FDA and Innovative Designs: Case Study of a Missed Opportunity

André Rogatko, Galen Cook-Wiens, Mourad Tighiouart and Steven Piantadosi

Biostatistics and Bioinformatics Research Center, Cedars-Sinai Medical Center, 8700 Beverly Blvd., PACT, Suite 900C, Los Angeles, CA 90048

Abstract

The standard 3+3 or "modified Fibonacci" up-and-down (MF-UD) method of dose escalation is by far the most used design in dose-finding cancer early trials. MF-UD was the state-of-the-art in 1971 when it was first used in cancer clinical studies. Since then, over one hundred methodological papers proposing or evaluating dose escalation designs have been published and MF-UD has always shown inferior performance when compared with its competitors regarding maximizing number of patients treated at optimal doses and minimizing number of patients under or overdosed. A consequence of using less effective designs is that more patients are treated with doses outside the therapeutic window. We present a case study where the U S Food and Drug Administration (FDA) rejected the proposal to use Escalation with Overdose Control (EWOC), an established dose-finding method which has been extensively used in FDA-approved first in human trials and imposed a suboptimal design, a variation of the MF-UD, known as accelerated titration (AT) design. We show through extensive simulation studies that the AT design has poor operating characteristics relative to two versions of EWOC under several practical scenarios.

Key Words: dose finding, cancer early trials, phase I trials, escalation with overdose control, accelerated titration, EWOC

1. Introduction

Clinical trials of new anti-cancer therapies are widespread, critically important tools in the search for more effective cancer treatments. Cancer trials typically proceed through several distinct phases. The major objective in dose finding (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), and its definition depends on the treatment under investigation, the severity and reversibility of its side effects, and on clinical attributes of the target patient population. 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 [1],

$$P\{\text{DLT} \mid \text{Dose} = \gamma\} = \theta \tag{1}$$

One of the very few options of patients with advanced cancer refractory to available treatment is to participate in dose-finding (phase I) oncology trial. Horstmann et al.[2]

reviewed all non-pediatric phase I oncology trials sponsored by the Cancer Therapy Evaluation Program at the National Cancer Institute between 1991 and 2002. They analyzed 460 trials involving 11,935 participants, all of whom were assessed for toxicity and 10,402 of whom were assessed for a response to therapy. The overall response rate (i.e., for both complete and partial responses) was 10.6 percent, with considerable variation among trials. These results demonstrate that it is reasonable to expect a therapeutic intent from a phase I trial. Thus, to increase the chances that a patient will benefit from participating in a phase I trial one should improve the design of these trials so that the number of patients receiving optimal doses is maximized. Consequently, more patients would be treated with therapeutic doses of promising new agents, and fewer patients would have to suffer the deleterious effects of toxic doses.

Since there is an ethical motivation to not harm the patient, controlling this risk is a desirable goal for physician and patients. Escalation with Overdose Control (EWOC) was the first dose-finding procedure to directly incorporate the ethical constraint of minimizing the chance of treating patients at unacceptably high doses[3]. Its defining property is that the expected proportion of patients treated at doses above the MTD is equal to a specified value α, the feasibility bound. This value is selected by the clinician and reflects his/her level of concern about overdosing. The method is flexible enough to allow prior information about the drug from laboratory or animal studies to be incorporated in the model, makes use of all the information available at the time of each dose assignment, controls the probability of overdosing patients at each stage, allows the estimation of the precision of the MTD, is optimally Bayesian feasible[4], produces a sequence of doses that converges in probability to the MTD[4], is coherent[5], and accounts for patients' pre-treatment characteristics[6-8]. EWOC can be implemented with the user friendly software EWOC [9, 10] or WinBUGS [11] for general class of prior distributions[12], or ordinal toxicity grades[13]. EWOC allows flexible patient enrollment, and conforms with the ethical goal of maximizing the number of patients receiving optimal doses. At the time this article was written, we were aware of twenty peer-reviewed articles describing trials designed with EWOC.

In June 2012 a protocol accompanying the Investigational New Drug (IND) application "*A Phase 1 Clinical and Pharmacokinetic Study of OVI-117 Administered Once Weekly for Three Consecutive Weeks, with Cycles Repeated Every Four Weeks, to Patients with Advanced Solid Tumors*" designed with EWOC was submitted to the U.S. Food and Drug Administration (FDA) and was reviewed by Frank H. Cross, Jr., Chief, Project Management Staff, Division of Oncology Products 1, Office of Hematology and Oncology Products (OHOP), Center for Drug Evaluation and Research (CDER). The interaction led to the rejection of EWOC and the imposition of the accelerated titration (AT) design, which is a variation of the "modified Fibonacci" up-and-down (MF-UD) method of dose escalation.

The standard 3+3 or MF-UD is by far the most used design in dose-finding cancer early trials[14]. MF-UD was the state-of-the-art in 1971 when it was first used in cancer clinical studies. Since then, over one hundred methodological papers proposing or evaluating dose escalation designs have been published and MF-UD has always shown inferior performance when compared with its competitors regarding maximizing number of patients treated at optimal doses and minimizing number of patients under or

overdosed[14]. AT was first proposed by Storer[15] and popularized by Simon and colleagues[16].

The goal of this manuscript is to compare the operating characteristics regarding safety and efficiency between AT and EWOC designs through extensive simulations studies under several practical scenarios.

