Published in final edited form as:

Clin Trials. 2017 October ; 14(5): 491–498. doi:10.1177/1740774517722760.

Isotonic Designs for Phase I Trials in Partially Ordered Groups

Mark Conaway¹

¹The University of Virginia Health System

Abstract

Background/Aims—Dose finding trials can be conducted such that patients are first stratified into multiple risk groups before doses are allocated. The risk groups are often completely ordered in that, for a fixed dose, the probability of toxicity is monotonically increasing across groups. In some trials, the groups are only partially ordered. For example, one of several groups in a trial may be known to have the least risk of toxicity for a given dose, but the ordering of the risk among the remaining groups may not be known. The aim of the paper is to introduce a method for designing dose-finding trials of cytotoxic agents in completely or partially ordered groups of patients.

Methods—The paper presents a method for dose-finding that combines previously proposed mathematical models, augmented with results using order restricted inference. The resulting method is computationally convenient and allows for dose finding in trials with completely or partially ordered groups. Extensive simulations are done to evaluate the performance of the method, using randomly generated dose-toxicity curves where, within each group, the risk of toxicity is an increasing function of dose.

Results—Our simulations show that the hybrid method, in which order restricted estimation is applied to parameters of a parsimonious mathematical model, gives results that are similar to previously proposed methods for completely ordered groups. Our method generalizes to a wide range of partial orders among the groups.

Conclusions—The problem of dose-finding in partially ordered groups has not been extensively studied in the statistical literature. The proposed method is computationally feasible, and provides a potential solution to the design of dose finding studies in completely or partially ordered groups.

Keywords

partial order; multiple risk groups; dose-finding

Introduction

Yuan and Chappell proposed several methods for dose finding trials in which patients are first stratified into multiple risk groups¹. The risk groups are completely ordered in the sense of Robertson et al.² in that for a fixed dose, the probability of a dose-limiting toxicity (DLT) is monotonically increasing across groups. There are several examples of such trials in the

Reprints and permission: sagepub.co.uk/journalsPermissions.nav

Corresponding author: Mark Conaway, Division of Translational Research and Applied Statistics Department of Public Health Sciences, Box 800717, Charlottesville, VA, 22901, USA. mconaway@eservices.virginia.edu.

literature, including Ramanathan et al.³ and LoRusso et al.⁴, who stratify patients into ‘none,’ ‘mild,’ ‘moderate,’ or ‘severe’ liver dysfunction at baseline. Leal et al. divide patients into groups defined by increasing renal dysfunction at baseline based on a measure of 24-hour creatinine clearance.⁵ Kim et al. genotyped *UGT1A1* *28 and *6, and defined patient cohorts by the numbers of defective alleles: 0, 1, or 2.⁶ Satoh et al. also investigated dose-finding with groups defined by *UGT1A1* genotypes.⁷ In this study, patients were divided into three groups according to wild-type (*1*1), heterozygous (*28*1, *6*1) or homozygous (*28*28, *6*6, *28*6). Each of these studies is an example of “completely ordered” groups. For a given dose of the agent, we would anticipate a greater chance that a patient with a greater degree of organ dysfunction at baseline or a greater number of defective alleles would experience a DLT.

There are fewer published examples of partially ordered groups, but these would arise naturally when one of K groups in a trials may be known to have the greatest risk of toxicity for a given dose, but the ordering of the risk among the remaining $K - 1$ groups may not be known. Innocenti et al. provides an example of this in a genotype-directed dose-finding study of irinotecan.⁸ Based on the observation that the risk of severe neutropenia with irinotecan is related in part to *UGT1A1* *28, Innocenti et al.⁸ sought to identify the maximum-tolerated dose (MTD) of irinotecan in patients with advanced solid tumors stratified by the *1/*1, *1/*28, and *28/*28 genotypes. The greatest risk of a DLT was expected to be in patients with the *28/*28 genotype. In other trials, patients might be stratified into groups based on the presence or absence of two genetic mutations. For a fixed dose of an agent, those without either mutation might have the least risk and those with the both mutations might have the greatest risk. The ordering between the probability of toxicity at a given dose for patients with one mutation might not be known.

