

# A Review of Dose Finding Methods and Theory

Ying Kuen Cheung<sup>1,a</sup>

<sup>a</sup>Department of Biostatistics, Mailman School of Public Health, Columbia University, USA

---

## Abstract

In this article, we review the statistical methods and theory for dose finding in early phase clinical trials, where the primary objective is to identify an acceptable dose for further clinical investigation. The dose finding literature is initially motivated by applications in phase I clinical trials, in which dose finding is often formulated as a percentile estimation problem. We will present some important phase I methods and give an update on new theoretical developments since a recent review by Cheung (2010), with an aim to cover a broader class of dose finding problems and to illustrate how the general dose finding theory may be applied to evaluate and improve a method. Specifically, we will illustrate theoretical techniques with some numerical results in the context of a phase I/II study that uses trinary toxicity/efficacy outcomes as basis of dose finding.

**Keywords:** adaptive designs, coherence, consistency,  $h$ -sensitivity, optimality bounds, phase I trials, phase I/II trials, trinary outcomes

---

## 1. Introduction

Phase I and phase I/II clinical trials of a new drug are typically small dose finding studies, whose goal is to identify an acceptable dose, defined with respect to some pre-specified toxicity and efficacy criteria, for further clinical investigation. Because little is known about the new drug in the early phase investigation, these studies are conducted in an adaptive and *small-group-sequential* manner, characterized by an iterative process: (1) start the trial at a safe dose, (2) treat a small group of patients at the dose, and (3) make dose decisions for the next group based on the outcomes of the current group. This process continues until a pre-specified sample size is reached. This approach, motivated by ethical considerations, avoids randomizing patients to doses with excessive toxicity risks and aims to concentrate treatments around the “good” doses.

Storer and DeMets (1987) give one of the earliest discussions on the topic of dose finding in the context of standard phase I trials, where the primary objective is to find the maximum tolerated dose based on *binary* toxicity outcomes. The authors formulate phase I dose finding as a percentile estimation problem, and define the maximum tolerated dose  $\theta$  as a dose associated with a pre-specified toxicity probability  $p$ . Precisely,

$$\theta = \arg \min_k |\pi_1(k) - p|, \quad (1.1)$$

where  $\pi_1(k)$  is the probability of toxicity at dose level  $k$ . Note that for practical reasons, a typical dose finding study tests a drug at a discrete set of dose levels  $\{1, \dots, K\}$ ; thus, the estimand  $\theta$  lives in a discrete and finite parameter space  $\{1, \dots, K\}$ . This feature, together with the ethical considerations in

---

<sup>1</sup> Department of Biostatistics, Mailman School of Public Health, Columbia University, 722 West 168th Street, New York, 10032, USA. E-mail: [yc632@columbia.edu](mailto:yc632@columbia.edu)

human studies, differentiates the dose finding problem in clinical trials from the percentile estimation problem in bioassay (Finney, 1978). Since Storer and DeMets (1987), several authors have proposed a great variety of methods and statistical principles to estimate  $\theta$ . Important examples include the up-and-down and random walk designs (Storer, 1989; Durham *et al.*, 1997), model-based methods (O’Quigley *et al.*, 1990; Zacks *et al.*, 1998), stepwise procedures (Lin and Shih, 2001; Cheung, 2007) and the stochastic approximation (Anbar, 1984; Cheung and Elkind, 2010). The study of the theoretical framework for phase I dose finding, however, has not received much attention until about a decade ago. A recent review of the theoretical criteria in phase I dose finding is given in Cheung (2010).

As the dose finding objective in clinical trials becomes increasingly sophisticated, novel adaptive designs have been proposed to address estimation beyond the percentile estimation objective (1.1). Bekele and Thall (2004), Yuan *et al.* (2007), and Lee *et al.* (2011) consider phase I dose finding with non-binary toxicity outcomes. For phase I/II trials in which dose finding is based on bivariate efficacy and toxicity outcomes, several model-based methods have been proposed (e.g., Braun, 2002; O’Quigley *et al.*, 2001; Thall and Cook, 2004). The use of utility scores has been considered in combination therapy trials (Houede *et al.*, 2010). Most of these methods are model-based in that they sequentially treat the next patients at a dose estimated to be “good” based on a dose-outcome model. Because the dose-outcome model is usually elaborate, these methods are computer intensive, and their properties are not theoretically tractable. As a result, the theoretical study of these complex methods has been difficult and rare when compared to the standard phase I dose finding problem.

The main goal of this article is to provide an overview of dose finding methods in clinical trials and an up-to-date review of the theoretical study. We will first focus on the standard phase I dose finding problem as defined in (1.1), with Section 2 presenting important examples of statistical methods, and Section 3 reviewing key theoretical criteria. Section 4 presents a general formulation for dose finding methods in clinical trials, along with general optimality bounds. In Section 5, we will illustrate the applications of the theoretical criteria in a specific dose finding situation using trinary outcomes. This article is ended with some concluding remarks in Section 6.

## 2. Methods for Phase I Dose Finding Problem

### 2.1. Notation

Let  $x_i \in \{1, \dots, K\}$  denote the dose level that patient  $i$  receives, and let  $Y_i = Y_i(x_i)$  denote the patient’s binary toxicity outcome where  $Y_i(k)$  is the toxicity indicator when the patient receives dose level  $k$ . In a dose finding study, each patient receives only one dose (i.e.,  $x_i$ ) and thus only  $Y_i$  is observable. Further let  $\pi_1(k) = \Pr\{Y_i(k) = 1\}$  so that  $\{\pi_1(1), \pi_1(2), \dots, \pi_1(K)\}$  denote the dose-toxicity curve defined on the test doses. A dose finding method consists of two components. The first component is the *design* that pertains to the determination of the sequence  $\{x_i\}$ , possibly with respect to some ethical constraints due to the fact that the experimental units are human subjects. The second component is the *estimation* of  $\theta$  as defined in (1.1). This section reviews some important classes of phase I dose finding methods.

