The Treatment Versus Experimentation Dilemma in Dose Finding Studies

Azriel D., Mandel M. and Rinott Y.

Department of Statistics, The Hebrew University of Jerusalem

Abstract

Phase I clinical trials are conducted in order to find the maximum tolerated dose (MTD) of a given drug from a finite set of doses. For ethical reasons, these studies are usually sequential, treating patients or groups of patients with the optimal dose according to the current knowledge, with the hope that this will lead to using the true MTD from some time on. However, the first result proved here is that this goal is infeasible, and that such designs, and, more generally, designs that concentrate on one dose from some time on, cannot provide consistent estimators for the MTD unless very strong parametric assumptions hold. Allowing some non MTD treatment, we construct a randomized design that assigns the MTD with probability that approaches 1 as the size of the experiment goes to infinity and estimates the MTD consistently. We compare the suggested design with several methods by simulations, studying their performances in terms of correct estimation of the MTD and the proportion of individuals treated with the MTD.

Keywords: Isotonic regression, Maximum tolerated dose, Phase I trial, Stochastic approximation, Up-and-down design.

1. Introduction

Let x be a dose of a given drug and let y be a binary outcome, where $y = 1$ ($y = 0$) represents a toxic (non-toxic) response. Let $f(x) := P(y = 1|x)$ be the probability of a toxic response at dose x, where $f : \mathbb{R}^+ \to (0, 1)$ is an unknown strictly increasing function. Typically, the dose range D consists of only a few doses, $d_1 < d_2 < \ldots < d_K$, and one aims at finding the dose d_{j^*} having toxicity that is closest to a prescribed target toxicity level p^*, i.e.,

$j^* = \arg\min_j |f(d_j) - p^*|$ (assuming it is unique). The dose d_{j^*} is called the maximum tolerated dose (MTD).

MTD-finding studies, conducted as part of phase I clinical trials, are usually performed sequentially for reasons of efficiency and ethics, according to the following two, possibly conflicting, principles:

1. Treatment: ideally, treat each subject with the MTD d_{j^*}. Since it is unknown, subjects are often treated with an estimate of the MTD.

2. Experimentation: obtain a good estimate for the MTD at the end of the study.

The choice between the two is called the treatment versus experimentation dilemma in Bartroff and Lai [3]. The first point above reflects the ethical consideration of avoiding treatment with doses that may have a high toxicity rate or doses with low efficacy. In the words of Shu and O'Quigley [23]: "being optimal for anything other than the best estimated treatment for the next patient, or group of patients, to be included in the study is not acceptable". The second point is the core of MTD studies, but may require to treat subjects with non-optimal doses in order to find the MTD as fast as possible.

In the case of a continuous response that follows a simple linear regression model and a continuous dose space, Lai and Robbins [13] show that this dilemma can be resolved asymptotically by treating each subject with the estimated MTD based on a truncated version of the least squares estimators. The aim of the current paper is to examine if and how this dilemma can be resolved in the more common phase I framework of a finite dose space, under minimal assumptions on the dose-response curve.

Many sequential designs for MTD studies have been suggested in the literature. Some of them assume a functional parametric model for the dose-response curve, e.g., the continual reassessment method (CRM)(O'Quigley et al. [17], O'Quigley [16]), and escalation with overdose control (Babb et al. [1]). Others are non-parametric in nature, e.g., Gasparini and Eisele [7]; Whitehead et al. [24]; Leung and Wang [14]; Ivanova et al. [9]. All these methods focus on the 'Treatment' purpose, requiring that each subject be treated with the estimated MTD. We show in Section 2 that they fail to satisfy the 'Experimentation' purpose in the sense that such designs cannot yield consistent estimators for all response curves. Consequently, as we show in Section 2, they may also fail in the 'Treatment' goal.

In contrast, methods that do not require treatment at the current estimated MTD can yield consistent estimators. The classical non-parametric up and down methods, e.g., Dixon and Mood [6], and Derman [5], can provide consistent estimators (of the MTD), but use only a small part of the available data at each step to determine the next dose, and therefore have undesirable properties (O'Quigley and Zohar [18]). In particular, when considered asymptotically, these methods will continue to assign subjects to all doses, and will therefore fail to accomplish the 'Treatment' goal, and will be wasteful in the sense of assigning subjects to irrelevant doses.

Ivanova et al. [10] and Ivanova and Kim [11], suggest an improved up and down methods that estimate the dose-response curve by isotonic regression. These methods use all available data at each step and provide consistent estimator for the MTD for every increasing dose-response curve. Isotonic regression was considered earlier in the framework of stochastic approximation on a lattice by Mukerjee [15], who proves consistency of a method which eventually concentrates on two doses. All the methods described in this paragraph concentrate eventually on doses adjacent to the MTD, but have a non-vanishing probability of treating at a non-MTD dose.

In Section 2, we show that if treatment in a sequential experiment is according to the current estimator of the MTD, then this estimator cannot be strongly consistent, that is, it cannot coincide with the true MTD from some time on. Moreover, any design that concentrates on a single dose from some time on has a non negligible probability of concentrating on the wrong dose and cannot lead to strongly consistent estimation of the MTD. The implication is that a design having an optimal treatment from some time on is impossible, and the practice of assigning sequentially the current estimated MTD is statistically open to doubt. Though it is impossible to assign the MTD from a certain stage of the sequential experiment and on, it is possible to assign the MTD with probability approaching 1 as the experiment grows, or equivalently, have a vanishing probability of treating at a non-MTD dose. We introduce in Section 3 a design, which is based on Mukerjee [15], that accomplishes this. Properties of several designs for small and moderate sample sizes are studied via simulations in Section 4. Concluding remarks are given in Section 5 and the proofs appear in the Appendix.

2. An Impossibility Result.

We consider sequential designs and denote by x_n and y_n the dose assigned to the n'th subject and his response, respectively, and by $\mathcal{F}_{n-1} := \sigma\{(x_1, y_1), (x_2, y_2), \ldots, (x_{n-1}, y_{n-1})\}$ the available data prior to the decision on the n'th subject's dose. We assume for $n \geq 2$ that

$$x_n \in \mathcal{F}_{n-1} \quad , \quad y_n | \mathcal{F}_{n-1} \sim \text{Bernoulli}\left(f(x_n)\right). \tag{1}$$

The sequence $\{x_n\}_{n=1}^{\infty}$, is called a design; a sequence of estimators $\{\widehat{MTD}_n\}_{n=1}^{\infty}$ is said to be strongly consistent with respect to a given design if $\widehat{MTD}_n \to d_{j^*}$ almost surely for all increasing functions f or equivalently, by discreteness $P(\exists N \ s.t. \ \forall n \geq N : \ \widehat{MTD}_n = d_{j*}) = 1$.

