Published in final edited form as:

J Biopharm Stat. 2017 ; 27(1): 34–43. doi:10.1080/10543406.2016.1148713.

A curve-free Bayesian decision-theoretic design for two-agent Phase I trials

Bee L. Lee^a, Shenghua K. Fan^b, and Ying Lu^{c,d}

^aDepartment of Mathematics and Statistics, San José State University, San José, California, USA

^bDepartment of Statistics and Biostatistics, California State University at East Bay, Hayward, California, USA

^cDepartment of Health Research and Policy, Stanford University, Stanford, California, USA

^dThe Cooperative Studies Program Coordinating Center, VA Palo Alto Health Care System, Palo Alto, California, USA

Abstract

Although Bayesian statistical methods are gaining attention in the medical community, as they provide a natural framework for incorporating prior information, the complexity of these methods limited their adoptions in clinical trials. This article proposes a Bayesian design for two-agent Phase I trials that is relatively easy for clinicians to understand and implement, yet performs comparably to more complex designs, so that it is more likely to be adopted in actual trials. In order to reduce model complexity and computational burden, we choose a working model with conjugate priors so that the posterior distributions have analytical expressions. Furthermore, we provide a simple strategy to facilitate the specification of priors based on the toxicity information accrued from single-agent Phase I trials. The proposed method should be useful in terms of the ease of implementation and the savings in sample size without sacrificing performance. Moreover, the conservativeness of the dose-finding algorithm renders it a relatively safe method.

Keywords

Dose allocation; maximum tolerated dose; Phase I trial; stopping rules; two-agent trials; utility function

1. Introduction

In hopes of improving overall survival in cancer patients, many recent Phase III trials combined two well-known anticancer agents without prior toxicity study of the combinations. That is, instead of conducting Phase I trials of the combinations as a first step, the results of single-agent Phase I trials were extrapolated to define the maximum tolerated dose (MTD), the highest dose combination at which not more than a prespecified proportion of patients experience dose-limiting toxicities (DLTs). Many of these trials failed to establish

the efficacy of the combined agents because “dose modifications and discontinuations as a result of toxicity have contaminated the overall survival results” (Kelley and Venook, 2012). Therefore, Phase I combination trials are necessary even if the toxicity profiles of each agent (when administered alone) have been well established.

Phase I combination trials are far more challenging than single-agent trials. First, the ordering of the DLT probabilities between dose combinations cannot be fully predetermined even in the case of two agents (Wages et al., 2011). Second, the number of combinations grows rapidly as the number of dose levels of one or more agents is increased. In order to study many dose combinations with a small sample that is typical in Phase I trials, the prior information for the component agents, in addition to the observed data, should be utilized. The Bayesian paradigm provides a natural framework to do so. The continual reassessment method (CRM) developed by O’Quigley et al. (1990) was one of the first clinical applications of Bayesian methodology in a dose-finding trial. However, several practical issues have limited the use of the CRM and its variants, such as the difficulty in prior specification and computational complexity. Numerous Bayesian methods have been developed for single-agent trials since the CRM but extending these methods to two-agent trials is nontrivial. Wages et al. (2011) proposed a computationally intensive Bayesian method in which inference is based on a weighted average of the utilities of all possible orderings of the dose combinations. Braun and Jia (2013) adapted the original CRM to two-agent trials by generalizing the one-parameter dose–toxicity curve to a two-parameter curve, with the assumption that the agents have independent toxicities. A comprehensive review of these methods can be found in Harrington et al. (2013).

In this paper, we extend the Bayesian method developed by Fan et al. (2012) for single-agent trials to two-agent trials. The proposed method is curve-free in that no assumption is made about the functional form of the relationship between dose and toxicity. As demonstrated in Carlin et al. (2013), a simple parametric function may not readily fit the true probability of DLT, a drawback not suffered by our method. Because there exists a closed form for the posterior distribution of the DLT probabilities, the computational burden is substantially less than most of the existing Bayesian methods for combination trials. Furthermore, the specification of the priors is relatively straightforward when the marginal distributions of the DLT probabilities are available (estimated from single-agent trials). Although it is always desirable to involve a statistician in the design of a clinical trial, our method can readily be implemented by a clinician with a basic knowledge of statistics, using an R program that is available upon request from the authors.

