

Dose Finding with Escalation with Overdose Control (EWOC) in Cancer Clinical Trials

Mourad Tighiouart *,
André Rogatko

Emory University and Cedars-Sinai Medical Center

Abstract.

Traditionally, the major objective in phase I trials is to identify a working-dose for subsequent studies, whereas the major endpoint in phase II and III trials is treatment efficacy. The dose sought is typically referred to as the maximum tolerated dose (MTD). Several statistical methodologies have been proposed to select the MTD in cancer phase I trials. In this manuscript, we focus on a Bayesian adaptive design, known as escalation with overdose control (EWOC). Several aspects of this design are discussed, including large sample properties of the sequence of doses selected in the trial, choice of prior distributions, and use of covariates. The methodology is exemplified with real life examples of cancer phase I trials. In particular, we show in the recently completed ABR-217620 (naptumomab estafenatox) trial that omitting an important predictor of toxicity when dose assignments to cancer patients are determined results in a high percent of patients experiencing severe side effects and a significant proportion treated at sub-optimal doses.

Key words and phrases: Cancer phase I trials, Dose limiting toxicity, Escalation with overdose control, tolerated dose, Optimal Bayesian feasible.

1. INTRODUCTION

The main objective in cancer phase I clinical trials is to identify a tolerable dose of a cytotoxic or therapeutic agent for subsequent studies. Phase I trials represent the first testing of an investigational agent or combination of agents whose safety profile has been established individually. These trials typically enroll patients with advanced cancer stages and who have exhausted available standard treatment options [24].

Department of Biostatistics and Bioinformatics and Winship Cancer Institute, Atlanta, GA 30322, Samuel Oschin Comprehensive Cancer Institute, Los Angeles, CA 90048 (e-mail:

mourad.tighiouart@emory.edu; andre.rogatko@cshs.org)

*Supported in part by NIH/NCI Grants No. 1 P01 CA116676, 1 P30 CA138292-01, and 5 P50 CA128613

Cancer phase I trials are carried out sequentially, assigning dose levels to subjects based on the observed side effects of the previously treated patients. From a safety and therapeutic perspective, these trials should be designed to minimize the number of unacceptable toxic events and maximize the number of patients treated at an optimal dose. Ideally, the design should control the probability of overdosing patients at each stage of the trial, produce a sequence of doses that converge to the MTD, and should take into account the heterogeneity nature of cancer phase I trial patients [29].

Decisions to escalate or de-escalate dose levels in cancer phase I trials are made after one cycle of therapy to patients. The length of a cycle is usually between 3 to 6 weeks. Therefore, the target phase I dose is typically defined in terms of treatment related side effects, ignoring treatment efficacy. This is due to the fact that treatment efficacy, expressed as a reduction in tumor size or an increase in survival, requires months (if not years) of observation [21; 34], a length of time far greater than the length of one cycle of therapy. Thus, it can be stated that the main objective of a cancer phase I clinical trial is to determine a safe dose of a new drug or combination of drugs for subsequent clinical evaluation of efficacy. This dose is known as the maximum tolerated dose (MTD), or phase II dose. Specifically, the MTD, γ , is defined as the dose expected to produce some degree of medically unacceptable, dose-limiting toxicity (DLT) in a pre-specified proportion θ of patients [14],

$$(1) \quad P(\text{DLT} \mid \text{Dose} = \gamma) = \theta.$$

The target probability of DLT θ depends on the severity of the treatment attributable toxicity. It is set relatively high when the DLT is reversible or non-fatal condition, and low if it is life threatening [3]. Ting [33] and Rosenberger and Haines [27] give a good review of statistical methods for dose finding in cancer phase I trials. In particular, the widely used continual reassessment method (CRM) proposed by O’Quigley et al. [21] and its extensions by Faries [12]; Goodman et al. [15]; Möller [19]; Piantadosi et al. [23]; and Storer [28], and escalation with overdose control (EWOC) method proposed by Babb et al. [3]; Zacks et al. [36]; Babb and Rogatko [4]; Tighiouart et al. [31], and Rogatko et al. [25] are Bayesian adaptive and produce consistent sequences of doses under some model assumptions and regularity conditions. These designs can be easily implemented in practice using published tutorials and free interactive software, see e.g. Garrett [13], Zohar et al. [38], Xu et al. [35], and Rogatko et al. [26].

In this paper, we review several aspects of EWOC, including large sample properties, choice of prior distributions, and use of covariates. The methodology is exemplified with cancer phase I clinical trials we designed and conducted at Fox Chase Cancer Center in Philadelphia and the Winship Cancer Institute in Atlanta.

This article is organized as follows. In section 2, we introduce the phase I design known as EWOC and review its large sample properties. We illustrate its implementation using a real life example. An extension of this design to account for patients’ specific characteristics is described in Section 3 and the methodology is illustrated by a recently completed phase I cancer trial. Section 4 contains some concluding remarks and discussion.

2. ESCALATION WITH OVERDOSE CONTROL

Denote by Y the binary indicator of DLT for a patient given dose x . Assume that there exists x^* and x^{**} , $x^* < x^{**}$ such that

$$(2) \quad P(Y = 1|x = x^*) = 0$$

$$(3) \quad P(Y = 1|x = x^{**}) = 1 - \epsilon$$

where $0 < \epsilon < 1$ is known and $\theta < 1 - \epsilon$.