Under the above general framework we obtain the following impossibility result:

Theorem 1. *Assuming (1), there exists no design that satisfies for all increasing functions f*

$$P(\exists N \ s.t. \ \forall n \geq N : \ x_n = d_{j*}) = 1, \tag{2}$$

or equivalently that $P(x_n \neq d_{j^*} \ i.o. \ ) = 0$.

The crux of the above result is that a design that concentrates eventually on one dose, say d_j, can yield a consistent estimator for $f(d_j)$, but cannot estimate consistently $f(d_i)$ for $i \neq j$, and therefore may miss the MTD, and eventually treat patients with a non-optimal dose, so that asymptotically both treatment and estimation fail.

Results like Theorem 1 and its consequences holds also for parametric models where the dimension of the parameter space is two or larger. A design that concentrates eventually on one dose can yield a consistent estimator for the probability of a toxic response at that dose only, but knowing the response curve in a single dose cannot yield a consistent estimator for the unknown parameters, unless the dimension is one. Shen and O'Quigley [21] and Shu and O'Quigley [23], make a similar argument in favor of a one parameter model as a working model for the CRM.

Perhaps the most striking implication of the theorem is the following corollary:

Corollary 1 (Treatment versus experimentation dilemma). *Let* $\{\widehat{MTD}_n\}_{n=1}^{\infty}$ *be any sequence of estimators of the MTD. If for all* n, $x_{n+1} = \widehat{MTD}_n$ *then* $\{\widehat{MTD}_n\}_{n=1}^{\infty}$ *is not strongly consistent.*

Corollary 1 has direct implications for phase I studies because many designs, including the CRM and the non-parametric methods mentioned in Section 1, assign the estimated MTD at each stage. Such designs cannot yield consistent estimators for the MTD unless severe parametric assumptions on f are imposed. Hence, in our framework, the aforementioned ethical requirement of Shu and O'Quigley seems arguable. In fact, it is easy to see from the above results that designs that concentrates eventually on one dose cannot yield weakly consistent estimation of the MTD, that is, estimators that converge to the MTD in probability. In Section 3 we show that designs leading to strongly consistent estimation of the MTD, and treatment at the MTD with probability approaching 1, do exist.

3. An asymptotically optimal design

In this section, we construct a design that satisfies $P(x_n = d_{j^*}) \xrightarrow[n\to\infty]{} 1$ and a strongly consistent sequence of estimators of the MTD. Such a design 'almost' resolves the treatment versus experimentation dilemma.

We shall use the isotonic regression estimator of f, which maximizes the likelihood $\prod_{i=1}^{n} f(x_i)^{y_i}[1-f(x_i)]^{1-y_i}$ under the restriction that f is nondecreasing (Barlow et al. [2], p. 38). Specifically, for any $d_r, d_s \in D$ such that $r \le s$, define

$$\mathcal{N}_n(d_r, d_s) = \sum_{i=1}^{n} I(x_i \in [d_r, d_s]),$$

$$\bar{y}_n(d_r, d_s) = \begin{cases} \frac{1}{\mathcal{N}_n(d_r,d_s)} \sum_{i=1}^{n} y_i I(x_i \in [d_r, d_s]) & \mathcal{N}_n(d_r, d_s) > 0 \\ 0 & \mathcal{N}_n(d_r, d_s) = 0 \end{cases}.$$

The estimator for f at stage n is

$$\hat{f}_n(d_j) = \max_{r \le j} \min_{s \ge j} \bar{y}_n(d_r, d_s) \quad j = 1, \ldots, K.$$

The corresponding estimator of the MTD is defined as follows. Let j be the maximal element in $\{1, \ldots, K-1\}$ that satisfies $\hat{f}_n(d_j) \le p^*$ (if no j

satisfies this, set $j = 1$); our estimator of the MTD is

$$\widehat{MTD}_n = \begin{cases} d_j & p^* \leq \frac{\hat{f}_n(d_j)+\hat{f}_n(d_{j+1})}{2} \\ d_{j+1} & p^* > \frac{\hat{f}_n(d_j)+\hat{f}_n(d_{j+1})}{2} \end{cases}. \tag{3}$$

We suggest the following randomized allocation design (RAD), which is a randomized version of Mukerjee [15] stochastic approximation scheme:

1. If $p^* < \hat{f}_n(d_1)$, set $x_{n+1} = d_1$.
2. If $p^* > \hat{f}_n(d_K)$, set $x_{n+1} = d_K$.
3. If $B_n(d_j) := \{\hat{f}_n(d_j) \leq p^* \leq \hat{f}_n(d_{j+1})\}$ occurs, select x_{n+1} to be d_j or d_{j+1} according to the following rule:

$$\begin{aligned} &\text{if } p^* \leq (>) \frac{\hat{f}_n(d_j) + \hat{f}_n(d_{j+1})}{2} \quad \text{then} \\ &x_{n+1} = \begin{cases} d_j & \text{with probability } 1 - \frac{1}{k} \quad (\frac{1}{k}) \\ d_{j+1} & \text{with probability } \frac{1}{k} \quad (1 - \frac{1}{k}) \end{cases}, \end{aligned} \tag{4}$$

where $k := k(n, j) = \sum_{i=1}^{n} I\{B_i(d_j)\}$.

Note that if $p^* \leq (>) \frac{\hat{f}_n(d_j)+\hat{f}_n(d_{j+1})}{2}$ for some j, then d_j (d_{j+1}) is the estimated MTD and $x_{n+1} = d_j$ (d_{j+1}) with large probability. The design is constructed in such a way that if $B_n(d_j)$ occurs infinitely often then the probability of choosing the estimated MTD tends to one (if $f(d_j) \leq p^* \leq f(d_{j+1})$), and both $\{x_n = d_j\}$ and $\{x_{n+1} = d_{j+1}\}$ occur infinitely often. Thus, in the proposed design, non MTD treatment occurs asymptotically only rarely (though it occurs infinitely often) and the probability of treatment at the true MTD approaches one. The properties of this design are summarized in the following theorem.