2. Methods

In this section, we propose a Bayesian decision-theoretic design for two-agent Phase I trials. The key components of the design are a probability model defined on observable data (toxicity responses) and unknown parameters (DLT probabilities), described in Section 2.1; a utility function which measures the values of various dose allocation decisions and a dose selection rule, defined in Section 2.2; and a set of rules for deciding when to stop a trial, set forth in Section 2.3.

2.1. Probability model

Consider a two-agent trial in which agent A has I dose levels and agent B has J dose levels. Let (i, j) denote the combination of agent A at dose level i and agent B at dose level j , and let t_{ij} denote the binary outcome of a patient treated at (i, j) , where $t_{ij} = 1$ ($t_{ij} = 0$) represents the outcome where the patient experiences a DLT (no DLT). We model t_{ij} as a Bernoulli random variable with unknown success (or DLT) probability p_{ij} .

2.1.1. Specifying the prior distribution—The prior distribution of p_{ij} is chosen to be the flexible, conjugate Beta distribution with parameters a_{ij} and b_{ij} , denoted $\text{Beta}(a_{ij}, b_{ij})$. Suppose that $p_{ij} \leq p_{rs}$ whenever $(i, j) < (r, s)$, where $<$ denotes a partial order on the dose combinations. Plausible orderings include:

- $(i, j) < (r, s)$ if and only if $i \leq r, j \leq s$, and $i + j < r + s$, which we shall refer to as strict ordering;
- $(i, j) < (r, s)$ if and only if $i + j < r + s$, which we shall refer to as diagonal ordering.

The initial parameter values, denoted a_{ij}^0 and b_{ij}^0 , should be specified in a way that preserves the ordering of the mean DLT probabilities:

$$\mu_{ij}^0 = \frac{a_{ij}^0}{a_{ij}^0 + b_{ij}^0}; \quad (1)$$

that is, a_{ij}^0 and b_{ij}^0 are chosen such that $\mu_{ij}^0 \leq \mu_{rs}^0$ whenever $(i, j) < (r, s)$.

Before we describe our procedure for prior specification, we recall that a Beta distribution can be characterized by its mean and a credible interval. That is, if p_{ij} has a $\text{Beta}(a_{ij}, b_{ij})$ distribution, μ_{ij} denotes $E(p_{ij})$ and $\Pr(q_{ij}^L \leq p_{ij} \leq q_{ij}^U) = \gamma_0$ for some specified γ_0 , then there is a one-to-one correspondence between (a_{ij}, b_{ij}) and $(\mu_{ij}, q_{ij}^L, q_{ij}^U)$. The specification of prior is now described as follows:

1. *Specification of prior means.* Ideally, the clinician provides his or her best initial guess of μ_{ij} at each dose combination, denoted earlier as μ_{ij}^0 . In the event that the clinician is unable to do so and results from single-agent trials are available, we propose the following as an aid. Let $\tilde{p}_{A_i}$ and $\tilde{p}_{B_j}$ be the estimated marginal DLT probabilities of agent A at dose i and agent B at dose j , respectively, from single-agent trials. A preliminary estimate of μ_{ij}^0 is given by:

$$\tilde{\mu}_{ij}^0 = 1 - \Pr(\text{no toxicity from either agent}) = 1 - (1 - p_{A_i})(1 - p_{B_j}), \quad (2)$$

based on the simplifying assumption that the toxicity of the two agents are independent. The estimates in equation (2) are only preliminary because the order restrictions, $\tilde{\mu}_{ij}^0 \leq \tilde{\mu}_{rs}^0$ whenever $(i, j) < (r, s)$, may not hold for certain partial orders, such as the diagonal ordering. In this case, the clinicians should adjust the estimates to satisfy the order restrictions, and the resulting values will serve as μ_{ij}^0 .