2. Simulation Studies

2.1 Models, Responses and General Settings

Let $G = 0,1,\ldots,4$ be the maximum grade of toxicity experienced by a patient by the end of one cycle of therapy and define DLT as a maximum of grade 3 or 4 toxicity. Let

$$Y = \begin{cases} 0 \text{ if } G = 0 \text{ or } 1 \\ 1 \text{ if } G = 2 \\ 2 \text{ if } G = 3 \text{ or } 4 \end{cases} \tag{2}$$

The dose-toxicities relationship is modeled by

$$P(Y \geq j \mid x) = F(\alpha_j + \beta x),\ j = 1, 2, \tag{3}$$

where $F(\cdot)$ is a known strictly increasing c.d.f. This implies that $\alpha_2 \leq \alpha_1$. We assume that $\beta > 0$ so that the probability of DLT is an increasing function of dose. Doses are assumed to be continuous and standardized to the interval [0, 1]. The MTD, γ, is defined as the dose that is expected to produce DLT in a specified proportion θ of patients:

$$P(Y = 2 \mid x = \gamma) = F(\alpha_2 + \beta\gamma) = \theta. \tag{4}$$

We reparameterize model (2) in terms of ρ_0, the probability that a DLT manifests within the first cycle of therapy for a patient given dose $x = 0$

$$\rho_0 = P\left(Y = 2 \mid x = 0\right), \tag{5}$$

ρ_1, the probability that a grade 2 or more toxicity manifests within the first cycle of therapy for a patient given dose $x = 0$

$$\rho_1 = P(Y > 1 \mid x = 0), \tag{6}$$

and the MTD γ. This reparameterization is convenient to clinicians since γ is the parameter of interest, and both ρ_0 and ρ_1 are easier to interpret in practice. It can be shown that

$$\begin{aligned} &\alpha_1 = F^{-1}(\rho_1),\ \alpha_2 = F^{-1}(\rho_0)\ , \\ &\beta = \frac{1}{\gamma}\left(F^{-1}(\theta) - F^{-1}(\rho_0)\right). \end{aligned} \tag{7}$$

Five models that specify the dose-toxicities relationship in (3) were used in the simulation studies to generate responses:

(A) Proportional odds logistic with link function $F(w) = 1/(1+e^{-w})$,
(B) Normal link function with shape parameter $\boldsymbol{\sigma = 0.5}$,
(C) Normal link function with shape parameter with $\boldsymbol{\sigma = 2}$,
(D) Non-proportional odds with logistic link function used on (A) with $\boldsymbol{\rho_0 = 0.126}$, and
(E) Non-proportional odds with logistic link function used on (A) with $\boldsymbol{\rho_0 = 0.02}$.

For all simulations, the target probability of DLT is set to $\theta = 0.33$. Figure 1 shows the dose-toxicity relationship for P(DLT | dose) on the left hand side and for P(grade 2 or higher | dose) on the right hand side, for three selected values of true MTD, for five models considered in the simulation studies.

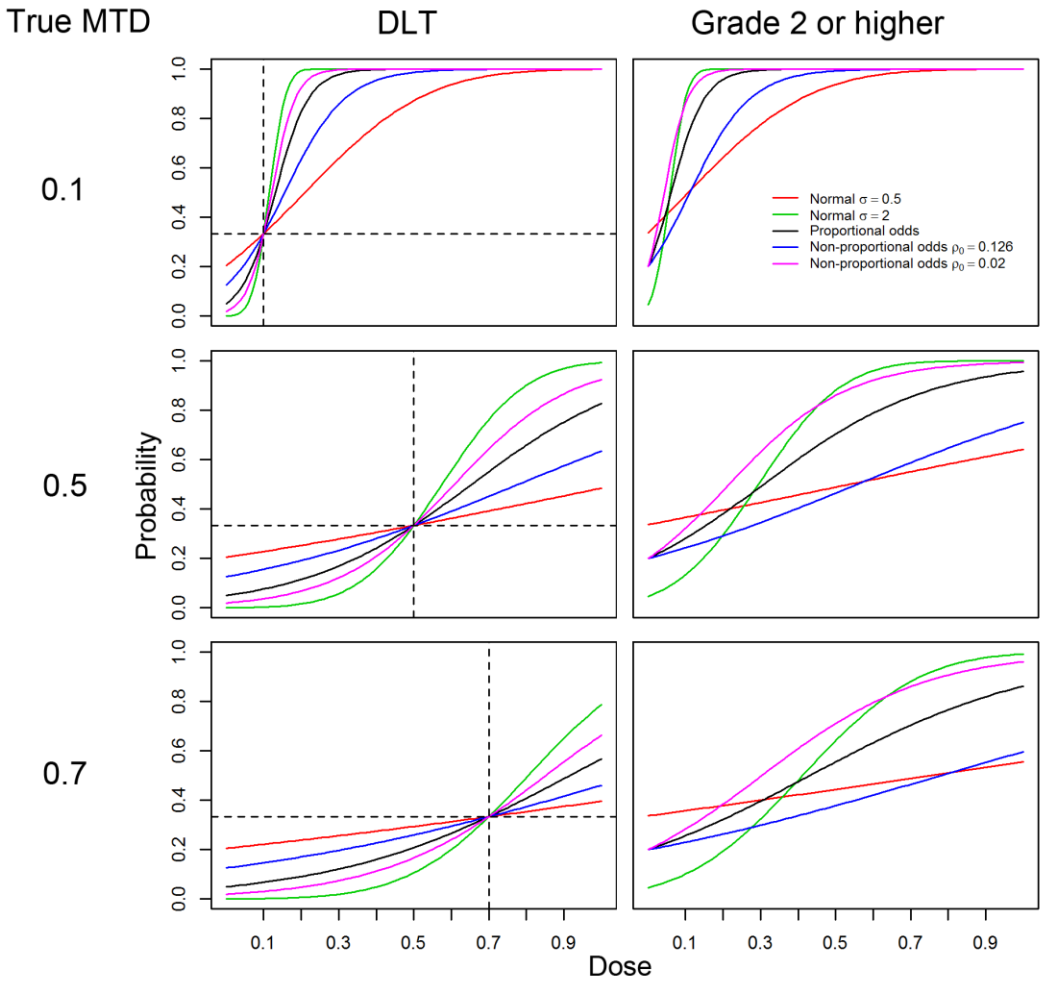

Figure 1: Dose-toxicity relationship for P(DLT | dose) on the left hand side and for P(grade 2 or higher | dose) on the right hand side, for three selected values of true MTD, for five models considered in the simulation studies: proportional odds logistic (black), normal with $\sigma = 0.5$ (red), normal with $\sigma = 2$ (green), non-proportional odds with $\rho_0 = 0.126$ (blue), and non-proportional odds with $\rho_0 = 0.02$ (magenta). The horizontal dashed lines represent the target probability of DLT θ=0.33 and the vertical lines correspond to the true values of the MTD γ.

