An ordinal categorization of solid tumour response is used commonly in oncology to compute descriptive statistics but almost never is used for decision making by dose finding designs. The

most common practice is to define  $Y_T$  and  $Y_E$  as binary variables. In the present setting, this would be done by defining  $Y_E$  to indicate a ‘response’ event, which could be PR/CR, {PR/CR or SD2} or {PR/CR, SD2 or SD1}. Most commonly,  $Y_T$  indicates a composite adverse event, ‘dose limiting toxicity’. These assumptions usually are made for phase I–II designs (see Braun (2002), Thall and Cook (2004), Bekele and Shen (2005), Zhang *et al.* (2006) and Yin *et al.* (2006)). A further reduction is the conventional approach of ignoring efficacy and conducting a phase I trial based on the probability of dose limiting toxicity as a function of dose (see Storer (1989), O’Quigley *et al.* (1990) and Babb *et al.* (1998)). Curve-free dose finding methods have been proposed by Gasparini and Eisele (2000) and Whitehead *et al.* (2010) for phase I trials, and by Whitehead *et al.* (2011) for phase I–II combination trials. Bekele and Shen (2005) and Zhou *et al.* (2006) proposed parametric model-based phase I–II methods to accommodate binary  $Y_T$  and continuous  $Y_E$ .

The utility-based two-agent phase I–II design of Houede *et al.* (2010) accounts for bivariate ordinal  $\mathbf{Y}$  and models marginal dose–dose interactions by using a generalization of the Aranda-Ordaz model (1981), which is given in Appendix A. Since this design deals with the same general problem as that addressed here, it is a natural comparator to our proposed methodology. The main differences between our methodology and that of Houede *et al.* (2010) are that

- (a) we account for joint effects of two doses on each marginal outcome distribution by using parametric dose standardization, and
- (b) we use adaptive randomization to reduce the probability of becoming stuck at a suboptimal dose pair.

Additionally, our motivating application has an outcome of dimension (4,4) whereas that in the application of Houede *et al.* (2010) is (3,3) dimensional. Our simulations comparing the methods show that our proposal has more consistent performance across a range of dose–outcome scenarios and in particular has better worst-case performance (Tables 5 and 6, and Fig. 5 in Section 4).

Medically, the trial that is considered here is similar to the trial motivating the phase I–II design of Riviere *et al.* (2015), in that both trials aim to find an optimal dose pair of a targeted agent and a chemotherapy agent. Key differences are that Riviere *et al.* (2015) addressed settings where toxicity is a binary variable and efficacy is a time-to-event variable and, assuming a proportional hazards model, the dose–efficacy curve may increase initially but then reach a plateau. The problem of optimizing the doses of a two-agent combination based on bivariate binary ( $Y_E, Y_T$ ) outcomes has been addressed in the phase I–II designs that were proposed by Yuan and Yin (YY) (2011) and Wages and Conaway (WC) (2014). Our computer simulations, which are reported in Section 4 and Table 6 there, show that defining efficacy and toxicity as ordinal variables with three or more levels is more informative than collapsing categories and defining two binary indicators, e.g. by dichotomizing  $Y_E$  in one of the ways noted above and defining binary  $Y_T = I(\text{high or severe toxicity})$ .

Formulating a probability model and decision rules that use a (4,4) dimensional bivariate ordinal outcome to choose dose pairs in a sequentially adaptive phase I–II trial is challenging. In this trial, a maximum of 60 patients will be accrued, treated in 20 cohorts of size 3, starting at  $\mathbf{d} = (4,60)$ . Denote an elementary outcome by  $\mathbf{y} = (y_E, y_T)$ , with the efficacy outcomes ordered from worst to best by  $y_E = 0, 1, \dots, L_E$ , and the toxicity outcomes ordered from least to most severe by  $y_T = 0, 1, \dots, L_T$ . Even if the trial’s 60 patients were distributed evenly among the 16 possible  $\mathbf{y}$  pairs at completion, there would be only about four patients per outcome. This sample size allocation is an unrealistic ideal, however, because the elementary outcomes are not equally likely for any  $\mathbf{d}$ , and moreover dose pairs are assigned in a sequentially adaptive

**Table 2.** Elicited numerical utilities of the 16 joint (efficacy, toxicity) outcomes

Toxicity	Utilities for the following disease statuses (efficacy):			
	PD	SD1	SD2	PR1CR
Mild	25	55	80	100
Moderate	20	35	70	90
High	10	25	50	70
Severe	0	10	25	40

manner. Unavoidably, in practice, the final distribution of patients among the 144 possible  $(\mathbf{d}, \mathbf{y})$  combinations will be very unbalanced. Consequently, a dose–outcome model  $\pi(\mathbf{y}, \mathbf{d}, \boldsymbol{\theta}) = \Pr(\mathbf{Y} = \mathbf{y} | \mathbf{d}, \boldsymbol{\theta})$ , parameterized by  $\boldsymbol{\theta}$ , must borrow strength across many possible  $(\mathbf{d}, \mathbf{y})$  values. We shall take the common practical approach of modelling the marginal probabilities  $\pi_k(y_k, \mathbf{d}, \boldsymbol{\theta}_k) = P(Y_k = y_k | \mathbf{d}, \boldsymbol{\theta}_k)$  for  $k = E$  and  $k = T$ , and using a bivariate copula (Nelsen, 2006) to induce association between  $Y_E$  and  $Y_T$  and to obtain  $\pi(\mathbf{y}, \mathbf{d}, \boldsymbol{\theta})$ .

Our goal in modelling the marginals is to obtain a dose finding design with desirable properties. Each marginal model must account for four outcome level main effects, two dose effects on each outcome level and possibly complex dose–dose interactions. The most difficult dose–outcome scenarios are those where the optimal pair  $\mathbf{d}$  is located in a middle portion of the two-dimensional domain, rather than at one of its four corners. To address these issues in a practical way, we assume a generalized continuation ratio (GCR) model (Fienberg, 1980; Cox, 1988) for each marginal. Our main departure from conventional approaches to constructing a dose finding model is that we standardize each agent’s dose parametrically in the linear term of each marginal. This gives a robust model that accounts for a wide variety of possible effects of  $\mathbf{d}$  on  $\pi_E(y_E, \mathbf{d}, \boldsymbol{\theta}_E)$  and  $\pi_T(y_T, \mathbf{d}, \boldsymbol{\theta}_T)$ .

Once the 16 possible elementary outcomes  $\mathbf{y} = (y_E, y_T)$  had been established, their numerical utilities  $U(\mathbf{y})$  were elicited from RGZ to quantify their relative desirability. These elicited utilities subsequently were reviewed by members of the Department of Investigational Cancer Therapeutics at the MD Anderson Cancer Center, and a consensus was obtained without changing any of the numerical values. In practice, utility elicitation may be carried out more formally by using the so-called ‘Delphi method’ (Dalkey, 1969; Brook *et al.*, 1986) or, for example, the methods that were described by Hunink *et al.* (2014) or Swinburn *et al.* (2010). Our elicited utilities are given in Table 2. A general admissibility criterion for any utility function  $U(y_E, y_T)$  in this setting is that it should increase as either  $y_E$  or  $y_T$  becomes more desirable on its ordinal scale, i.e. one should not use a utility function that does not make sense. These utilities are used during the trial as a basis for computing the posterior mean utility of each dose pair, which is the design’s optimality criterion. Adaptive randomization (AR) among nearly optimal dose pairs is used to avoid becoming stuck at a suboptimal pair (see Azriel *et al.* (2011) and Thall and Nguyen (2012)). Our simulations, which are given below in Section 4, show that parametrically standardizing the two doses and including them additionally in the model’s linear terms provides a robust basis for dose finding for a wide variety of  $\pi_k(\mathbf{y} | \mathbf{d})$  probability surfaces. In particular, our design’s performance compares favourably with what is obtained by assuming a more conventional model with multiplicative dose–dose interaction terms.

The dose–outcome model is given in Section 2. Decision criteria and algorithms for conduct of the trial are presented in Section 3. The methodology is applied to the motivating trial in Section 4, including a simulation study. We close with a discussion in Section 5.

## 2. Dose–response models

### 2.1. Parametric dose standardization

In phase I–II trials, a key issue is modelling the effects of intermediate doses on both  $\pi_E(y, \mathbf{d}, \theta_E)$  and  $\pi_T(y, \mathbf{d}, \theta_T)$ . First, consider a single-agent trial with lowest dose  $d_1$ , highest dose  $d_M$  and mean dose  $\bar{d}$ . For a given intermediate dose  $d_j$  between  $d_1$  and  $d_M$ , and each  $k = E, T$ , the actual value of  $\pi_k(d_j, \theta)$  may be, approximately, close to  $\pi_k(d_1, \theta)$ , midway between  $\pi_k(d_1, \theta)$  and  $\pi_k(d_M, \theta)$  or close to  $\pi_k(d_M, \theta)$ . If  $\pi_E$  and  $\pi_T$  both are defined by using the same standardized dose, say  $x = d - \bar{d}$  or  $d/\bar{d}$ , a problem arises from the facts that the shapes of the two curves  $\pi_E(x, \theta)$  and  $\pi_T(x, \theta)$  may be very different, and the desirabilities of an intermediate  $d_j$  in terms of  $\pi_E(x_j, \theta)$  and  $\pi_T(x_j, \theta)$  also may be very different. For example,  $d_j$  may have desirably low  $\pi_T(d_j, \theta)$  close to  $\pi_T(d_1, \theta)$ , and low, intermediate or high  $\pi_E(d_j, \theta)$ . An important case is one where  $\pi_T(d_j, \theta)$  is close to  $\pi_T(d_1, \theta)$  and  $\pi_E(d_j, \theta)$  is close to  $\pi_E(d_M, \theta)$ , so  $d_j$  is optimal for any reasonable criterion. If the model does not accurately reflect the different shapes of  $\pi_T(d, \theta)$  and  $\pi_E(d, \theta)$  as functions of  $d$ , the utility-based method may not select  $d_j$  with sufficiently high probability.

Next, consider a phase I–II combination trial. For each agent,  $a = 1, 2$ , denote the dose vector by  $d_a = (d_{a,1}, \dots, d_{a,M_a})$  with mean  $\bar{d}_a = (d_{a,1} + \dots + d_{a,M_a})/M_a$ . The modelling problem here is to characterize the joint effects of  $(d_{1,j}, d_{2,r})$  on both  $Y_E$  and  $Y_T$ . An intermediate dose pair is any  $\mathbf{d} = (d_{1,j}, d_{2,r})$  that is not located at one of the four corners of the rectangular dose pair domain, i.e.  $1 < j < M_1$  and  $1 < r < M_2$ . Standardizing each dose as  $x_{a,j} = d_{a,j} - \bar{d}_a$  or  $d_{a,j}/\bar{d}_a$  suffers from the same limitations as described above for an individual agent. Consequently, the problems that were described above for a single agent are more complex in that they now are elaborated in terms of the two probability surfaces  $\pi_E(d_{1,j}, d_{2,r})$  and  $\pi_T(d_{1,j}, d_{2,r})$ .

These problems motivate the use of two parametrically standardized versions of each dose: one with parameters corresponding to  $\pi_E$  and the other with parameters corresponding to  $\pi_T$ . For each outcome  $k = E, T$  and agent  $a$ , we define parametric dose standardization (PDS) for  $d_{a,j}$  to be

$$d_{k,a,j}^\lambda = \frac{d_{a,1}}{\bar{d}_a} + \left( \frac{d_{a,j} - d_{a,1}}{d_{a,M_a} - d_{a,1}} \right)^{\lambda_{k,a}} \frac{d_{a,M_a} - d_{a,1}}{\bar{d}_a} \quad (1)$$

where all entries of the dose standardization parameter vector  $\lambda = (\lambda_{E,1}, \lambda_{E,2}, \lambda_{T,1}, \lambda_{T,2})$  are positive valued. This construction gives two parametrically standardized versions of each dose of each agent: one for each outcome, mapping each  $d_{1,j}$  for agent 1 to  $(d_{E,1,j}^\lambda, d_{T,1,j}^\lambda)$ ,  $j = 1, \dots, M_1$ , and each  $d_{2,r}$  for agent 2 to  $(d_{E,2,r}^\lambda, d_{T,2,r}^\lambda)$ ,  $r = 1, \dots, M_2$ . Formula (1) is a two-agent version of that used by Thall *et al.* (2013) in the context of a design for optimizing the dose and schedule of one agent.