Theorem 2. *Assume that $f(d_{j'}) < p^* < f(d_{j'+1})$ for some $j' \in \{1, \ldots, K-1\}$. The RAD (4) satisfies*

I. $P(\exists N \, s.t. \, \forall n \geq N : x_n \in \{d_{j'}, d_{j'+1}\}) = 1$.

II. The sequence $\widehat{MTD}_n$ (3) is strongly consistent.

III. $P(x_n = d_{j^*}) \xrightarrow[n\to\infty]{} 1$.

The proof is given in the appendix. We note that the RAD guarantees strongly consistent estimation of the MTD, hence satisfies the Experimentation criterion, and from some stage on it treats with either the estimated MTD or an adjacent dose (with an increasing probability to treat with the former), and hence almost satisfies the Treatment purpose of MTD studies.

Remark 1. *By a similar argument as in the proof of Theorem 2, The RAD design (4) satisfies*

i. *If* $p^* = f(d_{j*})$ *then* $P(\exists N \, s.t. \, \forall n \geq N : \, x_n \in \{d_{j*-1}, d_{j*}, d_{j*+1}\}) = 1.$

ii. *If* $p^* < f(d_1)$ *then* $P(\exists N \; s.t. \; \forall n \geq N : \; x_n = d_1) = 1.$

iii. *If* $p^* > f(d_K)$ *then* $P(\exists N \; s.t. \; \forall n \geq N : \; x_n = d_K) = 1.$

In all three cases, $\widehat{MTD}_n$ *is strongly consistent and* $P(x_n = d_{j^*}) \xrightarrow[n\to\infty]{} 1$. *We assume that* $\arg\min_j |f(d_j) - p^*|$ *is unique (otherwise, the MTD is not well-defined). In the case that* p^* *is exactly in the middle of* $[f(d_{j'}), f(d_{j'+1})]$ *then for large enough* n *(with probability 1),* x_n *and* $\widehat{MTD}_n$ *will oscillate between* $d_{j'}$ *and* $d_{j'+1}$.

Remark 2. *Theorem 2 guarantees large probability of optimal treatment for large* n*, and treatment at one of the two closest doses to the MTD with probability 1 from some time on (Part I). However, for practical purposes (small* n*), we found that the algorithm performs better when the rate of choosing the estimated MTD is reduced. For example, replacing* k *in (4) with* $a \cdot k + 2$*, where* a *is a (small) constant, yields better small sample performance. Thus, the choice between* d_j *and* d_{j+1} *is done with probability* $\frac{1}{a \cdot k+2}$ *which is* $\approx 1/2$ *for small* a*. This modification does not change the asymptotic behavior of the estimator given in Theorem 2, while improving the learning rate of the response curve in early stages of the experiment.*

4. A Simulation study

In this section, we compare the small sample performances of several designs under the following dose-response curves:

A. $(f(d_1), \ldots, f(d_6)) = (0.1, 0.13, 0.15, 0.17, 0.25, 0.3)$;

B. $(f(d_1), \ldots, f(d_6)) = (0.07, 0.11, 0.23, 0.43, 0.84, 0.98)$.

Scenario B is taken from Table 4 of O'Quigley et al. [17]. Scenario A represents a dose-response curve with a much smaller slope. For scenario A, we considered two target probabilities: $p^* = 0.2$ and $p^* = 0.22$ (MTD=d_3 and MTD=d_4, respectively), and for scenario B, we considered $p^* = 0.2$ and $p^* = 0.3$ (MTD=d_3 in both cases). The two dose-response curves we consider do not satisfy the working model for the CRM. However, scenario B with $m^* = 0.2$ satisfies the conditions given by Cheung and Chappell [4], conjectured (and checked by simulations, but not proved) to be sufficient for consistency of the CRM, and therefore one may expect the CRM to perform well in this case.

We compare the randomized allocation design (RAD) (4) with three consistent methods: the up and down design of Ivanova et al. [10] (IVA), the design of Ivanova and Kim [11], with $\Delta = 0.01$ (IVA1) as recommended in Ivanova et al. [8], the design of Mukerjee (MUK), and the widely used one-parameter CRM design (Shen and O'Quigley [21]) with maximum likelihood as the estimation approach and the constants suggested by O'Quigley et al. [17]. Three different versions of the RAD are studied according to different choices of a (see Remark 2): $a = (10-2)/30$, $(10-2)/50$, and $(10-2)/100$, denoted RAD1, RAD2, and RAD3, respectively; these values correspond to 0.9 chance of assigning the estimated MTD for $k = 30$, 50, and 100.

All methods started with $x_1 = d_1$ and $x_{n+1} = d_{(n+1)\wedge K}$ until the first toxicity response was observed, and continued according to the specific rules described above. We conducted 10,000 replications from each scenario and ran the experiment for a maximum of 500 individuals. For better comparison, we coupled all designs in a manner that is akin to the notion of antithetic variables in the following way: when individual n in replication r had the outcome $y = 1$ in one design, then the same outcome was obtained in all designs that assigned the same or a higher dose to individual n in replication r. Similarly for the outcome $y = 0$.

The performances of the designs were measured according to the different purposes of MTD studies, that is, the probability of finding the true MTD at stage n and the proportion of subjects treated with the true MTD. For CRM, the MTD was estimated according to the maximum likelihood approach; for the other methods, (3) was used. The results for small sample sizes are given in Tables A.1 and A.2; Figures A.1 and A.2 present the results for all sample sizes.

Overall, the performances under scenario B are much better than under scenario A. This is expected, as the response-curve of the latter is much

flatter. Also of note is the small probability of correct estimation for $n \leq 50$, which are the typical sample sizes of MTD studies. This is a known problem in such studies – the probability of selecting the true MTD is more often than not smaller than 1/2.

When comparing the designs on the basis of the probability of finding the true MTD, we found that IVA, IVA1 and MUK outperformed the others for large n. The CRM preformed well only in scenario B with $p^* = 0.2$ as suggested by the consistency considerations mentioned above. The randomized designs and especially RAD3 estimated the MTD relatively well for small n. This is because the allocation probability is close to one half in these stages. For small n, the performances of all designs except the CRM are comparable, though, it seems that IVA, IVA1 and MUK are over all the best for the goal of MTD estimation. IVA and IVA1 do not get "stuck" in one dose but rather oscillate in some fashion around the MTD. In fact, these methods and MUK have a similar property as in Theorem 2 I: for large enough n, only doses near the MTD are assigned. Our simulation study shows that these three methods preform in a similar way in estimating the MTD for large n and they are generally better than the RAD's. However, these methods treat at non-MTD doses more often than the RAD's for large n.

