

A Parallel Phase I/II Clinical Trial Design for Combination Therapies

Xuelin Huang,^{1,*} Swati Biswas,² Yasuhiro Oki,³ Jean-Pierre Issa,³ and Donald A. Berry^{1,**}

¹Department of Biostatistics and Applied Mathematics, University of Texas,
M. D. Anderson Cancer Center, Houston, Texas 77030, U.S.A.

²Department of Biostatistics, School of Public Health,
The University of North Texas Health Science Center, Fort Worth, Texas 76107, U.S.A.

³Department of Leukemia, University of Texas, M. D. Anderson Cancer Center, Houston, Texas 77030, U.S.A.

* *email*: xlhuang@mdanderson.org

** *email*: dberry@mdanderson.org

SUMMARY. The use of multiple drugs in a single clinical trial or as a therapeutic strategy has become common, particularly in the treatment of cancer. Because traditional trials are designed to evaluate one agent at a time, the evaluation of therapies in combination requires specialized trial designs. In place of the traditional separate phase I and II trials, we propose using a parallel phase I/II clinical trial to evaluate simultaneously the safety and efficacy of combination dose levels, and select the optimal combination dose. The trial is started with an initial period of dose escalation, then patients are randomly assigned to admissible dose levels. These dose levels are compared with each other. Bayesian posterior probabilities are used in the randomization to adaptively assign more patients to doses with higher efficacy levels. Combination doses with lower efficacy are temporarily closed and those with intolerable toxicity are eliminated from the trial. The trial is stopped if the posterior probability for safety, efficacy, or futility crosses a prespecified boundary. For illustration, we apply the design to a combination chemotherapy trial for leukemia. We use simulation studies to assess the operating characteristics of the parallel phase I/II trial design, and compare it to a conventional design for a standard phase I and phase II trial. The simulations show that the proposed design saves sample size, has better power, and efficiently assigns more patients to doses with higher efficacy levels.

KEY WORDS: Bayesian; Dose selection; Efficacy; Logistic regression; Synergy; Toxicity.

1. Introduction

Combination therapies are now commonly used in medical practice and clinical trials, especially in the treatment for cancer. A literature search in PubMed, using the term “combination therapy,” identified 132,215 articles. Korn and Simon (1993) listed three reasons for the utility of combinations: (1) biochemical synergism, (2) differential susceptibility of tumor cells to different agents, and (3) higher achievable dose intensity by exploiting nonoverlapping toxicities to the host. Despite this growing popularity, caution is required when designing and evaluating trials for combination therapies. This is particularly true when combining toxic chemotherapeutic agents in the treatment of cancer. Investigators should not assume that the established dose level for an individual drug is automatically safe when it is used in combination with another drug.

The statistical design of combination therapy trials is worth special consideration. The traditional clinical trial design is based on the consideration of a single drug. The first phase of the trial is used to find the maximum tolerated dose (MTD) of the drug. A second phase is then conducted to determine whether the drug has promising effects at the MTD. The underlying theory is that a larger dose is associated with greater efficacy and greater toxicity. Although this theory may be

true for many drugs, it is false for some. In combination drug therapy, the MTDs form a curve in a two-dimensional plane. The location of the optimal combination dose on the curve is unknown. Therefore, a trial of drugs in combination requires a dose-selection procedure that is not needed for a single-drug trial.

Dose-finding methods for two agents in phase I trials have been proposed by Simon and Korn (1990, 1991), Thall et al. (2003), Conaway, Dunbar, and Peddada (2004), and Wang and Ivanova (2005), among others. They all focused on the identification of the MTD(s), and did not consider efficacy. For phase II trials, we are not aware of any published article that explicitly addresses the problem of statistical design for combination therapies. Traditional phase II designs for single-agent trials include the activity trial design (Gehan, 1961) and the optimal two-stage trial design (Simon, 1989). They did not consider dose selection.