For each agent  $a$ , the lowest and highest standardized doses in equation (1) are  $d_{k,a,1}^\lambda = d_{a,1}/\bar{d}_a$  and  $d_{k,a,M_a}^\lambda = d_{a,M_a}/\bar{d}_a$ . Thus, the parametrically standardized doses at the lower and upper limits of the dose domain are usual standardized doses and do not depend on either  $\lambda$  or the outcome  $k$ . These serve as anchors for the intermediate doses,  $1 < j < M_a$ , where the PDS involves  $\lambda$  and  $k$ , and  $d_{k,a,j}^\lambda$  is a parametric, outcome-specific modification of the commonly used form  $d_{a,j}/\bar{d}_a$ , which corresponds to  $\lambda_{k,a} \equiv 1$ . Exponentiating the proportion  $(d_{a,j} - d_{a,1})/(d_{a,M_a} - d_{a,1})$  by the

model parameter  $\lambda_{k,a}$  in equation (1) shifts each intermediate dose  $d_{a,j}/\bar{d}_a$  either up towards  $d_{a,M_a}/\bar{d}_a$  or down towards  $d_{a,1}/\bar{d}_a$ . Since  $\lambda$  is updated along with the other model parameters in the posterior, formulation (1) provides a data-driven refinement of dose effects on each outcome that is not obtained if we use the usual standardized values  $d_{a,j}/\bar{d}_a$  or  $d_{a,j} - \bar{d}_a$ .

## 2.2. Generalized continuation ratio models

Reviews of GCR models, and of copulas that are used to obtain bivariate distributions having given marginals, are given in Appendix A. Given the PDS form (1), we may stabilize numerical computations by using either  $x_{k,a,j}^\lambda = \log(d_{k,a,j}^\lambda)$  or  $x_{k,a,j}^\lambda = d_{k,a,j}^\lambda - 1$  in the model's linear component. For a given dose pair  $(d_{1,j}, d_{2,r})$ , when no meaning is lost we shall suppress the dose indices  $j = 1, \dots, M_1$  and  $r = 1, \dots, M_2$ , and use the generic notation  $\mathbf{d} = (d_{1,j}, d_{2,r})$  and  $\mathbf{x}_k^\lambda = (x_{k,1,j}^\lambda, x_{k,2,r}^\lambda)$ . Denote the conditional probabilities

$$\gamma_k(y, \mathbf{d}, \boldsymbol{\theta}_k) = P(Y_k \geq y | Y_k \geq y - 1, \mathbf{d}, \boldsymbol{\theta}_k), \quad \text{for } k = \text{E, T}, \quad y = 0, \dots, L_k. \quad (2)$$

To construct the GCR model with PDS, we define the linear components

$$\eta_k(y_k, \mathbf{x}_k^\lambda, \boldsymbol{\theta}_k) = \alpha_{k,y} + \beta_{k,y,1} x_{k,1,j}^\lambda + \beta_{k,y,2} x_{k,2,r}^\lambda, \quad \text{for } k = \text{E, T}, \quad y_k = 1, \dots, L_k. \quad (3)$$

To enhance robustness, we use the parametric link function of Aranda-Ordaz (1981), which defines a probability  $p$  in terms of a real-valued linear term  $\eta$  and parameter  $\phi > 0$  as

$$p = 1 - \{1 + \phi \exp(\eta)\}^{-1/\phi}. \quad (4)$$

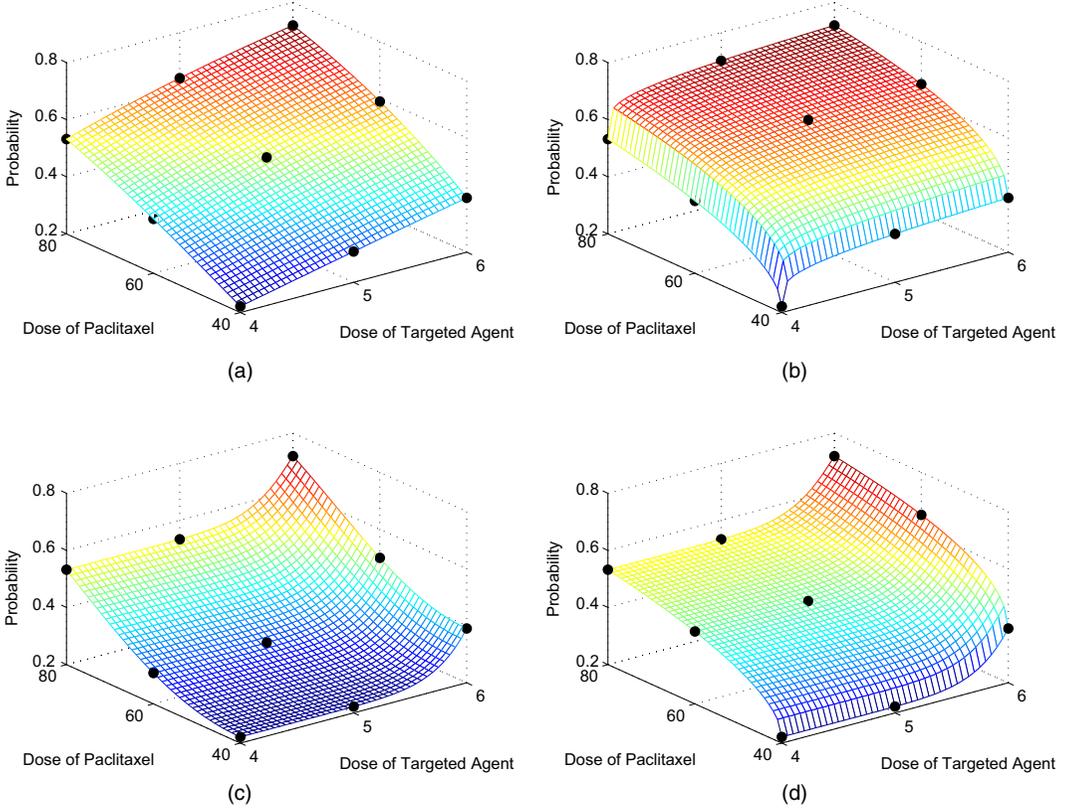
The Aranda-Ordaz link gives a very flexible model for  $p$  as a function of  $\eta$ , with  $\phi = 1$  corresponding to the logit link and the complementary log-log-link obtained as the limiting case when  $\phi \rightarrow 0$ . For the GCR model with PDS, we define the marginal of  $[Y_k | \mathbf{d}]$  by the equation

$$\gamma_k(y, \mathbf{x}_k^\lambda, \boldsymbol{\theta}_k) = 1 - [1 + \phi_k \exp\{\eta_k(y, \mathbf{x}_k^\lambda, \boldsymbol{\theta}_k)\}]^{-1/\phi_k} \quad \text{for } k = \text{E, T}, \quad y = 1, \dots, L_k,$$

i.e. we assume an Aranda-Ordaz link with PDS in the linear terms. We define  $\eta_k(0, \mathbf{x}_k^\lambda, \boldsymbol{\theta}_k) = \infty$  and  $\eta_k(L_k + 1, \mathbf{x}_k^\lambda, \boldsymbol{\theta}_k) = -\infty$  to ensure that  $\gamma_k(L_k + 1, \mathbf{x}_k^\lambda, \boldsymbol{\theta}_k) = 0$ . We require  $\beta_{k,y,1}, \beta_{k,y,2} > 0$  for each  $y \geq 1$  to ensure that  $\gamma_k(y, \mathbf{x}_k^\lambda, \boldsymbol{\theta}_k)$  increases with each dose. Writing  $\boldsymbol{\alpha}_k = \{\alpha_{k,y}, y = 1, 2, 3\}$  and  $\boldsymbol{\beta}_k = \{\beta_{k,y,a}, y = 1, 2, 3, a = 1, 2\}$ , the marginal parameter vector is  $\boldsymbol{\theta}_k = (\boldsymbol{\alpha}_k, \boldsymbol{\beta}_k, \lambda_{k,1}, \lambda_{k,2}, \phi_k)$ . The key components of the marginal model are that the linear components (3) include the doses of the two agents additively by using PDS (1), it has a GCR form (2) and it uses an Aranda-Ordaz link (4). In what follows, for brevity we shall abuse the notation slightly by identifying this model and the corresponding dose finding method by using the abbreviation ‘PDS’.

Since each intermediate standardized dose  $d_{k,a,j}^\lambda$  varies between the positive values  $d_{a,1}/\bar{d}_a$  and  $d_{a,M_a}/\bar{d}_a$ , we may consider 1 to be the middle numerical dose value. Mapping each  $d_{k,a,j}^\lambda$  to either  $x_{k,a,j}^\lambda = \log(d_{k,a,j}^\lambda)$  or  $x_{k,a,j}^\lambda = d_{k,a,j}^\lambda - 1$  has the same effect as centering the covariates at their means to reduce collinearity in conventional regression. Similarly, we define  $x_{k,a,j}^\lambda$  so that it varies around 0 rather than 1 to improve numerical stability. If, instead, we were to transform  $d_{k,a,j}^\lambda$  to maximize numerical stability at either the minimum or maximum of the dose domain, this would have the effect of destabilizing computations at the other end. Consequently, it is very desirable to transform  $d_{k,a,j}^\lambda$  to stabilize computations in the middle portion of the dose domain, and for values of  $\gamma_k(y, \mathbf{x}_k^\lambda, \boldsymbol{\theta}_k)$  near  $\frac{1}{2}$ . For  $x_{k,a,j}^\lambda = \log(d_{k,a,j}^\lambda)$ , this implies that

$$\exp\{\eta_k(y, \mathbf{x}_k^\lambda, \boldsymbol{\theta}_k)\} = \exp(\alpha_{k,y}) (d_{k,1,j}^\lambda)^{\beta_{k,y,1}} (d_{k,2,j}^\lambda)^{\beta_{k,y,2}},$$



**Fig. 1.** Illustration of the probability surface  $\gamma_E(1, \mathbf{d}, \theta_E)$  as a function of dose pair  $\mathbf{d}$ , for four different values of the dose standardization parameters  $(\lambda_{E,1}, \lambda_{E,2})$ : (a)  $\lambda_1 = \lambda_2 = 1$  (as a basis for comparison); (b)  $\lambda_1 = 0.2$ ,  $\lambda_2 = 0.5$ ; (c)  $\lambda_1 = 8$ ,  $\lambda_2 = 2$ ; (d)  $\lambda_1 = 8$ ,  $\lambda_2 = 0.5$

with  $\gamma_k(y, \mathbf{x}_k^\lambda, \theta_k) = \frac{1}{2}$  is obtained if  $\eta_k(y, \mathbf{x}_k^\lambda, \theta_k) = 0$  and  $\phi_k = 1$ , corresponding to a logit link in equation (2). In this case,  $\exp(\alpha_{k,y})(d_{k,1,j}^\lambda)^{\beta_{k,y,1}}(d_{k,2,j}^\lambda)^{\beta_{k,y,2}} = 1$ , and if  $d_{k,1,j}^\lambda = d_{k,2,j}^\lambda = 1$  then  $\alpha_{k,y} = 0$ . Thus, numerical stability is greatest in this dose pair neighbourhood, equivalently for  $x_{k,1,j}^\lambda = x_{k,2,j}^\lambda = 0$ . Alternatively, we could use  $x_{k,a,j}^\lambda = d_{k,a,j}^\lambda - 1$ .

Fig. 1 illustrates possible shapes of the probability surface  $\gamma_E(1, \mathbf{d}, \theta_E) = \Pr(Y_E \geq 1 | \mathbf{d}, \theta_E)$  as a function of the pair  $\mathbf{d} = (d_M, d_P) = (\text{dose of targeted agent}, \text{dose of paclitaxel})$ , for each of four different numerical dose standardization parameter pairs  $(\lambda_{E,1}, \lambda_{E,2})$ . The surface in Fig. 1(a) for  $\lambda_{E,1} = \lambda_{E,2} = 1$  corresponds to the additive model with linear term

$$\eta_E\{1, (d_{1,j}, d_{2,r}), \theta_k\} = \alpha_{E,1} + \beta_{E,1,1} \frac{d_{1,j}}{d_1} + \beta_{E,1,2} \frac{d_{2,r}}{d_2},$$

which may be used as a basis for visual comparison. Other probability surfaces as functions of  $\mathbf{d}$  may be drawn similarly, such as  $\gamma_E(y, \mathbf{d}, \theta_E)$ ,  $\gamma_T(y, \mathbf{d}, \theta_T)$ ,  $\pi_E(y, \mathbf{d}, \theta_E)$  or  $\pi_T(y, \mathbf{d}, \theta_T)$ , for integer  $y \geq 1$ . Fig. 1 shows that parametrically standardizing the doses in this way gives a very flexible model for the probabilities that are the basis for the dose finding design.