Looking at the proportion of subjects treated with the true MTD, we see that the RADs perform the best for large n for all scenarios except scenario B with $p^* = 0.2$ in which CRM is better. Generally, RAD3, which has the smallest a, is the best among the RADs, though, in scenario B with $p^* = 0.3$, RAD1 and RAD2 perform somewhat better. For small n, it seems that the best design depends on the specific dose-response curve and on p^*.

The CRM performs very well in scenario B, but performs poorly under scenario A. This demonstrates the potential benefit and risk of using parametric models. No single method among the nonparametric approaches outperforms the others. Further study is needed in order to understand the operating characteristics of the different designs under different scenarios.

5. Conclusions

The asymptotic point of view of this paper suggests that the estimation and treatment goals cannot be achieved simultaneously. This finding necessitates a second thought about the way doses should be assigned in phase I studies. In particular, this result may imply that one should try to learn the

responses in the two closest doses to the desired level, rather than the closest one.

The proposed randomized allocation design in which the MTD is assigned to the n'th subject with increasing probability, treats subjects 'almost' in an optimal way for large n. However, for small and moderate sample sizes, this design does not estimate the MTD as well as some of the other designs, as it aims mainly at the treatment part of the dilemma. This implies that, even tough the MTD can be consistently estimated, a price is being paid in the experimentation part, and it seems that further investigation on how to balance estimation and treatment is needed.

Acknowledgment We thank an anonymous referee for very helpful comments.

Appendix A. Proofs

Proof of Theorem 1. We will exhibit two response probabilities f and f' such that if (2) holds for $\{x_n\}_{n=1}^{\infty}$ under f then it does not hold under f'. Let f satisfy $f(d_{j^*}) < p^*$, and set, for example, $j^* = 2$. Then it is easy to construct f' that differs from f only at one value in D, say d_3, and such that d_3 is the MTD associated with f'. Consider two probability measures, P and P', generated by f and f', defined on the measurable space $(\Omega, \mathcal{F})$, where Ω is the sample space of the experiment and $\mathcal{F}$ is the sigma-field generated by the union of all $\mathcal{F}_n$.

Let $A_n = \{x_k = d_2, k \geq n\}$ be the event that from the n'th subject on we always choose d_2. If (2) holds for $\{x_n\}_{n=1}^{\infty}$ under f, that is, under P, there is an index n_0 such that $P(A_{n_0}) > 0$. Consequently, there exists a vector $(x_1^0, y_1^0, \ldots, x_{n_0-1}^0, y_{n_0-1}^0)$ such that

$$\tilde{A} := \{\omega \in \Omega : (x_1, y_1, \ldots, x_{n_0-1}, y_{n_0-1})(\omega) = (x_1^0, y_1^0, \ldots, x_{n_0-1}^0, y_{n_0-1}^0)\} \cap A_{n_0}$$

satisfies $P(\tilde{A}) > 0$. We also have $P'(\tilde{A}) > 0$, since the above $n_0 - 1$ outcomes that lead to A_{n_0} have positive probability also under P' by finiteness, and conditioned on them, the probability of having $x_k = d_2$ for all $k \geq n_0$ is the same under both P and P' by the relation between f and f'. Therefore, $P'(A_{n_0}) \geq P'(\tilde{A}) > 0$ so that on a set of P' positive measure, A_{n_0}, we sample d_2 from n_0 on, while the MTD under f' is d_3. □

We define the randomized x_{n+1} explicitly by

$$\text{if } p^* \le \frac{\hat{f}_n(d_j)+\hat{f}_n(d_{j+1})}{2} \text{ then } x_{n+1} = d_j I(U_{n+1} \le \frac{k-1}{k}) + d_{j+1} I(U_{n+1} > \frac{k-1}{k});$$
$$\text{if } p^* > \frac{\hat{f}_n(d_j)+\hat{f}_n(d_{j+1})}{2} \text{ then } x_{n+1} = d_j I(U_{n+1} \le \frac{1}{k}) + d_{j+1} I(U_{n+1} > \frac{1}{k}),$$

where $\{U_i\}_{i=1}^{\infty}$ are i.i.d Uniform$[0,1]$ random variables independent of the y's, j is such that $I\{B_n(d_j)\} = 1$, and $k = \sum_{i=1}^{n} I\{B_i(d_j)\}$.

For the proof of Theorem 2, we first need some lemmas.

Lemma 1. *For almost all ω in the sample space, if $B_n(d_j)$ occurs infinitely often for some $j \in \{1, ..., K-1\}$ then both $x_n = d_j$ and $x_n = d_j + 1$ occur infinitely often.*

Proof. Let $\{n_k\}_{k=1}^{\infty}$ be the (random) subsequence in which $I\{B_n(d_j)\} = 1$. The design implies $\{x_{n_k+1} = d_j\} \supseteq \{U_{n_k+1} < \frac{1}{k}\}$, and $\{U_{n_k+1}\}_{k=1}^{\infty}$ are independent and identically distributed. (since $n_k = \min\{n : I\{B_n(d_j)\} = 1, n > n_{k-1}\}$ is a stopping time for all k; see Lemma 2 below). As, $\sum_k P(U_{n_k+1} < \frac{1}{k}) = \sum_k \frac{1}{k} = \infty$, the second Borel-Cantelli lemma shows that $\{U_{n_k+1} < \frac{1}{k}\}$, and hence $\{x_{n_k+1} = d_j\}$ occur infinitely often. Similar arguments show that $\{x_n = d_{j+1}\}$ occur infinitely often.

Lemma 2. *Let $\{U_n\}_{n=1}^{\infty}$ be a sequence of independent and identically distributed random variables with distribution F, and let $\{\tau_n\}_{n=1}^{\infty}$ be an increasing sequence of finite stopping times with respect to $\mathcal{G}_n \supseteq \sigma(U_1, \ldots, U_n)$. Assume that $U_{n+k}|\mathcal{G}_n \sim F$ for all $k \ge 1$ and n. Then $\{U_{\tau_n+1}\}_{n=1}^{\infty}$ is also a sequence of independent and identically distributed random variables with distribution F.*

Lemma 2 is quite standard, we include a proof for completeness.