Three selected values for the true MTD: {0.1, 0.5, 0.7} combined with three chosen values for the true ρ_1: {0.2, 0.5, 0.8} produced nine scenarios used in the simulations. One thousand trials were simulated for each scenario.

The designs compared in the study are described next.

2.2 AT Designs

Two starting doses were selected: 0.01 and 0.1. The AT design comprises two phases. In the first phase (accelerated phase), only one patient is treated at each dose level and doses increase at a faster rate than in the second phase. The second phase enrolls three patients and dose escalation proceeds using the standard MF-UD algorithm. Three sets of dose increments were chosen:

(1) Double the dose in the accelerated phase and increase by 50% in the MF-UD phase (recommended by FDA);
(2) Increase by 69% in the accelerated phase and 30% in the MF-UD phase; and
(3) Increase by 96% in the accelerated phase and 40% in the MF-UD phase.

Six versions of AT design were studied by combining two starting doses and three sets of dose increments. Thus, each version is identified by a vector with three elements

(starting dose, accelerated phase increase rate, MF-UD increase rate) (8)

resulting in: (0.01, 2, 1.5), (0.1, 2, 1.5), (0.01, 1.69, 1.3), (0.1, 1.69, 1.3), (0.01, 1.96, 1.4), and (0.1, 1.96, 1.4).

Each trial starts in the accelerated phase, and the first patient at the starting dose is given a safe response (grade 0 or 1). The accelerated phase proceeds as follows: the algorithm checks whether there have been any DLTs or grade 2 toxicities. If there have been no DLTs and no grade 2 responses then the accelerated phase is continued, the dose is increased according to the accelerated increment, and a response is simulated at that dose. If there have been any DLTs or grade 2 toxicities, the accelerated phase is ended, the dose is kept at the previous dose, and two more patients and responses are simulated at that dose. This gives three patients at the end of the accelerated phase that were treated at the dose that gave the first grade 2 or DLT response. The simulation then moves to the MF-UD phase. If in the accelerated phase the dose becomes larger than 1 and the previous dose was less than 1, then the maximum dose of 1 is assigned to the next patient and a response is simulated. If in the accelerated phase the dose becomes larger than 1 and the previous dose was equal to 1, then the trial is ended and no estimated MTD is assigned.

The MF-UD phase proceeds as follows: the algorithm counts the number of DLTs in the previous three patients. There are three cases that follow: 0 DLTs in the last three patients, 1 DLT, or more than 1 DLT. If there were *no DLTs* then the dose is increased, three patients are added at this higher dose and three responses are simulated, unless any of the following situations are true: there have been patients already treated at this higher dose, or more than three patients have already been treated at the previous dose. If patients were already treated at the higher dose then, as long as more than three patients were treated at the previous dose level, that previous level is the estimated MTD and the trial ends. If more than three patients have already been treated at the previous dose level, then that is the estimated MTD and the trial ends. Otherwise three more patients are

added at the previous dose so that a total of more than three patients are treated at that level for testing. As in the accelerated phase, if the higher dose is greater than 1 and the previous dose was less than 1, then the next dose assigned is 1. However if the higher dose is greater than one and the previous dose was 1, then the trial ends and no estimated MTD is assigned.

If there was *more than 1 DLT* in the previous three patients, then the dose is decreased by the MF-UD increment, three patients are added at this lower dose and three responses are simulated, unless more than three patients have already been treated at the lower dose level. If that is the case then the lower dose level is the estimated MTD and the trial ends. It may be that the doses had increased to 1, in which case the lower dose level may not match any previous dose levels assigned and more than three patients will have to be tested at these new dose levels. For example, if the dose increment in the MF-UD phase were 40%, or 1.4, and the doses escalated as 0.54, 0.75, 1, and there were more than one DLT at dose level 1, then the next lower dose level would be 0.71 and not 0.75.

If there was *1 DLT* in the last three patients, then the dose remains the same as the previous dose, three patients are added at this dose and three responses are simulated. If more than three patients have been treated at this dose and the number of DLTs was less than or equal to 2 then this dose is the estimated MTD and the trial ends.

If no stopping rule has been reached by the time the maximum number of patients have been added, then the trial ends and no estimated MTD is given.

The maximum number of patients in the AT trial is 62 patients.

2.4 EWOC Designs

Two designs were used: the standard EWOC [3] that follows a binary logistic model and calculates the recommended dose for the next patient based on the occurrence of DLT (toxicity grades 3 or 4) or no DLT (toxicity grades 0, 1, or 2) response from previous patients; and an extension of the standard EWOC that uses a proportional odds model assumption of dose-toxicity relationship (EWOC-PO) and allows information on grade 2 toxicities to be incorporated as well [13]. The sample size for each trial in both versions of EWOC was fixed at 30. The feasibility bound α was fixed at 0.25 for every trial.

We designed an MCMC sampler based on the Metropolis-Hastings algorithm ([17, 18] to implement EWOC and EWOC-PO designs. We also used WinBUGS [11] to estimate features of the posterior distribution of the MTD.