Index patients by  $i = 1, \dots, n$  for interim sample size  $n \leq N$ , and denote the dose pair given to the  $i$ th patient by  $\mathbf{d}_{[i]}$ . The likelihood is the product

**Table 3.** Elicited prior mean marginal outcome probabilities, for each dose pair

$(d_M, d_P)$	Probabilities for efficacy				Probabilities for toxicity			
	PD	SD1	SD2	PR/CR	Mild	Moderate	High	Severe
(4,40)	0.70	0.10	0.10	0.10	0.70	0.20	0.05	0.05
(5,40)	0.50	0.10	0.20	0.20	0.60	0.20	0.10	0.10
(6,40)	0.30	0.20	0.20	0.30	0.50	0.20	0.15	0.15
(4,60)	0.50	0.10	0.20	0.20	0.60	0.20	0.10	0.10
(5,60)	0.30	0.20	0.20	0.30	0.50	0.20	0.15	0.15
(6,60)	0.20	0.20	0.20	0.40	0.30	0.20	0.30	0.20
(4,80)	0.30	0.20	0.20	0.30	0.50	0.20	0.15	0.15
(5,80)	0.20	0.20	0.20	0.40	0.30	0.20	0.30	0.20
(6,80)	0.10	0.20	0.20	0.50	0.20	0.20	0.30	0.30

$$\mathcal{L}(\text{data}_n | \boldsymbol{\theta}) = \prod_{i=1}^n \pi(Y_{i,E}, Y_{i,T}, \mathbf{d}_{[i]}, \boldsymbol{\theta})$$

and the posterior is

$$p(\boldsymbol{\theta} | \text{data}_n) \propto \mathcal{L}(\text{data}_n | \boldsymbol{\theta}) p(\boldsymbol{\theta} | \tilde{\boldsymbol{\theta}}),$$

where  $p(\boldsymbol{\theta} | \tilde{\boldsymbol{\theta}})$  denotes the prior with fixed hyperparameters  $\tilde{\boldsymbol{\theta}}$ . Collecting terms, for  $k = E, T$ ,  $y = 1, 2, 3$  and  $a = 1, 2$ , the model parameters are  $\boldsymbol{\lambda} = \{\lambda_{k,a}\}$  for parametric dose standardization, the intercepts  $\boldsymbol{\alpha} = \{\alpha_{k,y}\}$ , the dose effects  $\boldsymbol{\beta} = \{\beta_{k,y,a}\}$ , the Aranda-Ordaz link parameters  $\boldsymbol{\phi} = \{\phi_E, \phi_T\}$  and the copula's association parameter  $\rho$ . Thus  $\boldsymbol{\theta} = (\boldsymbol{\lambda}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\phi}, \rho)$ .

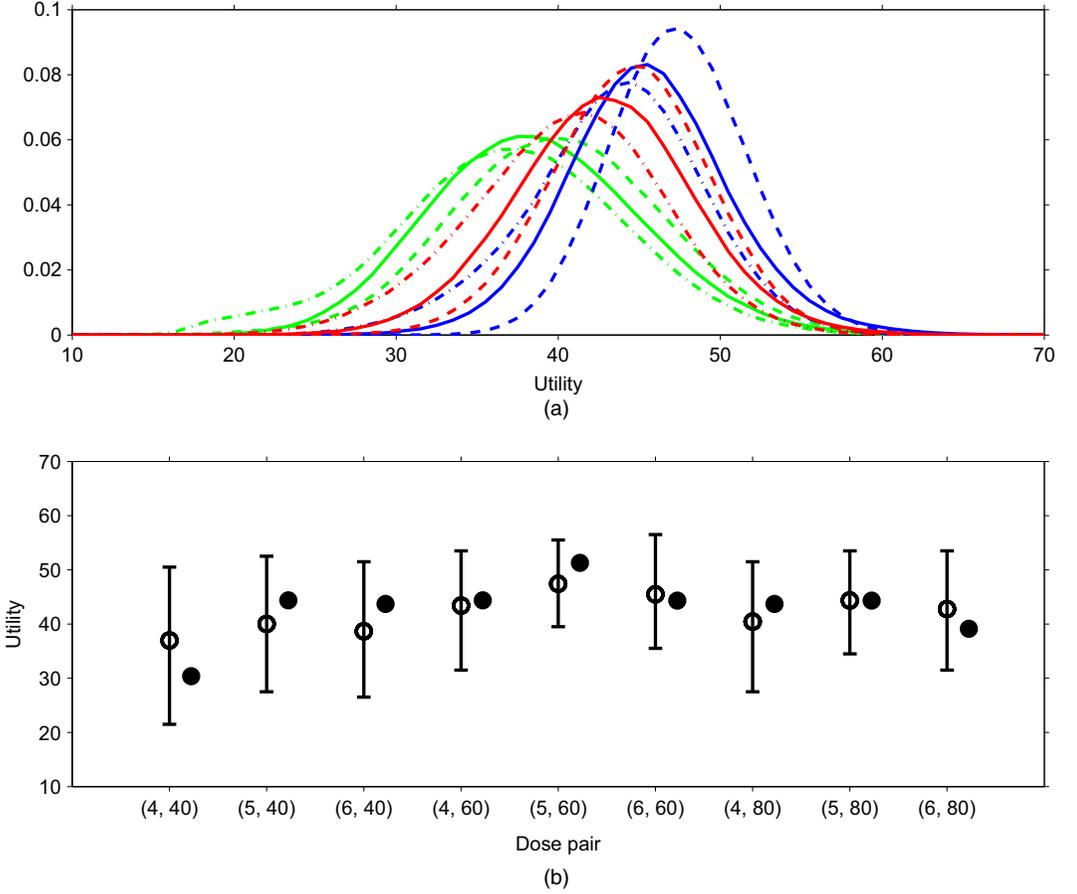
### 2.3. Establishing priors

Normal priors were assumed for the real-valued parameters  $\{\alpha_{k,y}\}$ , the positive-valued dose main effect coefficients  $\{\beta_{k,y,a}\}$  were assumed to follow normal priors truncated below at 0, the copula association parameter was assumed to be uniform on  $[-1, 1]$  and each  $\lambda_{k,a}$  and the Aranda-Ordaz link parameter  $\phi$  were assumed to follow log-normal priors. Prior means were estimated from the elicited probabilities given in Table 3 by using the pseudosampling method that was described in Thall *et al.* (2011), section 4.2, and Thall and Nguyen (2012), section 4.3. Prior variances were calibrated to make the effective sample size, as defined by Morita *et al.* (2008, 2010), of the prior of each marginal probability  $\pi_k(y, \mathbf{d}, \boldsymbol{\theta}_k)$  suitably small, and to give a design with good operating characteristics over a diverse set of scenarios. The effective sample size of each prior was approximated by equating the prior mean and variance of  $\pi_k(y, \mathbf{d}, \boldsymbol{\theta}_k)$  to the mean  $\tilde{\mu} = a/(a+b)$  and variance  $\tilde{\sigma}^2 = \tilde{\mu}(1-\tilde{\mu})/(a+b+1)$  of a beta( $a, b$ ) distribution. Thus,  $a+b$  was used to approximate the effective sample size of the prior of  $\pi_k(y, \mathbf{d}, \boldsymbol{\theta}_k)$ . The overall mean of these effective sample size values was 0.09 for the selected prior standard deviation of 20. Detailed descriptions of the prior parameters are given in the on-line supplementary Table S1.

## 3. Posterior decision criteria and trial design

### 3.1. Utility-based decision criteria

Given the Bayesian dose–outcome model and elicited numerical utilities  $U(\mathbf{y})$  in Table 2, the



**Fig. 2.** Posterior distributions of the mean utilities  $u(\theta|\text{data}_{60})$  for each of the nine dose pairs, based on a selected 60-patient data set obtained from one trial simulated under scenario 5: (a) posterior utilities (---, (4,40); - - -, (5,40); —, (6,40); - · - ·, (4,60); - - -, (5,60); —, (6,60); - · - ·, (4,80); - - -, (5,80); —, (6,80)); (b) posterior mean utilities with 95% credible intervals (●, true mean utility)

trial is conducted by using the following decision criteria. Given  $\theta$ , the mean utility of dose pair  $\mathbf{d}$  is

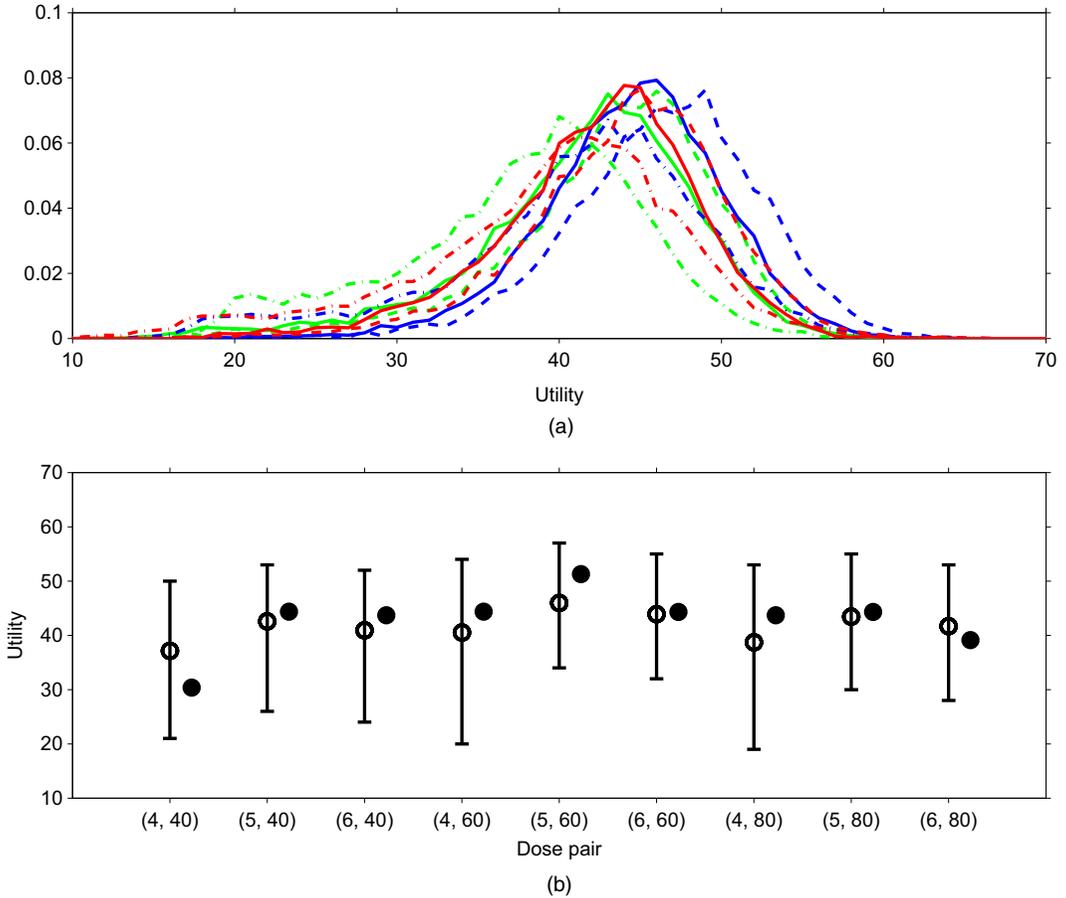
$$\bar{U}(\mathbf{d}, \theta) = \sum_{\mathbf{y}} U(\mathbf{y}) \Pr(\mathbf{Y} = \mathbf{y} | \mathbf{d}, \theta),$$

where the sum is over all  $\mathbf{y}$ -pairs in the support of  $\mathbf{Y}$ . Since  $\theta$  is not known, we compute each dose pair's posterior mean utility

$$u(\mathbf{d} | \text{data}_n) = \int_{\theta} \bar{U}(\mathbf{d}, \theta) p(\theta | \text{data}_n) \quad (5)$$

given the data on  $n$  patients available when an interim decision must be made. This integral is approximated by generating a posterior sample  $\theta^{(1)}, \dots, \theta^{(M)}$  by using Markov chain Monte Carlo sampling (Robert and Casella, 1999) and computing the sample mean of  $\bar{U}(\mathbf{d}, \theta^{(1)}), \dots, \bar{U}(\mathbf{d}, \theta^{(M)})$ .