To our knowledge, the only paper that addresses dose-finding in more than two groups is Yuan and Chappell,¹ although there have been several methods proposed for dose-finding in two groups of patients. O’Quigley et al. introduced the two-sample continual reassessment method, which allowed for the identification of the appropriate MTD’s for two groups simultaneously.⁹ O’Quigley and Paoletti proposed a two-parameter continual reassessment method for ordered groups that utilizes known differences between the groups.¹⁰ Ivanova and Wang incorporate isotonic regression estimates into designs for ordered groups that take into account both toxicity and efficacy endpoints.¹¹ Wages et al. describe the design of a dose-finding trial that explicitly uses the known ordering in the probabilities associated with ‘good’ and ‘poor’ prognosis patients.¹² Their design is based on the shift model^{13,14} that generalizes the continual reassessment method to two ordered groups. Conaway and Wages¹⁵ propose two methods based on the estimation method of Hwang and Peddada.¹⁶

In this paper, we first propose a simple modification of the “isotonic regression continual reassessment method” proposed by Yuan and Chappell¹ for completely ordered groups, and show that the properties of this modification are the same as those of the original method. The advantage of this modification is that it generalizes naturally to partially ordered groups.

Complete Orders

Review of the original Yuan-Chappell method

In Yuan and Chappell,¹ there are J dose levels under consideration, and the goal of the study is to find the “maximum tolerated dose” (MTD) in each of G risk groups. By definition, the MTD in group g is highest dose under consideration that can be administered to patients in that group with an “acceptable level” of toxicity. “Acceptable” toxicity means limiting the proportion of patients in that group who experience a sufficiently severe, protocol-specified adverse event, usually called a ‘dose-limiting toxicity’ (DLT). With cytotoxic agents, it is generally assumed that regardless of the group, the greater the dose administered, the greater the probability that a patient will experience a DLT.

The isotonic regression continual reassessment method proposed in Yuan and Chappell¹ is built on the continual reassessment method¹⁷. Originally proposed for a single risk group, the continual reassessment method assigns doses adaptively, with the dose for the current patient determined by the observed toxicities and dose allocations from the prior patients on the trial. Specifically, at any point in the trial, we have observed n_j patients on dose level j , $j = 1, \dots, J$ and among these patients, we have observed Y_j patients experiencing a DLT. The number of patients with a DLT on dose level j , Y_j , are distributed as independent $\text{Bin}(n_j, \pi_j)$, where π_j is the probability that a patient will experience a DLT given dose level j . The continual reassessment method uses a working model (1) for the π_j

$$\pi_j = \psi_j^{\exp(a)} \quad (1)$$

where $0 < \psi_1 < \psi_2 < \dots < \psi_J < 1$ are pre-specified constants, often referred to as the ‘skeleton’ values and a is a scalar parameter. The parametrization $\exp(a)$ ensures that the probability of toxicity is increasing in dose for all $-\infty < a < \infty$. Advice on choosing skeleton values is given in Lee et al.¹⁸

With these specifications, at any point in the trial, the likelihood from the previously observed DLTs and dose allocations is

$$L(a; n_j, y_j) = \prod_{j=1}^J \psi_j^{y_j \exp(a)} (1 - \psi_j^{\exp(a)})^{n_j - y_j} \quad (2)$$

Current estimates of the DLT probabilities are $\hat{\pi}_j = \psi_j^{\exp(\hat{a})}$, where $(\hat{a})$ is found by maximizing (2) as a function of a . Based on the current estimates, the next patient is assigned to the dose level j^* with the smallest value of $(\hat{\pi}_j, \theta)$, where (u, v) is a pre-specified measure of distance, usually quadratic distance, between u and v . The parameter θ is a target toxicity probability specified prior to the trial. The process of updating of a and allocating patients to doses continues until a pre-specified number of patients have been observed. At the end of the study, the MTD is taken to be the dose that the next patient would have received had the trial not ended.