Let $F(z)$ be a strictly increasing cumulative distribution function (c.d.f) having probability density function $f(z)$. We consider a dose-toxicity relationship of the form

$$(4) \quad P(Y = 1|x) = F\left(F^{-1}(1 - \epsilon) + \beta \log\left(\frac{x - x^*}{x^{**} - x^*}\right)\right),$$

where β is unknown, and $0 < \beta^* \leq \beta \leq \beta^{**}$ for some positive real numbers β^* and β^{**} . This model assumes that the quantiles of F are linear in the log-standardized dose $z = \log[(x - x^*)/(x^{**} - x^*)]$. An example of F that is commonly used in practice is the logistic model $F(z) = e^z/(1 + e^z)$. It is easy to verify that model (4) satisfies the constraints (2) and (3). The condition $\beta > 0$ implies that the probability of DLT is an increasing function of dose. Let $\phi = F^{-1}(1 - \epsilon) - F^{-1}(\theta)$. Using (4), it can be shown that the MTD γ defined in (1) is

$$(5) \quad \gamma = x^* + (x^{**} - x^*)e^{-\phi/\beta}.$$

This also shows that $\gamma \in [x^*, x^{**}]$. Let $\gamma' = \log((\gamma - x^*)/(x^{**} - x^*))$ be the MTD on the log-standardized scale. Then (5) implies that $\gamma' = -\phi/\beta$.

2.1 Dose Escalation Based on Bayesian Estimates

Let $G(u) = F(F^{-1}(\theta) + \phi + u)$, $g(u) = G'(u)$, and $z_1 = -\phi/\beta^*$ be the level assigned to the first patient. Then,

$$\begin{aligned} G(\beta z_1) &= F(F^{-1}(\theta) + F^{-1}(1 - \epsilon) - F^{-1}(\theta) + \beta z_1) = F(F^{-1}(1 - \epsilon) + \beta z_1) \\ &= P(Y = 1|z_1) \leq F(F^{-1}(1 - \epsilon) + \beta^* z_1) = F(F^{-1}(\theta)) = \theta, \end{aligned}$$

since $z_1 < 0$ and $F(z)$ is strictly increasing. This shows that this log-standardized dose z_1 is safe in the sense that the probability of DLT at this level does not exceed θ . Let $D_n = \{(z_i, Y_i), i = 1, \dots, n\}$ be the data after enrolling n patients to the trial where Y_i is the observed DLT status of the patient getting level z_i , $z_i \in L^* = [-\frac{\phi}{\beta^*}, -\frac{\phi}{\beta^{**}}]$.

Let $h(\beta)$ be a prior density function for the parameter β on $[\beta^*, \beta^{**}]$ and $\Pi_n(\beta) = \Pi(\beta|D_n)$ the posterior c.d.f given the data D_n . Let $0 < \alpha < 1$. A sequence of dose levels z_n such that

$$(6) \quad P(z_n \leq -\phi/\beta|D_{n-1}) \geq 1 - \alpha$$

for all $n \geq 2$ is called Bayesian-feasible at level $(1 - \alpha)$; see Zacks et al. [36]. Let

$$(7) \quad z_n^{(\alpha)} = -\frac{\phi}{\Pi_{n-1}^{-1}(\alpha)}, n \geq 2.$$

Then, it is easy to verify that for all $n \geq 2$, $z_n^{(\alpha)}$ is Bayesian-feasible at level $(1 - \alpha)$. The choice of $z_n^{(\alpha)}$ as the log-standardized dose levels in the trial implies that the posterior probability of exceeding the MTD is equal to the feasibility bound α . Let $\mathcal{F}_n = \sigma(D_n)$ be the sigma-field generated by D_n and $\psi^{(\alpha)}$ be the class of all Bayesian-feasible sequences $z_n \in \mathcal{F}_n$ of level $(1 - \alpha)$.

DEFINITION 2.1. *A sequence of levels $\{z_n^*, n \geq 1\} \in \psi^{(\alpha)}$ is called optimal Bayesian-feasible at level $(1 - \alpha)$, if for all $N \geq 1$,*

$$\sum_{n=1}^N E_h\{(\gamma' - z_n^*)^+\} = \inf_{\{z_n\} \in \psi^{(\alpha)}} \sum_{n=1}^N E_h\{(\gamma' - z_n)^+\},$$

where $z^+ = zI(z > 0)$ denotes the positive part of a random variable.

This means that z_n^* minimizes the average amount by which patients are underdosed. Using the law of total expectation, Zacks et al. [36] showed that $z_n^{(\alpha)}$ is optimal Bayesian feasible. Conditions under which this sequence converges to the true MTD in probability are stated in the next theorem.

THEOREM 2.1. *Suppose that for $\beta_0 \in [\beta^*, \beta^{**}]$,*

1. $0 < \epsilon_1 < G(-\beta_0\phi/\beta^*) \leq G(-\beta_0\phi/\beta^{**}) \leq 1 - \epsilon$
2. $0 < \epsilon_2 < \inf\{g(\beta_0x) : x \in L^*\} \leq \sup\{g(\beta_0x) : x \in L^*\} \leq g^*$
3. $g(x)$ is continuously differentiable
4. $-\infty < \inf\{g'(\beta_0x) : x \in L^*\} \leq \sup\{g'(\beta_0x) : x \in L^*\} < \infty$.
5. $h(\beta)$ is uniform on $[\beta^*, \beta^{**}]$. Then,

$$z_n^{(\alpha)} \xrightarrow{p} -\phi/\beta_0 \text{ as } n \rightarrow \infty.$$

PROOF. See Zacks et al. [36]. \square

2.2 Coherence of EWOC

Coherence of adaptive designs was introduced by Cheung [7] in the context of cancer phase I clinical trials. Due to ethical concerns, the dose of a cytotoxic agent for the next patient in a trial should not be higher than the current allocated dose if the current patient exhibits DLT. Likewise, the dose for the next patient should not be lower than the current one if the current patient does not exhibit DLT. This desirable property is known as coherence and Cheung [7] showed that CRM is coherent. The author also showed how the coherence property can be lost when ad-hoc modifications are introduced to CRM. In this section, we show that EWOC as described in Section (2.2) is coherent.

