ROYAL STATISTICAL SOCIETY
DATA | EVIDENCE | DECISIONS

Journal of the Royal Statistical Society
Applied Statistics
Series C

Appl. Statist. (2019)
68, *Part* 2, *pp.* 331–345

Phase I–II trial design for biologic agents using conditional auto-regressive models for toxicity and efficacy

Daniel G. Muenz, Jeremy M. G. Taylor and Thomas M. Braun

University of Michigan, Ann Arbor, USA

[Received June 2017. Final revision August 2018]

Summary. A traditional assumption in the design of chemotherapy phase I–II trial designs is that dose increases lead to both more toxicity as well as more efficacy. This assumption of monotonic rates of toxicity and efficacy has come into question as potential cancer treatments are less likely to be chemotherapy and are instead biologic agents. These biologic agents tend to have mechanisms of action that act as 'on–off' switches for cancer growth, so giving more of the biologic agents will not necessarily provide any more benefit (and possibly no further risk) to the patient. We propose the use of a conditional auto-regressive (CAR) model as a way to estimate adaptively the rates of dose limiting toxicities (DLTs) and efficacy by smoothing the data collected for all doses in such a way that allows for non-increasing rates of either outcome with dose. We present the study design for our CAR model approach and compare, via simulation, the operating characteristics of our design with two existing contemporary published approaches. We demonstrate that our CAR model approach is a viable design for an adaptive phase I–II trial that can accommodate a variety of toxicity–dose and efficacy–dose patterns.

Keywords: Biologically optimal dose; Clinical trial; Dose finding; Immunotherapy; Molecularly targeted agent; Randomization

1. Introduction

Classical phase I trial designs that are used in cancer research, such as the 3+3 design (Storer, 1989) and the continual reassessment method (CRM) (O'Quigley *et al.*, 1990; Faries, 1994; Goodman *et al.*, 1995), assume that the probability of dose limiting toxicity (DLT) strictly increases with dose, with this assumption also made implicitly for the probability of efficacy. These assumptions were generally acceptable for cytotoxic agents, with the implication that, among doses with acceptable rates of DLTs, the largest dose would have the highest probability of efficacy; this dose is known as the maximum tolerated dose (MTD). Thus, the 3+3 method, CRM and other phase I trial designs were created that made dose assignments and selected the MTD based solely on DLTs.

However, with the advent of biologic agents for the treatment of cancer, the assumption of monotonic toxicity and/or efficacy has come into question. To understand this issue, we focus on a specific class of biologic agents known as monoclonal antibodies; specific members of this class have names with the suffix 'mab'. These agents can be used to target a specific protein in cancer cells or to block specific pathways, such as immune checkpoints, that are believed to be important for cancer cell growth. Through these actions, monoclonal antibodies enable the immune system

Address for correspondence: Thomas M. Braun, Department of Biostatistics, University of Michigan, 1415 Washington Heights, Ann Arbor, MI 48109, USA.
E-mail: tombraun@umich.edu

 0035–9254/19/68331

to recognize and develop a response to the cancer; well-known examples include ipilimumab (Lipson and Drake, 2011), nivolumab (Rizvi *et al.*, 2015) and pembrolizumab (Hamid *et al.*, 2013).

Because of their mechanisms of action, biologic agents are viewed as an 'on–off' switch for cancer growth. Once a sufficient amount of the biologic agent has caused a pathway to be turned off or a specific protein to be targeted, giving more of the biologic agent will not necessarily provide any more benefit, and possibly no further risk, to the patient. Thus, it has been suggested that the patterns of toxicity and efficacy probabilities may plateau (increase and then level off) or peak (increase and then decrease) (Postel-Vinay *et al.*, 2009; Jain *et al.*, 2010). As a result, a phase I trial design for a biologic agent should

(a) incorporate both DLTs and efficacy outcomes when determining dose assignments and
(b) allow the dose patterns for DLTs and/or efficacy to vary flexibly with dose; see chapter 2 of Yuan *et al.* (2016) for a comprehensive discussion.

Since the MTD is no longer assumed to have maximum efficacy, the optimal dose that is selected at the end of the trial is instead viewed as the optimal biological dose (OBD). A visual demonstration of the difference between the MTD and the OBD can be found in Riviere *et al.* (2016).

When the goal is to identify the MTD by assuming increasing rates of DLT and efficacy with dose, there are several methods to model both the probability of DLT and the probability of efficacy for a single agent, including Thall and Russell (1998), Gooley *et al.* (1994), Braun (2002), Thall and Cook (2004), Thall and Nguyen (2012) and Liu and Johnson (2016). More recent approaches for identifying the OBD by removing the assumption of increasing rates of DLT or efficacy with dose include the methods of Zang *et al.* (2014), Wages and Tait (2015) and Riviere *et al.* (2016).

To complement these approaches of identifying the OBD, we propose a method in which we borrow information across doses without imposing a strict parametric form for either DLT rates or efficacy rates. Separately for each of DLT and efficacy, each dose level has its own parameter for the probability of the outcome, and these probabilities for all doses are linked via a covariance matrix that borrows information across doses. In particular, we model the log-odds of DLT and efficacy with a conditional auto-regressive (CAR) model, which has been used in geospatial analysis of lattice data (Besag, 1974; Cressie, 2015; Wall, 2004). We introduce our method in Section 2, which describes our statistical model and dose finding algorithm and outlines the specific steps that are needed for a clinical trial using our design. In Section 3, we compare the operating characteristics of our design with two existing designs, and we conclude with a discussion in Section 4.

2. Model and methods

2.1. Defining the model

Consider J candidate dose levels, ordered by increasing dose from 1 to J. We let n_j denote the number of patients who are assigned to dose $j = 1, 2, \ldots, J$, of whom Y_j^{D} patients have experienced DLTs and Y_j^{E} patients have experienced efficacy. We assume that $Y_j^{\mathrm{D}} \sim \mathrm{Bin}(n_j, \pi_j^{\mathrm{D}})$ and $Y_j^{\mathrm{E}} \sim \mathrm{Bin}(n_j, \pi_j^{\mathrm{E}})$. We induce correlation between the elements of $\boldsymbol{\pi}^{\mathrm{D}} = \{\pi_1^{\mathrm{D}}, \pi_2^{\mathrm{D}}, \ldots, \pi_J^{\mathrm{D}}\}$ and between the elements of $\boldsymbol{\pi}^{\mathrm{E}} = \{\pi_1^{\mathrm{E}}, \pi_2^{\mathrm{E}}, \ldots, \pi_J^{\mathrm{E}}\}$ by using a CAR covariance structure. Both $\boldsymbol{\pi}^{\mathrm{D}}$ and $\boldsymbol{\pi}^{\mathrm{E}}$ will have the same structure; for clarity, we first describe the structure for $\boldsymbol{\pi}^{\mathrm{D}}$.

If we define $\mathrm{logit}(x) = \log(x) - \log(1-x)$, we assume a multivariate normal distribution for $\mathrm{logit}(\boldsymbol{\pi}^{\mathrm{D}}) = \{\mathrm{logit}(\pi_1^{\mathrm{D}}), \mathrm{logit}(\pi_2^{\mathrm{D}}), \ldots, \mathrm{logit}(\pi_J^{\mathrm{D}})\}$, with mean equal to $\mathrm{logit}(\boldsymbol{\pi}_0^{\mathrm{D}})$, where $\boldsymbol{\pi}_0^{\mathrm{D}} =$

$\{\pi_{01}^{\mathrm{D}} \ldots \pi_{0J}^{\mathrm{D}}\}$ is a set of 'skeleton' values for the DLT rates of the doses. Conditionally on given values of λ and σ^2, $\mathrm{logit}(\boldsymbol{\pi}^{\mathrm{D}})$ has covariance matrix $\boldsymbol{\Sigma}^{\mathrm{D}} = \sigma^2(\mathbf{I}_J - \lambda\mathbf{W})^{-1}$, in which $\mathbf{I}_J$ is a $J \times J$ identity matrix and $\mathbf{W}$ is a matrix that is used to define the neighbourhood for each dose.