Vague priors for the model parameters were used:

$$\begin{aligned} \gamma &\sim \text{Unif}[0,1] \\ \rho_0 &\sim \text{Unif}[0, 0.33] \\ \rho_1 \mid \rho_0 &\sim \text{Unif}[\rho_0, 1]. \end{aligned} \tag{9}$$

Dose levels in the trial are selected in the interval [0, 1]. The adaptive design proceeds as follows. The first patient receives a dose $x_1 = 0$ and a safe response (grade 0 or 1). Denote by $\Pi_k(\gamma) = \Pi(\gamma \mid D_k)$ the marginal posterior cdf of the MTD, $k = 1,\ldots,n\text{-}1$. The $(k+1)$-st patient receives the dose $x_{k+1} = \Pi_k^{-1}(\alpha)$ so that the posterior probability of exceeding the

MTD is equal to the feasibility bound α.=0.25. This is the overdose protection property of EWOC, where at each stage of the design, we seek a dose to allocate to the next patient while controlling the posterior probability of exposing patients to toxic dose levels. The trial proceeds until 30 patients are enrolled to the trial. At the end of the trial, we estimate the MTD as $\hat{\gamma} = \Pi_n^{-1}(\alpha)$.

2.5 Efficiency and Safety Comparisons

Simulated trials following AT, EWOC and EWOC-PO designs were compared according to the average bias,

$$\text{bias}_{\text{ave}} = M^{-1}\left(\sum_{i=1}^{M}\hat{\gamma}_i - \gamma_{\text{true}}\right), \tag{10}$$

and the square root of the estimated mean square error,

$$\text{MSE} = M^{-1}\left(\sum_{i=1}^{M}\left(\hat{\gamma}_i - \gamma_{\text{true}}\right)^2\right), \tag{11}$$

where $\hat{\gamma}_i$ is the estimate of the MTD at the end of the *i*-th trial and M is the total number of simulated trials. In addition, models were compared with respect to the average proportion of patients exhibiting DLT, the proportion of trials for which the DLT rate exceeds θ + 0.05 = 0.38, the proportion of trials with estimated MTD within γ-0.15γ and γ+0.15γ (referred to as optimal dose), and the proportion of patients receiving optimal doses. It is important to highlight that, from the perspective of a patient participating in a dose finding trial, the best design is the one with the highest proportion of patients receiving optimal doses.

3. Results

Figures 2 and 3 show results of eight types of designs (EWOC, EWOC-PO, and six versions of AT) simulated under nine scenarios of true MTD and true ρ_1. Responses were simulated according to the proportional odds logistic model (model A, section 2.1). Responses for EWOC were generated in the same way as for EWOC-PO and AT, however the grade 2 toxicities were considered as the same level of response as the lower grade toxicities, leaving two response levels for DLT or no DLT. It is worth noting that, for these figures, responses were simulated with the same link function assumed by EWOC and EWOC-PO designs.

Regarding bias, all designs perform well when the true MTD=0.1. On the other hand, when the true MTD>0.1, AT designs display higher bias than EWOC and EWOC-PO. Similarly, √MSE is small for all designs when the true MTD=0.1, while AT designs display higher √MSE than EWOC and EWOC-PO when the true MTD>0.1. Regarding the average proportion of patients exhibiting DLT, EWOC and EWOC-PO are uniformly safe, being on target when the true MTD=0.1 while AT designs tend to be mostly safe, except when the true MTD=0.1 for three versions with starting dose=0.1; in these instances the average proportion of patients exhibiting DLT nears 50%. Similar safety concerns are shown regarding the proportion of trials with DLT rate > 0.38 for

three AT versions with starting dose=0.1 when the true MTD=0.1. Regarding the proportion of trials with estimated MTD around the optimal dose, except for AT (0.1, 1.96, 1.4) when the true MTD=0.1, EWOC and EWOC-PO show higher proportions than AT. It is worth noting that AT (0.1, 1.96, 1.4) had poor safety characteristics. Regarding the proportion of patients receiving optimal doses, EWOC and EWOC-PO uniformly display higher proportions than the AT designs.

Figure 3 shows the median of the number of patients per trial for EWOC, EWOC-PO and AT according to nine scenarios. While the number of patients per EWOC and EWOC-PO trial was set to 30, the number of patients per AT trial varied considerably. Considering all AT simulated trials, the overall minimum 5th percentile was 5 patients and the overall maximum 95th percentile was 52.