Let $F(x, \gamma) = P(Y = 1|x)$ be the model given in (4) reparameterized in terms of the MTD γ . Let $D_n = \{(x_1, Y_1), \dots, (x_n, Y_n)\}$ be the data generated using EWOC scheme described in Section (2.2). This design is said to be coherent in escalation if for all $n \geq 2$, $x_n \geq x_{n-1}$ whenever $Y_{n-1} = 0$. The design is said to

be coherent in de-escalation if for all $n \geq 2, x_n \leq x_{n-1}$ whenever $Y_{n-1} = 1$. The design is said to be coherent if it is coherent in both escalation and de-escalation.

THEOREM 2.2. *Suppose that $F(x, \gamma)$ is non-increasing in gamma for fixed dose x . Then EWOC scheme described in Section (2.2) is coherent.*

The proof of Theorem (2.2) is given in the Appendix. It is easy to verify that the monotonicity condition on $F(x, \gamma)$ is satisfied by model (4), and in particular, the logistic function.

2.3 Two-Parameter Logistic Model

Denote by X_{\min} and X_{\max} the minimum and maximum dose levels available for use in the trial. One chooses these levels in the belief that X_{\min} is safe when administered to humans. Babb et al. [3] considered a two-parameter logistic model for the dose-toxicity relationship:

$$(8) \quad P(Y = 1 | Dose = x) = \frac{\exp(\beta_0 + \beta_1 x)}{1 + \exp(\beta_0 + \beta_1 x)},$$

where we assume that $\beta_1 > 0$ so that the probability of DLT is a monotonic increasing function of dose. Model (8) is reparameterized in terms of the MTD γ and the probability of DLT at the starting dose ρ_0 , parameters clinicians can easily interpret. This might be advantageous since γ is the parameter of interest and one often conducts preliminary studies at or near the starting dose so that one can select a meaningful informative prior for ρ_0 . Using the definition of the MTD in (1) and (8), it can be shown that

$$(9) \quad \begin{aligned} \beta_0 &= \frac{X_{\min} \text{logit}(\theta) - \gamma \text{logit}(\rho_0)}{x_{\min} - \gamma} \\ \beta_1 &= \frac{\text{logit}(\rho_0) - \text{logit}(\theta)}{x_{\min} - \gamma} \end{aligned}$$

The second equation in (9) shows that the assumption that $\beta_1 > 0$ implies $0 < \rho_0 < \theta$.

2.3.1 Trial Design After specifying a prior distribution $h(\rho_0, \gamma)$ for (ρ_0, γ) , denote by $\Pi_n(\gamma)$ the marginal posterior c.d.f of γ given D_n . EWOC can be described as follows. The first patient receives the dose $x_1 = X_{\min}$ and conditional on the event $\{y_1 = 0\}$, the $(n + 1)$ st patient receives the dose $x_{n+1} = \Pi_n^{-1}(\alpha)$ so that the posterior probability of exceeding the MTD is equal to the feasibility bound α . If $y_1 = 1$, we recommend that the clinician stops the trial. Calculation of the marginal posterior distribution of γ is performed using numerical integration, see [3]. Often in practice, phase I clinical trials are typically based on a small number of prespecified dose levels d_1, \dots, d_r . In this case, the $(n + 1)$ st patient receives the dose

$$\hat{d}_{n+1} = \max_{1 \leq i \leq r} \{d_i : d_i - x_{n+1} \leq T_1 \text{ and } \Pi_n(x_{n+1}) - \alpha \leq T_2\}$$

where T_1, T_2 are nonnegative numbers we refer to as tolerances. We note that this design scheme does not require that we know all patient responses before we

can treat a newly accrued patient. Instead, we can select the dose for the new patient on the basis of the data currently available. At the conclusion of the trial, the MTD is estimated by minimizing the posterior expected loss with respect to some suitable loss function l . One should consider asymmetric loss functions since underestimation and overestimation have very different consequences. Indeed, the dose x_n selected by EWOC for the n th patient corresponds to the estimate of γ having minimal risk with respect to the asymmetric loss function

$$l_\alpha(x, y) = \begin{cases} \alpha(\gamma - x) & \text{if } x \leq \gamma \text{ that is, if } x \text{ is an underdose} \\ (1 - \alpha)(x - \gamma) & \text{if } x > \gamma \text{ that is, if } x \text{ is an overdose} \end{cases}$$

Note that the loss function l_α implies that for any $\delta > 0$, the loss incurred by treating a patient at δ units above the MTD is $(1 - \alpha)/\alpha$ times greater than the loss associated with treating the patient at δ units below the MTD. This interpretation might provide a meaningful basis for the selection of the feasibility bound. The above methodology can be implemented using the user friendly software of Rogatko et al. [26].

2.3.2 Correlated Priors on ρ_0 and γ In models (4) and (8), we assumed that the support of the MTD was strictly contained in $[x^*, x^{**}]$ and $[X_{\min}, X_{\max}]$, respectively. The assumption that γ is bounded from above may be too restrictive. In the absence of toxicity, this assumption causes the dose escalation rate to slow down and in general, the target MTD will never be achieved if it lies outside the support of γ . Furthermore, since the support of the probability of DLT at the initial dose ρ_0 is $[0, \theta]$ and γ is a function of θ , the assumption of prior independence between ρ_0 and γ may not be realistic. Intuitively, the closer is ρ_0 to θ , the closer the MTD is to X_{\min} . Tighiouart et al. [31] introduced a class of correlated priors for $h(\rho_0, \gamma)$ on $[0, \theta] \times [X_{\min}, \infty)$ using truncated normal distributions for the parameter γ . They showed that a candidate joint prior for (ρ_0, γ) with negative a priori correlation structure results in a safer trial than the one that assumes independent priors for these two parameters while keeping the efficiency of the estimate of the MTD essentially unchanged.