2. *Specification of quantiles.* We propose specifying the first and third quartiles of the Beta priors since quartiles are easily understood by clinicians. That is, the pairs (q_{ij}^L, q_{ij}^U) satisfy:

$$\Pr(p_{ij} \leq q_{ij}^L) = \Pr(p_{ij} \geq q_{ij}^U) = 0.25. \quad (3)$$

To facilitate the specification, several pairs of (q_{ij}^L, q_{ij}^U) are generated for each dose combination as follows. If $\mu_{ij}^0 \geq 0.5$, we fix $a_{ij}^0 = 1, 2, \dots$ and solve for the corresponding b_{ij}^0 using equation (1); otherwise, we fix $b_{ij}^0 = 1, 2, \dots$ and solve for the corresponding a_{ij}^0 using equation (1). The resulting (a_{ij}^0, b_{ij}^0) are used to solve for (q_{ij}^L, q_{ij}^U) . A numerical illustration is given in Table 1. For each dose combination (i, j) , the clinician will select a pair of (q_{ij}^L, q_{ij}^U) that best approximates his or her belief. The reason for generating the (q_{ij}^L, q_{ij}^U) pairs according to whether $\mu_{ij}^0 \geq 0.5$ or $\mu_{ij}^0 < 0.5$ is to ensure that a_{ij}^0 and b_{ij}^0 are small so that the prior is not overly informative.

2.1.2. Computing the posterior distribution—Although a patient is normally treated at a single dose combination in a clinical trial, his or her response can be extrapolated to certain dose combinations based on the assumption that $p_{ij} = p_{rs}$ whenever $(i, j) < (r, s)$. Specifically, patients who experienced a DLT at (i, j) are expected to do so when treated at (r, s) for $(r, s) > (i, j)$; likewise, patients who tolerated (i, j) are expected to do so when treated at (r, s) for $(r, s) < (i, j)$. For example, consider a study of combinations of $I=3$ doses of agent A and $J=4$ doses of agent B. Under a strict ordering of dose combinations, if a patient suffers a DLT at (2, 3), the working data for this patient consist of a DLT at (2, 3), (2, 4), (3, 3), (3,4), (4, 3), and (4,4); in the absence of DLT, the working data consist of no DLT at (2, 3), (2, 2), (2,1), (1, 3), (1, 2), and (1,1).

Suppose that a cohort of n_{ij} patients were treated at (i, j) and t of them experienced DLTs. After extending the data in the manner described above, the updated distribution is Beta (a_{rs}^*, b_{rs}^*) , where a_{rs}^* and b_{rs}^* are given as follows:

1. $a_{ij}^* = a_{ij} + t$ and $b_{ij}^* = b_{ij} + n_{ij} - t$.

2. For $(r, s) \succ (i, j)$, $a_{rs}^* = a_{rs} + t$ and $b_{rs}^* = b_{rs}$.
3. For $(r, s) \prec (i, j)$, $a_{rs}^* = a_{rs}$ and $b_{rs}^* = b_{rs} + n_{ij} - t$.
4. For all incomparable dose combinations, $a_{ij}^* = a_{ij}$ and $b_{ij}^* = b_{ij}$.

It is straightforward to show that the posterior means satisfy the order restrictions.

We hasten to point out that the use of data extrapolation is not intended for an accurate estimation of DLT probabilities. For ease of reference, we shall refer to dose combinations for which the response of a given patient is not extrapolated as missing cells. Missing cells occur at lower and incomparable dose combinations when a DLT is observed, and at higher and incomparable dose combinations when there is no DLT. Consequently, there is a tendency for the mean DLT probabilities to be underestimated at the lower dose combinations and overestimated at the higher dose combinations. Such bias can be advantageous, for it enables a fast escalation away from the low and potentially subtherapeutic dose combinations, yet impedes an aggressive escalation to the high and possibly overly toxic dose combinations. This advantage is based on the assumption that, in a two-agent trial where each agent has been thoroughly investigated when administered alone, the target combination is surrounded by dose combinations with lower as well as higher DLT probabilities.

2.2. Dose allocation

It is recognized that Bayesian methods are difficult to apply at the beginning of a trial because of limited information (Liu and Ning, 2013). To obtain preliminary data, we treat the first cohort of patients at the lowest dose combination, and then randomly escalate the dose of one agent for subsequent cohorts until the first toxicity is observed or the highest dose combination is reached, whichever occurs earlier. Afterward, we switch to the following decision-theoretic approach.