Proof. For any subset of indices $n_1 < n_2 < \ldots < n_l$ and any measurable sets $A_1, \ldots, A_l$,

$$P(\bigcap_{1\le k\le l} \{U_{\tau_{n_k}+1} \in A_k\}) = E\{\prod_{1\le k\le l} I(U_{\tau_{n_k}+1} \in A_k)\} = E[E\{\prod_{1\le k\le l} I(U_{\tau_{n_k}+1} \in A_k)|\mathcal{G}_{\tau_{n_l}}\}]$$
$$= E[\prod_{1\le k\le l-1} I(U_{\tau_{n_k}+1} \in A_k) E\{I(U_{\tau_{n_l}+1} \in A_l)|\mathcal{G}_{\tau_{n_l}}\}]$$
$$= E\{\prod_{1\le k\le l-1} I(U_{\tau_{n_k}+1} \in A_k)\} P(U_1 \in A_l) = \ldots = \prod_{1\le k\le l} P(U_1 \in A_k);$$

hence the random variables $\{U_{\tau_n+1}\}_{n=1}^{\infty}$ are independent and identically distributed. □

Lemma 3. *The RAD design (4) satisfies*

I. $\bar{y}_n(d_j) \to f(d_j)$ *almost surely on* $\{\mathcal{N}_n(d_j) \to \infty\}$, $j = 1, \ldots, K$, *where* $\bar{y}_n(d_j) := \bar{y}_n(d_j, d_j)$ *and* $\mathcal{N}_n(d_j) := \mathcal{N}_n(d_j, d_j)$.

II. $\hat{f}_n(d_j) \to f(d_j)$ *almost surely on* $\{\mathcal{N}_n(d_j) \to \infty\}$, $j = 1, \ldots, K$.

Proof. *I.* $\sum_{i=1}^{n} I(x_i = d_j)(y_i - f(d_j))$ is a square integrable martingale with respect to the filtration $\mathcal{F}_n$, with quadratic variation

$$\sum_{i=1}^{n} [I(x_i = d_j)]^2 \cdot f(d_j) \cdot [1 - f(d_j)] = f(d_j) \cdot [1 - f(d_j)] \mathcal{N}_n(d_j),$$

since $\mathcal{N}_n(d_j) = \sum_{i=1}^{n} I(x_i = d_j)$. Therefore, by the strong law of large numbers for square integrable martingales, (Shiryaev [22] p. 519, Theorem 4)

$$\frac{1}{\mathcal{N}_n(d_j)} \sum_{i=1}^{n} I(x_i = d_j)\{y_i - f(d_j)\} \to 0 \quad a.s.\ on \quad \{\mathcal{N}_n(d_j) \to \infty\}.$$

Since $\bar{y}_n(d_j) = f(d_j) + \frac{1}{\mathcal{N}_n(d_j)} \sum_{i=1}^{n} I(x_i = d_j)\{y_i - f(d_j)\}$, the first part of the lemma follows.

II. We first consider the case $j \in \{2, \ldots, K-1\}$. The RAD design (4) satisfies, due to Lemma 1, that if $\mathcal{N}_n(d_j) \to \infty$ then either $\mathcal{N}_n(d_{j+1}) \to \infty$ or $\mathcal{N}_n(d_{j-1}) \to \infty$ (or both); without loss of generality, we assume that $\mathcal{N}_n(d_{j+1}) \to \infty$, and we condition on the event $\{\mathcal{N}_n(d_j) \to \infty\} \cap \{\mathcal{N}_n(d_{j+1}) \to \infty\}$.

We first show that for any $r \leq j$ and $s > j$ and almost all $\omega \in \Omega$ (where the measurable space is $(\Omega, \mathcal{F})$ defined in the proof of Theorem 1) there exists $N(\omega)$ such that

$$\bar{y}_n(d_r, d_s) > \bar{y}_n(d_r, d_j) \tag{A.1}$$

holds for all $n \geq N(\omega)$, when the random variables above are evaluated at ω. To see that, write

$$\bar{y}_n(d_r, d_j) = \sum_{k=r}^{j} \frac{\mathcal{N}_n(d_k)}{\mathcal{N}_n(d_r, d_j)} \bar{y}_n(d_k),$$

and recall that $\mathcal{N}_n(d_r, d_j) \geq \mathcal{N}_n(d_j) \to \infty$. If $\lim_n \mathcal{N}_n(d_k) < \infty$ then the corresponding term in the sum above has zero limit; if $\lim_n \mathcal{N}_n(d_k) = \infty$ then, by part *I* of the lemma, $\lim_n \bar{y}_n(d_k) = f(d_k)$ almost surely, and in particular, $\lim_n \bar{y}_n(d_j) = f(d_j)$ almost surely. Thus, $\limsup_n \bar{y}_n(d_r, d_j) \leq f(d_j)$ almost surely. A similar argument shows that $\liminf_n \bar{y}_n(d_{j+1}, d_s) \geq f(d_{j+1})$ almost surely, and therefore, for large enough n, with probability 1, $\bar{y}_n(d_{j+1}, d_s) > \bar{y}_n(d_r, d_j)$. The inequality (A.1) follows, as

$$\bar{y}_n(d_r, d_s) = \frac{\mathcal{N}_n(d_r, d_j)}{\mathcal{N}_n(d_r, d_s)} \bar{y}_n(d_r, d_j) + \frac{\mathcal{N}_n(d_{j+1}, d_s)}{\mathcal{N}_n(d_r, d_s)} \bar{y}_n(d_{j+1}, d_s) >$$
$$\frac{\mathcal{N}_n(d_r, d_j)}{\mathcal{N}_n(d_r, d_s)} \bar{y}_n(d_r, d_j) + \frac{\mathcal{N}_n(d_{j+1}, d_s)}{\mathcal{N}_n(d_r, d_s)} \bar{y}_n(d_r, d_j) = \bar{y}_n(d_r, d_j).$$

In view of (A.1) and the definition of $\hat{f}_n(d_j) = \max_{r \leq j} \min_{s \geq j} \bar{y}_n(d_r, d_s)$, we have for large enough n, $\hat{f}_n(d_j) = \max_{r \leq j} \bar{y}_n(d_r, d_j)$ with probability 1. Now, since $\limsup_n \bar{y}_n(d_r, d_j) \leq f(d_j)$ almost surely for $r < j$, and $\lim_n \bar{y}_n(d_j, d_j) = f(d_j)$ almost surely, the second part of the lemma follows for $j \in \{2, \ldots, K-1\}$.