2.4 EWOC with Varying Feasibility Bound

Many of the phase I cancer trials the authors designed at Fox Chase Cancer Center and Winship Cancer Institute used a variable feasibility bound α , see Babb and Rogatko [4; 2], Cheng et al. [6], Tighiouart and Rogatko [30; 29], and Xu et al. [35]. The rationale behind this approach is that uncertainty about the MTD is high at the onset of the trial and a small value of α offers protection against the possibility of administering dose levels much greater than the MTD. As the trial progresses, uncertainty about the MTD declines and the likelihood of selecting a dose level significantly above the MTD becomes significantly smaller. However, design operating characteristics were not studied. Chu et al. [9] compared the performance of different versions of CRM with EWOC with both constant and varying α . The design of EWOC with varying α was termed "hybrid design". The authors conducted extensive simulations to compare these designs in terms of (1) the proportion of patients given doses above the "true" MTD and (2) the proportion of times the recommended dose is the "true" MTD after each patient is enrolled in the trial and his or her DLT status is resolved. It was found in

general that both the hybrid and CRM designs had better convergence rate than EWOC with fixed α and that EWOC with fixed and varying feasibility bound α provide a better overdose protection than the CRM designs in the sense that fewer patients are given doses above the "true" MTD.

2.5 Example

EWOC was used to design a phase I clinical trial that involved the R115777 drug at Fox Chase Cancer Center in Philadelphia, USA in 1999. R115777 is a selective nonpeptidomimetic inhibitor of farnesyltransferase (FTase), one of several enzymes responsible for posttranslational modification that is required for the function of p21(ras) and other proteins. This was a repeated dose, single center trial designed to determine the MTD of R115777 in patients with advanced incurable cancer. The target probability of DLT was set to $\theta = 1/3$. The dose-escalation scheme was designed to determine the MTD of R115777 when drug is administered orally for 12 hours during 21 days followed by a 7- day rest. This constitutes one cycle of therapy. Toxicity was assessed by the National Cancer Institute (NCI) Common Toxicity Criteria [20]. DLT was determined by week 3 of cycle 1, as defined by Grade III non-hematological toxicity (with the exception of alopecia or nausea/vomiting) or hematological Grade IV toxicity with a possible, probable, or likely causal relationship to administration of R115777. Dosing continued until there was evidence of tumor progression or DLT leading to permanent discontinuation. The initial dose judged to be safe by the clinician for this study was $X_{\min} = 60 \text{ mg/m}^2$ and the maximum allowable dose was $X_{\max} = 600 \text{ mg/m}^2$. More details about the dosing regimen for this trial can be found Tighiouart and Rogatko [29]. Assuming vague priors for ρ_0 on $[0, \theta]$ and γ on $[60, 600]$, the prior probability density of (ρ_0, γ) is

$$h(\rho_0, \gamma) = \begin{cases} 1/180 & \text{if } (\rho_0, \gamma) \in [0, 1/3] \times [60, 600] \\ 0 & \text{otherwise} \end{cases}$$

Thus, ρ_0 and γ are independent a priori, uniformly distributed over their corresponding interval. Figure (1) shows the posterior distributions of the MTD as the trial progressed and Figure (2) shows the posterior density of the MTD after 33 patients have been treated. The posterior mode is 323 which correspond to the 47th percentile of the distribution. In this trial, we used a variable feasibility bound α , starting with $\alpha = 0.3$, this value being a compromise between the therapeutic aspect of the agent and its toxic side effects. As the trial progressed, α increased in small increments until $\alpha = 0.5$ so that, by the end of the trial, the given dose corresponds to the 50th percentile, i.e., the median of the marginal posterior probability density function. Thus, the dose to be given to the 34th patient is 328. The 95% highest posterior density interval is [160.5, 536.1].

3. ADJUSTING FOR PATIENTS' BASELINE COVARIATES

A key assumption implied by the definition of the phase I target dose (MTD) is that every subgroup of the patient population has the same MTD. That is, it is assumed that the patient population is homogeneous in terms of treatment tolerance and every patient should be treated at the same dose. As a result, no allowance is made for individual patient differences in susceptibility to treatment [11].

Babb and Rogatko [4] extended EWOC to allow the incorporation of information concerning individual patient differences in susceptibility to treatment. The method adjusts doses according to patient-specific characteristics while safeguarding against overdosing.

3.1 Model

Let W be a p -dimensional baseline covariate vector. We consider the dose-toxicity model

$$(10) \quad P(Y = 1 | Dose = x, W = w) = \frac{\exp(\beta_0 + \beta_1 x + \eta'w)}{1 + \exp(\beta_0 + \beta_1 x + \eta'w)},$$

where $\eta \in \mathbb{R}^p$ is the effect of the baseline covariate vector on DLT. Let $p_x(w) = P(Y = 1 | Dose = x, W = w)$. We assume that $\beta_1 > 0$ so that $p_x(w)$ is an increasing function of dose x for fixed w . The MTD for a patient with baseline covariate value w is defined as the dose $\gamma(w)$ that results in a probability equal to θ that a DLT will manifest. It follows from model (10) that

$$(11) \quad \gamma(w) = \beta_1^{-1} \left[\log \left(\frac{\theta}{1 - \theta} \right) - \beta_0 - \eta'w \right].$$

As in Section (2.3), we reparameterize this model in terms of $(\gamma(w^*), \rho)$ for a selected value of the baseline covariate vector $w = w^*$ and ρ is a $p+1$ dimensional nuisance parameter.