Central to the decision-theoretic approach is a utility function which measures the values of various dose allocation decisions. A dose allocation is considered optimal if it maximizes the expected utility with respect to the posterior distribution of the unknown DLT probabilities, given all data observed at the time of decision. In our proposal, the data include extrapolated responses described in Section 2.1. Formally, let θ_0 be the target DLT probability and $u_k(p, \theta_0)$ be the utility of assigning the k th patient to a dose combination with DLT probability p . A (total) utility function, in the simplest form, can be expressed as:

$$u(p, \theta_0) = \sum_{k=1}^n u_k(p, \theta_0),$$

where n is the number of patients in the trial. In our simulation studies presented in Section 3, we consider:

$$u_k(p, \theta_0) = \begin{cases} -\alpha_0(\theta_0 - p) & \text{if } p \leq \theta_0 \\ -\eta_0(p - \theta_0) & \text{if } p > \theta_0, \end{cases}$$

where α_0 and η_0 are specified positive constants. This function assigns a negative value to the selection of dose combinations with a DLT probability different from θ_0 ; furthermore, it has the flexibility of allowing for the utility of overdosing to decrease faster than that of underdosing (by setting $\eta_0 > \alpha_0$) when patient safety is the primary concern, or the reverse when assigning patients to subtherapeutic doses is the primary concern.

In practice, finding the optimal dose allocation is computationally expensive, if not infeasible. For this reason, myopic strategies, such as one-step-look-ahead (OSLA), are often used instead. With the OSLA strategy, we act as if the next patient were the last patient of the study, so that the next dose allocation is based on the current posterior distribution. If the next patient is the $(k+1)$ th patient, the expected utility function corresponding to the OSLA strategy is proportional to the conditional expectation of $u_{k+1}(p, \theta_0)$, given all data collected up to the k th patient. This results in a considerable simplification of the maximization problem. To illustrate, we used cohort sizes of one in our simulation studies, so that $u(p, \theta_0) = u_{k+1}(p, \theta_0)$. Let $f(p; a_{ij}^*, b_{ij}^*)$ denote the density function of a Beta a_{ij}^*, b_{ij}^* distribution, where a_{ij}^* and b_{ij}^* are the parameters of the posterior distribution for p_{ij} given all data collected up to the k th patient. Then:

$$\begin{aligned} E[u(p, \theta_0) | a_{ij}^*, b_{ij}^*] &= \int_0^{\theta_0} -\alpha_0(\theta_0 - p) f(p; a_{ij}^*, b_{ij}^*) dp + \int_{\theta_0}^1 -\eta_0(p - \theta_0) f(p; a_{ij}^*, b_{ij}^*) dp \quad (4) \\ &= -(\alpha_0 + \eta_0) \left[\theta_0 F(\theta_0; a_{ij}^*, b_{ij}^*) - \left(\frac{a_{ij}^*}{a_{ij}^* + b_{ij}^*} \right) F(\theta_0; a_{ij}^* + 1, b_{ij}^*) \right] \\ &\quad - \eta_0 \left[\left(\frac{a_{ij}^*}{a_{ij}^* + b_{ij}^*} \right) - \theta_0 \right], \end{aligned}$$

Where $F(\theta_0; a_{ij}^*, b_{ij}^*) = \int_0^{\theta_0} f(p; a_{ij}^*, b_{ij}^*) dp$ is the distribution function of a Beta (a_{ij}^*, b_{ij}^*) distribution. The next dose assignment is given by the combination (i, j) which yields the largest value of $E[u(p, \theta_0) | a_{ij}^*, b_{ij}^*]$. This can be easily evaluated using statistical softwares such as R. In real trials, clinicians may prefer a conservative dose escalation approach, in which dose skipping is not allowed. In this case, if the current dose combination is (r, s) , the next dose allocation is limited to dose combinations (i, j) such that $i \geq r$ and $j \geq s$, as well as the combinations $(r+1, s)$ and $(r, s+1)$.