The posterior mean utilities that are given by equation (5) are the basis for the design's



**Fig. 3.** Distributions of the final posterior mean utility  $u(\theta|\text{data}_{60})$  (---, (4,40); - - -, (5,40); —, (6,40); ···, (4,60); - - -, (5,60); —, (6,60); - · - ·, (4,80); - - -, (5,80); —, (6,80)) and (b) 95% probability intervals (●, true mean utility) for each of the nine dose pairs for the PDS-model-based method proposed, based on a sample of 10000 trials, each of size  $n = 60$ , with the data for each trial generated under scenario 5

sequential decision rules to select dose pairs during the trial. It is very important to bear in mind that each posterior mean utility is a statistic that can be quite variable. This is illustrated by Fig. 2, which plots the distributions of  $u(\mathbf{d}|\text{data}_{60})$  and corresponding 95% probability intervals for each of the nine dose pairs, based on one 60-patient data set from a trial simulated under scenario 5. To illustrate how such final utility distributions may vary across trials, Fig. 3 provides similar plots based on a sample of 10000 trials, each of size  $n = 60$ , with the data generated under scenario 5. From a Bayesian perspective, the randomness of each distribution in Fig. 2 is due to posterior uncertainty about  $\theta$ , whereas the randomness of each distribution in Fig. 3 is due to the random variation in the data. It is also important to bear in mind that, for the smaller sample sizes that are the basis for interim decisions during the trial, the variability of  $u(\mathbf{d}|\text{data}_n)$  for each  $\mathbf{d}$  is greater than that shown by Fig. 2 for the final data of  $n = 60$  patients. In general, the substantial variability of each  $u(\mathbf{d}|\text{data}_n)$  also would be the case for any statistic that is used as a decision criterion in this or similar small-scale trial settings using any other adaptive design. These considerations motivate, in part, our use of adaptive randomization

between nearly optimal dose pairs in the trial design. The general point is that, in early phase trials, decision making must be done under great uncertainty.

### 3.2. Dose acceptability criteria and adaptive randomization

To ensure that the trial is ethically acceptable, rather than simply choosing  $\mathbf{d}$  from the nine pairs to maximize  $u(\mathbf{d}|\text{data}_n)$ , we impose additional constraints to ensure that any dose pair that is used to treat patients is both acceptably safe and acceptably efficacious. This follows the approach that was used by Thall and Cook (2004) and many others. We use the following two posterior acceptability criteria. For each  $k = E$  or  $k = T$ , denote  $\bar{\pi}_k(y, \mathbf{d}, \boldsymbol{\theta}_k) = \Pr(Y_k \geq y | \mathbf{d}, \boldsymbol{\theta}_k)$ . Indexing the toxicity levels by  $y = 0, 1, 2, 3$  for mild, moderate, high and severe,  $\bar{\pi}_T(2, \mathbf{d}, \boldsymbol{\theta})$  is the probability of high or severe toxicity with  $\mathbf{d}$ . A dose pair  $\mathbf{d}$  is considered *unacceptably toxic* if

$$\Pr\{\bar{\pi}_T(2, \mathbf{d}, \boldsymbol{\theta}) > 0.45 | \text{data}_n\} > 0.90, \quad (6)$$

i.e.  $\mathbf{d}$  is not acceptable if, on the basis of the current data, it is likely that  $\mathbf{d}$  has a probability of high or severe toxicity that is above 0.45. For the efficacy rule, we similarly index the outcomes  $\{\text{PD}, \text{SD1}, \text{SD2}, \text{PR/CR}\}$  by 0, 1, 2 and 3, so that  $\bar{\pi}_E(2, \mathbf{d}, \boldsymbol{\theta})$  is the probability of SD2 or better with dose pair  $\mathbf{d}$ . A dose pair  $\mathbf{d}$  is considered *unacceptably inefficacious* if

$$\Pr\{\bar{\pi}_E(2, \mathbf{d}, \boldsymbol{\theta}) < 0.40 | \text{data}_n\} > 0.90. \quad (7)$$

This says that  $\mathbf{d}$  is not acceptable if, given the current data, it is likely that achieving SD2 or better occurs at a rate below 40%. A dose pair  $\mathbf{d}$  is considered *acceptable* if it has both acceptable toxicity and acceptable efficacy, and we denote the set of acceptable dose pairs based on  $\text{data}_n$  by  $\mathcal{A}_n$ . As data are acquired during the trial and the posterior becomes more reliable,  $\mathcal{A}_n$  may change, so a given  $\mathbf{d}$  that is not in  $\mathcal{A}_n$  may be in  $\mathcal{A}_{n+k}$ , or conversely. The events that are used to define conditions (6) and (7) and the corresponding numerical probabilities 0.45 and 0.40 are specific to the solid tumour trial. These particular values were determined by RGZ in collaboration with oncologist colleagues who are involved in planning the trial. In other trials, different toxicity and efficacy events and probability cut-offs should be chosen as appropriate.

Given the acceptability criteria, it may seem that we simply may choose the  $\mathbf{d} \in \mathcal{A}_n$  that maximizes  $u(\mathbf{d}|\text{data}_n)$ . This may lead to a design with undesirable properties, in some cases, due to the well-known ‘optimization-versus-exploration’ dilemma in sequential decision making (see Gittins (1979) and Sutton and Barto (1998)). The problem is that, given some optimality criterion, a ‘greedy’ sequential decision rule that always takes the empirically optimal action on the basis of the current data carries a risk of becoming stuck at a truly suboptimal action. The problem that greedy sequential algorithms are ‘sticky’ in this sense only recently has been discussed in the context of dose finding trials, by Azriel *et al.* (2011), Thall and Nguyen (2012), Oron and Hoff (2013), Braun *et al.* (2013) and Thall *et al.* (2014).

We address the problem of stickiness by applying AR among  $\mathbf{d}$  having  $u(\mathbf{d}|\text{data}_n)$  close to the maximum, similarly to Thall and Nguyen (2012). Denote the acceptable dose pair maximizing  $u(\mathbf{d}|\text{data}_n)$  by  $\mathbf{d}_n^{\text{opt}}$ . Although nominally this dose pair is ‘optimal’, it is only empirically optimal on the basis of the most recent data, and it may not be the truly optimal pair that would maximize  $\bar{U}(\mathbf{d}, \boldsymbol{\theta})$  if  $\boldsymbol{\theta}$  were known. In practice the truly optimal dose pair cannot be known, but in a simulation study all assumed  $\pi^{\text{true}}(y|\mathbf{d})$  are specified, so the  $\mathbf{d}^{\text{opt}}$  under this assumed state of nature is known, and design performance can be evaluated accordingly. Although this distinction may seem obvious, the difference between an empirically optimal action and the truly optimal action is at the heart of the optimization-versus-exploration dilemma. A general form for AR probabilities for dose pair  $\mathbf{d}^*$  based on the posterior mean utilities of the acceptable dose pairs is

$$r_n(\mathbf{d}^*) = \frac{u(\mathbf{d}^* | \text{data}_n)}{\sum_{\mathbf{d} \in \mathcal{A}_n} u(\mathbf{d} | \text{data}_n)}.$$

We studied several modified versions of AR, called AR( $m$ ), which is limited to randomizing between only the best  $m$  dose pairs on the basis of their current posterior mean utilities, for  $m = 1$  (a greedy design with no AR), 2, 3, 4, 9. The results are summarized in on-line supplementary Table S5. Additionally, we studied the required difference between the subsample sizes of the empirically best and other acceptable dose pairs, to ensure that an adequate number of patients have been treated at  $\mathbf{d}_n^{\text{opt}}$  before applying any AR rule. On the basis of this preliminary study, for the actual trial design, we used AR(2), with AR applied only if at least three or more patients have been treated at the current  $\mathbf{d}_n^{\text{opt}}$  than at any other acceptable  $\mathbf{d}$ . Denote the empirically second-best acceptable dose pair by  $\mathbf{d}_n^{\text{second}}$ , i.e.  $u(\mathbf{d}_n^{\text{second}} | \text{data}_n)$  is the second-largest posterior mean utility. For our implementation of AR(2), the next cohort of patients are treated with dose pair  $\mathbf{d}_n^{\text{opt}}$  with probability

$$r_n = \frac{u(\mathbf{d}_n^{\text{opt}} | \text{data}_n)}{u(\mathbf{d}_n^{\text{opt}} | \text{data}_n) + u(\mathbf{d}_n^{\text{second}} | \text{data}_n)},$$

and treated with dose pair  $\mathbf{d}_n^{\text{second}}$  with probability  $1 - r_n$ .

### 3.3. Trial conduct

Using the above decision criteria, the trial is conducted as follows. Recall that the maximum sample size is  $N = 60$ , and the cohort size is  $c = 3$ .

*Step 1:* the first cohort is treated at  $\mathbf{d} = (d_M, d_P) = (4, 60)$ .

*Step 2:* for each cohort after the first, the posterior decision criteria (5), (6) and (7) are computed on the basis of the most current data.

*Step 3:* when escalating, an untried dose of either agent may not be skipped.

*Step 4:* if no  $\mathbf{d}$  is acceptable, the trial is terminated with no  $\mathbf{d}$  selected.

*Step 5:* if exactly one  $\mathbf{d}$  is acceptable, the next cohort is treated at that dose pair.

*Step 6:* for cohort size  $c$ , if two or more  $\mathbf{d}$ s are acceptable and the number of patients treated at  $\mathbf{d}_n^{\text{opt}}$  minus the largest number of patients treated at any other acceptable dose is

- (a)  $c$  or greater, then apply AR(2) to choose randomly between  $\mathbf{d}_n^{\text{opt}}$  and  $\mathbf{d}_n^{\text{second}}$ , or
- (b) less than  $c$ , then treat the next cohort at  $\mathbf{d}_n^{\text{opt}}$ .

## 4. Simulations

### 4.1. General design performance evaluation

The trial design was simulated under each of 12 dose–outcome scenarios, given in the on-line supplementary Table S2, assuming an accrual rate of 1.5 patients per month. Each scenario is specified in terms of fixed true four-level marginal efficacy and toxicity probabilities, which are not based on the design's model or any other model. Association was induced by assuming a Gaussian copula with true association parameter 0.10. Additional simulations were conducted by using alternative models, or different cohort size or maximum sample size. For each case studied, the trial was replicated 3000 times, and all posterior quantities were computed by using Markov chain Monte Carlo with Gibbs sampling.

We use the following summary statistics, given by Thall and Nguyen (2012), to quantify overall design performance. For given  $\mathbf{d}$  and assumed true outcome probabilities  $\{\pi^{\text{true}}(\mathbf{y} | \mathbf{d})\}$ ,

we define the *true mean utility of  $\mathbf{d}$*  to be

$$\bar{U}^{\text{true}}(\mathbf{d}) = \sum_{\mathbf{y}} U(\mathbf{y}) \pi^{\text{true}}(\mathbf{y}|\mathbf{d}).$$

Thus,  $\bar{U}^{\text{true}}(\mathbf{d})$  is analogous to, but different from, the mean utility  $\bar{U}(\mathbf{d}, \boldsymbol{\theta})$  based on the unknown parameter  $\boldsymbol{\theta}$ , and the posterior mean utility  $u(\mathbf{d}|\text{data}_n)$ , which is a statistic. Let  $\bar{U}_{\text{max}}^{\text{true}}$  and  $\bar{U}_{\text{min}}^{\text{true}}$  denote the largest and smallest possible true mean utilities among all dose pairs. To quantify the method's reliability for selecting a dose pair with high true utility, which benefits future patients, denoting the final selected dose pair by  $\mathbf{d}_{\text{select}}$ , we use the statistic

$$R_{\text{select}} = 100 \left\{ \frac{\bar{U}^{\text{true}}(\mathbf{d}_{\text{select}}) - \bar{U}_{\text{min}}^{\text{true}}}{\bar{U}_{\text{max}}^{\text{true}} - \bar{U}_{\text{min}}^{\text{true}}} \right\}.$$

To quantify benefit to the patients enrolled in the trial, we use the statistic

$$R_{\text{treat}} = 100 \left\{ \frac{(1/N) \sum_{i=1}^N \bar{U}^{\text{true}}(\mathbf{d}_{[i]}) - \bar{U}_{\text{min}}^{\text{true}}}{\bar{U}_{\text{max}}^{\text{true}} - \bar{U}_{\text{min}}^{\text{true}}} \right\},$$

where  $\mathbf{d}_{[i]}$  is the dose pair given to the  $i$ th patient, and  $N$  is the final sample size. For both statistics, a larger value in the domain  $[0, 100]$  corresponds to better performance. We report also the selection percentage of the best acceptable  $\mathbf{d}$ , denoted by %Best.

Simulation results for six selected scenarios are summarized in Table 4. The results for all 12 scenarios are given in the on-line supplementary Table S3 and supplementary Fig. S1. In terms of true utilities and selection percentages of the nine dose pairs, Table 4 shows that the design does a reliable job of selecting acceptable dose pairs having true mean utility at or near the maximum, while also reliably avoiding unacceptable dose pairs. Fig. 4 illustrates how the utility function  $U(\mathbf{y})$  maps the eight assumed true outcome probability pairs ( $\pi_{\text{E}}^{\text{true}}(y_{\text{E}}, \mathbf{d})$ ,  $\pi_{\text{T}}^{\text{true}}(y_{\text{T}}, \mathbf{d})$ ) for  $y_{\text{E}} = 0, 1, 2, 3$  and  $y_{\text{T}} = 0, 1, 2, 3$  to  $\bar{U}^{\text{true}}(\mathbf{d})$  for each  $\mathbf{d}$ , in scenario 5. For each outcome, the assumed probabilities  $\pi_k^{\text{true}}(y, \mathbf{d})$  for  $y = 0, 1, 2, 3$  are represented by successively darker shades of red for  $k = \text{T}$  and green for  $k = \text{E}$ . Fig. 4 shows, for the PDS-model-based design, how the dose pair selection probabilities follow the magnitudes of the true mean utilities. A key point is that, if we wish to compare dose pairs, inevitably a one-dimensional criterion is needed. The utility function provides this in a way that makes sense medically, provided that we accept the particular numerical utilities that are given in Table 2.