Specifically, $\mathbf{W}$ is a $J \times J$ matrix with element (i, j) equal to

$$w_{ij} = \begin{cases} 1 & \text{if } |i-j| = 1, \\ 0 & \text{if } i = j, \\ 0 & \text{otherwise,} \end{cases}$$

so that $\mathbf{W}$ has 1s just off the main diagonal, and 0s everywhere else.

The matrix $\boldsymbol{\Sigma}^{\mathrm{D}}$ is a function of two parameters:

(a) σ^2, which controls the prior standard deviation of each $\mathrm{logit}(\pi_j^{\mathrm{D}})$ and thus of each π_j^{D}, and
(b) λ, which controls the amount of information that is borrowed across neighbouring doses.

Furthermore, as shown in Wall (2004), the CAR model leads to element j of $\mathrm{logit}(\boldsymbol{\pi}^{\mathrm{D}})$ having a conditional normal distribution with mean

$$E\{\mathrm{logit}(\pi_j^{\mathrm{D}}) \mid \boldsymbol{\pi}_{(j)}^{\mathrm{D}}\} = \mathrm{logit}(\pi_{0j}^{\mathrm{D}}) + \lambda \sum_{k=1}^{J} w_{jk}\{\mathrm{logit}(\pi_k^{\mathrm{D}}) - \mathrm{logit}(\pi_{0k}^{\mathrm{D}})\}, \tag{1}$$

and variance

$$\mathrm{var}\{\mathrm{logit}(\pi_j^{\mathrm{D}}) \mid \mathrm{logit}(\boldsymbol{\pi}_{(j)}^{\mathrm{D}})\} = \sigma^2, \tag{2}$$

in which $\boldsymbol{\pi}_{(j)}^{\mathrm{D}}$ is $\boldsymbol{\pi}^{\mathrm{D}}$ with element j removed.

In this form, we see that the DLT rate of each dose, given its first-order neighbours ($w_{jk} = 1$), is conditionally independent of the other doses. We also see that λ controls how much the conditional mean of $\mathrm{logit}(\pi_j^{\mathrm{D}})$ moves away from the *a priori* value $\mathrm{logit}(\pi_{0j}^{\mathrm{D}})$ as a function of the rates of its first-order neighbours. From equation (1), we also see that negative values of λ correspond to negative correlations between all directly neighbouring doses, whereas positive values correspond to positive correlations between all doses. As negative correlation lacks plausibility in our setting, we shall not consider values of $\lambda < 0$.

More importantly, to ensure that $\boldsymbol{\Sigma}^{\mathrm{D}}$ is positive definite, λ is restricted to the interval $(1/w_{\min}, 1/w_{\max})$, where $w_{\min} < 0$ and $w_{\max} > 0$ are the minimum and maximum eigenvalues of $\mathbf{W}$ (Wall, 2004). Thus, in our methods, we shall assume a uniform prior distribution on λ over the range $[0, 1/w_{\max}]$. We shall consider σ^2 to be fixed, treating it as a tuning parameter whose value is selected by the user, which is a common approach to variance parameters in many phase I trial designs.

The marginal CAR model for $\mathrm{logit}(\boldsymbol{\pi}^{\mathrm{D}})$ just described is also used to model the marginal distribution of $\mathrm{logit}(\boldsymbol{\pi}^{\mathrm{E}})$, with every superscript 'D' replaced with superscript 'E'. The CAR model for $\mathrm{logit}(\boldsymbol{\pi}^{\mathrm{E}})$ also requires specification of a skeleton $\boldsymbol{\pi}_0^{\mathrm{E}} = \{\pi_{01}^{\mathrm{E}} \ldots \pi_{0J}^{\mathrm{E}}\}$. In their current formulation, the toxicity and efficacy rate CAR models use the same values for the parameters σ^2 and λ in their corresponding covariance matrices, although different values among the two models could be used, if desired.

We note that our approach models the marginal distributions of $\mathrm{logit}(\boldsymbol{\pi}^{\mathrm{D}})$ and $\mathrm{logit}(\boldsymbol{\pi}^{\mathrm{E}})$ and assumes that they are independent so that no information is shared between them. Although we could consider modelling the joint distribution of $\mathrm{logit}(\boldsymbol{\pi}^{\mathrm{D}})$ and $\mathrm{logit}(\boldsymbol{\pi}^{\mathrm{E}})$ by including an additional association parameter, we refrain from doing so on the basis of the recommendation of Cunanan and Koopmeiners (2014). They investigated various copula models to link DLT

and efficacy outcomes within subject and found that copula association parameters are difficult to estimate with binary outcomes, especially with the small sample sizes that are used in phase I trials. Thus, assuming within-subject independence of DLT and efficacy outcomes leads to suitable operating characteristics even if there is within-subject dependence of the outcomes; see Liu and Johnson (2016) for a similar discussion.

2.2. Estimation of model parameters and event probabilities

After m subjects have been observed for DLT and efficacy, we have three vectors of data:

(a) the number of subjects assigned to each dose, $\mathbf{n} = \{n_1, n_2, \ldots, n_J\}$, such that $m = \Sigma_{j=1}^{J} n_j$,
(b) the number of DLTs observed for each dose, $\mathbf{Y}^{\mathrm{D}} = \{Y_1^{\mathrm{D}}, Y_2^{\mathrm{D}}, \ldots, Y_J^{\mathrm{D}}\}$, and
(c) the number of efficacy outcomes that are observed for each dose, $\mathbf{Y}^{\mathrm{E}} = \{Y_1^{\mathrm{E}}, Y_2^{\mathrm{E}}, \ldots, Y_J^{\mathrm{E}}\}$.

Because we choose to model the marginal distributions of $\boldsymbol{\pi}^{\mathrm{D}}$ and $\boldsymbol{\pi}^{\mathrm{E}}$, each can be estimated separately by using Bayesian methods. For clarity, we first focus on estimation of $\boldsymbol{\pi}^{\mathrm{D}}$ and then explain how the same estimation procedure is applied to $\boldsymbol{\pi}^{\mathrm{E}}$.

We wish to find the posterior distribution of $\boldsymbol{\eta}^{\mathrm{D}} = \mathrm{logit}(\boldsymbol{\pi}^{\mathrm{D}})$, conditionally on $\mathbf{n}$ and $\mathbf{Y}^{\mathrm{D}}$, in which $\eta_j^{\mathrm{D}} = \mathrm{logit}(\pi_j^{\mathrm{D}})$. Because of the nature of the CAR model, it is straightforward to use the Metropolis–Hastings algorithm (Minh and Minh, 2015) to generate draws from the posterior distribution of $\boldsymbol{\pi}^{\mathrm{D}}$, i.e. using the JAGS software (Plummer, 2003). The Metropolis–Hastings algorithm provides us with draws from the joint posterior distribution $f(\boldsymbol{\pi}^{\mathrm{D}} | \mathbf{Y}^{\mathrm{D}}, \mathbf{n})$, which also gives us draws from each marginal posterior $f(\pi_j^{\mathrm{D}} | \mathbf{Y}^{\mathrm{D}}, \mathbf{n})$. The mean of each of these J vectors of marginal samples gives us the posterior mean DLT rate for dose j, which we denote as $\hat{\pi}_j^{\mathrm{D}}$. An analogous approach for computing the posterior mean efficacy rate of each dose j, denoted $\hat{\pi}_j^{\mathrm{E}}$, is found by replacing the 'D' superscripts with an 'E' in the posterior computations just described. These two estimators can be interpreted as the respective posterior expected probabilities of DLT and efficacy for a new patient assigned to dose j. We let $\hat{\boldsymbol{\pi}}_j = (\hat{\pi}_j^{\mathrm{E}}, \hat{\pi}_j^{\mathrm{D}})$.