2.3. Stopping rules

Before the trial begins, the minimum and maximum sample sizes, denoted $n_{\min}$ and $n_{\max}$ respectively, should be specified. In addition, let $\delta_0 \geq 0$ be specified such that $\theta_0 + \delta_0$ is an

upper bound on the acceptable DLT probabilities. The following rules are imposed in our design for deciding when to stop a trial.

S1. To mitigate the effects of prior misspecification, early stopping is prohibited before the sample size reaches $n_{\min}$.

S2. The trial is stopped when the sample size reaches $n_{\max}$.

S3. The trial is stopped if $\Pr(p_{11} > \theta_0 + \delta_0 | \text{data}) > r_1$ for some prespecified probability r_1 , that is, if the lowest (and hence all) dose combinations are likely to be overly toxic.

S4. If (i, j) is the current selected dose combination, the trial is stopped if $\min_{(i, j) < (r, s)} \Pr(p_{rs} > \theta_0 + \delta_0 | \text{data}) > r_2$ for some prespecified probability r_2 , that is, if all dose combinations of higher partial order than the current dose combination are likely to be overly toxic.

S5. S1 supersedes S3 and S4.

If a trial is stopped based on S3, no MTD is recommended; otherwise, the current dose combination is recommended as an MTD.

3. Simulation study

To examine the operating characteristics of the proposed design, we performed two sets of simulations, denoted LFL^{NS} and LFL; the former does not allow the skipping of doses, as described at the end of 2.2. In each set of simulations, 10,000 trials with cohort sizes of one were generated under all seven scenarios (Table 2) previously considered by Braun and Jia (2013), in which the target DLT probability, θ_0 , was 20%. A strict ordering of the dose combinations (as defined in Section 2.1) was used in all the simulated trials. In both LFL^{NS} and LFL, the true DLT probabilities p_{ij} were used as the prior means. Because it is not feasible to select the first and third quartiles individually for the simulated trials, automated selections were made among the generated (a_{ij}^0, b_{ij}^0) pairs such that $a_{ij}^0 \approx c\mu_{ij}^0$ and $b_{ij}^0 \approx c(1 - \mu_{ij}^0)$, where $c = 4$; note that a larger value of c corresponds to a more informative prior. For dose allocation, we chose the combination that maximizes the expected utility function given in equation (4), with $\alpha_0 = 1.2$ and $\eta_0 = 1$. We chose a_0 to be slightly greater than η_0 in view of the bias introduced by data extrapolation (last paragraph of Section 2.1) and stopping rule S4 (Section 2.3), the combination of which would render our method too conservative. For the stopping rules, we chose $n_{\min} = 10$, $n_{\max} = 50$, $\delta_0 = 0.05$, $r_1 = 0.5$, and $r_2 = 0.95$; the maximum sample size was chosen to match that of Braun and Jia (2013).

The results are summarized in Table 3, the format of which is similar to Table 2 of Braun and Jia (2013), which compares the operating characteristics of their gCRM model to the models of Yuan and Yin (2009) and Yin and Yuan (2009). We shall refer to these three designs collectively as “the other designs”. First of all, the operating characteristics of LFL^{NS} and LFL are nearly identical in terms of the selection probabilities of dose combinations and the average sample size. A plausible explanation is that, under the

monotonic assumption, which allows for data extrapolation as described in Section 2.1, the proposed algorithm narrows down the search range quickly; that is, the subtherapeutic and overly toxic dose combinations were eliminated from the search early in the trials. This characteristic has been previously demonstrated in Fan et al. (2012) for single-agent Phase I trials. As a result, the candidates for an MTD tend to be reduced to adjacent dose combinations early in the trials. Next, our design compares favorably to the other designs in all scenarios except B, where all dose combinations have DLT probabilities that are well below 20%. As pointed out earlier, our design assumes that θ_0 lies within the range of DLT probabilities of the dose combinations under study, an assumption that is violated in scenario B. A related scenario is D, where all dose combinations are overly toxic; however, despite a violation of the aforementioned assumption, our design was able to stop the trials early with no recommendation of MTD in nearly all the trials, similar to the other designs. This is due to the conservative stopping rule S3, which stops a trial when the lowest dose combination is deemed more likely than not to be overly toxic. An important advantage of our design is seen in the average sample size required to identify an MTD: except in scenario B, our design compares favorably to the other designs at almost half the sample size. The reason for this advantage is the same as explained earlier, that our algorithm takes full advantage of the partial ordering of the dose combinations. Note that both Yuan and Yin (2009) and Yin and Yuan (2009) restricted dose escalations to adjacent dose combinations, as with LFL^{NS}, whereas Braun and Jia (2013) allowed for the possibility of escalation of both doses at the same time.