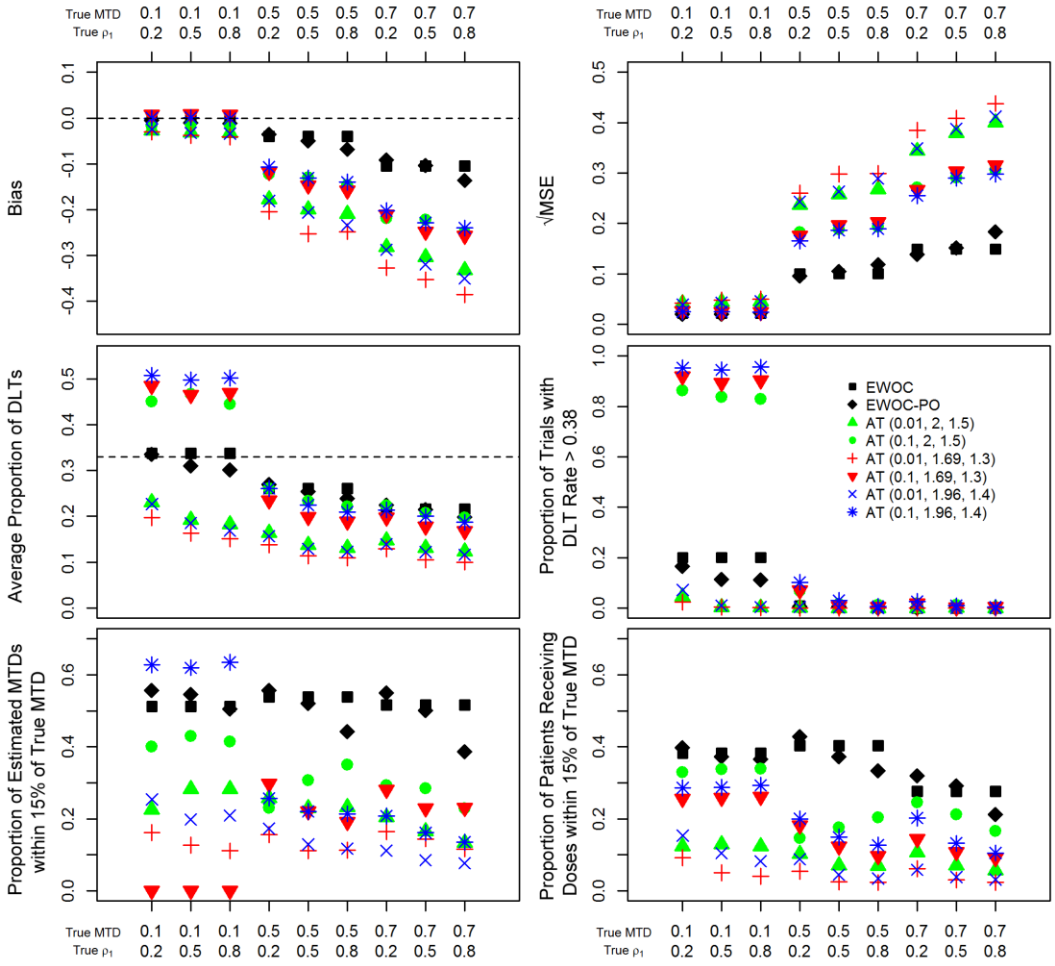

Figure 2: Summary statistics for trial efficiency and safety for EWOC, EWOC-PO and AT (starting dose, accelerated phase increase rate, MF-UD increase rate) under nine scenarios. Each symbol represents 1000 simulated trials.

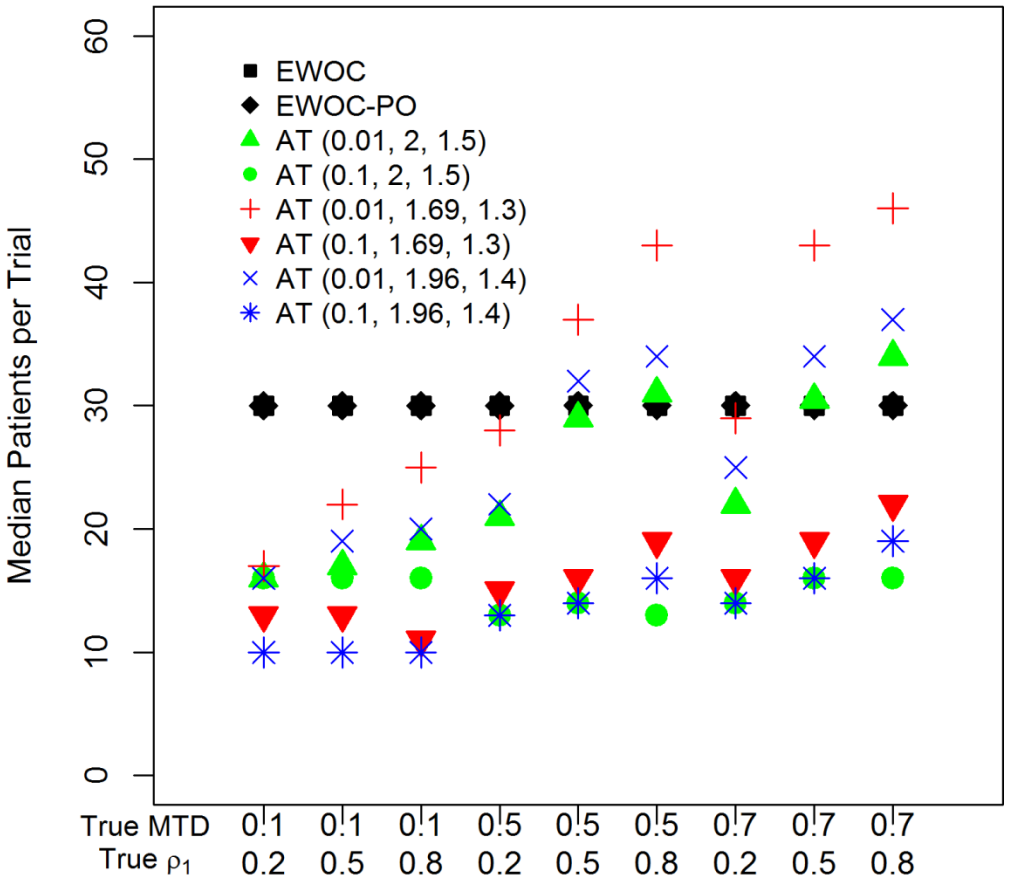

Figure 3: Median of the number of patients per trial for EWOC, EWOC-PO and AT (starting dose, accelerated phase increase rate, MF-UD increase rate) under nine scenarios. Each symbol represents 1000 simulated trials. The number of patients per EWOC and EWOC-PO trial is set to 30. The maximum number of patients per AT trial is 62 patients.