Instead of conducting a phase I trial for toxicity and a separate phase II trial for efficacy, we propose to integrate these two phases into one. We still require an initial dose-escalation period for the combination doses to designate the admissible doses, but we assess the toxicity rates during this period as well as throughout the trial. This is because the toxicity profile from a small number of patients is not reliable. Also, to

make full use of the data, for every patient enrolled in this dose-escalation period, his/her efficacy data will be used in the next dose-selection stage, provided he/she meets the eligibility criteria for the next phase. After completing the initial dose-escalation period and designating the admissible doses, we adaptively randomize patients to the admissible combination doses with unequal probabilities. The goal is to assign more patients to doses with higher efficacy. The trial is monitored at each accrual of a cohort of, say, five patients. During these interim analyses, doses with lower efficacy are temporarily closed and those with intolerable toxicity are eliminated.

In the following section, we describe the parallel phase I/II design. In Section 3, we apply the proposed design to a trial evaluating combination chemotherapeutic agents in the treatment of acute myelogenous leukemia (AML) and myelodysplastic syndrome (MDS). We report the operating characteristics of the design, and show its advantages compared to a conventional design. We conclude with a brief discussion in Section 4.

2. A Parallel Phase I/II Design

2.1 Dose Escalation

Throughout this article, we assume that the outcomes for both toxicity and efficacy are binary. We use “response” and “no response” to describe efficacy. Suppose two drugs, A and B, have been tested separately in their previous phase I, II, or even III trials. The MTDs for them are D_a and D_b , respectively. An investigator now wants to evaluate the combined effect of A and B, testing the combination of doses d_1^a, \dots, d_m^a for drug A, and d_1^b, \dots, d_k^b for drug B. Often, but not necessarily, $d_m^a = D_a$ and $d_k^b = D_b$. For easy illustration, we take $m = k = 3$, corresponding to the search over some low, medium, and high doses. The dose combinations are labeled d_1 through d_9 , as illustrated in Figure 1. They are zoned. The first zone has a single combination dose $d_1 = d_1^a \times d_1^b$. The

second zone has two combination doses: $d_2 = d_2^a \times d_1^b$ and $d_3 = d_1^a \times d_2^b$. The dose escalation starts with the first zone, and follows with the second zone, then the third zone, and so on. The combination doses in a zone closer to the coordinate origin of the plane are tested earlier. Combination doses in the same zone are tested simultaneously using blocked randomization to control the number of patients for each open dose. Some combination doses may close, in which case blocked randomization is applied to the remaining open combination doses in the same zone. We use a modified “3 + 3” design (Storer, 1989) for the initial dose escalation. Other designs for phase I mentioned in Section 1 can also be used.

To find acceptable dose levels, the investigator specifies the dose-limiting toxicity (DLT) for the drug(s) under consideration. DLT is usually defined as side effects sufficiently morbid that they constitute a practical limitation to the delivery of the treatment. The modified “3 + 3” design proceeds as described below. Three patients receive dose d_1 . Their responses determine the subsequent steps, as described in the following decision rules.

- (1) If two patients (2/3) experience DLT, then we close d_1 and all larger doses to testing.
- (2) If no patient experiences DLT, then we designate d_1 as an admissible dose, and open doses in zone 2 for testing.
- (3) If only one of the first three patients (1/3) experiences DLT, then we assign three more patients to d_1 . In this case, we assign up to six patients to this dose, and the following decision rules apply:
 - (a) If only one of the six patients (1/6) experiences DLT, then we designate d_1 as an admissible dose, and open doses in zone 2 for testing.
 - (b) If two of the six patients (2/6) experience DLT, then we designate d_1 as an admissible dose, and doses in zone 2 and higher are closed to testing.
 - (c) If more than two of the six patients (>2/6) experience DLT, then we close d_1 to testing, and we also close doses in zone 2 and higher.