2.3. Dose finding algorithm

As is typical for most phase I–II trial designs, we choose to assign the first cohort of patients to the lowest dose; all future patients are assigned to a dose that is determined by the following algorithm. First, we define a set of acceptable doses as those that meet both of the following criteria:

$$\Pr(\hat{\pi}_j^{\mathrm{D}} > \bar{\pi}^{\mathrm{D}} \mid \mathbf{Y}^{\mathrm{D}}, \mathbf{n}) < p^{\mathrm{D}}, \tag{3}$$

$$\Pr(\hat{\pi}_j^{\mathrm{E}} > \underline{\pi}^{\mathrm{E}} \mid \mathbf{Y}^{\mathrm{E}}, \mathbf{n}) > p^{\mathrm{E}}, \tag{4}$$

where $\underline{\pi}^{\mathrm{E}}$ and $\bar{\pi}^{\mathrm{D}}$ are lower and upper bounds respectively, for the rates of efficacy and DLT, and p^{E} and p^{D} are respective thresholds for the posterior cumulative probabilities. These four quantities are fixed at the beginning of the trial, and this is similar to the approach that was used in both Wages and Tait (2015) and Liu and Johnson (2016). On the basis of the data from m subjects, if no dose meets these acceptability criteria, the trial is terminated and no dose is identified as the OBD.

Otherwise, we define $\mathcal{S}$ to be the set of safe doses, i.e. those that satisfy equation (3). To promote exploration early in the trial, we shall randomize the assignments of the first half of the cohorts (or just over half if an odd number of cohorts), using the posterior mean efficacy rates to randomize among doses in $\mathcal{S}$, i.e. dose $k \in \mathcal{S}$ is assigned to the next cohort of patients with probability

$$p_k^{\text{rand}} = \frac{\hat{\pi}_k^{\text{E}}}{\sum_{l \in \mathcal{S}} \hat{\pi}_l^{\text{E}}}.$$

Randomization is not used with the remaining latter cohorts; instead, we assign each cohort to the member of $\mathcal{S}$ with largest posterior probability of efficacy: a so-called 'greed' strategy. When the trial has finished collecting data on the desired number of patients N, the OBD is the member of $\mathcal{S}$ with highest posterior mean efficacy.

2.4. Designing a trial

We now outline the steps that are necessary to implement our design successfully in an actual clinical trial.

Step 1: identify the number of doses J, the maximum number of patients N and the cohort size c.
Step 2: for each dose j, select a skeleton value π_{0j}^{D} for the probability of DLT and π_{0j}^{E} for the probability of efficacy.
Step 3: select values for $\bar{\pi}^{\text{D}}$, the maximum acceptable probability of DLT, and p^{D}, the maximum acceptable amount of posterior mass above $\bar{\pi}^{\text{D}}$, used in equation (3).
Step 4: select values for $\underline{\pi}^{\text{E}}$, the minimum acceptable probability of efficacy, and p^{E}, the minimum acceptable amount of posterior mass above $\bar{\pi}^{\text{E}}$, used in equation (4).
Step 5: select a value for the prior standard deviation σ used in the CAR model covariance.
Step 6: enrol the first cohort of c patients on the lowest dose.
Step 7: use the methods in Sections 2.2 and 2.3 to determine whether the study should continue and, if so, which dose to assign to the next cohort of c patients.
Step 8: repeat step 7 after each successive cohort has been followed for DLT and efficacy.
Step 9: once all N subjects have been followed for DLT and efficacy, make a final determination about which dose is the OBD.

We note that suitable parameter values in steps 3–5 will require calibration through small simulation studies and grid searches over plausible values of each parameter, as is done in other phase I trial designs.

3. Simulation studies

We now compare the operating characteristics of our CAR model design with two existing designs: one which does not assume monotonicity of efficacy rates with dose and one that does make this assumption. The methods of Wages and Tait (2015) apply the traditional CRM to model toxicity rates, whereas they use a Bayesian model averaging approach for efficacy by incorporating several vectors of *a priori* (skeleton) efficacy rates for each dose. The skeletons vary by the location of the highest efficacy rate, whether the efficacy rates increase or plateau with dose, and, if a plateau exists, at which dose the plateau occurs. By averaging over all these possible skeletons, it is hoped that the methods are sufficiently flexible to identify both monotonic and non-monotonic patterns of efficacy with dose.

The methods of Liu and Johnson (2016) assume monotonically increasing patterns of both toxicity and efficacy with dose. They create a Markov structure for the toxicity rates, whereby the toxicity rate of each dose is assumed to be equal to the toxicity rate of the next-lowest dose plus a random positive quantity, such that all the toxicity rates are bounded above 0 and below 1. A similar structure is assumed for the efficacy rates. Skeleton vectors of toxicity and efficacy

rates are used to define prior distributions for the differences in toxicity and efficacy rates of neighbouring doses. No regression model is used to model the DLT and efficacy rates, so it is hoped that this design is more flexible than many designs like the CRM that adopt a parametric regression model to enforce monotonicity.

Our methods can be viewed somewhat as a hybrid of those of Wages and Tait (2015) and Liu and Johnson (2016). First, like Wages and Tait (2015), we choose to allow efficacy rates to be possibly non-monotonic in dose, but we adopt a model that allows for non-monotonicity directly through its parameterization, rather than averaging over many linear and piecewise linear skeletons. Unlike Wages and Tait (2015), we also allow for possibly non-monotonic patterns of toxicity with dose. Like the methods of Liu and Johnson (2016), our approach is a conditional Markov model, in which the toxicity and efficacy rates of each dose are correlated with the toxicity and efficacy rates respectively of neighbouring doses. However, unlike Liu and Johnson (2016), our model does not enforce strict increases in toxicity and efficacy rates with dose.

For our method and our two comparators, we examine a group of settings that define the actual rates of DLT and efficacy for each dose. In each setting, if the true values of DLT and efficacy for dose j are $\pi_j^{D^*}$ and $\pi_j^{E^*}$, we simulate the binary DLT and efficacy outcomes for subjects who are assigned to dose j from respective Bernoulli distributions with probabilities $\pi_j^{D^*}$ and $\pi_j^{E^*}$. The specific values that are used for the parameters in our CAR model will be defined in each of the following sections. Unless explicitly stated otherwise, we use parameter values for the designs of Wages and Tait (2015) and Liu and Johnson (2016) as they originally proposed.

We note that the methods of Liu and Johnson (2016) combined the DLT and efficacy rates of each dose into a utility score and then sought to find the dose that was most likely to be the dose with greatest posterior mean utility. In contrast, our design and that of Wages and Tait (2015) are not utility based, which is the more common approach that is adopted in phase I trial designs; see Thall and Cook (2004) for an exception. Thus, in the tables that follow, the results for Liu and Johnson (2016) are based on simulations that used their Markov model to estimate the posterior mean DLT and efficacy rates. However, dose assignments and final selection of the OBD are based on the same rules as that of the CAR model explained in Section 2.3.

We have two reasons for doing so. First, we can directly compare how the operating characteristics of designs using the CAR model and the Markov model of Liu and Johnson (2016) differ from each other. Second, we found that the utility measure that was proposed in Liu and Johnson (2016) was somewhat subjective regarding what the OBD should be. For example, in scenario 1 of Liu and Johnson (2016) the first two doses have respective DLT rates of 0.15 and 0.32, and respective efficacy rates of 0.28 and 0.30. Their utility metric placed a large penalty on doses with DLT rates above 0.30. As a result, the lowest dose had greater utility than the second dose. However, in many settings, it would seem that the second dose is the OBD because it is more effective than the lowest dose and its DLT rate is only two points above the targeted DLT rate.