### 2.2. Random walk

Storer (1989) describes a number of up-and-down escalation designs. For example, according to his Design B, a trial starts at a low and safe dose (e.g.,  $x_1 = 1$ ), and a single patient is treated and observed for a toxicity outcome. The next patient will receive the next lower dose level if a toxicity is observed, and the next higher dose otherwise. Mathematically, Storer’s Design B can be represented as

$$x_{i+1} = Y_i \max(x_i - 1, 1) + (1 - Y_i) \min(x_i + 1, K), \quad (2.1)$$

where the starting dose  $x_1$  is specified by the clinical investigators. By using Markov chain representation, the sequence  $\{x_i\}$  according to (2.1) can be shown to sample around a dose that causes toxicity with 0.50 probability. For situations where a lower toxicity tolerance is desired, Storer (1989) proposes alternative up-and-down schemes that escalate after a small number of consecutive non-toxic outcomes, and suggests using a combination of these schemes in a study. At the end of the study, the maximum tolerated dose  $\theta$  can be estimated by using logistic regression or isotonic regression.

Durham *et al.* (1997) introduce a randomized generalization of the up-and-down designs for any target rate  $p \leq 0.50$ . Like Storer's Design B, the design de-escalates dose level for the next patient if the current patient has a toxicity. When no toxicity is observed, however, the method will escalate based on the outcome of a biased coin with probability  $p/(1-p)$ . This biased coin design can be described as a random walk:

$$x_{i+1} = Y_i \max(x_i - 1, 1) + (1 - Y_i)\{x_i(1 - c_i) + \min(x_i + 1, K)c_i\}, \quad (2.2)$$

where  $c_i$  is a Bernoulli variable with probability  $p/(1-p)$  independent of  $Y_i$ . A starting dose  $x_1$  needs to be specified to initiate the recursion (2.2). The biased coin design makes dose assignments with respect to the target rate  $p$ . Exploiting the Markov property of random walk, we can derive analytically that the sequence  $\{x_i\}$  will sample around a dose with toxicity probability about  $p$ . Thus, it is natural to estimate  $\theta$  with the mode of  $x_i$ s, although logistic and isotonic regression may also be used for estimation purposes.

Because the escalation rules of (2.1) and (2.2) for any given dose can be written down in advance without regards to the outcomes at other doses, they can be implemented in a study without interim statistical calculations. These methods, often called algorithm-based designs, are simple to implement and are transparent to non-statistical investigators in terms of the dosing decisions.

### 2.3. Model-based designs

The continual reassessment method (CRM; O'Quigley *et al.*, 1990) makes dose assignments based on a parametric dose-toxicity model, postulating  $\pi_1(k) = F(k, \beta)$  for some parameter  $\beta$ . The main idea of the CRM is to treat the next subject at the dose with toxicity probability estimated to be closest to the target  $p$ . Precisely, the next dose is determined by

$$x_{i+1} = \hat{\theta}_i := \arg \min_k \left| F(k, \hat{\beta}_i) - p \right|, \quad (2.3)$$

where  $\hat{\beta}_i$  is the posterior mean of  $\beta$  based on the observations  $\{(x_j, Y_j) : j \leq i\}$ . That is,

$$\hat{\beta}_i = \frac{\int_{-\infty}^{\infty} \beta L_i(\beta) dG(\beta)}{\int_{-\infty}^{\infty} L_i(\beta) dG(\beta)},$$

where  $L_i(\beta)$  is the binomial likelihood of  $\beta$  given  $\{(x_j, Y_j) : j \leq i\}$  and  $G(\beta)$  is the prior distribution of  $\beta$ . Under a Bayesian framework, the CRM sets the starting dose  $x_1 = \hat{\theta}_0$  in accordance with (2.3) with  $\hat{\beta}_0$  being the prior mean of  $\beta$ . The choice of  $G$  is discussed in Cheung (2011) and Lee and Cheung (2011) who study the use of least informative prior.

The CRM combines the design and estimation components of a dose finding study by setting  $x_{i+1}$  at  $\hat{\theta}_i$ . At the end of a study with sample size  $n$ , the maximum tolerated dose  $\theta$  is estimated by  $\hat{\theta}_n$ , a dose that would have been given to the next patient if the study continued. This coupling of design and estimation motivates a subsequent literature of model-based methods that adopt similar

ideas. In a nutshell, a model-based method makes dose decisions and estimation based on a dose-toxicity model that is being updated repeatedly throughout a study. The main difference between the various model-based methods lies in how the model  $F$  is specified and how the dose-toxicity curve  $\pi_1$  is estimated. Zacks *et al.* (1998), for instance, estimate  $\theta$  using an asymmetric loss function that penalizes over-dosing more than under-dosing, while (2.3) places equal penalty on the two erroneous dosing decisions. Haines *et al.* (2003) consider Bayesian  $c$ -optimal designs under parametric models, whereas Gasparini and Eisele (2000) and Cheung (2002) discuss the use of Bayesian nonparametric model, and Leung and Wang (2001) use a frequentist analogue that continually estimates  $\pi_1$  using nonparametric isotonic regression. Meanwhile, Shen and O'Quigley (1996) and Cheung and Chappell (2002) consider robustness of parametric estimation of  $\pi_1$ .

#### 2.4. Hybrid two-stage designs

Most model-based methods are Bayesian and assume a prior distribution on the dose-toxicity curve. At the start of a trial when there is few or no data, these Bayesian methods will rely on the prior belief about  $\theta$  in prescribing doses. As an alternative, O'Quigley and Shen (1996) propose a version of the CRM that estimates  $\beta$  using maximum likelihood estimation, that is, replacing  $\hat{\beta}_i$  in (2.3) with  $\arg \max_{\beta} L_i(\beta)$ . The likelihood CRM avoids the needs for choosing a prior. However, since maximum likelihood estimate does not exist until there is heterogeneity in the toxicity outcomes, an initial dose escalation sequence  $\{x_{i0}\}$  is required before the first toxicity occurs. This leads to the concept of a hybrid two-stage design: define  $x_i = x_{i0}$  if  $Y_j = 0$  for all  $j < i$ , and  $x_i$  is defined by the model-based recommendation such as (2.3) if  $Y_j = 1$  for some  $j < i$ . Generally, a hybrid design consists of three parts: a model-based design, an initial set of escalation rules  $\{x_{i0}\}$ , and a transition rule (e.g., when the first toxicity occurs). Section 3.1 gives additional discussion on the choice of  $\{x_{i0}\}$ .