Figures 4 to 9 show results of AT, EWOC and EWOC-PO designs for different models: (A) proportional odds logistic, (B) normal with σ = 0.5, (C) normal distribution with σ = 2, (D) non-proportional odds with ρ_0 = 0.126, and (E) non-proportional odds with ρ_0 = 0.02. In these series, the simulation results of two AT versions: (0.01, 2, 1.5) and (0.1, 2, 1.5) were combined. Also combined were the results of the nine simulation scenarios for of true MTD and true ρ_1. Therefore, AT summarizes the results of 18,000 trials while EWOC and EWOC-PO condenses the results of 9,000 trials each. As in the previous simulations, EWOC and EWOC-PO designs assume a proportional odds model for the dose-toxicity relationship. Since responses simulated for models B, C, D, and E were generated by different distributions that these assumed by EWOC and EWOC-PO, they allow us to assess model robustness.

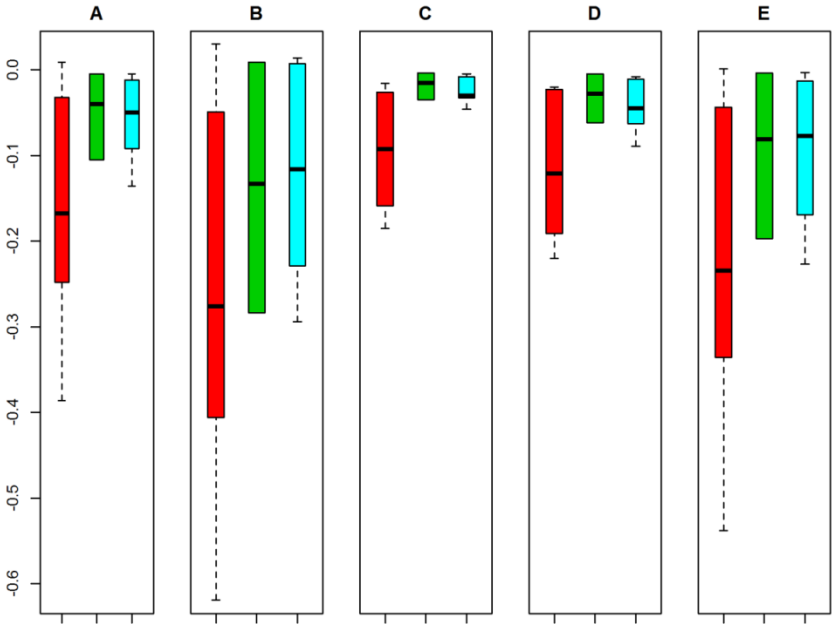

Figure 4: Average bias for AT (red), EWOC (green), and EWOC-PO (blue) designs for different models: (A) proportional odds logistic, (B) normal with σ = 0.5, (C) normal distribution with σ = 2, (D) non-proportional odds with ρ_0 = 0.126, and (E) non-proportional odds with ρ_0 = 0.02.

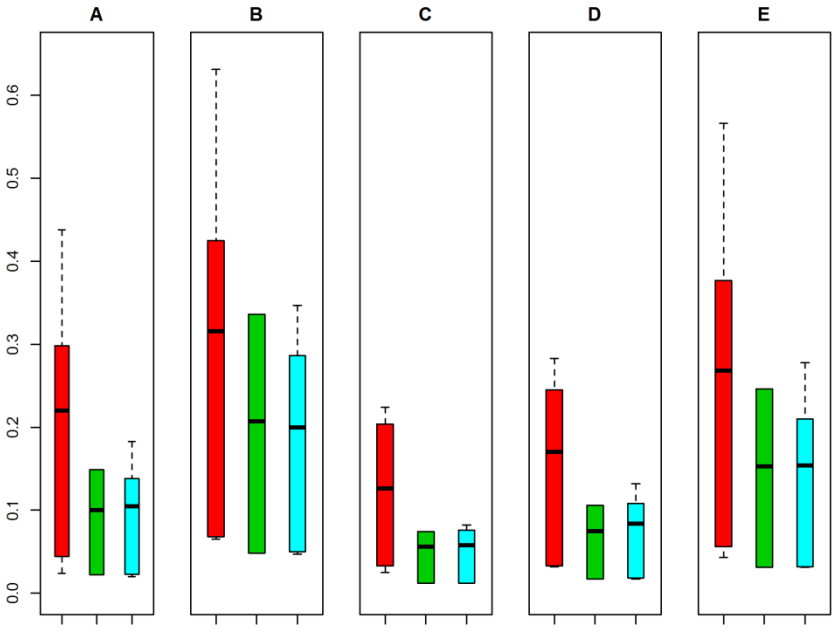

Figure 5: Average √MSE for AT (red), EWOC (green), and EWOC-PO (blue) designs for different models: (A) proportional odds logistic, (B) normal with σ = 0.5, (C) normal distribution with σ = 2, (D) non-proportional odds with ρ_0 = 0.126, and (E) non-proportional odds with ρ_0 = 0.02.