3.1. Comparison in settings explored in Liu and Johnson (2016)

We first compare the operating characteristics of the three methods in eight settings that were described in Liu and Johnson (2016). Each of the settings has toxicity and efficacy rates that increase monotonically with dose. Thus, we use these settings to examine how our CAR model design and the design of Wages and Tait (2015), neither of which requires monotonicity, perform when monotonicity exists. The location of the OBD varies among the eight settings, with one of the settings having no OBD among the doses examined. The actual rates of toxicity and efficacy in each setting are demonstrated in Fig. 1.

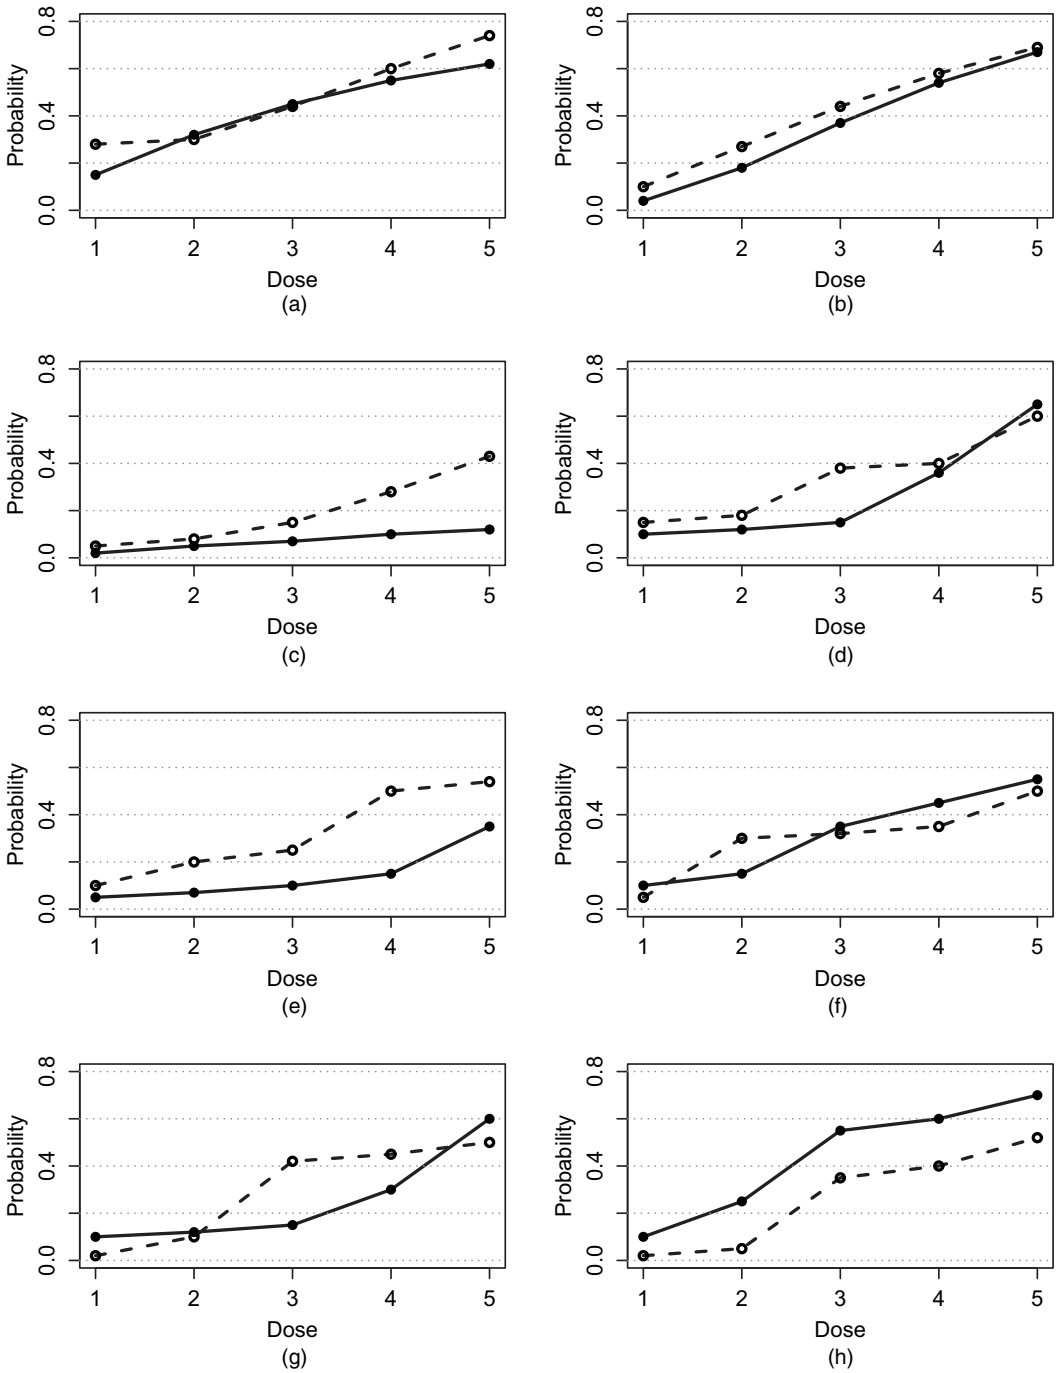

Fig. 1. DLT (●) and efficacy (○) rates for each of the eight settings examined in Section 3.1 (the exact values for the DLT and efficacy rates can be found in Table 1): (a) setting 1; (b) setting 2; (c) setting 3; (d) setting 4; (e) setting 5; (f) setting 6; (g) setting 7; (h) setting 8

In each setting, we have a study of five doses that enrols a total of $N = 48$ patients in cohorts of size $c = 3$. In all three designs, we allocated randomized dose assignments to the first $N/2 = 24$ patients, with 'greedy' assignments given to the remaining 24 patients. We performed simulations of 1000 trials in each setting to obtain the operating characteristics of each method. For our CAR design, we used the same skeleton values for toxicity and efficacy as proposed by Liu and Johnson (2016). The same toxicity skeleton was used in the method of Wages and Tait (2015). For efficacy, we defined a set of nine skeletons, five that were unimodal and four with plateaus, modified from those in Wages and Tait (2015) to have only five doses.

Also following directly from Liu and Johnson (2016), we used the value $\bar{\pi}^{\mathrm{D}} = 0.30$ for the maximal allowed toxicity rate and $\bar{\pi}^{\mathrm{E}} = 0.20$ for the minimal allowed efficacy rate; these values were also used for the corresponding values in the design of Wages and Tait (2015). We also used the values $p^{\mathrm{D}} = p^{\mathrm{E}} = 0.20$ for the definition of acceptable doses in steps 3 and 4 in Section 2.4. The association parameter λ in the CAR model variance matrices is uniformly distributed over $[0.00, 0.58]$. Finally, through a grid search over a range of values, we selected the standard deviation in the CAR model covariance to be $\sigma = 0.75$. We found that lower values of σ made the *a priori* skeleton rates too informative, leading to larger doses always being preferred in all settings, whereas larger values of σ led to much higher rates of early termination because so little *a priori* information was available when few patients had been enrolled.

In Table 1, we summarize the performance of the three designs in the eight settings with two operating characteristics:

(a) the proportion of simulations in which each dose was selected as the OBD at the end of the study, and
(b) the simulationwide average number of subjects assigned to the true OBD during the study.

In setting 1, we see that the two lowest doses have efficacy rates that are closest to the targeted DLT rate of 0.30; the CAR model design prefers the second dose to the first dose, whereas the design of Wages and Tait (2015) moderately prefers the first dose to the second dose and the design of Liu and Johnson (2016) strongly prefers the first dose. In terms of patient assignments, the design of Wages and Tait (2015) assigns the most patients to the first two doses, whereas the CAR model design assigns the least. Thus, in this setting, the CAR model is not competitive with the other designs.