If $j = 1$ (the case $j = K$ is similar), then if $\mathcal{N}_n(d_k) \to \infty$ for some $k \geq 2$ the proof is the same (however, Theorem 2 shows that for $k > 2$ this is impossible). If d_1 is the only dose that is assigned infinitely often then for every j

$$\bar{y}_n(d_1, d_j) = \sum_{k=1}^{j} \frac{\mathcal{N}_n(d_k)}{\mathcal{N}_n(d_r, d_j)} \bar{y}_n(d_k) \to f(d_1) \quad a.s.$$

because for $k = 1$ the limit of the k'th term above is $f(d_1)$, and for $k > 1$, this limit is zero. Since $\hat{f}_n(d_1) = \min_{j \geq 1} \bar{y}_n(d_1, d_j)$ the lemma follows. □

Proof of Theorem 2. *I.* Denote by $B_n(d_K)$ the event that $\hat{f}_n(d_K) < p^*$. We show by contradiction that the only j satisfying that $\{B_n(d_j)\}$ occurs infinitely often is $j = j'$. The event that $\hat{f}_n(d_1) > p^*$ is treated similarly.

Assume that j satisfies $j > j'$ and $\{B_n(d_j)\}$ occurs infinitely often. By Lemmas 1 and 3, there exists N such that $\hat{f}_n(d_j) > p^*$ for all $n \geq N$, with probability 1. Then, $I\{B_n(d_j)\} = 0$ for $n \geq N$ in contradiction to $\{B_n(d_j)\}$ occurring infinitely often. A similar argument shows that $\{B_n(d_j)\}$ does not occur infinitely often for $j < j'$; hence, for large enough n, with probability 1, $I\{B_n(d_{j'})\} = 1$.

II. The first part of the theorem ensures that the design will concentrate

eventually on the two closest doses to $f^{-1}(p^*)$ and both of these doses will be chosen infinitely often due to Lemma 1. By Lemma 3, $\hat{f}_n(d_{j'})$ and $\hat{f}_n(d_{j'+1})$ are strongly consistent. This implies that for large enough n, with probability 1, $\widehat{MTD}_n = d_{j^*}$.

III. The MTD, d_{j^*}, is either $d_{j'}$ or $d_{j'+1}$. Assume that $d_{j^*} = d_{j'}$ (the argument is symmetric). For any $n \geq k_0 \geq 1$, define the set

$$A_{n,k_0} = B_n(d_{j'}) \cap \{\widehat{MTD}_n = d_{j'}\} \cap \{k(n, j') \geq k_0\}.$$

where $k(n, j)$ is defined after (4). For any fixed k_0 and large enough n (so that A_{n,k_0} is not null),

$$P(x_{n+1} = d_{j'}) \geq P(x_{n+1} = d_{j'}|A_{n,k_0})P(A_{n,k_0}) \geq \frac{k_0 - 1}{k_0} P(A_{n,k_0}),$$

since $P(x_{n+1} = d_{j'}) = P(U_{n+1} \leq \frac{k-1}{k}) \geq P(U_{n+1} \leq \frac{k_0-1}{k_0})$ on A_{n,k_0}. For any fixed k_0, $P(A_{n,k_0}) \to 1$ by Parts *I* and *II*, so that $\liminf_n P(x_n = d_{j'}) \geq \frac{k_0-1}{k_0}$. As this is true for all k_0, the claim follows. □

References

[1] 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.

[2] Barlow R.E., Bartholomew D.J., Bremner J.M., Brunk H.D. (1972). *Statistical inference under order restrictions*, London ; New York : Wiley.

[3] Bartroff J., Lai T.L. (2010). Approximate dynamic programming and its applications to the design of phase I cancer trials, *Statistical Science*, **25**, 245 – 257.

[4] Cheung Y.K., Chappell R. (2002). A simple technique to evaluate model sensitivity in the continual reassessment method, *Biometrics*, **58**, 671 – 674.

[5] Derman C., (1957). Non-parametric up-and-down experimentation, *Annals of Mathematical Statistics*, **28**, 795 – 798.

[6] Dixon W.J., Mood A.M. (1948). A method for obtaining and analyzing sensitivity data, *Journal of the American Statistical Association*, **43**, 109 – 126.

[7] Gasparini, M., Eisele, J. (2000). Curve-free method for phase I clinical trials, *Biometrics*, **56**, 609 - 615.

[8] Ivanova A., Bolognese J.A., Perevozskaya I. (2008). Adaptive design based on t-statistic for dose-response trials. *Statistics in Medicine*, **27**, 1581 - 1592.

[9] Ivanova A., Flournoy N. ,Chung Y. (2007). Cumulative cohort design for dose-finding. *Journal of Statistical Planning and Inference*, **137**, 2316 - 2327.

[10] Ivanova A., Haghighi A., Mohanty S., Durham S. (2003). Improved up-and-down designs for phase I trials, *Statistics in Medicine*, **22**, 69 - 82.

[11] Ivanova A., Kim S.H. (2009). Dose finding for continuous and ordinal outcomes with a monotone objective function: a unified approach, *Biometrics*, **65**, 307 - 315.

[12] Ivanova, A., Wang, K. (2004). A nonparametric approach to the design and analysis of two-dimensional dose-finding trials, *Statistics in Medicine*, **23**, 1861 - 1870.

[13] Lai T.L., Robbins H. (1982). Iterated least squares in multiperiod control, *Advances in Applied Mathematics*, **3** , 50 – 73.

[14] Leung, D.H., Wang, Y.G. (2001). Isotonic designs for phase I trials, *Controlled Clinical Trials*, **22**, 126 - 138.

[15] Mukerjee H.G. (1981). A stochastic approximation by observation on a discrete lattice using isotonic regression, *The Annals of Statistics*, **9**, 1020 – 1025.

[16] O'Quigley, J. (2006). Theoretical study of the continual reassessment method, *Journal of Statistical Planning and Inference*, **136**, 1765 – 1780.

[17] O'Quigley, J., Pepe, M., Fisher, L. (1990). Continual reassessment method: A practical design for phase 1 clinical trials in cancer, *Biometrics*, **46**, 33 – 48.