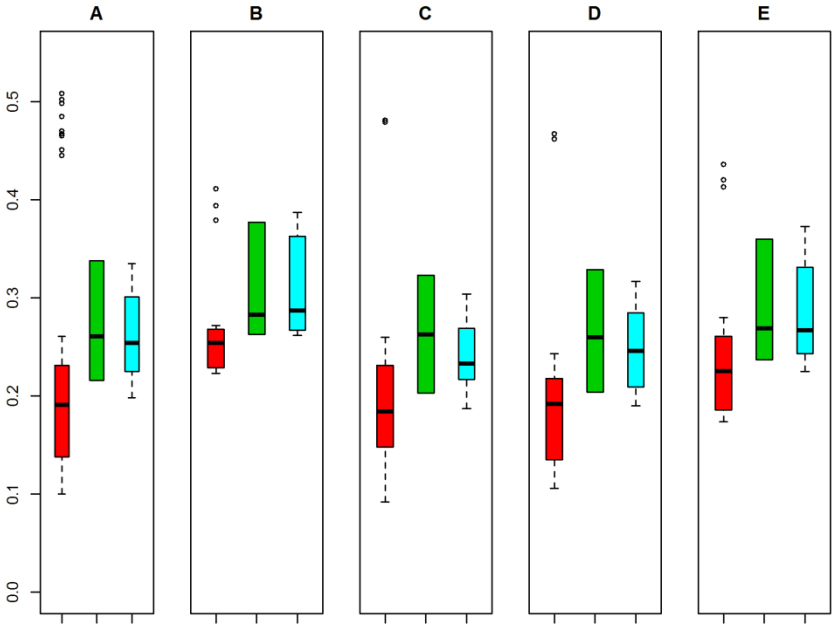

Figure 6: Average proportion of DLTs for AT (red), EWOC (green), and EWOC-PO (blue) designs for different models: (A) proportional odds logistic, (B) normal with σ = 0.5, (C) normal distribution with σ = 2, (D) non-proportional odds with ρ_0 = 0.126, and (E) non-proportional odds with ρ_0 = 0.02.

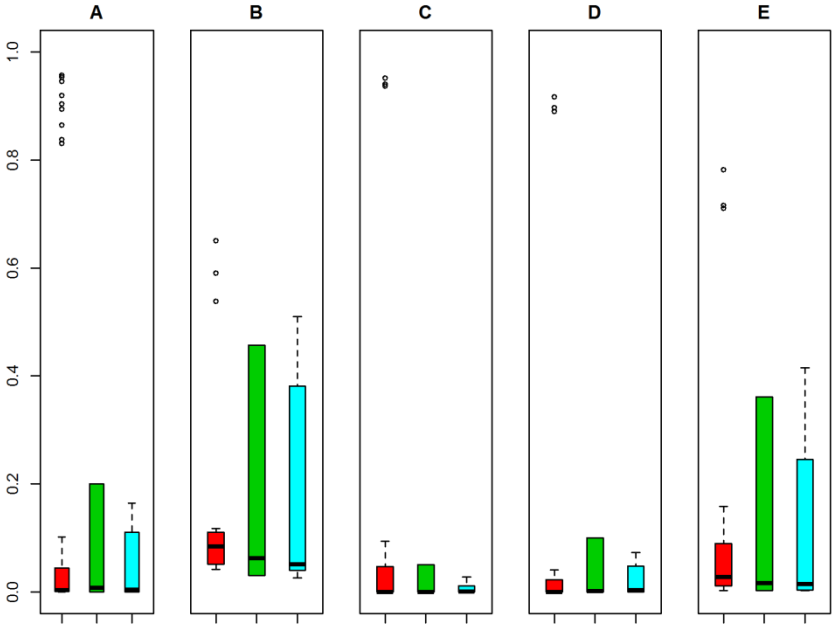

Figure 7: Proportion of trials with DLT rate above 0.38 (θ+0.05) for AT (red), EWOC (green), and EWOC-PO (blue) designs for different models: (A) proportional odds logistic, (B) normal with σ = 0.5, (C) normal distribution with σ = 2, (D) non-proportional odds with ρ_0 = 0.126, and (E) non-proportional odds with ρ_0 = 0.02.

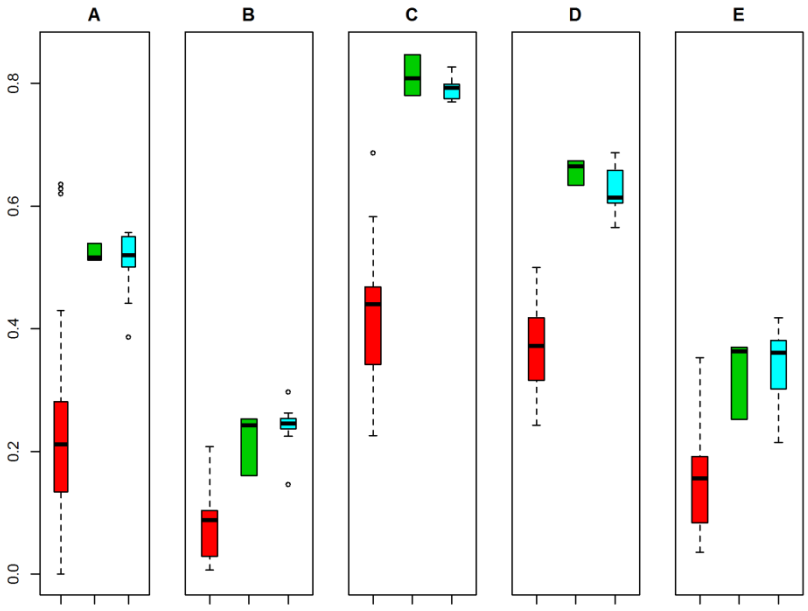

Figure 8: Proportion of estimated MTDs within 15% of True MTD for AT (red), EWOC (green), and EWOC-PO (blue) designs for different models: (A) proportional odds logistic, (B) normal with $\sigma = 0.5$, (C) normal distribution with $\sigma = 2$, (D) non-proportional odds with $\rho_0 = 0.126$, and (E) non-proportional odds with $\rho_0 = 0.02$.