#### 2.5. Stochastic approximation and virtual observations

Stochastic approximation is perhaps the earliest method proposed for the purpose of dose finding. Anbar (1984) consider using the Robbins-Monro method (1951) in phase I trials. Precisely, the stochastic approximation determines the next dose by

$$x_{i+1} = x_i - \frac{1}{ib}(Y_i - p), \quad (2.4)$$

for some constant  $b > 0$  and a pre-specified starting dose  $x_1$ . The theoretically properties of recursion (2.4) have been extensively studied, and numerous variations have been proposed. However, a few practical reasons prevent its use in phase I trials. First, while most dose finding studies are conducted on discrete dose levels, the stochastic approximation assumes a continuum of dosages. To circumvent this issue, O'Quigley and Chevret (1991) consider

$$x_{i+1} = C \left\{ x_i - \frac{1}{ib}(Y_i - p) \right\}, \quad (2.5)$$

where  $C(x)$  is the rounded value of  $x$  for  $x \in [0.5, K + 0.5)$ , and is set to 1 and  $K$  respectively if  $x < 0.5$  and  $> K + 0.5$ . However, since  $|(ib)^{-1}(Y_i - p)| < 0.5$  when  $i$  is large enough, the sequence  $\{x_i\}$  according to (2.5) will eventually stay at one dose level, i.e.,  $x_{i+1} = x_i = k'$ , which may not be  $\theta$ . Cheung (2010) describes the concept of virtual observations on which a stochastic approximation recursion is based: Set  $x_1^* = x_1$ , update

$$x_{i+1}^* = x_i^* - \frac{1}{ib}(V_i - p), \quad (2.6)$$

where  $V_i = Y_i + b(x_i^* - x_i)$  is the virtual observation of subject  $i$ , and treat the next patient at  $x_{i+1} = C(x_{i+1}^*)$ . Because the cumulative increment  $\sum_i (ib)^{-1}(V_i - p)$  in (2.6) is unbounded, following the standard theory of stochastic approximation under some regularity conditions, we can show that  $\{x_i^*\}$  is asymptotically normal with mean  $\phi$  where  $C(\phi) = \theta$ . In other words, the actual assigned dose sequence  $\{x_i\}$  is consistent for  $\theta$ .

Second, the choice of  $b$  is crucial to the success of the method. While one may choose adaptively based on interim data to guarantee consistency, the convergence rate is too slow to be relevant in small samples. Finally, Wu (1985) shows that the stochastic approximation is inefficient when using binary outcomes in small samples. However, an exception is when the binary outcome  $Y_i$  is defined by dichotomizing a continuous outcome that is observable. For this specific situation, Cheung and Elkind (2010) show that stochastic approximation using the continuous outcomes leads to substantial improvement in efficiency, when compared to any phase I methods using the dichotomized  $Y_i$ s only, and offer a practical way to choose  $b$  so as to control the bias in estimation at a pre-specified level.

### 3. Theoretical Criteria for Phase I Methods

#### 3.1. Coherence

Coherence is a dose finding principle that stipulates that no dose escalation (de-escalation) should take place for the next patient if the current subject experiences a toxic (non-toxic) outcome (Cheung, 2005). Mathematically, a dose finding design  $\mathcal{D}$  is coherent in escalation if with probability one

$$P_{\mathcal{D}}(x_i - x_{i-1} > 0 | Y_{i-1} = 1) = 0, \quad (3.1)$$

for all  $i$ ; and is coherent in de-escalation if with probability one

$$P_{\mathcal{D}}(x_i - x_{i-1} < 0 | Y_{i-1} = 0) = 0, \quad (3.2)$$

for all  $i$ , where  $P_{\mathcal{D}}(\cdot)$  denotes probability computed under the design  $\mathcal{D}$ .

Coherence is an ethical consideration for the design component of a dose finding method, and as a result, it may not correspond to efficient estimation of  $\theta$  with regard to the estimation objective (1.1). Coherence is a built-in feature of many algorithm-based designs such as (2.1) and (2.2), because they would not be clinically acceptable otherwise. It is also clear that coherence holds for the stochastic approximation (2.4). In contrast, for model-based designs or hybrid designs, coherence needs to be established on a method-by-method basis. Cheung (2005) establishes the coherence of a specific version of the CRM, and shows that a hybrid two-stage design is not necessarily coherent.

Specifically, a hybrid design is not coherent if the initial sequence  $\{x_{i0}\}$  is over-conservative and escalates very slowly (Cheung, 2005). Based on this interplay between the coherence principle and the initial sequence, Jia *et al.* (2014) subsequently show the unique existence of a most conservative, coherent initial sequence  $\{x_{i0}\}$  for a given CRM model, and prescribe a conservative choice of  $\{x_{i0}\}$ .

#### 3.2. Optimality bounds

While there are many phase I dose finding methods (cf. Section 2), it is common to compare their performance *relative to each other* by simulation. However, such comparison does not allow ascertainment of their efficiency in an absolute sense. O'Quigley *et al.* (2002) describe a nonparametric benchmark design to assess a method's performance relative to an objective optimality bound. To calculate the benchmark, we postulate that patient  $i$  carries a toxicity tolerance  $U_i$  that is uniformly distributed, so that  $Y_i(k) = I\{U_i \leq \pi_1(k)\}$ , where  $I(E)$  is an indicator function of the event  $E$ . Further

suppose that  $U_i$ s are observable for all  $n$  patients in a trial, implying that we observe a complete toxicity profile  $\{Y_i(1), Y_i(2), \dots, Y_i(K)\}$  for each patient and we can estimate the dose-toxicity curve  $\pi_1$  based on these complete profiles. Precisely, set  $\pi_1^*(k) = n^{-1} \sum_{i=1}^n Y_i(k)$  for all  $k$ , and estimate  $\theta$  with  $\theta^* = \arg \min_k |\pi_1^*(k) - p|$ .