[18] O'Quigley J., Zohar S. (2006). Experimental designs for phase I and phase I/II dose-finding studies, *British Journal of Cancer*, **94**, 609 - 613.

[19] Potter, D. M. (2006). Phase I studies of chemotherapeutic agents in cancer patients: a review of the designs, *Journal of Biopharmaceutical Statistics*, **16**, 579 - 604.

[20] Rosenberger, W. F. and Haines, L. M. (2002). Competing designs for phase I clinical trials: a review, *Statistics in Medicine*, **21**, 2757 - 2770.

[21] Shen L.Z., O'Quigley J. (1996). Consistency of continual reassessment method under model misspecification, *Biometrika*, **83** , 395 – 405.

[22] Shiryaev A.N. (translated by R.P. Boas.) (1996). *Probability (2nd. ed.)*, New York : Springer.

[23] Shu, J., O'Quigley, J. (2008). Dose-escalation designs in oncology: ADEPT and the CRM, *Statistics in Medicine*, **27**, 5345 – 5353.

[24] Whitehead, J., Thygesen H., Whitehead A. (2010). A Bayesian dose-finding procedure for phase I clinical trials based only on the assumption of monotonicity, *Statistics in Medicine*, **29**, 1808 - 1824.

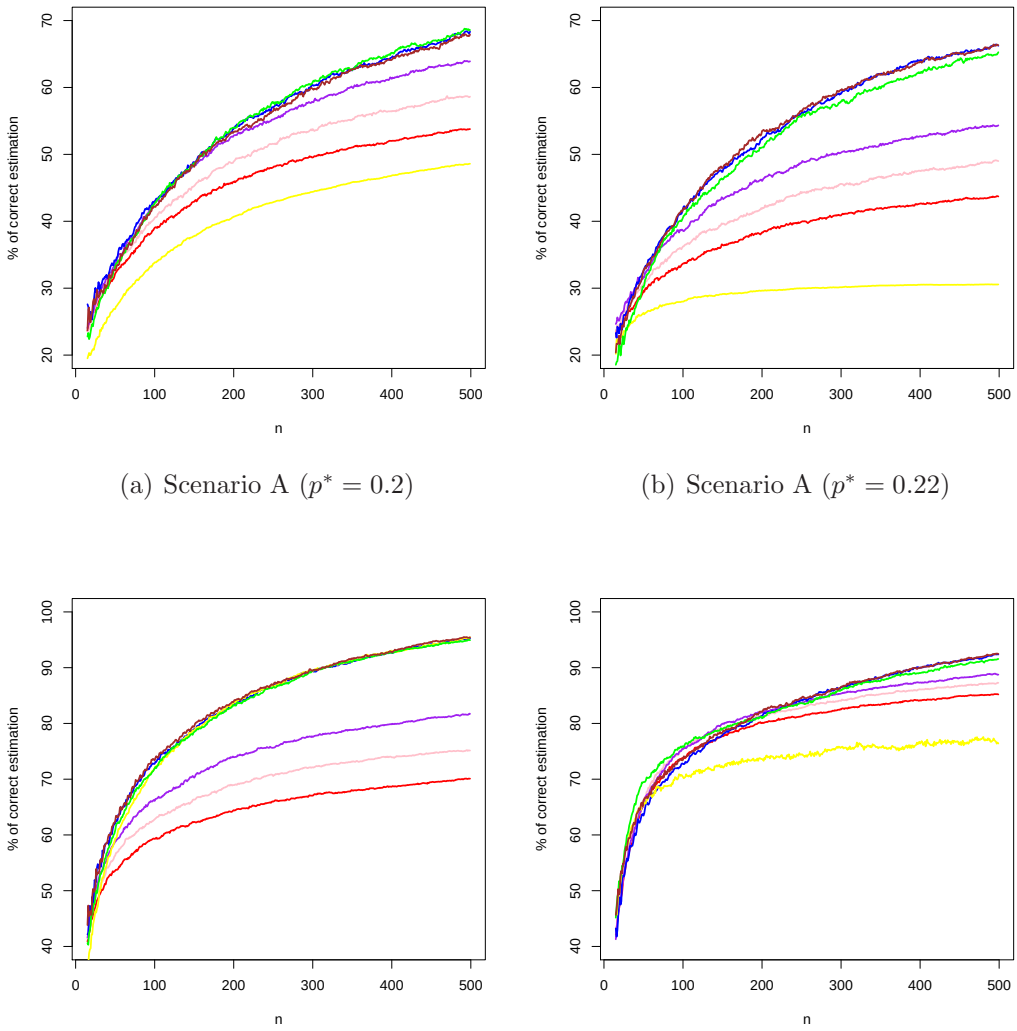

(a) Scenario A ($p^* = 0.2$)

(b) Scenario A ($p^* = 0.22$)

(c) Scenario B ($p^* = 0.2$)

(d) Scenario B ($p^* = 0.3$)

Figure A.1: The percent of finding the true MTD at stage n for $n = 20 \ldots 500$ based on 10,000 replications. The following designs were compared: RAD1 (red), RAD2 (pink), RAD3 (purple), MUK (blue), IVA (green), IVA1 (brown) and CRM (yellow).

Table A.1: The percent of correct estimation (standard errors) at stage n for $n = 20, 30, 40, 50$ based on 10,000 replications; the best designs are shown in bold. The simulation standard error is about 0.5 for all designs and all sample sizes.