Yuan and Chappell generalize the continual reassessment method to G completely ordered groups, where by convention, group 1 is the lowest risk group and patient in group G have the highest risk of a DLT.¹ At any point in the trial, within group g , we have observed n_{gj} patients allocated to dose level j and of these, Y_{gj} patients have experienced a DLT. Assuming skeleton values $0 < \psi_1 < \psi_2 < \dots < \psi_J < 1$, the working model for the probability π_{gj} that a patient in risk group g , given dose level j , experiences a DLT is given by

$$\pi_{gj} = \psi_j^{\exp(a_g)} \quad (3)$$

for $j = 1, \dots, J$ and $g = 1, \dots, G$. The likelihood as a function of $(a_1, a_2, \dots, a_G)$ is

$$L(\{a_g\}; n_{gj}, y_{gj}) = \prod_{g=1}^G \prod_{j=1}^J \psi_j^{y_{gj} \exp(a_g)} (1 - \psi_j^{\exp(a_g)})^{n_{gj} - y_{gj}} \quad (4)$$

Estimates $(\hat{a}_1, \hat{a}_2, \dots, \hat{a}_G)$ can be found by maximizing (4), and from these, we can obtain the current estimates of π_{gj} ,

$$\hat{\pi}_{gj} = \psi_j^{\exp(\hat{a}_g)} \quad (5)$$

The skeleton values and the working model guarantee that within groups, the estimated probabilities of toxicity are increasing with dose. Across groups, however, there may be “reversals”, meaning that at a given dose, the probability of a DLT is estimated to be greater in a lower risk group than in a higher risk group. To eliminate these reversals, Yuan and Chappell¹ create estimates $\hat{\pi}_{gj}^*$ by applying the 2-way isotonic regression procedure of Robertson et al.² to the model-based estimates, $\hat{\pi}_{gj}$, from (5). The 2-way isotonic regression estimation is an iterative process of adjusting the model-based estimates within rows and within columns to guarantee that the estimated probabilities of a DLT have no reversals, i.e., the estimates are non-decreasing within rows and within columns. As in the original continual reassessment method, if the next patient is from group g , that patient is assigned the dose level with the “isotonized” estimate closest to the target toxicity probability, θ , using a quadratic distance measure. The process of estimation and allocation continues until a pre-specified number of patients have been observed in each group.

A modification

As in the original Yuan and Chappell¹ method, we obtain the estimate $\hat{a} = (\hat{a}_1, \hat{a}_2, \dots, \hat{a}_G)$ by maximizing (4), and we let $v_j(\hat{a})$ denote the j^{th} diagonal element of the observed information matrix based on the log of the likelihood (4) evaluated at $\hat{a}$. Instead of estimating the group-dose DLT probabilities from the model and using 2-way isotonic regression, we apply the estimation procedure of Hwang and Peddada¹⁶ to $\hat{a} = (\hat{a}_1, \hat{a}_2, \dots, \hat{a}_G)$ using weights $\frac{1}{v_j(\hat{a})}$ to obtain estimates subject to the following restriction,

$$\hat{a}_1 \geq \hat{a}_2 \geq \dots \geq \hat{a}_G \quad (6)$$

We denote the resulting estimates by $\hat{a}^{HP} = (\hat{a}_1^{HP}, \hat{a}_2^{HP}, \dots, \hat{a}_G^{HP})$, which satisfy

$$\hat{a}_1^{HP} \geq \hat{a}_2^{HP} \geq \hat{a}_G^{HP} \quad (7)$$

After “isotonizing” the estimates, we compute estimates of the probabilities of a DLT for dose d in group g ,

$$\hat{\pi}_{gj}^{HP} = \psi_j^{\exp(\hat{a}_{gj}^{HP})} \quad (8)$$

Within groups, the estimated probabilities of a DLT increase with dose and for a fixed dose, the probabilities increase across groups. While Hwang and Peddada¹⁶ provide a general estimation procedure for parameters subject to partial orders, in the case of a complete order, the estimation procedure is identical to using isotonic regression on $-\hat{a} = (-\hat{a}_1, -\hat{a}_2, \dots, -\hat{a}_G)$ using weights $\frac{1}{v_j(\hat{a})}$, as in Robertson et al.². We would argue that this is a simpler method

than the original to avoid reversals in the DLT probability estimates. In general, the number of groups is small, and the modification requires isotonic regression of only a small number of parameters. In order to make sure we have finite estimates $\hat{a}_g$, we use a pseudo-data prior as in Whitehead and Brunier¹⁹ or Legedza et al.²⁰ with total weight equal to 1 observation per group.