Note that in a dose finding study, only  $Y_i = Y_i(x_i)$  is observed for patient  $i$  instead of the complete profile. Thus, intuitively, by using additional information that is not available in practice, the “estimator”  $\theta^*$  is expected to perform better than any methods that use only  $Y_i$ s for estimation. Therefore, while  $\theta^*$  cannot be implemented in practice, it may serve as a yardstick for other dose finding methods.

Importantly, the properties of  $\theta^*$  are theoretically tractable and can be easily computed. Cheung (2013) shows that the distribution of  $\theta^*$  can be approximated by

$$\Pr(\theta^* \geq k) \approx \Phi \left[ \frac{\sqrt{n} \{2p - \pi_1(k-1) - \pi_1(k) + 0.5n^{-1}\}}{\sigma_k} \right], \quad (3.3)$$

for  $k \geq 2$ , where  $\sigma_k^2 = \pi_1(k-1)\{1 - \pi_1(k-1)\} + \pi_1(k)\{1 - \pi_1(k)\} + 2\pi_1(k-1)\{1 - \pi_1(k)\}$ , and  $\Phi$  is the standard normal distribution function. The expression (3.3) is then used to derive a lower bound of the sample size required for a certain accuracy constraint, and also a sample size formula for a version of the CRM.

### 3.3. Consistency and $h$ -sensitivity

An estimator  $\tilde{\theta}_n$  is consistent for  $\theta$  if and only if  $\tilde{\theta}_n = \theta$  with probability one as  $n \rightarrow \infty$ . For model-based designs that couple design and estimation, Azriel *et al.* (2011) prove that consistency cannot be achieved under all dose-toxicity curves. Rather, a method will achieve consistency if the design visits every dose level infinitely often (e.g., up-and-down designs). This presents a dilemma between experimentation and estimation in the context of dose finding studies, in that this “best intention” approach is at odds with the estimation objective. In other words, consistency is an infeasible goal under the ethical constraints of phase I trials.

As a middle ground, Cheung (2010) discuss the concept of  $h$ -sensitivity and indifference interval in dose finding. Briefly, the interval  $p \pm h$  is called an indifference interval of a dose finding design  $\mathcal{D}$  if there is  $N > 0$  and  $h \in (0, p)$  such that

$$P_{\mathcal{D}} \{ \pi_1(\tilde{\theta}_n) \in (p - h, p + h), \text{ for all } n \geq N \} = 1.$$

For brevity of discussion, we confine our attention to the class of dose-toxicity curves such that  $\pi_1(k) \in p \pm h$  for some  $k$ . A design  $\mathcal{D}$  with half-width  $h$  is called a  $h$ -sensitivity design. In words, a design is  $h$ -sensitive if it will select a dose with toxicity probability on the interval  $p \pm h$ . An easy corollary is that any dose finding method that provides consistent estimation of  $\theta$  is also  $h$ -sensitive for any arbitrarily small  $h$ . The converse is not true: the CRM, while not consistent in the sense of Azriel *et al.* (2011), can be calibrated to be  $h$ -sensitive for any given  $h$  under all dose-toxicity curves (Lee and Cheung, 2009).

## 4. General Dose Finding Methods

### 4.1. Estimation objective function

This section extends the notation in Section 2 to cover the general dose finding settings. We consider multinomial outcome  $Y$  that takes on  $L + 1$  distinct values  $\{\omega_0, \omega_1, \dots, \omega_L\}$ . Without loss of generality,

Table 1: Weights of toxicity types and grades in Lee *et al.* (2012)

	Toxicity grade				
	0	1	2	3	4
Neuropathy	0.00	0.19	0.64	1.03	2.53
Low platelet count	0.00	0.17	0.17	0.40	0.85

assume that  $\omega_0 < \omega_1 < \dots < \omega_L$  so that the distribution of the outcome  $Y_i(k)$  of patient  $i$  at dose level  $k$  can be described by the tail function

$$\pi_l(k) = \Pr\{Y_i(k) \geq \omega_l\}, \quad \text{for } l = 1, \dots, L. \quad (4.1)$$

We shall also use  $\pi = \{\pi_l(k) : l = 1, \dots, L; k = 1, \dots, K\}$  to denote the entire collection of dose-outcome probabilities. Generally, the estimation component of a dose finding trial is to estimate a dose  $d(\pi) \in \{1, \dots, K\}$ , where  $d(\cdot)$  is a known objective function specified in the planning stage of a trial.

**Example 1.** (Dose finding with binary toxicity outcomes). The outcome  $Y$  in a phase I trial is binary, i.e.,  $L = 1$ ,  $\omega_0 = 0$  and  $\omega_1 = 1$ , so that  $\pi_1(k) = \Pr\{Y_i(k) = 1\}$  denotes the probability of a toxic outcome at dose  $k$ . The estimation objective function is defined as  $d_1(\pi_1) = \arg \min_k |\pi_1(k) - p|$  as displayed in (1.1).

**Example 2.** (Dose finding with total toxicity burden). In most clinical situations, a patient may experience more than one type of toxicity at different severity grades. Instead of summarizing this toxicity profile into a binary toxicity outcome, Bekele and Thall (2004) propose to measure the total toxicity burden of a patient using a weighted sum of grades and types of toxicities: each toxicity type and grade will have a pre-determined weight, and the total burden  $Y$  of a patient will be calculated by adding the weights of toxicity grade/type the patient experiences. To illustrate the concept of total burden, consider a clinical study described in Lee *et al.* (2012) where neuropathy and low platelet count are the potential toxicities of a drug, with the severity weights given in Table 1. Based on these weights, a patient with a grade 2 neuropathy and a grade 1 low platelet will have a total toxicity burden  $Y = 0.64 + 0.17 = 0.81$ . While  $Y$  is a numerical score, it only takes on a finite number of possible values depending on the toxicity grades/types considered. We can enumerate from Table 1 that the burden score  $Y$  takes on 18 possible values (i.e.,  $L = 17$ ), which are listed in the second column of Table 2. Lee *et al.* (2011) study the method of Bekele-Thall that aims to estimate  $d_{BT}(\pi) = \arg \min_k |E_\pi\{Y(k)\} - \tau|$ , where  $\tau$  is a pre-specified level. As  $\tau$  is set to 0.72 in the paper, the true target dose  $d_{BT}(\pi) = 3$ , with  $E_\pi\{Y(3)\} = 0.81$ , in the dose-outcome scenario given in Table 2.