3.1.1 Trial design Let $h(\rho, \gamma(w^*))$ be a prior distribution for $(\rho, \gamma(w^*))$ and denote by $\Pi_{n,w^*}(\gamma(w^*))$ the marginal posterior c.d.f of $\gamma(w^*)$ given the data $D_n = \{(x_1, Y_1, w_1), \dots, (x_n, Y_n, w_n)\}$. The first patient receives the dose $x_1 = X_{\min}$ and conditional on the event $\{y_1 = 0\}$, the $(n+1)$ st patient with covariate vector value w_{n+1} receives the dose $x_{n+1} = \Pi_{n,w_{n+1}}^{-1}(\alpha)$ so that the posterior probability of exceeding the MTD is equal to the feasibility bound α . Note that here, $\Pi_{n,w_{n+1}}^{-1}(\cdot)$ is the inverse c.d.f of $\Pi_{n,w_{n+1}}(\gamma(w_{n+1}))$.

For a binary covariate W , Tighiouart et al. [32] studied operating characteristics of this model with extensive simulations under different scenarios for the underlying true MTDs. They found that if the two MTDs are different and the design does not adjust for this heterogeneity, then the trial will result in more patients being overdosed. If the two MTDs are different and parallel trials are used, then the estimates of the MTDs are less efficient. Finally, if the two MTDs are the same and the design adjusts for patients' heterogeneity, then few more patients can be overdosed if the true MTD is low relative to a design with no covariate but the difference is not practically important. Thus, we stand to lose little if we do include a statistically non-significant covariate in the model. This conclusion is in agreement with the findings in O'Quigley et al. [22].

3.2 Example

ABR-217620 (naptumomab estafenatox) is a recombinant fusion protein that consists of the 5T4Fab moiety genetically fused to the engineered superantigen variant SEA/E-120. This fusion protein is a new generation tumor-targeted superantigen based on the previously described ABR-214936 (anatumomab mafenatox). ABR-217620 was designed to reduce antigenicity and toxicity. We use

model (10) with $W = C$ representing the Anti SEA/E120 covariate to design a phase I study for non-small cell lung cancer (NSCLC) patients. The goal is to determine the MTD of ABR-217620 as a function of patients' baseline Anti SEA/E120 and test whether the neutralizing effect of Anti SEA/E120 on the cytotoxic agent which was observed by Babb and Rogatko [4] has been reduced or eliminated with this new agent. The modeling approach is similar to the PNU trial described in [4]; the target probability of DLT θ was set to 0.2. The feasibility bound α was set at 0.25 for the first nine patients, then was increased to a maximum value of 0.5 by increments of 0.05 every time a new patient was enrolled in the trial and a DLT assessment was resolved. Based on preliminary clinical data, the minimum and maximum allowable dose for ABR-217620 set by the clinicians are $x_{\min} = 1$ ug/kg and $x_{\max} = 100$ ug/kg. The minimum and maximum values of Anti SEA/E120 anticipated in the trial are $c_1 = 0$ pmol/ml, and $c_2 = 200$ pmol/ml, respectively. As in the PNU trial [4], we reparameterize model (10) in terms of $\gamma_{\max} = \gamma(c_2)$, $\rho_1 = p_{x_{\min}}(c_1)$, $\rho_2 = p_{x_{\min}}(c_2)$ with $(\gamma_{\max}, \rho_1, \rho_2)$ uniformly distributed on $\{(x, y, z) : y \in (0, \theta], z \in (0, y), x \in [1, 100]\}$ a priori.

Figure (3) shows the doses allocated to all 39 patients as a function of their pre-treatment Anti SEA/E120. The solid line is the estimated conditional MTD, obtained by taking the posterior median of the marginal distribution of the MTD conditional on the covariate Anti SEA/E120. The dashed lines delimit the 95% Bayesian credible region. Six patients experience DLT (15.4%) and the MTD seems to indicate that the neutralizing effect of Anti SEA/E120 has been reduced considerably. The protocol was amended to include patients with renal cell (RCC) and pancreatic cancer (PC). Figures (4) and (5) show the doses allocated to NSCLC & PC patients and RCC patients, respectively. The solid line represents the conditional MTD obtained after fitting the data in each group to model (10), reparameterized in terms of $(\gamma_{\max}, \rho_1, \rho_2)$. This shows that NSCLC & PC patients were treated at sub-optimal doses and RCC patients were overdosed, with 36.4% experiencing DLT, way above the target probability of DLT $\theta = 0.2$.

The effects of Anti SEA/E120 and type of cancer were tested by fitting model (10) with $W = (C, Z)$, where C is the baseline Anti SEA/E120 and Z is a binary covariate representing the cancer type, $Z = z_1 = 1$ for NSCLC and PC patients and $Z = z_2 = 0$ for RCC patients. To be consistent with the priors used to design the trial, we reparameterized the model in terms of $\gamma_{\max} = \gamma(c_2, z_1)$, $\rho_1 = p_{x_{\min}}(c_1, z_1)$, $\rho_2 = p_{x_{\min}}(c_2, z_1)$, and $\rho_3 = p_{x_{\min}}(c_1, z_2)$. Independent uniform priors are placed on these parameters. It can be shown that this induces priors centered at 0 for the Anti SEA/E120 and cancer type effect parameters η_1 and η_2 . We used WinBUGS [18] to fit this model and the 95% HPD intervals for the parameters η_1 and η_2 were $(-0.14, 0.24)$ and $(-4.6, 0.6)$, respectively. We conclude that the agent ABR-217620 was successful in reducing the neutralizing capacity of Anti SEA/E120 and that the phase II dose should be carefully tailored to account for patients' cancer type and hence avoid excessive overdosing and underdosing patients.