#### 4.2. Comparison with models with qualitatively different dose–dose effects

The generalized Aranda-Ordaz (GAO) model that was used by the two-agent phase I–II design of Houede *et al.* (2010) to account for dose–dose interactions is given in Appendix A. As noted earlier, because this design addresses the same problem of choosing optimal  $\mathbf{d}$  on the basis of ordinal  $(Y_{\text{E}}, Y_{\text{T}})$ , it is a natural comparator to the PDS-model-based design that is proposed here. Another comparator may be obtained from the more conventional model formulation in which all  $\lambda_{k,a} = 1$  in the PDS linear components and a multiplicative dose–dose interaction term is included in the linear term, using the usual standardized doses  $x_{a,j} = \log(d_{a,j}/\bar{d}_a)$ . The linear components then would take the commonly assumed form

$$\eta_k(y, \mathbf{d}, \boldsymbol{\theta}_k) = \alpha_{k,y} + \beta_{k,y,1}x_{1,j} + \beta_{k,y,2}x_{2,r} + \beta_{k,12}x_{1,j}x_{2,r}, \quad k = \text{E}, \text{T}.$$

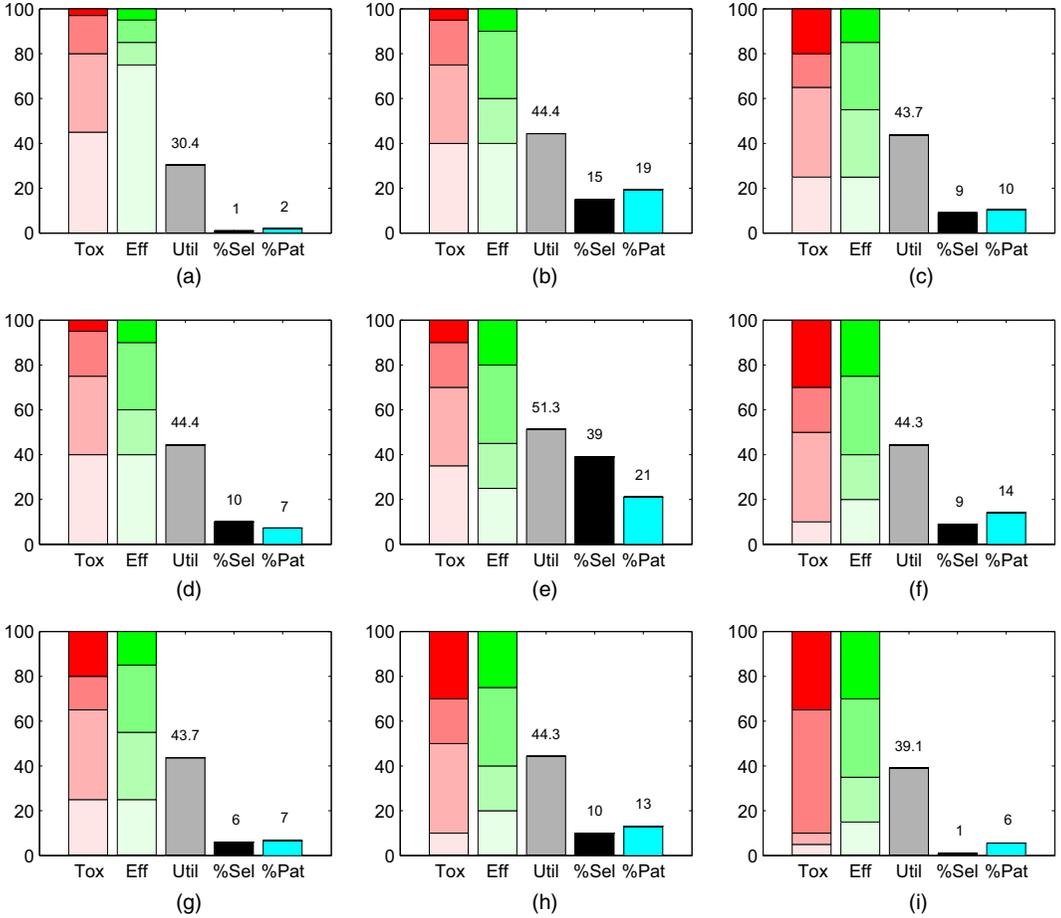
**Table 4.** Simulation results for the PDS–GCR partial orders model-based design†

	$d_M$	Results for the following values of $d_P$ :			Results for the following values of $d_P$ :		
		40	60	80	40	60	80
		<i>Scenario 1</i>			<i>Scenario 2</i>		
$\bar{U}^{\text{true}}(\mathbf{d})$	4	<b>56.0</b>	51.8	48.3	43.2	44.4	37.9
Selection percentage		32	25	9	8	17	3
Number of patients		9.4	14.8	8.4	3.5	12.2	5.3
$\bar{U}^{\text{true}}(\mathbf{d})$	5	51.8	47.2	44.7	<b>49.7</b>	46.8	38.9
Selection percentage		15	7	2	28	24	2
Number of patients		7.4	6.8	3.2	8.7	10.5	5.2
$\bar{U}^{\text{true}}(\mathbf{d})$	6	48.3	44.7	39.4	45.5	39.7	33.6
Selection percentage		8	2	0	10	5	0
Number of patients		6.0	2.6	0.9	6.7	5.6	1.7
% none selected			2			2	
		<i>Scenario 3</i>			<i>Scenario 5</i>		
$\bar{U}^{\text{true}}(\mathbf{d})$	4	39.4	40.1	36.7	30.4	44.4	43.7
Selection percentage		1	3	1	1	15	9
Number of patients		1.1	8.6	3.2	1.2	11.7	6.3
$\bar{U}^{\text{true}}(\mathbf{d})$	5	48.9	47.6	42.7	44.4	<b>51.3</b>	44.3
Selection percentage		12	17	3	10	39	9
Number of patients		6.1	9.1	5.7	4.4	12.7	8.5
$\bar{U}^{\text{true}}(\mathbf{d})$	6	<b>52.6</b>	49.8	44.6	43.7	44.3	39.1
Selection percentage		32	27	3	6	10	1
Number of patients		10.0	11.2	4.7	4.0	7.8	3.4
% none selected			1			1	
		<i>Scenario 8</i>			<i>Scenario 9</i>		
$\bar{U}^{\text{true}}(\mathbf{d})$	4	33.8	45.4	48.2	43.8	50.8	52.6
Selection percentage		1	9	15	1	2	4
Number of patients		0.8	10.0	7.6	0.5	8.3	5.3
$\bar{U}^{\text{true}}(\mathbf{d})$	5	37.2	48.9	<b>53.2</b>	50.8	52.6	58.4
Selection percentage		2	21	33	1	4	17
Number of patients		2.6	9.7	12.6	1.6	5.9	11.5
$\bar{U}^{\text{true}}(\mathbf{d})$	6	41.3	45.9	45.6	52.6	58.4	<b>64.0</b>
Selection percentage		3	9	6	4	18	48
Number of patients		2.5	6.2	7.9	3.4	8.6	14.7
% none selected			1			0	

†For each dose pair  $\mathbf{d} = (d_M, d_P)$ , selection percentage and number of patients treated. Utilities of unacceptable doses are in italics. The highest utility among acceptable doses is given in bold.

The  $\beta_{k,12s}$  are real valued and assumed to have normal priors. The element  $\beta_{k,12x_1,jx_2,r}$  of this linear term is widely considered to be an ‘interaction’ between two covariates in their joint effect on the outcome in a regression model. Here, the interaction is the joint effect of  $d_1$  and  $d_2$  on the marginal probability distribution of  $Y_k$ . We shall refer to this as the conventional multiplicative interaction (CMI) model.

Table 5 summarizes how well the design performs assuming each of these three alternative models, for (4,4) dimensional bivariate ordinal outcomes. All three designs reliably stop the trial early if no  $\mathbf{d}$ -pairs are acceptable, in scenarios 11 and 12. For scenarios 1–10, Fig. 5 shows the comparative  $R_{\text{select}}$ -results graphically. In the five scenarios  $\{2, 4, 5, 6, 8\}$  where  $\mathbf{d}^{\text{opt}}$  is a middle dose pair, not located at one of the four corners of the  $3 \times 3$  matrix of  $\mathbf{d}$ -pairs, the PDS



**Fig. 4.** Illustration of true marginal outcome probabilities  $\{\pi_T^{\text{true}}(y, \mathbf{d}), \pi_E^{\text{true}}(y, \mathbf{d}), y = 0, 1, 2, 3\}$ , the resulting true mean utility  $U^{\text{true}}(\mathbf{d})$ , and simulation results percentage selection %Sel and percentage of patients treated %Pat in the trial for each dose pair, using the PDS-model-based method proposed, under scenario 5 ( $\pi_k^{\text{true}}(y, \mathbf{d})$  for  $y = 0, 1, 2, 3$  are represented by successively darker shades of red for  $k = T$  and green for  $k = E$ ): (a) dose pair (4,40); (b) dose pair (4,60); (c) dose pair (4,80); (d) dose pair (5,40); (e) dose pair (5,60); (f) dose pair (5,80); (g) dose pair (6,40); (h) dose pair (6,60); (i) dose pair (6,80)

model gives much larger  $R_{\text{select}}$ -values than the other two designs. The differences  $R_{\text{select}}(\text{PDS}) - R_{\text{select}}(\text{GAO})$  vary from 5 to 13 (7–20%), whereas  $R_{\text{select}}(\text{PDS}) - R_{\text{select}}(\text{CMI})$  vary from 9 to 16 (13–26%). In the four scenarios  $\{1, 3, 7, 9\}$  where  $\mathbf{d}^{\text{opt}}$  is located at one of the four corners of the matrix of  $\mathbf{d}$ -pairs, the GAO and the CMI model give  $R_{\text{select}}$ -values that are larger than those of the PDS model by the smaller differences 5–7 (6–9%). The  $R_{\text{treat}}$  and %Best  $\mathbf{d}$  selected values also follow these general patterns. Scenario 10 corresponds to the prior and has three acceptable dose pairs all having the same maximum true utility. An important property of the PDS method is that it gives much more stable behaviour across scenarios 1–10, with  $R_{\text{select}}$ -values in the range [76, 89], compared with ranges [65, 88] for the GAO method and [62, 92] for the CMI model. Similarly, the %Best  $\mathbf{d}$  selected values have range [28, 48] for the PDS method *versus* ranges [8, 67] for the GAO model and [2, 68] for the CMI model. It thus appears that using parametrically standardized doses gives much more stable behaviour across a range of scenarios and provides

**Table 5.** Comparison of design performance by using three alternative GCR models for (4,4) dimensional bivariate ordinal outcomes†