**Example 3.** (Dose finding with multiple toxicity constraints). For the same clinical setting as in Example 2, Lee *et al.* (2011) propose an alternative estimation objective that chooses the smallest dose among doses that satisfy individual toxicity constraints. To be precise, they define  $d_{MC}(\pi) = \min_{j=1, \dots, J} \theta_j$  where  $\theta_j = \arg \min_k |\Pr\{Y(k) \geq t_j\} - p_j|$ , for pre-specified toxicity thresholds  $t_1 < \dots < t_J$  and toxicity rates  $p_1 > \dots > p_J$ , and  $J$  is the number of toxicity constraints. In the special case with  $J = 1$ , the objective function reduces to (1.1). To illustrate using Table 2, suppose  $J = 2$  and that  $t_1 = 1$ ,  $t_2 = 1.5$ ,  $p_1 = 0.25$ , and  $p_2 = 0.10$ . Under the scenario described in the table, we obtain  $\theta_1 = 3$ ,  $\theta_2 = 2$ , and hence  $d_{MC}(\pi) = 2$ .

**Example 4.** (Dose finding with bivariate binary outcomes). In a phase I/II trial, dose finding is based on both efficacy and toxicity outcomes. Generally, the outcome  $Y$  takes on four possible

Table 2: A dose-outcome distribution of toxicity burden score

$l$	$\omega_l$	$\pi_l(1)$	$\pi_l(2)$	$\pi_l(3)$	$\pi_l(4)$	$\pi_l(5)$
0	0.00	1.00	1.00	1.00	1.00	1.00
1	0.17	0.53	0.63	0.70	0.89	0.94
2	0.19	0.40	0.48	0.53	0.78	0.86
3	0.36	0.28	0.39	0.45	0.72	0.82
4	0.40	0.25	0.35	0.41	0.66	0.76
5	0.59	0.21	0.29	0.35	0.59	0.70
6	0.64	0.20	0.28	0.33	0.55	0.66
7	0.81	0.09	0.22	0.32	0.54	0.65
8	0.85	0.06	0.19	0.31	0.52	0.63
9	1.03	0.05	0.16	0.25	0.45	0.55
10	1.04	0.03	0.14	0.25	0.44	0.54
11	1.20	0.02	0.12	0.23	0.38	0.47
12	1.43	0.01	0.11	0.23	0.37	0.46
13	1.49	0.01	0.11	0.23	0.36	0.45
14	1.88	0.01	0.10	0.23	0.35	0.43
15	2.53	0.01	0.10	0.23	0.34	0.42
16	2.70	0.00	0.04	0.12	0.24	0.34
17	2.93	0.00	0.02	0.05	0.14	0.21
18	3.38	0.00	0.01	0.02	0.07	0.13
	$E_\pi\{Y(k)\}$	0.25	0.51	0.81	1.28	1.6

values. Using a particular convention, we may define  $Y = 0$  for a patient without no toxicity and no response,  $Y = 1$  for response without toxicity,  $Y = 2$  for toxicity without response, and  $Y = 3$  for both response and toxicity. Let  $F_T(k)$  and  $F_R(k)$  denote respectively the toxicity probability and response probability at dose  $k$ . Using the notation (4.1), these probabilities can be expressed in terms of  $\pi$  as follows:  $F_T(k) = \Pr\{Y_i(k) \geq 2\} = \pi_2(k)$  and  $F_R(k) = \Pr\{Y_i(k) = 1 \text{ or } Y_i(k) = 3\} = \pi_1(k) - \pi_2(k) + \pi_3(k)$ . The desirability  $\delta_k$  of dose  $k$  is then defined as a function of the response and toxicity rates, that is,  $\delta_k = \delta\{F_R(k), F_T(k)\}$ , where  $\delta(r, s)$  is increasing in  $r$  and decreasing  $s$ . The estimation objective function is defined as  $d_{ET}(\pi) = \arg \max_k \delta_k$ .

#### 4.2. General design strategy and optimality bound

As illustrated in the previous subsection, the estimation component of a dose finding study is highly study-specific. Examples 2 and 3, both dealing with situations with multiple toxicities, take very different objective functions and hence lead to different target doses; recall that  $d_{BT}(\pi) = 3$  and  $d_{MC}(\pi) = 2$  for Table 2. Therefore, it is very crucial to work with the clinical investigators to decide the estimation objective at the planning stage of a trial.

In contrast, the design component has predominantly taken the form of a model-based or hybrid strategy: start a trial at a safe dose, treat and observe a few patients, obtain some model-based estimate of  $\pi$  using on the most recent data, and treat the next group of patients based on this estimate. While this strategy is applicable to very general outcome types and estimation objective functions, the actual implementations are computer intensive and specific to the estimation objective and the assumed dose-outcome model. The performance of the dose finding method also depends on this assumed dose-outcome model and the underlying dose-outcome relationship, which can be quite complicated. Therefore, it is important to be able to derive some theoretical optimality bound for these methods as a way to ascertain their efficiencies and perform design diagnosis. Cheung (2014) extends the concept of nonparametric benchmark to the general dose finding settings. To calculate a benchmark for a trial with  $Y$  taking on  $L + 1$  possible values, we postulate that each patient carries a tolerance profile  $(U_{i1}, U_{i2}, \dots, U_{iL})$  that is uniformly distributed on  $[0, 1]^L$ , so that under a given  $\pi$ , the outcome



of patient  $i$  is determined by

$$Y_i(k) = \omega_l \text{ if } U_{i,l+1} > \frac{\pi_{l+1}(k)}{\pi_l(k)} \quad \text{and} \quad U_{ij} \leq \frac{\pi_j(k)}{\pi_{j-1}(k)}, \quad \text{for all } j \leq L, \quad (4.2)$$

where we set  $\pi_0(k) \equiv 1$  and  $\pi_{L+1}(k) \equiv 0$  by convention. Consequently, we may estimate  $\pi_l(k)$  by  $\pi_l^*(k) = n^{-1} \sum_{i=1}^n I\{Y_i(k) \geq \omega_l\}$ ; hence  $d(\pi^*)$  can be calculated as the benchmark for the objective  $d(\pi)$ . Some theoretical properties of  $\pi^*$  are given in Cheung (2014). For each given  $d(\cdot)$  and  $\pi$ , the sampling distribution of  $\theta^*$  can be derived analytically using asymptotic approximation in principle, analogously to (3.3). In practice, the sampling properties of a benchmark can be easily approximated by generating the outcomes (4.2) using simulation.