Comparing the original and the modification using a family of dose-toxicity curves

We denote the true MTD in group g by MTD_g . With G groups and J dose levels under consideration, complete ordering of the groups implies that

$$J \geq MTD_1 \geq MTD_2 \geq \dots \geq MTD_G \geq 1 \quad (9)$$

We generated 2 scenarios for each of the possible configurations of the MTDs consistent with (9). Table 1 lists the number of such configurations for 3 or 4 groups and 4 or 6 dose levels.

Letting $\gamma = (\gamma_1, \gamma_2, \gamma_3)$ denote one of these configurations, we generated group-dose toxicity probabilities π_{gj} , $g = 1, \dots, G$ and $j = 1, \dots, J$ starting from group 1. We define $\pi_{g,0} = 0$ and $\pi_{0,j} = 0$. In group g , for $j = \gamma_g$, we set $\pi_{g,j} = \theta$. To generate probabilities for $j < \gamma_g$, we define ‘Lower’ to equal $\max\{\pi_{g-1,j}, \pi_{g,j-1}\}$. For $j < \gamma_g$, we define ‘Upper’ to equal the target, θ and for $j > \gamma_g$, we define Upper to equal 1. With these definitions, we sequentially generated the true group-dose toxicity probabilities, $\pi_{g,j}$ as

$$\pi_{g,j} = \text{Lower} + (\text{Upper} - \text{Lower})\beta(1,4)$$

where $\beta(1, 4)$ are independently generated *beta* distributed random variables with mean 0.2 and variance 0.0267. This method of generating the group-dose toxicity probabilities guarantees that the resulting probabilities will follow the complete order.

For each scenario, we simulated 1000 trials with either a total of 36 or 54 patients, and either either equal or unequal group probabilities. For 3 groups, the group probabilities were equal 0.45, 0.40, 0.15 to match the approximate group proportions in Innocenti et al.⁸ For 4 groups, the group probabilities were chosen to equal 0.4, 0.3, 0.2, 0.1. For each combination of sample size and group probabilities, we computed for both the original and the modified method, the ‘percent correct selection’, defined as the proportion of times that the method correctly identifies the MTD in each group. We also computed the “accuracy index” of Cheung²¹ within groups. Figure 1 displays the distribution of differences in the percent correct selection, modified minus original and Figure 2 displays the distribution of differences in the accuracy index. The figures show a small difference between the original and modified methods with, in general, the percent correct selection and the accuracy index with the original method greater than for the modification.

Partial Orders

The results of the previous section indicate that for complete orders, applying order restricted estimation methods to $\hat{a} = (\hat{a}_1, \hat{a}_2, \dots, \hat{a}_G)$, found by maximizing (4) has similar properties to the original design proposed by Yuan and Chappell¹. By using $\hat{a}$, however, the extension to partially ordered groups is straightforward. Table 2 lists a number of partial orders that could be encountered in 3 or 4 groups, along with the corresponding order restrictions satisfied by the Hwang and Peddada¹⁶ estimates.

We use the study in Innocenti et al.⁸ to illustrate the performance of this method. In this study, there were 5 dose levels of irinotecan, and patients were stratified by the $*1/*1$, $*1/*28$, and $*28/*28$ genotypes. This is an example of the first ordering in Table 2; patients with the $*28/*28$ genotype were expected to have a greater risk of severe neutropenia than either of the other genotypes. Of the 68 patients enrolled, 31 (46%) were genotype $*1/*1$, 28 (41%) had genotype $*1/*28$, and 9 (13%) were in the $*28/*28$ group. In the complete order setting, we had randomly generated 2 scenarios for each of the 35 possible MTD configurations consistent with the complete order. To convert this to a partial order, we randomly chose one of the two scenarios for each of the possible MTD configurations, and interchanged the probabilities associated with the $*1/*1$ and $*1/*28$ groups. The other scenario in the pair was left unchanged. For each scenario we simulated 1000 trials of 36, 54 and 72 patients under the assumption that the genotypes each occur with probability 1/3 in the population. We also simulated trials under the assumption that the population proportions are 45% for genotype $*1/*1$, 40% for genotype $*1/*28$ and 15% for genotype $*28/*28$, similar to those observed in Innocenti et al.⁸. The average PCS and AI values are given in Table 3. We do not have a method with which to compare these results, but we can note that these values are comparable to the averages for the 3×5 case displayed in Figure 1