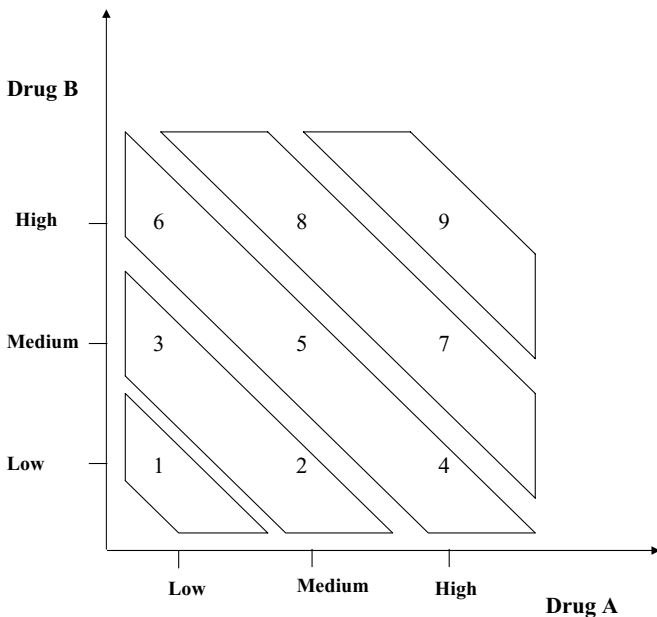


Figure 1. An illustration of combination doses.

We then evaluate the higher open doses using similar rules. The results of $d_2 (= d_2^a \times d_1^b)$ determine the opening or closure of $d_4 (= d_3^a \times d_1^b)$, as do the results of $d_3 (= d_1^a \times d_2^b)$ for $d_6 (= d_1^a \times d_3^b)$. However, dose $d_5 (= d_2^a \times d_2^b)$ is opened for testing only if the results of both d_2 and d_3 indicate that it can be opened. That is to say, the decision to open a combination dose to testing depends upon the results of both its horizontal and vertical neighbors in the lower dose zone.

2.2 Dose Selection Using Adaptive Randomization

After the initial dose-escalation period, the design adaptively randomizes patients to all the admissible dose levels. There are many approaches to determine the adaptive randomization assignment probability for each dose level. We denote the estimated response rates for dose levels d_1 to d_9 by P_1 to P_9 , respectively. For convenience, we call each dose level a treatment arm. One approach is to let the assignment probability for arm i be proportional to its response rate P_i estimated by the data accumulated so far. Another approach, the play-the-winner design, does not explicitly compute the response rates, but achieves a similar property asymptotically (Zelen, 1969;

Wei and Durham, 1978). A disadvantage of these approaches is that it does not account for the reliability of the estimated values of P_i . The difference between, say, $P_i = 0.4$ and $P_j = 0.5$ results in the same difference in assignment probabilities throughout the trial. Consequently, these designs are too liberal and unstable at the beginning of the trial when the number of patients is small. Then, at a late stage of the trial when the number of patients is large, these designs may be too conservative, and thus cannot take full advantage of the adaptive randomization.

Another approach is to compare the response rates with a target rate, say p , and let the assignment probabilities be proportional to the posterior probabilities $R_i = \Pr[P_i > p | \text{Data}]$. This approach takes into consideration the amount of information in the data. However, adaptive randomization may not work well with this approach when all of the true response rates are much higher or lower than p , because all the values of R_i are then close to 1 or 0.

We propose a new approach. We let the assignment probabilities be proportional to the credibility of the supposition that one treatment arm is superior to the other arms. That is, the assignment probability for arm i is proportional to the posterior probability $\Pr[P_i > \max\{P_j, j \neq i\} | \text{Data}]$. This approach does not have the disadvantages mentioned above. The assignment probabilities depend solely on internal comparisons between different treatment arms, rather than on the comparison with an arbitrarily chosen target response rate. It also distinguishes between early-stage differences (based on a small sample and thus unreliable) and late-stage differences (based on a large sample and thus reliable). For example, given estimates $P_i = 0.4$ and $P_j = 0.5$, the design will assign roughly the same number of patients to arms i and j at the beginning of the trial, because the results are based on data from a small number of patients. However, at a late stage of the trial, when the results are based on data from a relatively large number of patients, the design will assign many more patients to arm j than to arm i .