However, in setting 2, we see that, although all three designs equally prefer the second dose, the design of Liu and Johnson (2016) is predisposed towards the lowest dose, whereas the CAR model design and that of Wages and Tait (2015) skew towards the third dose, which, in this setting, seems preferable to the lowest dose. These results support the findings in setting 1 where the design of Liu and Johnson (2016) appears to skew selection towards lower doses. Furthermore, in setting 6, we once again see that the design of Liu and Johnson (2016) prefers lower doses, whereas the CAR model design and that of Wages and Tait (2015) prefer higher doses, although the differences in the operating characteristics among the three designs is modest.

In setting 3, where the highest dose is the OBD, all three designs have similar distributions of dose assignments, but the CAR model design and that of Liu and Johnson (2016) both correctly identify the OBD more than the design of Wages and Tait (2015). Furthermore, in settings 4, 5 and 7, the CAR model design has the best operating characteristics of the three designs, with the design of Liu and Johnson (2016) again predisposed to lower doses more than the others. Setting 8 has none of the doses being an OBD, as no dose simultaneously is both safe and effective; again all three designs have comparable operating characteristics.

We emphasize that each of the designs could be made to be more preferential to lower or higher doses by changing some of their corresponding tuning parameters. Thus, there is no

Table 1. Operating characteristics of the proposed CAR model design, the design of Wages and Tait (2015) (WT) and the design of Liu and Johnson (2016) (LJ) in the eight settings shown in Fig. 1†

Setting	*Method*	*Results for OBD selection*						*Results for patients assigned*				
		D1	*D2*	*D3*	*D4*	*D5*	*None*	*D1*	*D2*	*D3*	*D4*	*D5*
1		**(15,28)**	**(32,30)**	(45,44)	(55,60)	(62,74)						
	CAR	**36.2**	**43.1**	8.2	0.2	0.1	12.2	**8.7**	**15.5**	12.4	6.3	4.4
	WT	**51.2**	**43.7**	1.8	0.0	0.0	3.3	**27.1**	**16.0**	2.9	0.9	0.6
	LJ	**71.9**	**10.9**	2.1	1.0	0.8	13.3	**21.2**	**13.9**	5.4	2.3	1.0
2		(4,10)	**(18,27)**	**(37,44)**	(54,58)	(67,69)						
	CAR	3.2	**62.5**	**25.4**	0.8	0.0	8.1	6.0	**11.9**	**17.9**	7.5	4.3
	WT	0.9	**55.6**	**37.2**	0.8	0.0	5.5	10.5	**19.0**	**13.7**	2.6	1.6
	LJ	20.1	**58.7**	**9.8**	0.7	0.0	10.7	11.4	**18.2**	**10.8**	3.6	1.0
3		(2,5)	(5,8)	(7,15)	(10,28)	**(12,43)**						
	CAR	0.0	0.0	0.6	10.4	**88.3**	0.7	5.4	5.7	5.8	6.0	**25.0**
	WT	1.3	2.3	7.1	18.5	**64.8**	6.0	6.2	4.8	7.0	10.0	**19.5**
	LJ	0.1	0.6	2.1	7.3	**88.4**	1.5	3.2	3.5	4.1	9.9	**26.8**
4		(10,15)	(12,18)	**(15,38)**	**(36,40)**	(65,60)						
	CAR	1.1	7.3	**74.5**	**15.6**	0.0	1.5	5.9	6.2	**13.0**	**17.3**	5.6
	WT	9.4	10.9	**50.8**	**19.8**	1.1	8.0	13.3	10.7	**14.7**	**5.9**	2.4
	LJ	5.2	18.3	**68.9**	**3.4**	0.3	3.9	5.8	10.6	**17.7**	**9.7**	3.1
5		(5,10)	(7,20)	(10,25)	**(15,50)**	**(35,54)**						
	CAR	0.2	2.0	9.6	**67.5**	**20.4**	0.3	5.6	5.9	6.3	**12.1**	**18.1**
	WT	2.6	6.5	11.0	**47.2**	**30.3**	2.4	8.0	6.8	8.8	**13.0**	**11.1**
	LJ	0.7	3.0	15.6	**71.1**	**8.4**	1.2	3.6	4.7	8.7	**16.6**	**13.9**
6		(10,5)	**(15,30)**	**(32,35)**	(45,35)	(55,50)						
	CAR	0.3	**54.7**	**29.3**	5.5	0.4	9.8	5.8	**10.1**	**14.6**	10.6	6.4
	WT	0.0	**58.0**	**34.4**	1.4	0.0	6.2	10.5	**20.1**	**12.4**	2.4	1.6
	LJ	6.8	**75.6**	**6.7**	0.3	0.4	10.2	8.8	**18.8**	**11.3**	4.1	1.9
7		(10,2)	(12,10)	**(15,42)**	**(30,45)**	(60,50)						
	CAR	0.0	2.2	**63.5**	**32.1**	0.0	2.2	5.4	5.8	**10.7**	**19.1**	6.8
	WT	0.0	2.8	**50.2**	**29.2**	2.1	15.7	9.1	9.6	**15.4**	**8.3**	3.3
	LJ	2.4	6.9	**75.4**	**10.4**	0.0	4.9	4.6	8.2	**17.5**	**12.0**	4.0
8		(10,2)	(25,2)	(55,35)	(60,40)	(70,52)						
	CAR	1.0	2.7	1.2	0.1	0.0	**95.0**	6.3	12.3	9.2	4.9	3.3
	WT	0.0	5.2	1.1	0.0	0.0	**93.7**	12.5	14.9	4.2	1.1	1.1
	LJ	10.9	3.2	0.6	0.1	0.3	**84.9**	8.8	8.7	5.0	1.6	0.6

†Each setting has five doses D1, D2,..., D5. For each setting, the first row contains a pair $(100x, 100y)$ for each dose, in which x is the DLT rate and y is the efficacy rate. The six columns under OBD selection present the percentage of simulations (times 100) in which each dose was selected as the MTD at the end of the study. The five columns under patients assigned present the average number of subjects (out of 48) who were assigned to each dose. The blocks in bold for each setting highlight the doses that have DLT rates closest to the threshold of 0.30 and have minimally acceptable efficacy, making them the best candidates for the OBD.

absolute 'best' design among the three in all settings. The overall conclusion from this set of simulations is that all three designs are worthwhile designs for studies where monotonic rates of DLT and efficacy are expected.

3.2. Comparison in settings motivated by Wages and Tait (2015)

We next compare the operating characteristics of the three methods in eight settings motivated by settings in Wages and Tait (2015). Each of the settings is a study of six doses; the values of the true DLT and efficacy rates are shown in Fig. 2. The values of the true DLT and efficacy rates were selected to cover settings where DLT and/or efficacy rates that would monotonically