4. DISCUSSION

In this paper, we described EWOC, a Bayesian dose finding design for cancer phase I clinical trials. The method is flexible enough to allow prior information about the drug from laboratory or animal studies to be incorporated in the model,

is coherent, makes use of all the information available at the time of each dose assignment, and controls the probability of overdosing patients at each stage. EWOC can be implemented with the user friendly software EWOC 2.1 [26] or WinBUGS [18] for general class of prior distributions [31]. The two-parameter model described in section 2.3 accounts for the uncertainty regarding the probability of DLT at the initial dose by placing a vague prior distribution on ρ_0 . If expert opinion about this parameter is available, then it should be incorporated in the prior for ρ_0 . In particular, if the clinician strongly believes that this prior can be approximated by a point mass distribution, then the one-parameter model described in Section 2.1 may be used. In any case, design operating characteristics should be performed with a sensitivity analysis about the parameter ρ_0 when designing the trial. Our own experience in designing dose-finding studies in cancer is that the uncertainty of the clinicians regarding the probability of DLT at the initial dose is large. Thus, in more than ten years of designing trials with EWOC, the use of a one-parameter model was never chosen by the clinical researchers we worked with.

It is worth highlighting that the values of α and θ are chosen independently when the trial is designed. They have distinct meanings and functions. For example, taking a value of α greater than θ only affects the loss function used to estimate the next dose and the MTD at the conclusion of the trial. It does not mean that patients are given doses at a rate above the target probability of DLT θ . When $\alpha = 0.5$, the method differs from CRM in the sense that the loss functions are different. The loss function for EWOC is taken with respect to the parameter γ , the MTD. The overprotection property of EWOC is with respect to the posterior distribution of the MTD, given the data. The overprotection property states that the posterior probability of exceeding the MTD given the current data is bounded by α . This overprotection is as good as the posterior distribution of the MTD at each stage of the trial. For instance, if we used a flat prior on the MTD and the true MTD turns out to be very close to the initial dose, then it would take many patients for the median of the posterior distribution to cluster around the true MTD.

Another aspect of cancer phase I clinical trials not discussed here is the choice of the number of patients to enroll. Most sample size recommendations in the literature are based on pre-specified stopping rules, see for example, Zohar and Chevret [37] on selecting the number of patients by considering different stopping rules using the CRM. Lin and Shih [17] and Ivanova [16] describe sample size recommendations based on the expected number of patients allocated to each dose selected from a set of prespecified dose levels. However, these methods apply to a prespecified set of discrete doses and it is not clear how they can be applied to continuous doses. Unlike the frequentist approach, there is no consensus on a specific Bayesian method for the sample size determination problem, see Adcock [1] for a review of Bayesian approaches. We conducted extensive simulation studies in order to estimate the sample size based on a desired accuracy of the Bayes estimate on the average. Specifically, we determined the minimum number of patients so that the posterior variance of the MTD on the average over all possible trials is no more than a specified margin. Tabulated values of the average mean posterior standard deviation, length of 90% and 95% HPD intervals for different values of the target probability of DLT θ are available from the authors upon

request.

The methodology described in this article assumes that DLT status is binary and does not account for patients' time to toxicity. Information on time to DLT is crucial to clinicians in that it permits a dynamic updating of the posterior distribution of the MTD based on the number of patients who experienced DLT and the ones who are still at risk. If new patients are eligible to enter the clinical trial while the DLT status of currently enrolled patients are still being resolved, then the new patients are allocated to the current established dose because it is not ethical to resolve DLT status at the expense of treatment delay. In this case, there is no adaptation to the most current information and the design will not be efficient. Time to DLT was first investigated by Cheung and Chapell [8] and later adapted to estimation of a maximum cumulative dose by Braun et al. [5]. These methods are extensions of the CRM and incorporate information on partially observed patients using weighted binomial likelihoods. EWOC can be adapted to this framework by modeling time to DLT as a Cox [10] type model with a parametric or non-parametric baseline risk of toxicity $h_0(t)$. We are currently investigating the performance of a large class of models within this framework via extensive simulations.

ACKNOWLEDGMENT

We would like to thank the reviewer and executive editor for valuable comments and suggestions in writing the manuscript.

5. APPENDIX

Proof of Theorem 2.2.

Let $\Pi_n(t)$ be the posterior c.d.f of γ given $D_n, n \geq 2$. Then, it suffices to show that

1. $\Pi_n(t) \leq \Pi_{n-1}(t)$ for all t whenever $Y_n = 0$
2. $\Pi_n(t) \geq \Pi_{n-1}(t)$ for all t whenever $Y_n = 1$

Let

$$L_n(\gamma|D_n) = \prod_{i=1}^n (F(x_i, \gamma))^{Y_i} (1 - F(x_i, \gamma))^{1-Y_i}$$

be the likelihood function and $h(\gamma)$ be a proper prior distribution for γ . To simplify notation, let $L_n(\gamma|D_n) = L_n(\gamma)$, $F_i(\gamma) = F(x_i, \gamma)$ and suppose that $x^* = 0, x^{**} = 1$. Using Bayes rule, the posterior c.d.f $\Pi_n(t)$ given D_n is

$$\Pi_n(t) = \frac{\int_0^t L_n(\gamma)h(\gamma)d\gamma}{\int_0^1 L_n(\gamma)h(\gamma)d\gamma}.$$