In practice, the high-dimensional integration involved in the computation of the above quantities can be obtained using simulation. However, we can circumvent this still considerable amount of computational burden by using a modification that can give essentially the same assignment probabilities, but with much less computation expense. The modification is as follows. For $i \geq 2$, we define $R_i = \Pr[P_i > P_1 | \text{Data}]$. For $i = 1$, we let $R_1 = 0.5$. Then, we let the assignment probabilities be proportional to the values of R_i , or to their squares if we want to be more aggressive. This approach allows the quantity of interest, that is, the comparison results, to be built directly into the adaptive randomization scheme. It is reliable and efficient. The choice of $R_1 = 0.5$ is fair and is not arbitrary. If we suppose that some arm is independent of and identical to the first arm, then the chance for this arm to be superior to the first arm is 0.5.

We may update the estimate of the response rate for each dose separately. However, it is better to use a model to characterize the relationship between the response rates of all the combination doses, and update all of them together. The latter approach can borrow strength between doses, and thus is more efficient. In the next section, we apply the proposed design to a cancer trial. We monitor the efficacy and toxicity of

each dose level throughout the trial. Doses with lower efficacy are temporarily closed and those with intolerable toxicity are eliminated. To do the efficacy monitoring, a logistic regression model is used to estimate the response rates of combination dose levels. Bayesian methods are used because of their inherent convenience for sequential updating and monitoring. We also would like to achieve desirable frequentist operating characteristics for the design. This is done through simulations to calibrate the boundary values used in the interim monitoring and final decision.

3. Application of the Design to a Clinical Trial

3.1 Background

We design a parallel phase I/II trial for a clinical cancer study. The study aims to evaluate the safety and efficacy of low-dose decitabine combined with Ara-C in the treatment of relapsed/refractory AML or high-risk MDS. AML is a clonal myelopoietic stem cell disorder characterized by the accumulation of neoplastic cells in the bone marrow and peripheral circulation. MDS comprises a heterogeneous group of hematopoietic stem cell disorders that may evolve into AML in 10% to 70% of individuals with this disorder. We use complete remission (CR) as the response criterion. The historical CR rate for patients with relapsed/refractory AML or high-risk MDS is as low as 5%. The target CR rate for the experimental drug combination is 20%. We evaluate the response and toxicity rates at 6 weeks after treatment. The investigators define DLT as grade 3 or higher myelosuppression or nonhematologic toxicity, and estimate the accrual rate at 5 to 10 patients per month, which can be reduced if necessary. We determine that a maximum of 100 patients will be enrolled in the trial, and that the study will run for 1.5 to 2 years.

The MTD of decitabine as a single agent has been identified, and myelosuppression is the primary DLT (Kantarjian et al., 2003). Other trials have indicated that decitabine is most efficacious at doses much lower than the MTD (Koshy et al., 2000; Issa et al., 2004). The toxicity and efficacy profiles of Ara-C have also been established. The primary goal of this study is to find the most effective, safe dose levels for decitabine and Ara-C when used in combination. We investigate two dose levels of decitabine that are much lower than its MTD (low-dose levels 1 and 2), and two dose levels of Ara-C (low and high). We also investigate both sequential and concurrent schedules of decitabine and Ara-C. This leads to a total of eight combination dose levels, as denoted 1 to 8 in Figure 2. The four sequential combination dose levels (labeled 1 to 4) are zoned similarly as in Figure 1, as are the four concurrent combination dose levels (labeled 5 to 8).