To examine the robustness of our design to misspecifications of prior means, we performed three additional sets of simulations. In the first set of simulations, denoted LFL⁻, we randomly subtracted 1–5 percentage points from the true DLT probabilities and used the resulting values as the prior means. That is, the prior mean for combination (i, j) in the s th simulated trial was specified as $p_{ij} - e_{ij}^{(s)}$, where $e_{ij}^{(s)}$ are independent and uniformly distributed between 1–5 percentage points. In the event that $p_{ij} - e_{ij}^{(s)} < 0.01$, the prior mean for p_{ij} was specified to be 0.01. In the second set of simulations, denoted LFL⁺, we randomly added 1–5 percentage points to the true DLT probabilities in a similar fashion; and in the last set of simulations, denoted LFL*, we randomly added or subtracted 1–5 percentage points to the true DLT probabilities in a similar fashion. In each of these simulations, the prior means were not constrained to satisfy a strict ordering; the dose allocation and stopping rules were the same as before. The results are summarized in Table 4, along with the results of LFL to facilitate comparisons. The results suggest that our design is generally robust to misspecifications of prior means that are within 5 percentage points of the true means. Because of the conservativeness of our design, the impact is more apparent in LFL⁺, especially in scenarios F and G. Even then, there were hardly any changes in the percentage of recommendation of dose combinations with a DLT probability more than 10 percentage points below or above the target. Moreover, the results are still comparable to those of the other designs (Table 3) at nearly half the sample size.

4. Concluding remarks

Practitioners have long recognized the conceptual elegance and simplicity of the Bayesian framework, but a wider adoption of Bayesian designs in clinical trials has been hindered by two major issues: the complexity of prior specification and the computational effort required to implement Bayesian procedures. With this in mind, we developed a Bayesian design for two-agent Phase I trials that is relatively easy to understand and implement. We agree with Braun and Jia (2013) that “no one design for identifying the [MTD] will be best in all scenarios”, and the simulation results presented in Section 3 only serve to demonstrate that the proposed design is a viable alternative to more complex designs when there are significant time and sample size constraints. A limitation of our design is the dependence on the assumption of monotonicity, that the probability of toxicity increases monotonically with increasing dose. According to Wages and Conaway (2014), “situations in which the monotonicity assumption, for both toxicity and efficacy, may fail are becoming increasingly common in cancer research practice.” Additionally, the simulation results for scenario B suggest that our design may perform less than satisfactory when the MTD is the lowest or highest dose combination, due to the assumption that the target combination is surrounded by dose combinations with lower as well as higher DLT probabilities.

The proposed design can be further refined by choosing, for example, a different utility function or stopping rules. To facilitate such refinements, the source code for performing the simulations described in Section 3, written in the R language, is available upon request from the authors. Finally, the importance of performing a comprehensive simulation study before adopting any design in actual trials cannot be overemphasized. The simulation study should include a multitude of possible scenarios and different choices of priors, especially misspecified ones; this can be easily done since, on a laptop equipped with a 2.4 GHz Intel Core i7 processor, each set of 10,000 simulated trials can be completed in less than 30 seconds.