% of correct estimation					
Scenario	Design	n=20	n=30	n =40	n=50
A($p^* = 0.2$)	RAD1	25.6	28.3	30.1	32.2
	RAD2	25.4	28.6	31.1	33.1
	RAD3	25.9	29.3	31.4	33.4
	MUK	25.6	**30.1**	**31.5**	**34.1**
	IVA	24.2	27.8	30.2	33.2
	IVA1	**26.0**	28.8	30.4	32.6
	CRM	20.5	22.8	25.6	26.9
A($p^* = 0.22$)	RAD1	24.3	26.3	27.8	29.6
	RAD2	24.6	27.5	29.1	30.7
	RAD3	**25.4**	**28.8**	30.0	31.7
	MUK	23.8	26.6	**30.1**	32.5
	IVA	21.4	24.4	27.1	30.1
	IVA1	21.6	26.1	29.5	**32.7**
	CRM	22.9	24.1	25.4	26.3
B($p^* = 0.2$)	RAD1	44.8	49.2	52.2	53.6
	RAD2	45.6	50.5	54.2	56.3
	RAD3	**46.3**	52.6	55.8	58.8
	MUK	45.9	**53.9**	**58.6**	62.0
	IVA	46.3	51.9	56.9	60.8
	IVA1	45.6	53.6	**58.6**	**62.4**
	CRM	42.1	49.9	55.1	59.3
B($p^* = 0.3$)	RAD1	47.9	55.8	61.7	65.5
	RAD2	48.7	56.3	62.0	66.1
	RAD3	47.7	55.9	61.2	65.5
	MUK	48.8	55.4	59.9	63.47
	IVA	**51.2**	**60.2**	**65.4**	**69.5**
	IVA1	49.7	58.6	62.5	66.2
	CRM	50.6	58.3	63.0	65.2

Table A.2: The proportion of subjects treated with the true MTD for $n = 20, 30, 40, 50$ based on 10,000 replications; The simulation standard error is about 0.002 for all designs and all sample sizes.

Proportion of subjects treated with the true MTD					
Scenario	Design	n=20	n=30	n =40	n=50
A($p^* = 0.2$)	RAD1	0.190	0.212	0.229	0.244
	RAD2	0.187	0.209	0.227	0.243
	RAD3	0.187	0.207	0.225	0.240
	MUK	0.198	0.214	0.227	0.239
	IVA	0.202	0.220	0.234	0.246
	IVA1	**0.218**	**0.231**	**0.239**	**0.248**
	CRM	0.173	0.188	0.202	0.214
A($p^* = 0.22$)	RAD1	0.192	0.217	0.233	0.245
	RAD2	0.194	0.221	0.239	0.252
	RAD3	0.202	0.230	0.249	0.263
	MUK	**0.207**	**0.235**	**0.254**	**0.269**
	IVA	0.157	0.183	0.203	0.218
	IVA1	0.196	0.212	0.226	0.238
	CRM	0.170	0.191	0.205	0.216
B($p^* = 0.2$)	RAD1	0.325	0.361	0.389	0.411
	RAD2	0.324	0.361	0.389	0.413
	RAD3	0.314	0.351	0.381	0.405
	MUK	0.312	0.339	0.360	0.375
	IVA	0.333	**0.367**	0.391	0.409
	IVA1	**0.340**	0.362	0.380	0.3927
	CRM	0.310	0.360	**0.402**	**0.437**
B($p^* = 0.3$)	RAD1	0.335	0.382	0.422	0.456
	RAD2	0.330	0.376	0.414	0.446
	RAD3	0.318	0.359	0.393	0.423
	MUK	0.312	0.337	0.356	0.371
	IVA	0.362	0.403	0.433	0.456
	IVA1	0.348	0.366	0.382	0.394
	CRM	**0.372**	**0.434**	**0.478**	**0.512**

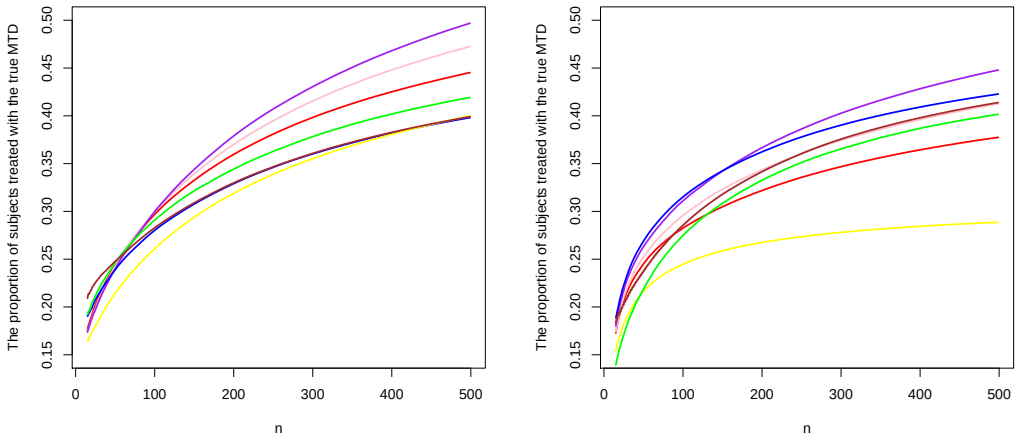

(a) Scenario A ($p^* = 0.2$)

(b) Scenario A ($p^* = 0.22$)

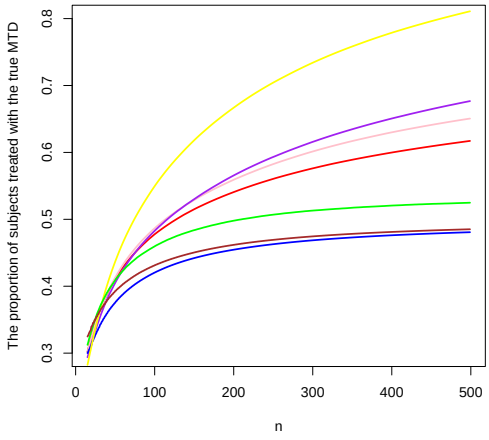

(c) Scenario B ($p^* = 0.2$)

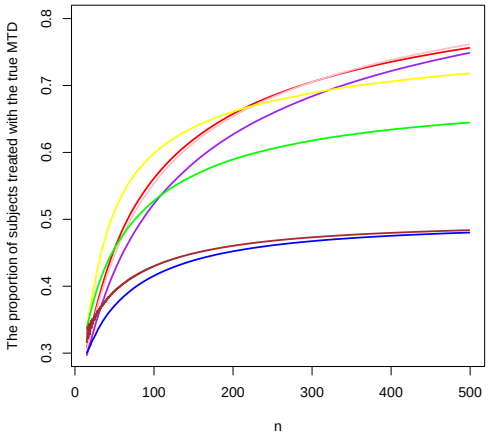

(d) Scenario B ($p^* = 0.3$)

Figure A.2: The proportion of subjects treated with the true MTD for $n = 20 \ldots 500$ based on 10,000 replications. The following designs were compared: RAD1 (red), RAD2 (pink), RAD3 (purple), MUK (blue), IVA (green), IVA1 (brown) and CRM (yellow).