3.2 Interim Monitoring for Efficacy and Toxicity

We conduct dose escalations for the sequential and concurrent therapies separately and simultaneously. Toxicities occurring during either the concurrent or sequential application of the therapies do not affect the admissibility of the dose levels under the other scheme. The modified "3 + 3" algorithm described in Section 2 is applied to both schemes. First, we assign three patients to each of dose levels 1 and 5, using blocked randomization. If results from both of them indicate dose escalation, then blocked randomization is applied to dose levels 2, 3, 6, and 7. If only dose level 1 leads to escalation

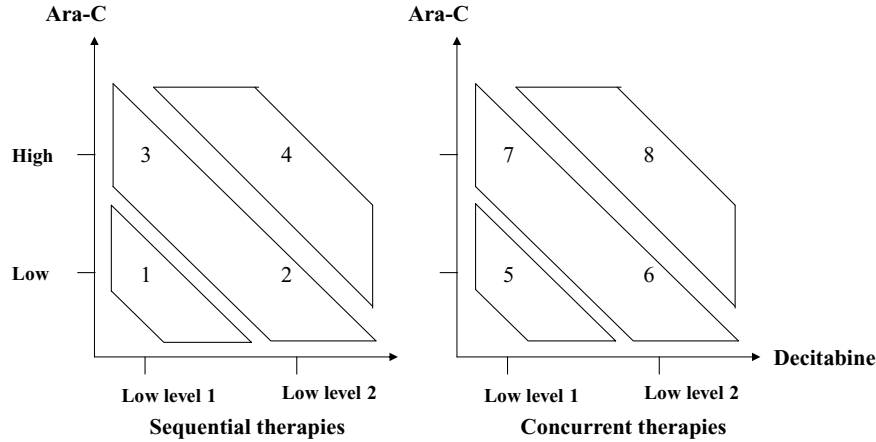


Figure 2. Simultaneous dose escalation for sequential and concurrent therapies.

but not dose level 5, then blocked randomization is applied to dose levels 2 and 3 only. If neither of dose levels 1 and 5 leads to escalation, but at least one of them is admissible, the trial proceeds to further evaluate the efficacy and toxicity of the admissible dose(s). If both dose levels 1 and 5 are too toxic, then we terminate the trial.

After completing the dose-escalation procedure, we begin to adaptively randomize patients among the admissible dose levels. We model the response rates of the combination dose levels as follows, assuming all of the eight dose levels are admissible. Modification is straightforward for situations where not all dose levels are admissible.

$$\begin{cases} X = -0.5, & \text{for decitabine at low-dose level 0,} \\ X = 0.5, & \text{for decitabine at low-dose level 1;} \\ \\ Y = -0.5, & \text{for low-dose Ara-C,} \\ Y = 0.5, & \text{for high-dose Ara-C;} \\ \\ Z = -0.5, & \text{for sequential therapy,} \\ Z = 0.5, & \text{for concurrent therapy.} \end{cases}$$

Denoting the CR rate for arm i by P_i , we use the following logistic regression model to estimate the values of P_i for all combination dose levels,

$$\text{Logit}(P_i) = \beta_0 + \beta_1 X + \beta_2 Y + \beta_3 Z. \tag{1}$$

We assume that the prior distribution of β_0 is Normal $(-2.944, 10)$, which is equivalent to a response rate of 5%, and that the prior distributions of $\beta_1, \beta_2,$ and β_3 are Normal $(0, 10)$. We choose values of -0.5 or 0.5 for $X, Y,$ and Z instead of the conventional 0 or 1, in order to give equal variance to the prior distributions of the eight arms.

For the adaptive randomization, we define

$$R_i = \Pr[P_i > P_1 | \text{data}], \quad i = 2, \dots, 8.$$

We also define $R_1 = 0.5$. We compute R_i for each arm using the above logistic regression model. The probability that the current patient is assigned to treatment arm i is proportional to the current estimate of R_i .