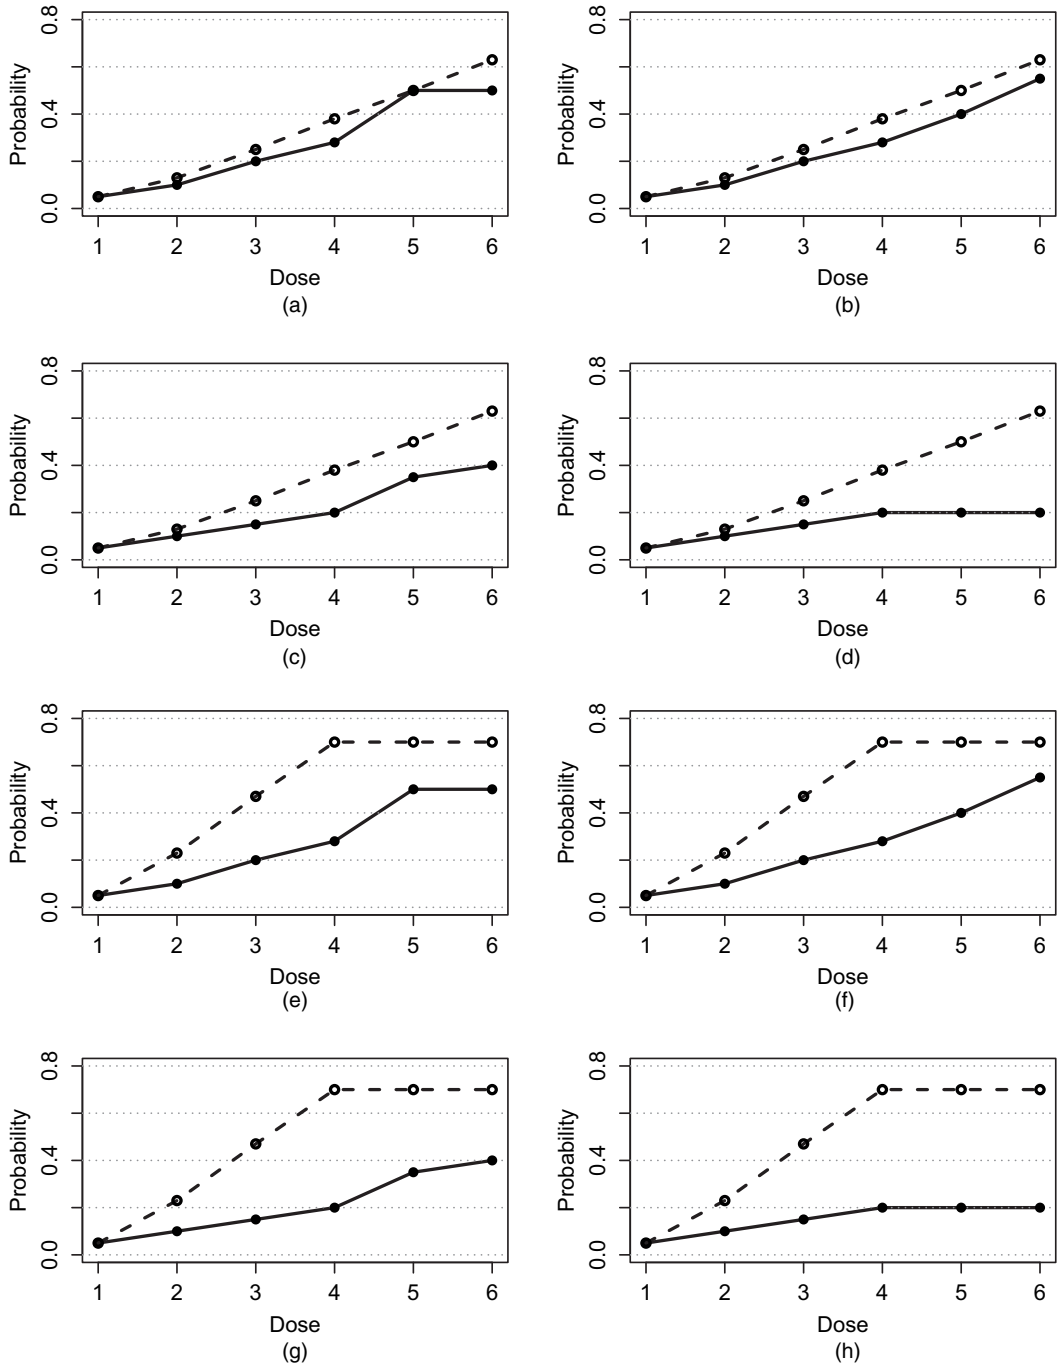

Fig. 2. DLT (●) and efficacy (○) rates for each of the eight settings examined in Section 3.2 (the exact values for the DLT and efficacy rates can be found in Table 2): (a) setting 1; (b) setting 2; (c) setting 3; (d) setting 4; (e) setting 5; (f) setting 6; (g) setting 7; (h) setting 8

increase with dose or plateau before the last dose. Thus, we expect that the methods of Liu and Johnson (2016) will struggle to identify the OBD correctly in the presence of a plateau in the toxicity and/or efficacy rates. We note that settings 1, 2, 3, 5, 6 and 7 in Fig. 2 are from Wages and Tait (2015), whereas settings 4 and 8 are new settings that were not investigated in Wages and Tait (2015).

We performed 1000 simulations for each setting to obtain operating characteristics of all three methods when the six doses were studied with a sample of $N = 64$ patients who were enrolled as singleton cohorts ($c = 1$). Our CAR design and that of Liu and Johnson (2016) use a skeleton vector of DLT rates identical to that proposed in Wages and Tait (2015). We use a skeleton vector of efficacy rates $\boldsymbol{\pi}_0^{\mathrm{E}} = (0.05, 0.18, 0.36, 0.54, 0.60, 0.67)$, which is the average value for each dose among the efficacy patterns that are shown in Fig. 2. In all three methods, randomized dose assignments were used for the first $N/2 = 32$ patients, with the remaining 32 patients receiving greedy dose assignments.

Following directly from Wages and Tait (2015), we used the value $\bar{\pi}^{\mathrm{D}} = 0.33$ for the maximal allowed toxicity rate and $\bar{\pi}^{\mathrm{E}} = 0.05$ for the minimal allowed efficacy rate. We also used the values $p^{\mathrm{D}} = p^{\mathrm{E}} = 0.20$ for the definition of acceptable doses in steps 3 and 4 in Section 2.4. The association parameter λ in the CAR model covariance is uniformly distributed over $[0.00, 0.55]$. As in Section 3.1, we selected the standard deviation in the CAR model covariance to be $\sigma = 0.75$.

In Table 2, we summarize the same two operating characteristics as in Table 1, namely

(a) the proportion of simulations in which each dose was selected as the OBD at the end of the study and
(b) the simulationwide average number of subjects assigned to the true OBD during the study.

For settings 1–3, which were originally examined in Wages and Tait (2015), we see that the CAR model can identify the OBD better than the design of Wages and Tait (2015), and both designs identify the OBD better than the design of Liu and Johnson (2016). The CAR model and the design of Wages and Tait (2015) also assign more patients to the OBD than does the model of Liu and Johnson (2016). Note that these results cover settings in which the DLT rates plateau at an unacceptable toxicity level (setting 1), as well as monotonically increase (settings 2 and 3). Setting 4 was a setting that we included to examine a situation when DLT rates plateau at an acceptable toxicity level, so that the largest of the doses after the plateau would be the OBD. In this setting, we see that the monotonic pattern of efficacy rates enables the design of Liu and Johnson (2016) to identify the OBD better than the others, with the design of Wages and Tait (2015) appearing to prefer all doses 4–6 equally.

Settings 5–7 have the same DLT rates as settings 1–3 respectively but now have efficacy rates that plateau at dose 4. As a result, the monotonic assumption of the design of Liu and Johnson (2016) is too strong, causing it to be much less likely to identify the OBD correctly compared with the other designs, with the CAR model design outperforming the design of Wages and Tait (2015). Setting 8 is a setting in which both DLT rates and efficacy rates plateau together, so the OBD exists simultaneously at doses 4–6; here the three designs have similar operating characteristics. Thus, across all eight settings, we have strong evidence for the benefit of the CAR model over the other two designs in settings where DLT rates and/or efficacy rates are expected to plateau.

We also ran simulations using these eight settings to assess the sensitivity of our design to

(a) the skeleton values for DLT and efficacy rates and
(b) the value of the variance parameter σ^2 (the results are not shown).