References

- Braun TM, Jia N (2013). A generalized continual reassessment method for two-agent phase I trials. *Statistics in Biopharmaceutical Research* 5:105–115. [PubMed: 24436776]
- Carlin BP, Zhong W, Koopmeiners JS (2013). Discussion of small-sample behavior of novel phase I cancer trial designs, by Assaf P. Oron and Peter D. Hoff. *Clinical Trials* 2013; 10:81–92.
- Fan SK, Lu Y, Wang YG (2012). A simple Bayesian decision-theoretic design for dose-finding trials. *Statistics in Medicine* 31:3719–3729. [PubMed: 22763943]
- Harrington JA, Wheeler GM, Sweeting MJ, Mander AP, Jodrell DI (2013). Adaptive designs for dual-agent phase I dose-escalation studies. *Nature Reviews Clinical Oncology* 10:277–288.
- Kelley RK, Venook AP (2012). Novel therapeutics in hepatocellular carcinoma. In Dizon DS (Ed.), *American Society of Clinical Oncology Educational book/ASCO, Vol. 2013 American Society of Clinical Oncology*, pp. 137–142.
- Liu S, Ning J (2013). A Bayesian dose-finding design for drug combination trials with delayed toxicities. *Bayesian Analysis* 8:703–722. [PubMed: 27924182]
- O’Quigley J, Pepe M, Fisher L (1990). Continual reassessment method: A practical design for phase I clinical trials in cancer. *Biometrics* 46:33–48. [PubMed: 2350571]
- Wages NA, Conaway MR (2014). Phase I/II adaptive designs for drug combination oncology trials. *Statistics in Medicine* 30:1990–2003.

- Wages NA, Conaway MR, O'Quigley J (2011). Dose-finding design for multi-drug combinations. *Clinical Trials*; 8:380–389. [PubMed: 21652689]
- Yuan Y, Yin G (2009). A latent contingency table approach to dose finding for combinations of two agents. *Biometrics* 65:866–875. [PubMed: 18759848]
- Yin G, Yuan Y (2009). Bayesian dose finding in oncology for drug combinations by copula regression. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 58:211–224.

Table 1.

An example of pairs of first and third quartiles, (q_{ij}^L, q_{ij}^U) , generated according to whether the prior mean $\mu_{ij}^0 \geq 0.5$ or $\mu_{ij}^0 < 0.5$. Also shown are the values of the hyperparameters, a_{ij}^0 and b_{ij}^0 , corresponding to each pair of quartiles.

$\mu_{ij}^0 = 0.65$			$\mu_{ij}^0 = 0.35$		
a_{ij}^0	b_{ij}^0	(q_{ij}^L, q_{ij}^U)	a_{ij}^0	b_{ij}^0	(q_{ij}^L, q_{ij}^U)
1	0.54	(0.41, 0.92)	0.54	1	(0.08, 0.59)
2	1.08	(0.48, 0.85)	1.08	2	(0.15, 0.52)
3	1.68	(0.51, 0.81)	1.62	3	(0.19, 0.49)
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$

Table 2.DLT probabilities ($\times 100$) for the seven scenarios examined.

Scenario	Agent A	Agent B			
		1	2	3	4
A	1	4	10	16	22
	2	8	14	20	26
	3	12	18	24	30
	4	16	22	28	34
B	1	2	5	8	11
	2	4	7	10	13
	3	6	9	12	15
	4	8	11	14	17
C	1	10	25	40	55
	2	20	35	50	65
	3	30	45	60	75
	4	40	55	70	85
D	1	44	50	56	62
	2	48	54	60	66
	3	52	58	64	70
	4	56	62	68	74
E	1	8	9	10	11
	2	18	19	20	21
	3	28	29	30	31
	4	29	30	31	41
F	1	12	16	44	50
	2	13	18	45	52
	3	14	20	46	54
	4	15	22	47	55
G	1	1	4	6	10
	2	2	10	15	30
	3	3	15	30	50
	4	4	20	45	80

Table 3.

Operating characteristics of the proposed design (denoted LFL) and the designs proposed by Braun and Jia (2013, denoted gCRM), Yuan and Yin (2009, YY09a), and Yin and Yuan (2009, denoted YY09b).