and Figure 2. As expected, the values are lower for the partial order, where there is less information about the groups than for the complete order.

An example of the dose allocations is given in Table 4 and is plotted in Figure 3 as in O’Quigley²². In this plot, open circles are patients who do not experience a DLT; closed circles represent DLTs. This trial was simulated based on the results in Innocenti et al.⁸. Unequal group probabilities of 0.45, 0.40 and 0.15 were used, with toxicity probabilities of (0.005, 0.02, 0.03, 0.2, 0.33) in group 1, (.005, .02, 0.20, 0.30, 0.40) in group 2 and (0.20, 0.35, 0.50, 0.60, 0.70) in group 3. The final recommendation for this example trial is level 5 in group 1, level 3 in group 2 and level 1 in group 3.

Summary and Future Directions

This paper presents a modification of the Yuan and Chappell¹ method that can be applied for completely or partially ordered groups. For complete orders, the simulation results suggest that the properties of the modified estimates are similar to those originally proposed by Yuan and Chappell. With a known ordering between 2 groups, Conaway and Wages¹⁵ show that the original method has excellent properties for 6 or 8 dose levels. The modification applies the Hwang-Peddada estimation method to the estimated parameters in the one parameter model continual reassessment method O’Quigley et al.¹⁷. While the method of Hwang and Peddada¹⁶ is computationally simple, and produces estimates with good statistical properties, it is not a well known algorithm. As an alternative, we are exploring the use of laying out multiple complete orders consistent with the partial order, and using model selection as in Wages et al.²³ to estimate which complete order is most consistent with the accumulated data. For example, if the partial order on the parameters were $\hat{a}_1^{HP} \geq \hat{a}_3^{HP}$ and $\hat{a}_2^{HP} \geq \hat{a}_3^{HP}$, we would postulate that either $\hat{a}_1^{HP} \geq \hat{a}_2^{HP} \geq \hat{a}_3^{HP}$ or $\hat{a}_2^{HP} \geq \hat{a}_1^{HP} \geq \hat{a}_3^{HP}$. There are several well-known algorithms, such as as the pool-adjacent-violators algorithm (PAVA) or the minimum lower sets algorithm Robertson et al.² that would produce estimates subject to each of these complete orders. Once the ordered estimates are computed, we would estimate the dose-toxicity probabilities using the set of ordered estimates that yields the larger value of (4). Further work is needed on the properties of this alternative method.

Additional research is also needed on calibrating the pseudo-data to ensure that dose escalation does not occur too rapidly early in the trial. An alternative is to do a 2-stage design as in O’Quigley and Paoletti¹⁰, where the first stage is a rule-based design that follows the known orderings among the groups. The modeling stage of our modification of Yuan and Chappell¹ would begin only after heterogeneity in responses is observed in each of the groups.

Acknowledgments

I would like to thank the two reviewers and the associate editor for their careful review of the manuscript; incorporating their comments greatly improved the manuscript. Research reported in this publication was supported by the National Cancer Institute of the National Institutes of Health under award number R01CA142859 and P30CA044579. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