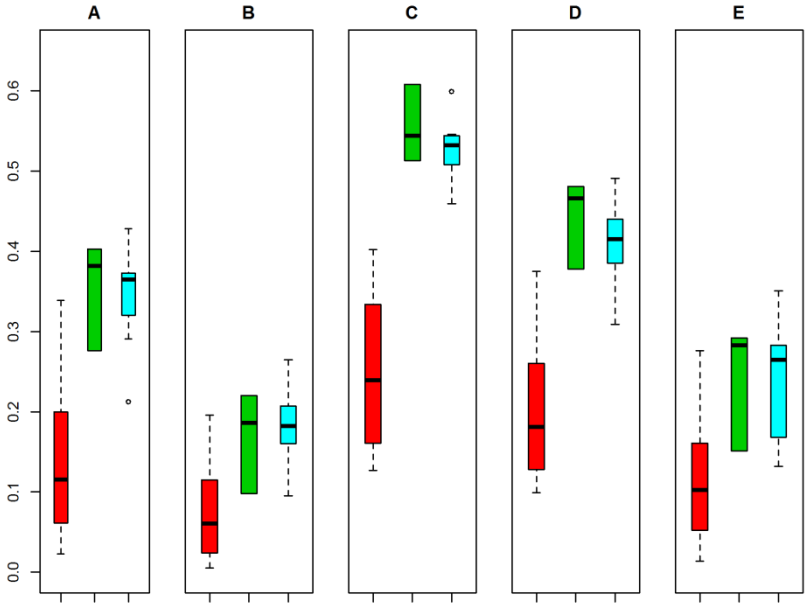

Figure 9: Proportion of patients receiving doses within 15% of True MTD for AT (red), EWOC (green), and EWOC-PO (blue) designs for different models: (A) proportional odds logistic, (B) normal with $\sigma = 0.5$, (C) normal distribution with $\sigma = 2$, (D) non-proportional odds with $\rho_0 = 0.126$, and (E) non-proportional odds with $\rho_0 = 0.02$.

4. Conclusion

Both EWOC and EWOC-POM designs tend to have smaller absolute bias than the AT designs, particularly when the true MTD is high. A higher proportion of trials with estimated MTD within an optimal range of the true MTD was obtained with the EWOC designs. The EWOC designs also tend to have lower MSE than the AT designs. The average proportion of DLTs for the AT designs were mostly between 0.1 and 0.25, although when the starting dose was equal to the true MTD the proportion was much higher. The EWOC designs targeted a DLT rate of 0.33 and the average proportion of DLTs in the simulations was usually a bit lower than 0.33, the exception being when the true MTD was low at 0.1. The proportion of patients receiving doses within 15% of true MTD (optimal doses) are uniformly higher with EWOC designs even when the model is misspecified. We conclude that EWOC designs are more efficient in estimating the MTD than AT designs and are safer under some scenarios.

Acknowledgements

Supported in part by the National Center for Research Resources, Grant UL1RR033176, and is now at the National Center for Advancing Translational Sciences, Grant UL1TR000124 (M.T and A.R), grant 5P01CA098912-05 (A.R) and P01 DK046763 (A.R). The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

References

1. Gatsonis C, Greenhouse JB. Bayesian methods for phase I clinical trials. *Stat Med* 1992; **11**: 1377-1389.
2. Horstmann E, McCabe MS, Grochow L, Yamamoto S, Rubinstein L, Budd T, Shoemaker D, Emanuel EJ, Grady C. Risks and benefits of phase 1 oncology trials, 1991 through 2002. *N Engl J Med* 2005; **352**: 895-904.
3. Babb J, Rogatko A, Zacks S. Cancer Phase I clinical Trials: efficient dose escalation with overdose control. *Stat Med* 1998; **17**: 1103-1120.
4. Zacks S, Rogatko A, Babb J. Optimal Bayesian-feasible dose escalation for cancer phase I trials. *Stat Prob Ltrs* 1998; **38**: 215-220.
5. Tighiouart M, Rogatko A. Dose Finding with Escalation with Overdose Control (EWOC) in Cancer Clinical Trials. *Statistical Science* 2010; **25**: 217-226.
6. Babb JS, Rogatko A. Patient specific dosing in a cancer phase I clinical trial. *Stat Med* 2001; **20**: 2079-2090.
7. Babb JS, Rogatko A. *Bayesian Methods for Cancer Phase I Clinical Trials.* Marcel Dekker: New York, 2004.
8. Tighiouart M, Cook-Wiens G, Rogatko A. Incorporating a Patient Dichotomous Characteristic in Cancer Phase I Clinical Trials Using Escalation with Overdose Control. *Journal of Probability and Statistics* 2012; **2012**: 10.
9. Rogatko A TM, Bresee C. Escalation with overdose control. User Guide Web-Based Portal. **2012**.
10. Rogatko A TM, Cook-Wiens G. Escalation with overdose control. User's guide. Version 3.1. **2012**.
11. Lunn DJ, Thomas A, Best N, Spiegelhalter D. WinBUGS - A Bayesian modelling framework: Concepts, structure, and extensibility. *Statistics and Computing* 2000; **10**: 325-337.

12. Tighiouart M, Rogatko A, Babb JS. Flexible Bayesian methods for cancer phase I clinical trials. Dose escalation with overdose control. *Stat Med* 2005; **24**: 2183-2196.
13. Tighiouart M, Cook-Wiens G, Rogatko A. Escalation with Overdose Control Using Ordinal Toxicity Grades for Cancer Phase I Clinical Trials. *Journal of Probability and Statistics* 2012; **2012**: 18.
14. Rogatko A, Schoeneck D, Jonas W, Tighiouart M, Khuri FR, Porter A. Translation of innovative designs into phase I trials. *Journal of Clinical Oncology* 2007; **25**: 4982-4986.
15. Storer BE. Design and analysis of phase I clinical trials. *Biometrics* 1989; **45**: 925-937.
16. Simon R, Freidlin B, Rubinstein L, Arbuck SG, Collins J, Christian MC. Accelerated titration designs for phase I clinical trials in oncology. *Journal of the National Cancer Institute* 1997; **89**: 1138-1147.
17. Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH, Teller E. Equations of State Calculations by Fast Computing Machines. *Journal of Chemical Physics* 1953; **21**: 1087:1092.
18. Hastings WK. Monte Carlo Sampling Methods Using Markov Chains and Their Applications. *Biometrika* 1970; **57**: 97:109.