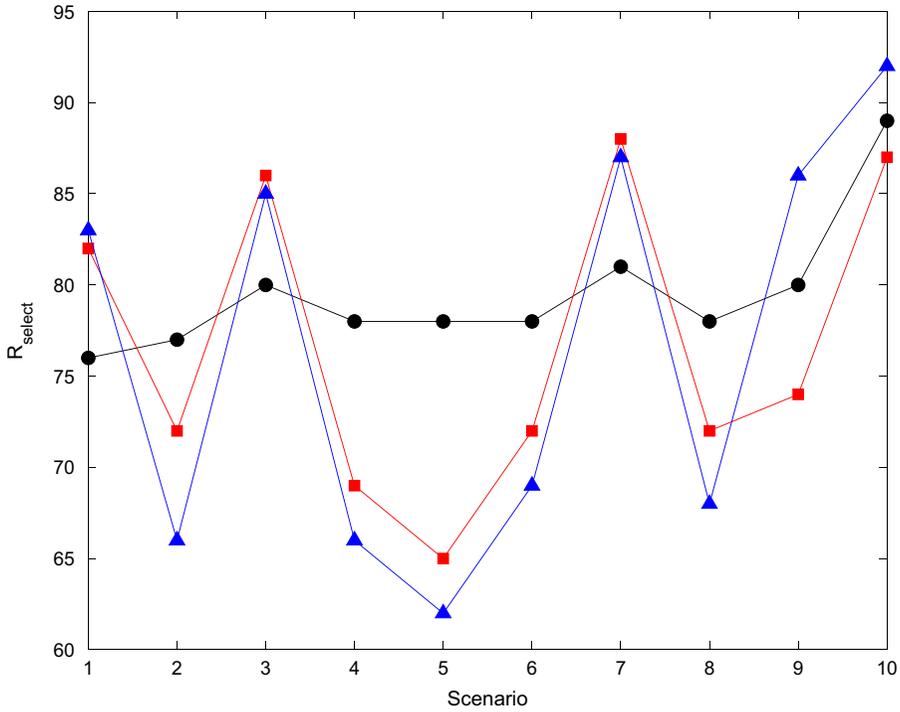
Parameter	Results for the following scenarios:											
	1	2	3	4	5	6	7	8	9	10	11	12
<i>PDS</i>												
$R_{\text{select}}$	76	77	80	78	78	78	81	78	80	89	93	94
$R_{\text{treat}}$	65	63	61	72	70	67	70	70	65	85	79	67
%None	2	2	1	0	1	7	1	1	0	0	95	96
%Best	32	28	32	46	39	34	39	33	48	44	—	—
Number of patients	59.4	59.4	59.6	59.9	59.8	57.9	59.7	59.8	59.9	59.9	26.5	27.7
Efficacy number	36.9	27.1	29.3	29.4	30.2	24.1	29.8	30.0	31.9	31.9	13.2	5.0
Toxicity number	28.6	25.2	23.6	19.2	22.9	20.8	22.1	22.1	19.8	22.6	18.1	7.7
<i>GAO model</i>												
$R_{\text{select}}$	82	72	86	69	65	72	88	72	74	87	93	95
$R_{\text{treat}}$	71	65	63	68	65	63	75	68	57	82	79	66
%None	2	3	1	0	1	6	1	1	1	0	94	94
%Best	49	25	64	31	8	19	67	22	47	54	—	—
Number of patients	59.5	59.3	59.7	60.0	59.8	58.4	59.7	59.7	59.8	59.9	27.1	32.2
Efficacy number	36.2	25.8	28.1	27.9	28.3	23.3	29.2	28.7	30.9	30.5	13.5	5.7
Toxicity number	25.7	22.9	21.2	17.6	22.1	20.6	19.7	20.8	19.0	20.9	18.5	8.9
<i>CMI</i>												
$R_{\text{select}}$	83	66	85	66	62	69	87	68	86	92	95	97
$R_{\text{treat}}$	69	60	61	67	62	64	73	68	65	85	80	68
%None	1	3	1	0	1	6	1	1	1	0	91	95
%Best	56	14	62	20	2	9	65	11	68	50	—	—
Number of patients	59.7	59.3	59.7	59.9	59.7	58.3	59.7	59.7	59.8	60.0	29.3	29.0
Efficacy number	36.8	26.0	28.8	29.4	28.9	24.0	29.9	29.9	31.9	32.2	14.6	5.3
Toxicity number	26.9	24.8	23.1	19.6	24.2	21.6	22.0	23.0	19.7	22.9	19.9	8.1

†Scenarios 11 and 12 have no acceptable dose, so  $R_{\text{select}}$ -values are less relevant and thus are in italics.

insurance against very poor performance in some scenarios. The PDS model gives substantially larger  $R_{\text{select}}$ -values in the more difficult cases where  $\mathbf{d}^{\text{opt}}$  is a middle dose pair, with the price being smaller  $R_{\text{select}}$ -values in the easier cases where  $\mathbf{d}^{\text{opt}}$  is at a corner of the rectangular dose pair domain.

### 4.3. Comparison with designs that reduce the ordinal outcomes

We next compare our proposed method, based on the (4,4) dimensional ordinal outcome  $\mathbf{Y} = (Y_E, Y_T)$ , with alternative designs that reduce this outcome by combining categories. The first two comparators are versions of the PDS and GAO designs based on (3,3) ordinal outcomes that are obtained by combining SD2 and CR/PR for  $Y_E$  and combining high and severe events for  $Y_T$ . We obtained a (2,2) outcome by also combining the  $Y_E$ -events PD and SD1 so that  $Y_E$  became the binary indicator of [CR/PR or SD2], and combining the  $Y_T$ -events mild and moderate so that  $Y_T$  became the binary indicator of [high or severe]. For each of these (3,3) and (2,2) cases, in each scenario the outcome probabilities were obtained from those in on-line supplementary Table S2 by summing the corresponding elementary event probabilities. For the (2,2) case, in addition to the reduced version of the PDS design, we also included as comparators the phase I–II designs of Yuan and Yin (2011), YY, and Wages and Conaway (2014), WC, both of which rely on bivariate binary  $\mathbf{Y}$ . A final comparator is the partial orders CRM of Wages *et al.* (2011), which uses only a binary version of  $Y_T$  to choose optimal  $\mathbf{d}$ .



**Fig. 5.**  $R_{\text{select}}$ -values for designs based on three different bivariate ordinal outcome models (●, PDS; ■, GAO; ▲, CMI) that account for dose–dose interactions differently: all three designs determine an optimal dose pair based on (4,4) dimensional bivariate ordinal outcomes under GCR models for the marginals

The YY design uses a copula to model the probability of toxicity as a function of  $\mathbf{d}$  in phase I and chooses a set of admissible  $\mathbf{d}$  for subsequent efficacy evaluation in parallel treatment arms in phase II. The design applies AR on the basis of the probability of a binary efficacy outcome in phase II, assuming a hierarchical binomial–beta–gamma model. At the end of phase II, the YY design selects the dose pair with acceptable toxicity that has highest posterior mean efficacy. Since the YY design allows us to vary the cohort size  $c$  and subsample sizes  $n_{\text{I}}$  and  $n_{\text{II}}$  in phases I and II, for comparison with the PDS-model-based design, we first simulated versions of the YY design with  $(c, n_{\text{I}}, n_{\text{II}}) = (3, 30, 30), (1, 30, 30), (1, 20, 40)$ , given in on-line supplementary Table S8. Since the YY design with  $(c, n_{\text{I}}, n_{\text{II}}) = (1, 30, 30)$  has slightly better overall performance than the other two, this version is used for comparison with the PDS design.

The WC design is based on partial orderings of  $\mathbf{d}$ . Like the YY design, the WC design also chooses the dose pair  $\mathbf{d}$  with acceptable toxicity that maximizes the probability of efficacy. We simulated both the YY and the WC designs by using the same toxicity probability acceptability upper limit, 0.45, and efficacy probability lower limit, 0.40, as those used by the PDS design. Since the total number of possible partial orderings in the rectangle of  $\mathbf{d}$ -pairs is impractically large, a subset must be chosen. For comparison with the PDS-model-based design, we first simulated versions of the WC design with either six partial orderings, starting the trial at  $\mathbf{d} = (1, 2)$  as in our design, or 26 partial orderings, starting the trial at either  $\mathbf{d} = (1, 2)$  or  $\mathbf{d} = (1, 1)$ , summarized in on-line supplementary Table S9. Since the version with 26 partial orderings, starting the trial at  $\mathbf{d} = (1, 2)$ , had slightly better overall performance than the other two, it is included in Table 6.

An important point is that both the YY and the WC designs choose  $\mathbf{d}$  that has acceptably

low toxicity and maximum efficacy, whereas the PDS design chooses  $\mathbf{d}$  that has acceptably low toxicity, acceptably high efficacy and maximum posterior mean utility, i.e. the criteria are qualitatively different. The three designs have the same ‘best’  $\mathbf{d}$  in scenarios 5, 6, 8, 9 and 10, and different best  $\mathbf{d}$  in scenarios 1, 2, 3, 4 and 7. To compare the methods, we used the same utility-based criteria, namely  $R_{\text{select}}$ ,  $R_{\text{treat}}$  and true mean utility  $\bar{U}^{\text{true}}(\mathbf{d})$ , to define %Best  $\mathbf{d}$  selected.

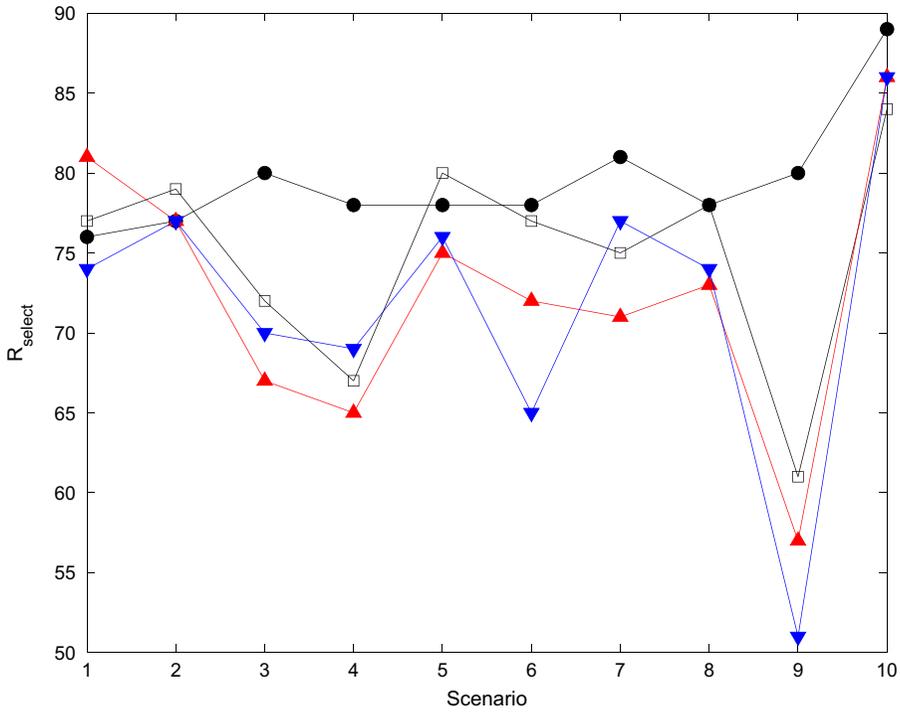
The results are given in Table 6. Comparing the PDS-model-based design with (4,4) versus (3,3) dimensional outcomes shows that the  $R_{\text{select}}$ -values differ by at most  $\pm 3$  for scenarios 1–8 and 10, but in scenario 9 using a (3,3) outcome greatly reduces  $R_{\text{select}}$  from 80 to 66. A similar pattern is seen for  $R_{\text{treat}}$  and %Best  $\mathbf{d}$  selected. Comparison of the PDS model, with either (4,4) or (3,3) outcomes, with the GAO model with (3,3) outcomes shows that the latter has much larger variability between scenarios in terms of  $R_{\text{select}}$ ,  $R_{\text{treat}}$  and %Best. Thus, as in Table 5, it appears that the PDS model provides a much more stable design, and in particular protects against very poor performance in some cases, as seen in scenarios 4, 5 and 9 with the GAO model.

Simulation results for four designs in Table 6 are illustrated graphically for  $R_{\text{select}}$  in Fig. 6, which shows that, in general, dichotomizing the ordinal outcomes substantively decreases

**Table 6.** Comparison of summary statistics for two-agent dose finding designs, with (4,4), (3,3) or (2,2) dimensional bivariate outcomes†

Design	Parameter	Results for the following scenarios:											
		1	2	3	4	5	6	7	8	9	10	11	12
4 efficacy and 4 toxicity levels, PDS	$R_{\text{select}}$	76	77	80	78	78	78	81	78	80	89	93	94
	$R_{\text{treat}}$	65	63	61	72	70	67	70	70	65	85	79	67
	%None	2	2	1	0	1	7	1	1	0	0	95	96
	%Best	32	28	32	46	39	34	39	33	48	44	—	—
3 efficacy and 3 toxicity levels, PDS	$R_{\text{select}}$	77	75	83	75	76	81	84	77	66	88	91	92
	$R_{\text{treat}}$	66	63	63	70	69	67	71	69	59	85	79	67
	%None	2	2	1	0	1	6	1	1	1	0	94	95
	%Best	35	28	38	36	33	40	47	35	25	46	—	—
3 efficacy and 3 toxicity levels, GAO	$R_{\text{select}}$	84	71	87	67	63	72	89	73	60	87	93	97
	$R_{\text{treat}}$	72	64	64	67	64	63	76	68	51	81	77	67
	%None	2	3	1	0	1	4	1	1	0	1	94	93
	%Best	53	25	67	24	5	19	70	26	25	61	—	—
2 efficacy and 2 toxicity levels, PDS	$R_{\text{select}}$	77	79	72	67	80	77	75	78	61	84	89	89
	$R_{\text{treat}}$	67	67	60	69	72	67	67	70	58	83	79	66
	%None	3	2	1	0	1	7	1	1	1	0	95	96
	%Best	34	30	19	13	43	30	23	32	15	46	—	—
2 efficacy and 2 toxicity levels, Yuan and Yin (2011) design	$R_{\text{select}}$	81	77	67	65	75	72	71	73	57	86	91	51
	$R_{\text{treat}}$	74	66	57	59	63	67	53	58	62	81	96	64
	%None	20	10	8	2	12	14	7	11	9	2	100	8
	%Best	31	29	20	8	26	22	23	21	18	48	—	—
2 efficacy and 2 toxicity levels, Wages and Conaway (2014) design	$R_{\text{select}}$	74	77	70	69	76	65	77	74	51	86	100	77
	$R_{\text{treat}}$	73	69	60	66	68	58	64	62	47	79	97	61
	%None	0	2	1	0	2	5	0	2	0	0	71	63
	%Best	12	34	21	11	29	13	33	23	9	62	—	—
No efficacy and 2 toxicity levels, partial orders CRM, target 0.35	$R_{\text{select}}$	79	73	55	64	71	64	60	61	61	90	100	81
	$R_{\text{treat}}$	78	68	49	62	63	57	56	55	55	84	99	71
	%None	0	0	0	0	0	0	0	0	0	0	0	0
	%Best	27	28	21	3	20	17	13	15	19	65	—	—

†Scenarios 11 and 12 have no acceptable dose, so  $R_{\text{select}}$ -values are less relevant and thus are in italics.