References

1. Yuan Z, Chappell R. Isotonic designs for phase I cancer clinical trials with multiple risk groups. *Clin Trials*. 2004; 1:499–508. [PubMed: 16279290]
2. Robertson, T., Wright, FT., Dykstra, R. Order Restricted Statistical Inference. New York: J. Wiley and Sons; 1988.
3. Ramanathan R, Egorin M, Takimoto C, Remick S, Doroshow J, LoRusso P, Mulkerin D, Grem J, Hamilton A, Murgo A, Potter D, Belani C, Hayes M, Peng B, Ivy P. Phase I and Pharmacokinetic Study of Imatinib Mesylate in Patients With Advanced Malignancies and Varying Degrees of Liver Dysfunction: A Study by the National Cancer Institute Organ Dysfunction Working Group. *J Clin Oncol*. 2008; 26:563–569. [PubMed: 18235115]
4. LoRusso P, Venkatakrishnan K, Ramanathan R, Sarantopoulos J, Mulkerin D, Shibata S, Hamilton A, Dowlati A, Mani S, Rudek M, Takimoto C, Neuwirth R, Esseltine D, Ivy P. Pharmacokinetics and Safety of Bortezomib in Patients with Advanced Malignancies and Varying Degrees of Liver Dysfunction: Phase I NCI Organ Dysfunction Working Group Study NCI-6432. *Clin Cancer Res*. 2012; 18:1–10.
5. Leal T, Remick S, Takimoto C, Ramanathan R, Davies A, Egorin M, Hamilton A, LoRusso P, Shibata S, Lenz H-J, Mier J, Sarantopoulos J, Mani S, Wright J, Ivy P, Neuwirth R, von Moltke L, Venkatakrishnan K, Mulkerin D. Dose-escalating and pharmacological study of bortezomib in adult cancer patients with impaired renal function: a National Cancer Institute Organ Dysfunction Working Group Study. *Cancer Chemoth Pharm*. 2011; 68:1439–1447.
6. Kim K-P, Kim H-S, Sym S, Bae K, Hong Y, Chang H-M, Lee J, Kang Y-K, Lee J, Shin JG, Kim T. A UGT1A1*28 and *6 genotype-directed phase I dose-escalation trial of irinotecan with fixed-dose capecitabine in Korean patients with metastatic colorectal cancer. *Cancer Chemoth Pharm*. 2013; 71:1609–1617.
7. Satoh T, Ura T, Yamada Y, Yamazaki K, Tsujinaka T, Munakata M, Nishina T, Okamura S, Esaki T, Sasaki Y, Koizuma W, Kakeji Y, Ishizuka N, Hyodo I, Sakata Y. Genotype-directed, dose-finding study of irinotecan in cancer patients with UGT1A1*28 and or UGT1A1*6 polymorphisms. *Cancer Sci*. 2011; 102:1868–1873. [PubMed: 21740478]
8. Innocenti F, Schilsky R, Ramirez J, Janisch L, Undevia S, House L, Das S, Wu K, Turcich M, Marsh R, Karrison T, Maitland M, Salgia R, Ratain M. Dose-Finding and Pharmacokinetic Study to Optimize the Dosing of Irinotecan According to the UGT1A1 Genotype of Patients With Cancer. *J Clin Oncol*. 2014; 32:2328–2334. [PubMed: 24958824]
9. O'Quigley J, Shen L, Gamst A. Two Sample Continual Reassessment Method. *J Biopharm Stat*. 1999; 9:17–44. [PubMed: 10091908]
10. O'Quigley J, Paoletti X. Continual Reassessment Method for Ordered Groups. *Biometrics*. 2003; 59:430–440. [PubMed: 12926728]
11. Ivanova A, Wang K. Bivariate isotonic design for dose-finding with ordered groups. *Stat Med*. 2006; 25:2018–2026. [PubMed: 16220476]
12. Wages N, Read P, Petroni G. A Phase I/II adaptive design for heterogeneous groups with application to a stereotactic body radiation therapy trial. *Pharm Stat*. 2015; 14:302–310. [PubMed: 25962576]
13. O'Quigley, J. Phase I and Phase I/II Dose Finding Algorithms Using Continual Reassessment Method. In: Crowley, J., Ankherst, D., editors. *Handbook of Statistics in Clinical Oncology*. 2. Boca Raton FL: Chapman and Hall/CRC Biostatistics Series; 2006. p. 31–58.
14. O'Quigley J, Iasonos A. Bridging Solutions in Dose-Finding Problems. *J Biopharm Stat*. 2014; 6:185–197.
15. Conaway M, Wages N. Designs for phase I trials in ordered groups. *Stat Med*. 2016; 36:254–265. [PubMed: 27624880]
16. Hwang J, Peddada S. Confidence Interval Estimation Subject to Order Restrictions. *Ann Stat*. 1994; 22:67–93.
17. O'Quigley J, Pepe M, Fisher L. Continual Reassessment Method: A Practical Design for Phase I Clinical Trials in Cancer. *Biometrics*. 1990; 46:33–48. [PubMed: 2350571]