### 5. An Application to Dose Finding with Trinary Outcomes

In this section, we illustrate how the theoretical criteria can be applied to build and evaluate dose finding methods in the context of a phase I/II trial where each patient will experience one of three possible outcomes: no toxicity and no response ( $Y = 0$ ), response without toxicity ( $Y = 1$ ), and toxicity ( $Y = 2$ ). It can be viewed as a special case of Example 4 above with  $\pi_3(k) \equiv 0$ , that is, a toxicity will preclude a response. That is, we have  $F_R(k) = \pi_1(k) - \pi_2(k)$  and  $F_T(k) = \pi_2(k)$ . For the estimation objective, we use the desirability function specified in Thall and Cook (2004):

$$\delta(r, s) = 1 - \frac{\sqrt{(1-r)^2 + s^2}}{\sqrt{(1-p_E)^2 + p_T^2}}, \quad (5.1)$$

where  $(p_E, p_T)$  solves  $(p_T + 0.045)p_E^2 - 0.347p_E + 0.147 = 0$  and  $s(p_E - r) + (1-r)(p_T - s) = 0$ . The function (5.1) attains a maximum of 1 when  $r = 1$  and  $s = 1$ , and can be negative. Thus, we consider a modified efficacy-toxicity estimation objective function defined as  $d_{ET}(\pi) = \arg \max_k \delta_k$  if  $\delta_k > 0$  for some  $k$ ;  $= 0$  otherwise.

Cheung (2015) introduces the concept of a tandem dose finding design, in which patients are enrolled in pairs and treated at adjacent doses, and dose escalation between pairs are conducted in tandem. Without loss of generality, the design points will be subject to the following constraints:

$$x_{2i} = x_{2i-1} + 1, \quad \text{for } i = 1, \dots, \frac{n}{2}, \quad (5.2)$$

for patient  $2i - 1$  and patient  $2i$ , in a trial with sample size  $n$ , i.e.,  $n/2$  pairs (when two patients are enrolled, we can pre-randomize which patient is the “first” patient in the pair). The tandem design (5.2) is intended as an escalation restriction that enables comparison of neighboring doses. For a phase I/II trial with trinary outcomes, there are 9 possible outcome configurations for each pair of subjects. And while there are many ways to specify the dose escalation rules corresponding to each of the outcome configurations, Table 3 prescribes a specific set of rules where the coherence principles are used to impose restrictions on certain rules. For example, suppose that the outcomes are  $Y_{2i-1} = 0$  and  $Y_{2i} = 2$  for the  $(2i - 1)^{th}$  and  $(2i)^{th}$  patients. Coherence in de-escalation (3.2) imposes that

$$x_{2i+1} \geq x_{2i-1} \quad (5.3)$$

following the outcome  $Y_{2i-1} = 0$ , and coherence in escalation (3.1) implies

$$x_{2i+2} \leq x_{2i} = x_{2i-1} + 1 \quad (5.4)$$

Table 3: Escalation in tandem design for phase I/II trial where  $x_{2i} = x_{2i-1} + 1$  for all  $i$

$(Y_{2i-1}, Y_{2i})$	Rule	Remarks
(0, 0)	$x_{2i+1} = x_{2i-1} + 1$	Escalate; stay if $x_{2i-1} = K - 1$
(0, 1)	$x_{2i+1} = x_{2i-1} + 1$	Escalate; stay if $x_{2i-1} = K - 1$
(0, 2)	$x_{2i+1} = x_{2i-1}$	Coherence restriction
(1, 0)	$x_{2i+1} = x_{2i-1}$	
(1, 1)	$x_{2i+1} = x_{2i-1}$	
(1, 2)	$x_{2i+1} = x_{2i-1} - 1$	De-escalate; stay if $x_{2i-1} = 1$
(2, 0)	$x_{2i+1} = x_{2i-1}$	Coherence restriction
(2, 1)	$x_{2i+1} = x_{2i-1}$	
(2, 2)	$x_{2i+1} = x_{2i-1} - 1$	De-escalate; stay if $x_{2i-1} = 1$

Table 4: Comparison of the tandem design, the Thall-Cook method (PO and CR), and the optimality bounds. Relative efficiency is calculated as the ratio of the accuracy index  $\mathcal{A}$  of a method to that of the benchmark

	Method	Percent of Selecting Dose					None	$\mathcal{A}$	Relative efficiency
		1	2	3	4	5			
		$(F_R, F_T)$	(.20, .02)	(.40, .03)	(.60, .04)	(.68, .06)	(.74, .20)		
		$\delta$	-0.48	-0.13	0.22	0.32	-0.26		
Scenario 3 in Thall and Cook (2004)	Tandem		0.1	0.2	12.6	76.1	9.1	2.0	.91
	PO		0.0	0.4	19.8	71.6	6.5	1.7	.92
	CR		0.0	1.6	32.2	49.4	15.7	1.0	.83
	$d_{ET}(\pi^*)$		0.0	0.0	13.0	84.8	1.2	0.0	.97
		$(F_R, F_T)$	(.52, .01)	(.62, .015)	(.71, .02)	(.79, .025)	(.86, .03)		
		$\delta$	0.12	0.29	0.45	0.58	0.69		
Scenario 4 in Thall and Cook (2004)	Tandem		1.3	0.4	2.5	28.8	67.0	0.0	.93
	PO		0.1	1.7	10.1	34.3	53.7	0.0	.90
	CR		0.0	0.1	1.1	4.6	94.0	0.1	.99
	$d_{ET}(\pi^*)$		0.0	0.0	0.4	4.5	95.1	0.0	.99