**Fig. 6.**  $R_{\text{select}}$ -values of competing phase I–II designs to choose an optimal dose pair, given (4,4) dimensional ordinal (efficacy, toxicity) outcomes: ●, PDS four-levels design proposed; □, PDS two-levels dichotomized outcomes design; ▲, Yuan and Yin (2011) design; ▼, Wages and Conaway (2014) design

$R_{\text{select}}$ -values in some scenarios, regardless of the design that is used. Fig. 6 also illustrates that the PDS-model-based design using the full (4,4) dimensional ordinal outcome is robust, in the sense that the  $R_{\text{select}}$ -values stay consistently high across all scenarios. In the special case of scenario 10, which corresponds to the prior, three of the nine  $\mathbf{d}$ -pairs are optimal; hence selecting an optimal  $\mathbf{d}$ -pair is much easier for all designs. On-line supplementary Table S10 shows that the partial orders CRM has greatly inferior performance compared with the PDS-based design. This may be attributed to the general fact that using binary toxicity alone for dose finding may ignore useful efficacy information.

#### 4.4. Additional sensitivity analyses

On-line supplementary Table S6 shows that the PDS-based design's behaviour is insensitive to cohort size  $c = 1, 2, 3$ . Supplementary Table S4 summarizes the PDS-model-based design's sensitivity to maximum sample sizes  $N = 30\text{--}300$ . The design's operating characteristic improves greatly as  $N$  increases. For example, in scenario 1, for  $N = 30, 60, 300$ , the corresponding  $R_{\text{select}}$ -values are 67, 76 and 95 and  $R_{\text{treat}}$ -values are 61, 76 and 78. The same pattern is seen for all other scenarios with acceptable dose pairs. In the two scenarios 11 and 12, where there is no acceptable  $\mathbf{d}$ , the simulated probability that no pair is selected is 1 for  $N \geq 120$ . These results provide an empirical validation of the method's consistency, in terms of both optimal dose pair selection and stopping the trial early for futility or safety in cases where this should be done. These numerical results also show that the maximum sample size 60 cannot reliably achieve  $R_{\text{select}}$ -values of 80 or larger across the scenarios that were studied, in this particular setting,

**Table 7.** Case-by-case example of a 60-patient trial†

Patient	Dose pair		Outcomes		Posterior mean utility $u(\mathbf{d} \text{data}_n)$ for each $(d_1, d_2)$ ‡								
	$d_1$	$d_2$	$Y_E$	$Y_T$	(4,40)	(4,60)	(4,80)	(5,40)	(5,60)	(5,80)	(6,40)	(6,60)	(6,80)
(Prior)	—	—	—	—	34.8 (22.3)	37.3 (25.5)	38.0 (26.4)	36.3 (24.2)	39.0 (27.3)	39.7 (28.1)	36.8 (24.8)	39.5 (27.9)	<b>40.2</b> (28.7)
1	4	60	1	0	49.2 (11.8)	54.2 (6.0)	54.8 (10.9)	51.0 (14.8)	55.9 (13.2)	56.5 (16.1)	51.6 (15.7)	56.4 (14.9)	<b>56.9</b> (17.6)
2	4	60	1	0	49.8 (11.1)	54.7 (3.5)	55.2 (9.9)	51.5 (14.2)	56.3 (12.3)	56.7 (15.5)	52.0 (15.2)	56.7 (14.2)	<b>57.1</b> (17.0)
3	4	60	0	2	33.6 (9.9)	35.5 (8.5)	<b>35.5</b> (12.3)	33.8 (14.5)	34.9 (14.4)	34.9 (16.6)	33.3 (15.5)	34.0 (15.3)	34.1 (17.3)
4	4	80	1	1	35.2 (9.4)	37.2 (7.2)	37.4 (8.1)	37.9 (12.5)	39.6 (11.6)	39.7 (12.5)	38.7 (13.2)	40.2 (12.5)	<b>40.3</b> (13.4)
5	4	80	0	1	31.7 (8.3)	33.7 (6.4)	31.8 (6.8)	35.3 (13.2)	37.4 (12.6)	35.5 (13.0)	36.3 (14.1)	<b>38.3</b> (13.5)	36.5 (14.0)
6	4	80	0	1	30.4 (6.9)	31.4 (5.7)	29.6 (5.9)	32.5 (11.4)	33.6 (11.4)	32.3 (11.7)	32.9 (12.3)	<b>34.0</b> (12.3)	33.0 (12.7)
7	5	60	1	1	30.4 (7.4)	31.8 (5.5)	30.0 (5.5)	34.2 (10.4)	35.5 (7.6)	34.8 (8.2)	34.6 (11.0)	<b>36.0</b> (8.7)	35.4 (9.3)
8	5	60	1	0	30.4 (7.4)	32.9 (5.3)	31.7 (5.5)	37.8 (10.6)	<b>40.0</b> (6.4)	38.7 (7.4)	37.6 (11.2)	39.4 (8.0)	38.3 (9.0)
9	5	60	0	2	31.5 (7.2)	31.6 (4.9)	30.7 (5.3)	32.3 (7.9)	31.9 (5.5)	31.1 (6.1)	<b>34.5</b> (10.6)	33.8 (9.5)	33.0 (10.1)
10	6	40	2	2	30.9 (6.4)	31.9 (5.2)	30.4 (5.5)	31.3 (7.5)	32.3 (6.6)	31.1 (7.1)	48.4 (16.7)	<b>50.2</b> (17.8)	49.6 (18.4)
11	6	40	2	2	31.7 (6.6)	32.2 (5.3)	29.5 (5.6)	30.9 (7.9)	31.4 (6.8)	29.0 (7.1)	52.7 (12.6)	<b>54.7</b> (14.1)	53.3 (15.4)
12	6	40	3	2	31.3 (6.5)	31.8 (5.6)	29.6 (5.9)	30.8 (7.7)	31.4 (7.0)	29.4 (7.4)	57.0 (13.9)	<b>62.9</b> (16.8)	62.7 (17.8)
15	6	60	1	2	32.0 (6.9)	32.6 (6.2)	30.4 (6.4)	32.1 (9.5)	32.8 (9.1)	31.2 (9.3)	56.0 (10.6)	56.9 (10.3)	<b>57.3</b> (12.0)
30	6	80	2	1	35.8 (7.4)	36.2 (6.3)	36.1 (6.7)	38.9 (10.4)	39.1 (9.6)	39.1 (10.0)	53.7 (5.7)	53.6 (5.9)	<b>53.9</b> (6.2)
45	6	40	1	0	33.6 (6.2)	33.6 (6.0)	28.2 (6.1)	34.6 (9.6)	34.3 (9.4)	28.5 (9.6)	<b>55.5</b> (4.7)	55.4 (4.7)	48.7 (6.1)
60	6	60	3	1	36.7 (6.9)	37.9 (6.4)	31.5 (7.7)	37.2 (10.3)	38.2 (9.9)	32.1 (10.5)	56.1 (4.0)	<b>57.0</b> (3.8)	50.4 (5.5)
Total number of patients assigned					0	3	3	0	3	0	21	24	6

†The largest current posterior mean utility is given in bold.

‡Standard deviations are given in parentheses.

and that  $N \geq 90$  is needed to achieve  $R_{\text{select}} \geq 80$ , and  $N$  roughly 200–240 is needed if  $R_{\text{select}} \geq 90$  is desired.

Supplementary Table S12 summarizes the behaviour of the PDS-based design when the trial is conducted by using each of three different numerical utilities. One is the elicited utility in Table 2, and two are hypothetical, given in supplementary Table S11, constructed to place greater value on either lower toxicity or greater efficacy. The simulations show that, across the 12 scenarios, the three resulting designs behave differently, but with no general pattern favouring one utility over the others. The utility favouring higher efficacy gives a design that escalates more aggressively and thus has greater observed toxicity and efficacy. Analogously, the utility favouring lower toxicity results in less toxicity but also less efficacy. A general conclusion is that the design behaves in a way that reflects the numerical values of  $U(\mathbf{y})$ , which is the intention.

Table 7 gives a patient-by-patient illustration of how the design may behave as the trial plays out, and what the interim estimates look like during the trial, for patients 1–12, 15, 30, 45 and 60. Since the maximum posterior mean utility after the first cohort is  $u(4, 80 | \text{data}_3) = 35.5$ , the pair  $\mathbf{d} = (4, 80)$  is used to treat cohort 2. Although  $u(6, 60 | \text{data}_6) = 34.0$  is largest for  $n = 6$ , the constraint that an untried dose may not be skipped when escalating results in  $\mathbf{d} = (5, 60)$  being used to treat cohort 3. The trial continues similarly, applying the AR method as described in step 6 of the design in Section 3.3. For each  $\mathbf{d}$ , the posterior variability of  $u(\mathbf{d} | \text{data}_n)$  decreases with sample size  $n$ , but not monotonically. At the end of the trial,  $\mathbf{d} = (6, 60)$  is optimal with  $u(6, 60 | \text{data}_{60}) = 57.0$ , but  $\mathbf{d} = (6, 40)$  also is a good choice since it is nearly optimal with  $u(6, 40 | \text{data}_{60}) = 56.1$ .

### 5. Discussion

Because the GCR model given by expression (2) links the conditional probability  $\gamma_k(y, \mathbf{x}_k^\lambda, \boldsymbol{\theta}_k)$  to the linear term  $\eta_k(y, \mathbf{x}_k^\lambda, \boldsymbol{\theta}_k)$ , it has the computational advantage that there are no order constraints on the intercept parameters  $\alpha_{k,1}, \dots, \alpha_{k,L_k}$ . An alternative model may be defined by

$$\bar{\pi}_k(y, \mathbf{x}_k^\lambda, \boldsymbol{\theta}_k) = 1 - [1 + \phi_k \exp\{\eta_k(y, \mathbf{x}_k^\lambda, \boldsymbol{\theta}_k)\}]^{-1/\phi_k}.$$

This generalizes the proportional odds model (McCullagh, 1980) by replacing the logit link with the Aranda-Ordaz link. Because this model links the unconditional probability  $\bar{\pi}_k(y, \mathbf{x}_k^\lambda, \boldsymbol{\theta}_k)$  rather than the conditional probability  $\gamma_k(y, \mathbf{x}_k^\lambda, \boldsymbol{\theta}_k)$  to the linear term, it requires the order constraints  $\alpha_{k,1} > \dots > \alpha_{k,L_k}$  for the probabilities to be well defined. Using this model for dose finding, the need to impose these constraints on each parameter vector  $\boldsymbol{\alpha}_k = (\alpha_{k,1}, \dots, \alpha_{k,L_k})$ ,  $k = E, T$  makes the Markov chain Monte Carlo computations to obtain posteriors much more difficult, especially for small amounts of data. This is one important motivation for our use of the GCR model.

Various special cases or alternative formulations of the PDS model can be obtained by changing one or more of its components. A natural question is whether adding a multiplicative dose–dose interaction term to the model with parametric dose standardization would improve the design’s behaviour. This model would have linear components

$$\eta_k(y, \mathbf{x}_k^\lambda, \boldsymbol{\theta}_k) = \alpha_{k,y} + \beta_{k,y,1} x_{k,1,j}^\lambda + \beta_{k,y,2} x_{k,2,r}^\lambda + \beta_{k,12} x_{k,1,j}^\lambda x_{k,2,r}^\lambda.$$

It may be considered a hybrid of the PDS and CMI model, in that it includes both parametric dose standardization and a conventional multiplicative interaction term. Supplementary Table S7 shows that, compared with the PDS model, the hybrid model gives a design with  $R_{\text{select}}$ -values 1–6 smaller in eight scenarios, 1–3 larger in two scenarios and slightly larger incorrect early stopping probabilities. Thus, on average, this more complex hybrid model produces a design with slightly worse performance than the PDS model.