18. Lee S, Cheung YK. Model calibration in the continual reassessment method. *Clin Trials*. 2009; 6:227–238. [PubMed: 19528132]
19. Whitehead J, Brunier H. Bayesian Decision Procedures for Dose Finding Studies. *Stat Med*. 1995; 14:885–893.
20. Legedza A, Ibrahim J. Heterogeneity in phase I clinical trials: prior elicitation and computation using the continual reassessment method. *Stat Med*. 2001; 20:867–882. [PubMed: 11252009]
21. Cheung, YK. Dose Finding by the Continual Reassessment Method. Boca Raton FL: Chapman and Hall/CRC Biostatistics Series; 2011.
22. O’Quigley J. Theoretical study of the continual reassessment method. *J Stat Plan Infer*. 2006; 136:1765–1780.
23. Wages N, Conaway M, O’Quigley J. Continual reassessment method for partial ordering. *Biometrics*. 2011; 67:1555–1563. [PubMed: 21361888]

Figure 1.
Comparison of Percent Correct Selection for Original and Modified Methods

Figure 3.
Plot of dose allocations in example trial

Table 1

Possible MTD configurations

Groups	Dose levels	Number
3	4	20
3	5	35
3	6	56
4	4	35
4	5	70
4	6	126

Table 2

Partial Order Examples

Ordering	Parameter Restriction
Group 3 has the greatest risk	$\hat{a}_1^{HP} \geq \hat{a}_3^{HP}$ and
Groups 1 and 2 unknown	$\hat{a}_2^{HP} \geq \hat{a}_3^{HP}$
Group 1 has the least risk	$\hat{a}_1^{HP} \geq \hat{a}_2^{HP}$ and
Groups 2 and 3 unknown	$\hat{a}_1^{HP} \geq \hat{a}_3^{HP}$
Group 1 has the least risk	$\hat{a}_1^{HP} \geq \hat{a}_2^{HP} \geq \hat{a}_4^{HP}$
Group 4 has the greatest risk	$\hat{a}_1^{HP} \geq \hat{a}_3^{HP} \geq \hat{a}_4^{HP}$
Groups 2 and 3 unknown	
Group 1 has the least risk	$\hat{a}_1^{HP} \geq \hat{a}_2^{HP} \geq \hat{a}_3^{HP}$
Group 2 has the next lower risk	$\hat{a}_1^{HP} \geq \hat{a}_2^{HP} \geq \hat{a}_4^{HP}$
Groups 3 and 4 are unknown	

Table 3Average PCS and AI for the 3×5 partial order

N	Equal Group Probabilities		Unequal Group Probabilities	
	PCS	AI	PCS	AI
36	0.465	0.548	0.462	0.545
54	0.513	0.616	0.515	0.615
72	0.555	0.666	0.552	0.664

Table 4Example Trial based on Innocenti et al.⁸

Patient	Group	Dose Level	Toxicity
1	3	1	0
2	3	2	0
3	1	2	0
4	2	3	0
5	1	3	0
6	3	3	0
7	1	4	0
8	3	4	1
9	3	2	0
10	2	3	0
11	1	4	0
12	2	4	1
13	1	5	0
14	3	2	1
15	3	1	0
16	3	1	0
17	1	5	0
18	2	2	0
19	1	5	1
20	3	2	1
21	2	2	0
22	1	4	0
23	3	1	0
24	3	1	1
25	3	1	0
26	1	5	0
27	3	1	0
28	2	3	0
29	1	5	0
30	3	1	0
31	3	1	0
32	1	5	0
33	2	3	0
34	2	3	0
35	3	1	1
36	2	4	1