Scenario	Design	Percentage of recommendation				Average Sample size
		At target	1–10 pts of target	< or > 10 pts of target	None recommended	
A	LFL	17	79	4	0	24
	LFL ⁻	11	85	4	0	23
	LFL ⁺	9	86	5	0	26
	LFL*	14	82	4	0	24
	gCRM	10	82	3	5	48
	YY09a	13	82	5	0	50
	YY09b	11	81	6	2	50
B	LFL	0	87	13	0	21
	LFL ⁻	0	88	12	0	20
	LFL ⁺	0	87	13	0	26
	LFL*	0	87	13	0	24
	gCRM	0	94	3	3	50
	YY09a	0	99	1	0	50
	YY09b	0	96	4	0	50
C	LFL	53	46	0	1	21
	LFL ⁻	43	56	0	1	20
	LFL ⁺	52	46	0	3	21
	LFL*	48	51	0	1	20
	gCRM	45	39	5	11	45
	YY09a	41	50	5	4	49
	YY09b	42	47	5	6	48
D	LFL	0	0	2	98	10
	LFL ⁻	0	0	4	96	10
	LFL ⁺	0	0	3	97	10
	LFL*	0	0	3	97	10
	gCRM	0	0	4	96	11
	YY09a	0	0	1	99	10
	YY09b	0	0	1	99	8
E	LFL	9	82	9	0	25
	LFL ⁻	12	81	7	0	24
	LFL ⁺	8	83	8	1	25
	LFL*	13	78	8	0	24
	gCRM	9	70	14	7	47
	YY09a	6	65	27	2	49
	YY09b	7	67	25	1	49
F	LFL	12	85	0	2	23

Scenario	Design	Percentage of recommendation			None recommended	Average Sample size
		At target	1–10 pts of target	< or > 10 pts of target		
G	LFL ⁻	18	81	0	1	21
	LFL ⁺	6	92	0	3	25
	LFL*	17	82	0	2	22
	gCRM	13	70	6	11	45
	YY09a	14	76	6	4	48
	YY09b	12	74	7	7	47
	LFL	47	50	3	0	22
	LFL ⁻	40	58	3	0	19
	LFL ⁺	24	72	4	0	25
	LFL*	38	59	3	0	22
	gCRM	25	68	5	2	50
	YY09a	12	76	12	0	50
	YY09b	15	72	13	0	50

In LFL, the priors are centered at the true DLT probabilities; in LFL⁻ and LFL⁺, the prior means underestimate and overestimate the true DLT probabilities by 1–5% in each of the 10,000 simulations, respectively; in LFL*, both under- and overestimation occur in the prior means. The percentages of recommendation may not add up to 100 due to rounding errors.

Table 4.

Sensitivity analysis.

Scenario	Design	At target	Percentage of recommendation			None recommended	Average sample size
			1–5 pts of target	6–10 pts of target	< or > 10 pts of target		
A	LFL	17	67	12	5	0	24
	LFL ⁻	11	63	22	4	0	23
	LFL ⁺	9	72	14	5	0	26
	LFL*	10	67	18	4	0	26
B	LFL	0	53	34	13	0	20
	LFL ⁻	0	57	32	11	0	20
	LFL ⁺	0	28	59	12	0	26
	LFL*	0	33	56	11	0	26
C	LFL	53	27	19	1	1	21
	LFL ⁻	43	37	19	0	1	20
	LFL ⁺	52	18	28	0	3	21
	LFL*	47	30	22	0	1	20
D	LFL	0	0	0	2	98	10
	LFL ⁻	0	0	0	4	96	10
	LFL ⁺	0	0	0	3	97	10
	LFL*	0	0	0	4	96	10
E	LFL	9	56	26	9	0	24
	LFL ⁻	12	45	35	7	0	24
	LFL ⁺	8	54	29	8	0	25
	LFL*	11	48	32	8	0	27
F	LFL	13	49	36	0	2	23
	LFL ⁻	18	51	30	0	1	21
	LFL ⁺	6	41	51	0	3	25
	LFL*	12	42	44	0	2	26
G	LFL	47	0	50	3	0	21
	LFL ⁻	40	0	57	3	0	19
	LFL ⁺	24	0	72	4	0	25
	LFL*	33	0	64	3	0	23

In LFL, the priors are centered at the true DLT probabilities; in LFL⁻ and LFL⁺, the prior means underestimate and overestimate the true DLT probabilities by 1–5%; in LFL*, both under- and overestimation of 1–5% occur in the prior means. The percentages of recommendation may not add up to 100 due to rounding errors.