We found that, when we increased the skeleton rates of both DLT and efficacy by 50%, correct selection of the OBD did generally decrease by a few percentage points, which was expected

Table 2. Operating characteristics of the proposed CAR model design, the design of Wages and Tait (2015) (WT) and the design of Liu and Johnson (2016) (LJ) in the eight settings shown in Fig. 2†

Setting	*Method*	*Results for OBD selection*							*Results for patients assigned*					
		D1	*D2*	*D3*	*D4*	*D5*	*D6*	*None*	*D1*	*D2*	*D3*	*D4*	*D5*	*D6*
1		(5,5)	(10,13)	(20,25)	**(28,38)**	(50,50)	(50,63)							
	CAR	0.0	2.4	28.6	**66.9**	1.7	0.3	0.1	2.2	4.6	9.3	**27.9**	11.3	8.6
	WT	0.9	4.4	29.0	**57.5**	8.0	0.2	0.0	6.8	9.5	17.4	**21.8**	6.7	1.9
	LJ	0.5	5.3	32.2	**48.0**	3.1	1.1	9.8	2.2	5.1	11.6	**22.9**	9.9	7.9
2		(5,5)	(10,13)	(20,25)	**(28,38)**	(40,50)	(55,63)							
	CAR	0.0	2.5	24	**61.1**	12.1	0.1	0.2	2.2	4.4	8.7	**21.9**	19.6	7.1
	WT	0.8	4.6	25.3	**51.8**	17.1	0.4	0.0	6.7	9.2	16.6	**19.8**	9.4	2.2
	LJ	0.2	5.0	26.9	**46.2**	12.0	0.4	9.3	2.2	4.9	10.1	**18.2**	15.5	9.2
3		(5,5)	(10,13)	(15,25)	**(20,38)**	**(35,50)**	(40,63)							
	CAR	0.0	0.5	8.4	**53.9**	**27.8**	9.4	0.0	2.1	3.9	6.1	**14.8**	**19.6**	17.6
	WT	1.0	2.6	11.7	**36.7**	**40.2**	7.8	0.0	5.8	7.6	11.4	**16.8**	**15.7**	6.7
	LJ	0.2	0.6	10.1	**40.3**	**27.1**	17.3	4.4	1.9	3.8	7.3	**13.8**	**14.2**	20.9
4		(5,5)	(10,13)	(15,25)	(20,38)	(20,50)	**(20,63)**							
	CAR	0.0	0.0	0.3	5.7	20.9	**73.1**	0.0	2.0	3.6	5.4	7.5	9.8	**35.8**
	WT	1.0	2.7	7.8	20.9	27.8	**39.8**	0.0	5.5	6.7	9.1	12.7	14.1	**16.0**
	LJ	0.1	0.3	1.7	1.0	2.9	**90.8**	3.2	1.8	3.4	5.8	7.3	8.1	**35.8**
5		(5,5)	(10,23)	(20,47)	**(28,70)**	(50,70)	(50,70)							
	CAR	0.1	0.9	24.3	**72.3**	1.7	0.7	0.0	2.1	4.4	10.0	**28.2**	11.0	8.3
	WT	0.0	2.6	24.0	**65.6**	7.7	0.1	0.0	7.2	9.5	15.6	**23.4**	6.6	1.7
	LJ	0.2	8.4	29.3	**51.3**	4.2	0.9	5.7	2.3	6.2	11.8	**22.6**	10.7	7.9
6		(5,5)	(10,23)	(20,47)	**(28,70)**	(40,70)	(55,70)							
	CAR	0.0	1.2	19.6	**68.8**	10.4	0.0	0.0	2.0	4.2	9.2	**22.8**	18.8	6.9
	WT	0.0	2.5	22.7	**59.5**	15.1	0.2	0.0	7.3	9.1	15.6	**21.0**	8.9	2.0
	LJ	0.2	6.0	25.7	**48.0**	13.6	0.6	5.9	2.1	5.5	10.7	**18.3**	16.1	8.7
7		(5,5)	(10,23)	(15,47)	**(20,70)**	**(35,70)**	(40,70)							
	CAR	0.0	0.0	5.9	**62.8**	**24.5**	6.8	0.0	2.0	3.9	6.5	**15.7**	**19.6**	16.3
	WT	0.2	0.7	6.9	**49.8**	**38.5**	3.9	0.0	6.1	7.2	9.9	**20.2**	**15.7**	4.9
	LJ	0.2	1.3	7.8	**43.9**	**27.6**	16.6	2.6	2.0	4.4	7.7	**14.2**	**14.3**	20.1
8		(5,5)	(10,23)	(15,47)	**(20,70)**	**(20,70)**	**(20,70)**							
	CAR	0.0	0.0	0.6	**23.3**	**29.9**	**46.2**	0.0	1.9	3.7	5.8	**8.1**	**10.1**	**34.4**
	WT	0.1	1.3	5.2	**34.0**	**34.9**	**24.5**	0.0	5.7	6.5	9.0	**15.6**	**15.4**	**11.8**
	LJ	0.1	0.6	1.4	**1.7**	**3.0**	**90.9**	2.3	1.9	3.9	6.5	**7.8**	**7.8**	**34.7**

†Each setting has six doses D1, D2,…, D6. For each setting, the first row contains a pair $(100x, 100y)$ for each dose, in which x is the DLT rate and y is the efficacy rate. The seven columns under OBD selection present the percentage of simulations (times 100) in which each dose was selected as the MTD at the end of the study. The six columns under patients assigned present the average number of subjects (out of 64) who were assigned to each dose. The blocks in bold for each setting highlight the doses that have DLT rates closest to the threshold of 0.33 and have minimally acceptable efficacy, making them the best candidates for the OBD.

because the DLT skeleton suggested that the doses were all more toxic than before. But, overall, the performance of our design was still as good as or better than those of the other designs.

When we halved the value of σ, we found that selection of the BOD was skewed towards higher doses, so that, although our design selected the correct dose more often, it was also slightly more likely to select overly toxic doses. Conversely, when we doubled the value of σ, we found that selection of the BOD was skewed towards lower doses, so dose 4, which was the BOD in most of the settings, was selected less often. Thus, our design is sensitive to the value of σ, which is common among most phase I–II designs, emphasizing the fact that calibration of σ^2 through simulation is an important exercise.

4. Discussion

Our proposed CAR model design is competitive with two contemporary designs and appears to be the optimal choice when DLT and/or toxicity rates are expected to plateau. The CAR model design is no more complex than the other two designs, with an equal number of tuning parameters required. Nonetheless, there are aspects of the CAR model design, which are common to the other designs, that we would like to explore further.

First, randomization is used in phase I trials assessing both toxicity and efficacy because it encourages greater exploration of the doses and increases a design's ability to identify the OBD correctly. However, there is little current direction on exactly how

(a) to restrict which doses to randomize among and
(b) to compute the randomization probabilities of the selected doses.

We chose to limit randomization among doses that were deemed safe on the basis of the data, but many other approaches could be used. Other designs consider randomizing only among doses that are neighbours of the most recently assigned dose or to doses that are deemed to be both safe and effective. An interesting avenue of research for any design is to compare how operating characteristics of that design vary with different randomization schemes.

We also chose to use randomization probabilities that were directly proportional to the value of each dose's posterior mean efficacy; many other approaches are possible. For example, we could have randomized using both the posterior rates of DLT and efficacy, or used transformed values, i.e. the square root or exponential of those means. The idea of exponentiating the posterior means is a specific application of Softmax learning, which is often used in multiarmed bandits (Luce, 1959). It also remains unclear whether randomization is needed during the entire trial or is only needed until a specific proportion of subjects has been enrolled, after which dose assignments are based on a greedy assignment algorithm.

As with many phase I–II trial designs, our method has tuning parameters whose values need to be selected before starting the trial. We emphasize that the values that we selected in our simulation studies are not necessarily the best values for general use. However, the number of tuning parameters in our design is no more than that of most other published phase I–II trials. The most crucial parameter to calibrate is the value of σ^2, which can be determined quite quickly through small simulation studies of a few selected values. In our experience, a good value of σ^2 across many settings is often in the interval [0.50, 1.50].