Suppose that $Y_n = 0$. Then, $L_n(\gamma) = L_{n-1}(\gamma)(1 - F_n(\gamma))$ and

$$\Pi_n(t) = \frac{\int_0^t L_{n-1}(\gamma)(1 - F_n(\gamma))h(\gamma)d\gamma}{\int_0^1 L_{n-1}(\gamma)(1 - F_n(\gamma))h(\gamma)d\gamma}.$$

It follows that

$$\begin{aligned}\Pi_n(t) - \Pi_{n-1}(t) &= \frac{\int_0^t L_{n-1}(\gamma)(1 - F_n(\gamma))h(\gamma)d\gamma}{\int_0^1 L_{n-1}(\gamma)(1 - F_n(\gamma))h(\gamma)d\gamma} - \frac{\int_0^t L_{n-1}(\gamma)h(\gamma)d\gamma}{\int_0^1 L_{n-1}(\gamma)h(\gamma)d\gamma} \\ &= A^{-1} \left[\int_0^t \int_0^1 L_{n-1}(\gamma)L_{n-1}(\gamma')h(\gamma)h(\gamma')[F_n(\gamma') - F_n(\gamma)]d\gamma'd\gamma \right]\end{aligned}$$

where

$$A = \int_0^1 \int_0^1 L_{n-1}(\gamma)(1 - F_n(\gamma))h(\gamma)L_{n-1}(\gamma')h(\gamma')d\gamma'd\gamma.$$

$$\Pi_n(t) - \Pi_{n-1}(t) = A^{-1} \int_0^t \int_0^t L_{n-1}(\gamma)L_{n-1}(\gamma')h(\gamma)h(\gamma')[F_n(\gamma') - F_n(\gamma)]d\gamma'd\gamma$$

$$+ A^{-1} \int_0^t \int_t^1 L_{n-1}(\gamma)L_{n-1}(\gamma')h(\gamma)h(\gamma')[F_n(\gamma') - F_n(\gamma)]d\gamma'd\gamma$$

$$= A^{-1} \int_0^t \int_t^1 L_{n-1}(\gamma)L_{n-1}(\gamma')h(\gamma)h(\gamma')[F_n(\gamma') - F_n(\gamma)]d\gamma'd\gamma \leq 0$$

since $F_n(\gamma)$ is non-increasing in γ . Hence, $\Pi_n(t) \leq \Pi_{n-1}(t)$, which implies that $\Pi_n^{-1}(\alpha) \geq \Pi_{n-1}^{-1}(\alpha)$, i.e., $x_{n+1} \geq x_n$. Using a similar argument, one can show that $\Pi_n^{-1}(\alpha) \leq \Pi_{n-1}^{-1}(\alpha)$ if $Y_n = 1$. This shows that EWOC is coherent. \square

REFERENCES

- [1] ADCOCK, C. J. (1997). Sample size determination: A review. *Statistician* **46** 261-283.
- [2] BABB, J. and ROGATKO, A. (2004). *Contemporary Biostatistical Methods in Clinical Trial Bayesian Methods for Cancer Phase I Clinical Trials* 1-39. Marcel Dekker: New York.
- [3] BABB, J., ROGATKO, A. and ZACKS, S. (1998). Cancer Phase I Clinical Trials: Efficient Dose Escalation With Overdose Control. *Statistics in Medicine* **17** 1103-1120.
- [4] BABB, J. S. and ROGATKO, A. (2001). Patient specific dosing in a cancer phase I clinical trial. *Statistics in Medicine* **20** 2079-2090.
- [5] BRAUN, T. M., LEVINE, J. E. and FERRARA, J. (2003). Determining a maximum tolerated cumulative dose: dose reassignment within the TITE-CRM. *Controlled Clinical Trials* **24** 669-681.
- [6] CHENG, J. D., BABB, J. S., LANGER, C., AAMDAL, S., ROBERT, F., ENGELHARDT, L. R., FERNBERG, O., SCHILLER, J., FORSBERG, G., ALPAUGH, R. K., WEINER, L. M. and ROGATKO, A. (2004). Individualized patient dosing in phase I clinical trials: The role of EWOC in PNU-214936. *Journal of Clinical Oncology* **22** 602-609.
- [7] CHEUNG, Y. K. (2005). Coherence principles in dose-finding studies. *Biometrika* **92** 863-873.
- [8] CHEUNG, Y. K. and CHAPPELL, R. (2000). Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics* **56** 1177-1182.
- [9] CHU, P. L., LIN, Y. and SHIH, W. J. (2009). Unifying CRM and EWOC designs for phase I cancer clinical trials. *Journal of Statistical Planning and Inference* **139** 1146-1163.