because  $Y_{2i} = 2$ . It can be shown that under the constraints (5.2), (5.3), and (5.4), the only possible dose assignment for the next pair is  $(x_{2i+1}, x_{2i+2}) = (x_{2i-1}, x_{2i})$ . This example illustrates how coherence can be applied to simplify clinical decisions that are ethically defensible. This type of clarity is particularly important in trials with complex dose-outcome structures. The tandem design described above is an algorithm-based design while most dose finding methods for phase I/II trials are model-based. For example, Thall and Cook (2004) consider two dose-outcome models to model  $F_T$  and  $F_R$  jointly: the proportional odds (PO) model and the continuation ratio (CR) model, and apply a model-based design that treats the next group of patients at a dose with the largest model-based estimate of desirability, subject to some acceptability constraint; see Thall and Cook (2004) for further model and algorithm details.

It is typical to use simulation to compare different dose finding methods. Table 4 gives the simulation results of the tandem design, the Thall-Cook methods with the PO and the CR models, along with the benchmark  $d_{ET}(\pi^*)$  for a trial with  $n = 72$ . The tandem design uses the escalation rules specified in Table 3 along with (5.2), and estimate  $d_{ET}(\pi)$  using logistic regression at the end of the trial. The results of the tandem design and the benchmark are based on 10,000 simulation replicates, and the results of the PO and the CR models are obtained from Thall and Cook (2004). The methods are compared based on an accuracy index  $\mathcal{A}$  calculated as a standardized, expected desirability of the selected dose:

$$\frac{\sum_{k=1}^K \delta_k \Pr(\hat{d} = k) - \min_k \delta_k}{\max_k \delta_k - \min_k \delta_k},$$

where  $\hat{d}$  denotes the selected dose by a method. The index  $\mathcal{A}$  is standardized so that it equals 0 (1) when the worst (best) dose is selected with probability one (Cheung *et al.*, 2015).

If we focus on the relative performance, we cannot tell apart the PO and the CR models because the former performs better than the latter under Scenario 3, and the converse holds under Scenario 4. This is possible because it is unlikely that a method or a model will perform *uniformly* better than another. However, when we look at the optimality bounds, the benchmark  $d_{ET}(\pi^*)$  has better accuracy (0.99) in Scenario 4 than in Scenario 3 (0.97), suggesting Scenario 4 is an “easier” scenario for estimation purposes. In this light, it is interesting to note that the PO model performs poorer in an easy scenario ( $\mathcal{A} = 0.90$ ) than in a less easy one ( $\mathcal{A} = 0.92$ ), which may be indicative of an intrinsic idiosyncrasy of the PO model in this context. This idea is analogous to unbiasedness in hypothesis testing where we expect a larger power under a greater effect size. In other words, the benchmark accuracy can be viewed as a numerical summary of the “effect size” of a dose-outcome scenario, thus providing a yardstick for what we should expect from a reasonable dose finding method.

## 6. Discussion

In this article, we review the theoretical concepts in dose finding with some details on coherence and optimality bounds. While these concepts were first developed for the standard phase I problem, we discuss their extension to general dose finding settings. Since we aim to cover theoretical developments after 2010, we have not covered all theoretical criteria in dose finding such as rigidity and unbiasedness, for which we refer readers to Cheung (2010).

In this article, we also give some important examples of “non-standard” dose finding problems including dose finding with trinary outcomes and multiple toxicity constraints. We, however, have not covered many important topics such as the use of time-to-event outcomes (Cheung and Chappell, 2000; Polley, 2011), testing drug combinations (Houede *et al.*, 2010), and dose finding with patient heterogeneity. However, the principles of estimation objective and design strategy can be applicable to these other situations.

Statistical methodology for dose finding trials is an application-oriented discipline. As a result, little attention has been given to studying the general theory of dose finding methods. This, however, falls short of the rigor necessitated by the ethics and scientific merits of a clinical study, and may miss the opportunity for developing powerful *theory-guided design diagnostic and calibration tools*. Examples abound for the standard phase I dose finding problems: Cheung (2013) uses the optimality bounds to derive a sample size formula for the CRM; Jia *et al.* (2014) provide tools based on the coherence principles to choose the initial sequence  $\{x_{i0}\}$  in a hybrid design; and Lee and Cheung (2009) develop an algorithm to fine-tune a CRM model to achieve  $h$ -sensitivity for any given  $h$ . In this article, we demonstrate the use of optimality benchmark in a quick diagnosis of the idiosyncrasy of the PO model. These are all tangible design tools that would not be possible without careful study of dose finding theory. As dose finding trials become increasingly sophisticated, complex, and study-specific, it is ever more important to be able to have a general theoretical framework that enhances the simplification and transparency of the dosing decisions by novel methods. The goal of this paper is to illustrate how these theoretical criteria can be useful beyond the standard phase I problem, and to inspire further research along this line.

## Acknowledgement

I would like to thank Mr. Jimmy Duong for proofreading of an early version of this article.