In our simulations, we assumed that the efficacy outcome of each subject occurred in temporal proximity to their DLT outcome so that the dose assignment of each successive subject could be determined from complete data. However, in many settings, the efficacy outcome may occur much later than the DLT outcome, so a new subject might be available for a dose assignment before currently enrolled subjects have completed follow-up for their efficacy end point. In these situations, we suggest that the posterior distributions of the efficacy rates be computed with a weight that reflects how much follow-up each subject has completed at the time that the new subject is to be enrolled. Such an approach is akin to the time-to-event CRM approach of Cheung and Chappell (2000); a latent variable approach to the problem is also possible (Jin *et al.*, 2014). The work of Liu and Johnson (2016) examined this approach with their design and found that the design had slightly reduced ability to identify the OBD but still had generally good operating characteristics. We would expect the same results for our, or any, method, that was applied to studies with later follow-up for efficacy relative to DLT.

Finally, because of the multifocal nature of cancer, many biologic agents are being studied in combination with an existing chemotherapy or with a second biologic agent. As a result, phase

I–II trial designs are needed that model the DLT and efficacy rates of two agents, as well as allow for non-increasing rates of DLT and/or efficacy for one or both agents. For that, some recent designs have been published (Cai *et al.*, 2014; Riviere *et al.*, 2015; Guo and Li, 2015). We are currently working on a sequel to our current work that applies our CAR model approach to study the toxicity and efficacy of combinations of doses of two agents, at least one of which is a biologic agent.

Acknowledgement

This research was supported by US National Institutes of Health grant CA 83654.

References

Besag, J. (1974) Spatial interaction and the statistical analysis of lattice systems (with discussion). *J. R. Statist. Soc.* B, **36**, 192–236.

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.

Cai, C., Yuan, Y. and Ji, Y. (2014) A Bayesian dose finding design for oncology clinical trials of combinational biological agents. *Appl. Statist.*, **63**, 159–173.

Cheung, Y. K. and Chappell, R. (2000) Sequential designs for Phase I clinical trials with late-onset toxicities. *Biometrics*, **56**, 1177–1182.

Cressie, N. A. C. (2015) *Statistics for Spatial Data*, vol. II, *Lattice Data*, ch. 6. Hoboken: Wiley.

Cunanan, K. and Koopmeiners, J. S. (2014) Evaluating the performance of copula models in Phase I-II clinical trials under model misspecification. *BMC Med. Res. Methodol.*, **14**, no. 51, 1–11.

Faries, D. (1994) Practical modifications of the continual reassessment method for Phase I cancer clinical trials. *J. Biopharm. Statist.*, **4**, 147–164.

Goodman, S. N., Zahurak, M. L. and Piantadosi, S. (1995) Some practical improvements in the continual re-assessment method for Phase I studies. *Statist. Med.*, **14**, 1149–1161.

Gooley, T. A., Martin, P. J., Fisher, L. D. and Pettinger, M. (1994) Simulation as a design tool for Phase I/II clinical trials: an example from bone marrow transplantation. *Contr. Clin. Trials*, **15**, 450–462.

Guo, B. and Li, Y. (2015) Bayesian dose-finding designs for combination of molecularly targeted agents assuming partial stochastic ordering. *Statist. Med.*, **34**, 859–875.

Hamid, O., Robert, C., Daud, A., Hodi, F. S., Hwu, W. J., Kefford, R., Wolchok, J. D., Hersey, P., Joseph, R. W., Weber, J. S., Dronca, R., Gangadhar, T. C., Patnaik, A., Zarour, H., Joshua, A. M., Gergich, K., Elassaiss-Schaap, J., Algazi, A., Mateus, C., Boasberg, P., Tumeh, P. C., Chmielowski, B., Ebbinghaus, S. W., Li, X. N., Kang, S. P. and Ribas, A. (2013) Safety and tumor responses with lambrolizumab (anti-PD-1) in melanoma. *New Engl. J. Med.*, **369**, 134–144.

Jain, R. K., Lee, J. J., Hong, D., Markman, M., Gong, J., Naing, A., Wheler, J. and Kurzrock, R. (2010) Phase I oncology studies: evidence that in the era of targeted therapies patients on lower doses do not fare worse. *Clin. Cancer Res.*, **16**, 1289–1297.

Jin, I., Liu, S., Thall, P. and Yuan, Y. (2014) Using data augmentation to facilitate conduct of Phase I/II clinical trials with delayed outcomes. *J. Am. Statist. Ass.*, **109**, 525–536.

Lipson, E. J. and Drake, C. G. (2011) Ipilimumab: an anti-ctla-4 antibody for metastatic melanoma. *Clin. Cancer Res.*, **17**, 6958–6962.

Liu, S. and Johnson, V. E. (2016) A robust Bayesian dose-finding design for phase I/II clinical trials. *Biostatistics*, **17**, 249–263.

Luce, D. (1959) *Individual Choice Behavior*. New York: Wiley.

Minh, D. D. L. and Minh, D. L. P. (2015) Understanding the Hastings algorithm. *Communs Statist. Simuln Computn*, **44**, 332–349.

O'Quigley, J., Pepe, M. and Fisher, L. (1990) Continual reassessment method: a practical design for Phase 1 clinical trials in cancer. *Biometrics*, **46**, 33–48.

Plummer, M. (2003) JAGS: a program for analysis of Bayesian graphical models using Gibbs sampling. In *Proc. 3rd Int. Wrkshp Distributed Statistical Computing* (eds K. Hornik, F. Leisch and A. Zeileis).

Postel-Vinay, S., Arkenau, H.-T., Olmos, D., Ang, J., Barriuso, J., Ashley, S., Banerji, U., De-Bono, J., Judson, I. and Kaye, S. (2009) Clinical benefit in Phase-I trials of novel molecularly targeted agents: does dose matter? *Br. J. Cancer*, **100**, 1373–1378.

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.

Riviere, M.-K., Yuan, Y., Jourdan, J.-H., Dubois, F. and Zohar, S. (2016) Phase I/II dose-finding design for molecularly targeted agent: plateau determination using adaptive randomization. *Statist. Meth. Med. Res.*, **27**, 466–479.

Rizvi, N. A., Mazières, J., Planchard, D., Stinchcombe, T. E., Dy, G. K., Antonia, S. J., Horn, L., Lena, H., Minenza, E., Mennecier, B., Otterson, G. A., Campos, L. T., Gandara, D. R., Levy, B. P., Nair, S. G., Zalcman, G., Wolf, J., Souquet, P. J., Baldini, E., Cappuzzo, F., Chouaid, C., Dowlati, A., Sanborn, R., Lopez-Chavez, A., Grohe, C., Huber, R. M., Harbison, C. T., Baudert, C., Lestini, B. J. and Ramalingam, S. S. (2015) Activity and safety of nivolumab, an anti-PD-1 immune checkpoint inhibitor, for patients with advanced, refractory squamous non-small-cell lung cancer (CheckMate 063): a Phase 2, single-arm trial. *Lancet Oncol.*, **16**, 257–265.

Storer, B. E. (1989) Design and analysis of Phase I clinical trials. *Biometrics*, **45**, 925–937.

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. and Russell, K. E. (1998) A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in Phase I/II clinical trials. *Biometrics*, **54**, 251–264.

Wages, N. A. and Tait, C. (2015) Seamless Phase I/II adaptive design for oncology trials of molecularly targeted agents. *J. Biopharm. Statist.*, **25**, 903–920.

Wall, M. M. (2004) A close look at the spatial structure implied by the CAR and SAR models. *J. Statist. Planng Inf.*, **121**, 311–324.

Yuan, Y., Nguyen, H. Q. and Thall, P. F. (2016) *Bayesian Designs for Phase I-II Clinical Trials*. Boca Raton: Chapman and Hall–CRC.

Zang, Y., Lee, J. J. and Yuan, Y. (2014) Adaptive designs for identifying optimal biological dose for molecularly targeted agents. *Clin. Trials*, **11**, 319–327.