- [10] COX, D. R. (1972). Regression Models and Life-Tables. *Journal of the Royal Statistical Society, Series B* **34** 187-220.
- [11] DILLMAN, R. O. and KOZIOL, J. A. (1992). Phase I cancer trials: limitations and implications. *Molecular Biotherapy* **4** 117-121.
- [12] FARIES, D. (1994). Practical modifications of the continual reassessment method for phase I cancer clinical trials. *Journal of Biopharmaceutical Statistics* **4** 147-164.
- [13] GARRETT-MAYER, E. (2006). The continual reassessment method for dose-finding studies: a tutorial. *Clinical Trials* **3** 57-71.
- [14] GATSONIS, C. and GREENHOUSE, J. (1992). Bayesian methods for phase I clinical trials. *Statistics in Medicine* **11** 1377-1389.
- [15] GOODMAN, S. N., ZAHURAK, M. L. and PIANTADOSI, S. (1995). Some practical improvements in the continual reassessment method for phase I studies. *Statistics in Medicine* **14** 1149-1161.
- [16] IVANOVA, A. (2006). Escalation, group and A + B designs for dose-finding trials. *Statistics in Medicine* **25** 3668-3678.
- [17] LIN, Y. and SHIH, W. J. (2001). Statistical properties of the traditional algorithm-based designs for phase I cancer clinical trials. *Biostatistics* **2** 203-215.
- [18] LUNN, D. J., THOMAS, A., BEST, N. and SPIEGELHALTER, D. (2000). WinBUGS – a Bayesian modelling framework: concepts, structure, and extensibility. *Statistics and Computing* **10** 325-337.
- [19] MOLLER, S. (1995). An extension of the continual reassessment methods using a preliminary up-and-down design in a dose finding study in cancer patients, in order to investigate a greater range of doses. *St* **14** 911-922.
- [20] NCI, (2003). Common Toxicity Criteria for Adverse Events v3.0 (CTCAE). Available at <http://ctep.cancer.gov/reporting/ctc.html>
- [21] 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.
- [22] O'QUIGLEY, J., SHEN, L. Z. and GAMST, A. (1999). Two-Sample Continual Reassessment Method. *Journal of Biopharmaceutical Statistics* **9** 17-44.
- [23] PIANTADOSI, S., FISHER, J. D. and GROSSMAN, S. (1998). Practical implementation of a modified continual reassessment method for dose-finding trials. *Cancer Chemotherapy and Pharmacology* **41** 429-436.
- [24] ROBERTS, T. G. J., GOULART, B., SQUITIERI, L., STALLINGS, S. C., HALPERN, E. F., CHABNER, B. A., GAZELLE, G. S., FINKELSTEIN, S. N. and CLARK, J. W. (2004). Trends in the risks and benefits to patients with cancer participating in phase I clinical trials. *Jama* **292** 2130-2140.
- [25] ROGATKO, A., GHOSH, P., VIDAKOVIC, B. and TIGHIOUART, M. (2008). Patient-Specific Dose Adjustment in the Cancer Clinical Trial Setting. *Pharmaceutical Medicine* **22** 345-350.
- [26] ROGATKO, A., TIGHIOUART, M. and XU, Z. (2005). EWOC 2.1 application software. Available at <http://winshipcancer.emory.edu/biostatistics/software.shtml>
- [27] ROSENBERGER, W. F. and HAINES, L. M. (2002). Competing designs for phase I clinical trials: a review. *Statistics in Medicine*. *Statistics in Medicine* **21** 2757-2770.
- [28] STORER, B. E. (2001). An evaluation of phase I clinical trial designs in the continuous dose-response setting. *Statistics in Medicine* **20** 2399-2408.
- [29] TIGHIOUART, M. and ROGATKO, A. (2006). *Dose Finding in Drug Development* Dose Finding in Oncology Parametric Methods 59-72. Springer: New York.
- [30] TIGHIOUART, M. and ROGATKO, A. (2006). *Statistical Methods for Dose-Finding Experiments* Dose Escalation with Overdose Control 173-188. John Wiley and Sons.
- [31] TIGHIOUART, M., ROGATKO, A. and BABB, J. S. (2005). Flexible Bayesian methods for cancer phase I clinical trials. Dose escalation with overdose control. *Statistics in Medicine* **24** 2183-2196.
- [32] TIGHIOUART, M., ROGATKO, A. and XU, Z. (2007). Incorporating patients characteristics in cancer phase I clinical trials using escalation with overdose control. Joint Statistical Meetings.
- [33] TING, N. (2006). *Dose Finding in Drug Development*. Springer: New York.
- [34] WHITEHEAD, J. (1997). Bayesian decision procedures with application to dose-finding studies. *International Journal of Pharmaceutical Medicine* **11** 201-208.
- [35] XU, Z., TIGHIOUART, M. and ROGATKO, A. (2007). EWOC 2.1: interactive software for dose escalation in cancer phase I clinical trials. *Drug Information Journal* **41** 221-228.




FIG 1. Posterior density of the MTD when the number of treated patients (from bottom to top) is 1, 5, 10, 15, 20, 25, 30, 33.

- [36] ZACKS, S., ROGATKO, A. and BABB, J. (1998). Optimal Bayesian-feasible dose escalation for cancer phase I trials. *Statistics and Probability Letters* **38** 215-220.
- [37] ZOHAR, S. and CHEVRET, S. (2001). The continual reassessment method: comparison of Bayesian stopping rules for dose-ranging studies. *Statistics in Medicine* **20** 2827-2843.
- [38] ZOHAR, S., LATOCHE, A., TACONNER, M. and CHEVRET, S. (2003). Software to compute and conduct sequential Bayesian phase I or II dose-ranging clinical trials with stopping rules. *Computer Methods and Programs in Biomedicine* **72** 117-125.




FIG 2. Posterior density of the MTD after 33 patients have been treated. The posterior mode is 323 (47th percentile) and the median and dose to be given to the 34th patient is 328. The 95% highest posterior density interval is [160.5, 536.1].




FIG 3. Dose allocation as a function of baseline Anti SEA/E120 during the trial for all 39 patients. The solid line is the MTD conditional on the covariate Anti SEA/E120 which corresponds to the posterior median of the conditional posterior distribution of the MTD and the dashed lines delimit the 95% Bayesian credible region.




FIG 4. Dose allocation as a function of baseline Anti SEA/E120 during the trial for all 28 NSCLC & PC patients. The solid line is the MTD conditional on the covariate Anti SEA/E120 which corresponds to the posterior median of the conditional posterior distribution of the MTD and the dashed lines delimit the 95% Bayesian credible region.




FIG 5. Dose allocation as a function of baseline Anti SEA/E120 during the trial for all 11 RCC patients. The solid line is the MTD conditional on the covariate Anti SEA/E120 which corresponds to the posterior median of the conditional posterior distribution of the MTD and the dashed lines delimit the 95% Bayesian credible region.