## References

- Anbar, D. (1984). Stochastic approximation methods and their use in bioassay and phase I clinical trials, *Communications in Statistics - Theory and Methods*, **13**, 2451–2467.
- Azriel, D., Mandel, M. and Rinott, Y. (2011). The treatment versus experimentation dilemma in dose finding studies, *Journal of Statistical Planning and Inference*, **141**, 2759–2768.
- Bekele, B. N. and Thall, P. F. (2004). Dose-finding based on multiple toxicities in a soft tissue sarcoma trial, *Journal of the American Statistical Association*, **99**, 26–35.
- Braun, T. M. (2002). The bivariate continual reassessment method: Extending the CRM to phase I trials of two competing outcomes, *Controlled Clinical Trials*, **23**, 240–256.
- Cheung, Y. K. (2002). On the use of nonparametric curves in phase I trials with low toxicity tolerance, *Biometrics*, **58**, 237–240.
- Cheung, Y. K. (2005). Coherence principles in dose-finding studies, *Biometrika*, **92**, 863–873.
- Cheung, Y. K. (2007). Sequential implementation of stepwise procedures for identifying the maximum tolerated dose, *Journal of the American Statistical Association*, **102**, 1448–1461.
- Cheung, Y. K. (2010). Stochastic approximation and modern model-based designs for dose-finding clinical trials, *Statistical Science*, **25**, 191–201.
- Cheung, Y. K. (2011). *Dose Finding by the Continual Reassessment Method*, CRC Press/Taylor & Francis Group, Boca Raton, FL.
- Cheung, Y. K. (2013). Sample size formulae for the Bayesian continual reassessment method, *Clinical Trials*, **10**, 852–861.
- Cheung, Y. K. (2014). Simple benchmark for complex dose finding studies, *Biometrics*, **70**, 389–397.
- Cheung, Y. K. (2015). *Dose Finding in Tandem to Identify the Maximum Response*, Unpublished manuscript.
- Cheung, Y. K., Chakraborty, B. and Davidson, K. W. (2015). Sequential multiple assignment randomized trial (SMART) with adaptive randomization for quality improvement in depression treatment program, *Biometrics*, **71**, 450–459.
- Cheung, Y. K. and Chappell, R. (2000). Sequential designs for phase I clinical trials with late-onset toxicities, *Biometrics*, **56**, 1177–1182.
- Cheung, Y. K. and Chappell, R. (2002). A simple technique to evaluate model sensitivity in the continual reassessment method, *Biometrics*, **58**, 671–674.
- Cheung, Y. K. and Elkind, M. S. (2010). Stochastic approximation with virtual observations for dose-finding on discrete levels, *Biometrika*, **97**, 109–121.
- Durham, S. D., Flournoy, N. and Rosenberger, W. F. (1997). A random walk rule for phase I clinical trials, *Biometrics*, **53**, 745–760.
- Finney, D. J. (1978). *Statistical Methods in Biological Assay*, Griffin, London.
- Gasparini, M. and Eisele, J. (2000). A curve-free method for phase I clinical trials, *Biometrics*, **56**, 609–615.
- Haines, L. M., Perevozskaya, I. and Rosenberger, W. F. (2003). Bayesian optimal designs for phase I clinical trials, *Biometrics*, **59**, 591–600.
- Houede, N., Thall, P. F., Nguyen, H., Paoletti, X. and Kramar, A. (2010). Utility-based optimization of combination therapy using ordinal toxicity and efficacy in phase I/II trials, *Biometrics*, **66**, 532–540.
- Jia, X., Lee, S. M. and Cheung, Y. K. (2014). Characterization of the likelihood continual reassessment method, *Biometrika*, **101**, 599–612.
- Lee, S. M., Cheng, B. and Cheung, Y. K. (2011). Continual reassessment method with multiple toxicity constraints, *Biostatistics*, **12**, 386–398.

- Lee, S. M. and Cheung, Y. K. (2009). Model calibration in the continual reassessment method, *Clinical Trials*, **6**, 227–238.
- Lee, S. M. and Cheung, Y. K. (2011). Calibration of prior variance in the Bayesian continual reassessment method, *Statistics in Medicine*, **30**, 2081–2089.
- Lee, S. M., Hershman, D. L., Martin, P., Leonard, J. P. and Cheung, Y. K. (2012). Toxicity burden score: A novel approach to summarize multiple toxic effects, *Annals of Oncology*, **23**, 537–541.
- Leung, D. H. Y. and Wang, Y. G. (2001). Isotonic designs for phase I trials, *Controlled Clinical Trials*, **22**, 126–138.
- 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.
- O’Quigley, J. and Chevret, S. (1991). Methods for dose finding studies in cancer clinical trials: A review and results of a Monte Carlo study, *Statistics in Medicine*, **10**, 1647–1664.
- O’Quigley, J., Hughes, M. D. and Fenton, T. (2001). Dose-finding designs for HIV studies, *Biometrics*, **57**, 1018–1029.
- O’Quigley, J., Paoletti, X. and Maccario, J. (2002). Non-parametric optimal design in dose finding studies, *Biostatistics*, **3**, 51–56.
- O’Quigley, J., Pepe, M. and Fisher, L. (1990). Continual reassessment method: A practical design for phase I clinical trials in cancer, *Biometrics*, **46**, 33–48.
- O’Quigley, J. and Shen, L. Z. (1996). Continual reassessment method: A likelihood approach, *Biometrics*, **52**, 673–684.
- Polley, M. Y. C. (2011). Practical modifications to the time-to-event continual reassessment method for phase I cancer trials with fast patient accrual and late-onset toxicities, *Statistics in Medicine*, **30**, 2130–2143.
- Robbins, H. and Monro, S. (1951). A stochastic approximation method, *Annals of Mathematical Statistics*, **22**, 400–407.
- Shen, L. Z. and O’Quigley, J. (1996). Consistency of the continual reassessment method under model misspecification, *Biometrika*, **83**, 395–405.
- Storer, B. E. (1989). Design and analysis of phase I clinical trials, *Biometrics*, **45**, 925–937.
- Storer, B. E. and DeMets, D. (1987). Current phase I/II designs: Are they adequate?, *Journal of Clinical Research and Drug Development*, **1**, 121–130.
- Thall, P. F. and Cook, J. D. (2004). Dose-finding based on efficacy-toxicity trade-offs, *Biometrics*, **60**, 684–693.
- Wu, C. J. (1985). Efficient sequential designs with binary data, *Journal of the American Statistical Association*, **80**, 974–984.
- Yuan, Z., Chappell, R. and Bailey, H. (2007). The continual reassessment method for multiple toxicity grades: A Bayesian quasi-likelihood approach, *Biometrics*, **63**, 173–179.
- Zacks, S., Rogatko, A. and Babb, J. (1998). Optimal Bayesian-feasible dose escalation for cancer phase I trials, *Statistics & Probability Letters*, **38**, 215–220.