A computer program named ‘U2OET’ for implementing this methodology is available from <https://biostatistics.mdanderson.org/SoftwareDownload>.

### Acknowledgements

This research was supported by National Institutes of Health National Cancer Institute grant RO1-CA-83932. We thank Nolan Wages and Mark Conaway for providing computer programs to simulate their designs. We also are grateful to two referees and the Associate Editor for their constructive comments and suggestions.

**Appendix A: Review of generalized continuation ratio models and copulas**

Recall that  $\gamma_k(y, \mathbf{d}, \boldsymbol{\theta}_k) = P(Y_k \geq y | Y_k \geq y - 1, \mathbf{d}, \boldsymbol{\theta}_k)$ , for  $y = 0, \dots, L_k$ . For given link function and linear term  $\eta(y, \mathbf{d}, \boldsymbol{\theta}_k)$ , a GCR model defines this conditional probability as

$$\gamma_k(y, \mathbf{d}, \boldsymbol{\theta}_k) = \text{link}\{\eta(y, \mathbf{d}, \boldsymbol{\theta}_k)\}.$$

The marginal probabilities of a GCR model are given by

$$\begin{aligned} \pi_k(0, \mathbf{d}, \boldsymbol{\theta}_k) &= 1 - \gamma_k(1, \mathbf{d}, \boldsymbol{\theta}_k), \\ \pi_k(y, \mathbf{d}, \boldsymbol{\theta}_k) &= \{1 - \gamma_k(y + 1, \mathbf{d}, \boldsymbol{\theta}_k)\} \prod_{r=1}^y \gamma_k(r, \mathbf{d}, \boldsymbol{\theta}_k), \quad \text{for } y = 1, \dots, L_k, \\ \bar{\pi}_k(y, \mathbf{d}, \boldsymbol{\theta}_k) &= \prod_{r=1}^y \gamma_k(r, \mathbf{d}, \boldsymbol{\theta}_k), \quad \text{for } y = 1, \dots, L_k. \end{aligned}$$

Since

$$P(Y_k \geq y | Y_k \geq y - 1, \mathbf{d}, \boldsymbol{\theta}_k) = 1 - \frac{P(Y_k = y - 1 | \mathbf{d}, \boldsymbol{\theta}_k)}{P(Y_k \geq y - 1 | \mathbf{d}, \boldsymbol{\theta}_k)},$$

the GCR model may be specified equivalently in the more commonly used form

$$\frac{\Pr(Y_k = y | \mathbf{d}, \boldsymbol{\theta}_k)}{\Pr(Y_k \geq y | \mathbf{d}, \boldsymbol{\theta}_k)} = 1 - \gamma_k(y + 1, \mathbf{d}, \boldsymbol{\theta}_k), \quad \text{for } y = 0, \dots, L_k - 1.$$

In general, the joint probability distribution of  $\mathbf{Y} = (Y_E, Y_T)$  given by a copula (Nelsen, 2006) can be defined in terms of the marginal cumulative distributions functions

$$F_k(y | \mathbf{d}, \boldsymbol{\theta}_k) = \Pr(Y_k \leq y | \mathbf{d}, \boldsymbol{\theta}_k) = 1 - \bar{\pi}_k(y + 1, \mathbf{d}, \boldsymbol{\theta}_k), \quad \text{for } y = 0, \dots, L_k - 1, \quad k = E, T,$$

by applying the formula

$$\Pr(Y_E = y_E, Y_T = y_T | \mathbf{d}, \boldsymbol{\theta}) = \sum_{a=1}^{y_E} \sum_{b=1}^{y_T} (-1)^{a+b} C_\rho(u_a, v_b)$$

where  $C_\rho(u_a, v_b)$  denotes the copula and  $u_1 = F_E(y_E | \mathbf{d}, \boldsymbol{\theta})$ ,  $v_1 = F_T(y_T | \mathbf{d}, \boldsymbol{\theta})$ ,  $u_2 = F_E(y_E - 1 | \mathbf{d}, \boldsymbol{\theta})$  and  $v_2 = F_T(y_T - 1 | \mathbf{d}, \boldsymbol{\theta})$ . To obtain a bivariate distribution under the PDS model, we assume a Farlie–Gumbel–Morgenstern copula

$$C_\rho(u, v) = uv\{1 + \rho(1 - u)(1 - v)\}, \quad \text{for } 0 \leq u, v \leq 1, \quad -1 \leq \rho \leq 1.$$

The GCR model that was given by Houede *et al.* (2010) accounts for the joint effects of the two doses on each ordinal outcome in a qualitatively different way. First, a conventional linear term for each agent  $a = 1, 2$ , level  $y$  of outcome  $Y_k$  for  $k = E, T$  and dose  $d^{(a)}$  is defined as

$$\eta_{k,y}^{(a)} = \alpha_{k,y,0}^{(a)} + \alpha_{k,y,1}^{(a)} d^{(a)}.$$

A GAO link is then defined as

$$\gamma_k(y, \mathbf{d}, \boldsymbol{\theta}_k) = 1 - [1 + \lambda_k \{\exp(\eta_{k,y}^{(1)}) + \exp(\eta_{k,y}^{(2)}) + \kappa \exp(\eta_{k,y}^{(1)} + \eta_{k,y}^{(2)})\}]^{-1/\lambda_k}$$

where  $\kappa > 0$  is a dose–dose interaction parameter. Houede *et al.* (2010) obtained bivariate distributions by assuming a Gaussian copula,

$$C_\rho(u, v) = \Phi_\rho\{\Phi^{-1}(u), \Phi^{-1}(v)\},$$

where  $\Phi_\rho$  denotes a bivariate normal cumulative distribution function with correlation  $\rho$  and  $\Phi$  denotes an  $N(0, 1)$  cumulative distribution function.

**References**

Aranda-Ordaz, F. J. (1981) On two families of transformations to additivity for binary response data. *Biometrika*, **68**, 357–363.

- Azriel, D., Mandel, M. and Rinott, Y. (2011) The treatment versus experimentation dilemma in dose-finding studies. *J. Statist. Plannng Inf.*, **141**, 2759–2768.
- Babb, J., Rogatko, A. and Zacks, S. (1998) Cancer phase I clinical trials: efficient dose escalation with overdose control. *Statist. Med.*, **17**, 1103–1120.
- Bekele, B. N. and Shen, Y. (2005) A Bayesian approach to jointly modeling toxicity and biomarker expression in a phase I/II dose-finding trial. *Biometrics*, **61**, 344–354.
- Braun, T. M. (2002) The bivariate continual reassessment method: extending the CRM to phase I trials of two competing outcomes. *Contr. Clin. Trials*, **23**, 240–256.
- Braun, T. M., Kang, S. and Taylor, J. M. G. (2013) A phase I/II trial design when response is unobserved in subjects with dose-limiting toxicity. *Statist. Meth. Med. Res.*, **22**, 1–15.
- Brook, R. H., Chassin, M. R., Fink, A., Solomon, D. H., Kosecoff, J. and Park, R. E. (1986) A method for the detailed assessment of the appropriateness of medical technologies. *Int. J. Technol. Assessmnt Hlth Care*, **2**, 53–63.
- Cox, C. (1988) Multinomial regression models based on continuation ratios. *Statist. Med.*, **7**, 435–441.
- Dalkey, N. C. (1969) An experimental study of group opinion. *Futures*, **1**, 408–426.
- Fienberg, S. E. (1980) *The Analysis of Cross-classified Categorical Data*, 2nd edn. Cambridge: Massachusetts Institute of Technology Press.
- Gasparini, M. and Eisele, J. (2000) A curve-free method for phase I clinical trials. *Biometrics*, **56**, 609–615.
- Gittins, J. C. (1979) Bandit processes and dynamic allocation indices (with discussion). *J. R. Statist. Soc. B*, **41**, 148–177.
- Houede, N., Thall, P. F., Nguyen, H., Paoletti, X. and Kramar, A. (2010) Utility-based optimization of combination therapy using ordinal toxicity and efficacy in phase I/II trials. *Biometrics*, **66**, 532–540.
- Hunink, M. G. M., Weinstein, M. C., Wittenberg, E., Pliskin, J. S., Drummond, M. F., Glasziou, P. P. and Wong, J. B. (2014) *Decision Making in Health and Medicine: Integrating Evidence and Values*. Cambridge: Cambridge University Press.
- McCullagh, P. (1980) Regression models for ordinal data (with discussion). *J. R. Statist. Soc. B*, **42**, 109–142.
- Morita, S., Thall, P. F. and Müller, P. (2008) Determining the effective sample size of a parametric prior. *Biometrics*, **64**, 595–602.
- Morita, S., Thall, P. F. and Müller P. (2010) Evaluating the impact of prior assumptions in Bayesian biostatistics. *Statist. Biosci.*, **2**, 1–17.
- Nelsen, R. B. (2006) *An Introduction to Copulas*, 2nd edn. New York: Springer.
- O’Quigley, J., Pepe, M. and Fisher, L. (1990) Continual reassessment method: a practical design for phase I clinical trials in cancer. *Biometrics*, **46**, 33–48.
- Oron, A. P. and Hoff, P. D. (2013) Small-sample behavior of novel phase I designs. *Clin. Trials*, **10**, 63–80.
- Riviere, M.-K., Yuan, Y., Dubois, F. and Zohar, S. (2015) A Bayesian dose finding design for clinical trials combining a cytotoxic agent with a molecularly targeted agent. *Appl. Statist.*, **64**, 215–229.
- Robert, C. P. and Casella, G. (1999) *Monte Carlo Statistical Methods*. New York: Springer.
- Storer, B. (1989) Design and analysis of phase I clinical trials. *Biometrics*, **45**, 925–937.
- Sutton, R. S. and Barto, A. G. (1998) *Reinforcement Learning: an Introduction*. Cambridge: Massachusetts Institute of Technology Press.
- Swinburn, P., Lloyd, A., Nathan, P., Choueiri, T. K., Cella, D. and Neary, M. P. (2010) Elicitation of health state utilities in metastatic renal cell carcinoma. *Curr. Med. Res. Opin.*, **26**, 1091–1096.
- Thall, P. F. and Cook, J. D. (2004) Dose-finding based on efficacy-toxicity trade-offs. *Biometrics*, **60**, 684–693.
- Thall, P. F. and Nguyen, H. Q. (2012) Adaptive randomization to improve utility-based dose-finding with bivariate ordinal outcomes. *J. Biopharm. Statist.*, **22**, 785–801.
- Thall, P. F., Nguyen, H. Q., Braun, T. M. and Qazilbash, M. H. (2013) Using joint utilities of the times to response and toxicity to adaptively optimize schedule-dose regimes. *Biometrics*, **69**, 673–682.
- Thall, P. F., Nguyen, H. Q., Zohar, S. and Maton, P. (2014) Optimizing sedative dose in preterm infants undergoing treatment for respiratory distress syndrome. *J. Am. Statist. Ass.*, **109**, 931–943.
- Thall, P. F., Szabo, A., Nguyen, H. Q., Amlie-Lefond, C. M. and Zaidat, O. O. (2011) Optimizing the concentration and bolus of a drug delivered by continuous infusion. *Biometrics*, **67**, 1638–1646.
- Wages, N. A. and Conaway, M. R. (2014) Phase I/II adaptive design for drug combination oncology trials. *Statist. Med.*, **33**, 1990–2003.
- Wages, N. A., Conaway, M. R. and O’Quigley J. (2011) Dose-finding design for multi-drug combinations. *Clin. Trials*, **8**, 380–389.
- Whitehead, J., Thygesen, H. and Whitehead, A. (2010) A Bayesian dose-finding procedure for phase I clinical trials based only on the assumption of monotonicity. *Statist. Med.*, **29**, 1808–1824.
- Whitehead, J., Thygesen, H. and Whitehead, A. (2011) Bayesian procedures for phase I/II clinical trials investigating the safety and efficacy of drug combinations. *Statist. Med.*, **30**, 1952–1970.
- Yin, G., Li, Y. and Ji, Y. (2006) Bayesian dose-finding in phase I/II clinical trials using toxicity and efficacy odds ratios. *Biometrics*, **62**, 777–787.
- Yuan, Y. and Yin, G. (2011) Bayesian phase I/II adaptively randomized oncology trials with combined drugs. *Ann. Appl. Statist.*, **5**, 924–942.

- Zhang, W., Sargent, D. J. and Mandrekar, S. (2006) An adaptive dose-finding design incorporating both efficacy and toxicity. *Statist. Med.*, **25**, 2365–2383.
- Zhou, Y., Whitehead, J., Bonvini, E. and Stevens, J. (2006) Bayesian decision procedures for binary and continuous bivariate dose-escalation studies. *Pharm. Statist.*, **5**, 125–133.

*Supporting information*

Additional ‘supporting information’ may be found in the on-line version of this article.