After every cohort of five patients, we update the prior distributions for the toxicity and efficacy parameters for the eight arms. We use all the data accumulated in the dose-escalation period and thereafter for updating. Let Q_i denote the toxicity rate for the i th arm and n_i the sample size used by this arm. Once we start the adaptive randomization, the following stopping rules apply.

- (1) Toxicity: Stop arm i if $\Pr[Q_i \geq 0.33 | \text{data}] > 0.8$, using a Beta $(0.1, 0.9)$ prior distribution for each Q_i .
- (2) Futility:
 - (a) Temporarily close any arm with $R_i < 0.01$. If newly accumulated data indicate $R_i > 0.01$ for that arm, then reopen it.
 - (b) Close the trial if $\max\{\Pr[P_i \geq 0.20 | \text{data}], i = 1, \dots, 8\} < 0.1$, and if for at least four arms $n_i \geq 5$. If the number of remaining arms is less than 4, then require all $n_i \geq 5$.
- (3) Efficacy: Stop the trial and select the i th arm as the best arm if the following conditions are met:
 - (a) $\Pr[P_i \geq 0.20 | \text{data}] > 0.90$,
 - (b) $\Pr[P_i > P_j | \text{data}] > 0.80$ for all $j \neq i, j = 1, \dots, 8$,
 - (c) At least four arms have $n_i \geq 5$. If the number of remaining arms is less than 4, then require all $n_i \geq 5$.

When the maximum sample size is reached, we apply the following decision rules.

- (1) Among the remaining arms in the trial, if arm i has the highest probability $\Pr[P_i \geq 0.05 | \text{data}]$ and this probability is greater than 0.95, then select arm i .
- (2) If none of the remaining arms in the trial satisfies $\Pr[P_i \geq 0.05 | \text{data}] > 0.95$, then select no arm, and conclude that none of the proposed experimental combinations is promising for further investigation.

3.3 Operating Characteristics and Advantages of the Parallel Phase I/II Design

We perform simulation studies to assess the operating characteristics of the proposed design in six different scenarios: increasing, constant, and decreasing response rates, each in

Table 1
Operating characteristics of the proposed parallel phase I/II design

Doses	Sequential therapies				Concurrent therapies				Total
	1	2	3	4	5	6	7	8	
	Scenario 1: Response rates increasing, toxicity rates increasing								
Response rate P_i	0.05	0.2	0.2	0.4	0.05	0.15	0.15	0.3	
Toxicity rate Q_i	0.05	0.10	0.15	0.20	0.10	0.15	0.20	0.25	
Pr(closed), toxicity	0	0.04	0.08	0.38	0.04	0.18	0.24	0.59	
Pr(selected), efficacy	0	0.16	0.13	0.49	0	0.03	0.03	0.13	
Sample size used (n_i)	7.3	13.5	12.7	8.5	6.4	8.6	8.4	5.4	70.9
	Scenario 2: Response rates increasing, toxicity rates constant								
Response rate P_i	0.05	0.2	0.2	0.4	0.05	0.15	0.15	0.3	
Toxicity rate Q_i	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	
Pr(closed), toxicity	0	0	0.01	0.04	0.01	0.03	0.03	0.08	
Pr(selected), efficacy	0	0.09	0.07	0.66	0	0.01	0.01	0.16	
Sample size used (n_i)	6.0	11.1	11.0	11.7	5.4	8.1	8.0	9.2	70.5
	Scenario 3: Response rates decreasing, toxicity rates increasing								
Response rate P_i	0.3	0.15	0.15	0.05	0.2	0.10	0.10	0.05	
Toxicity rate Q_i	0.05	0.10	0.15	0.20	0.10	0.15	0.20	0.25	
Pr(closed), toxicity	0	0.04	0.07	0.34	0.04	0.18	0.25	0.58	
Pr(selected), efficacy	0.72	0.04	0.03	0	0.13	0.01	0.02	0	
Sample size used (n_i)	28.7	9.5	9.4	4.3	13.8	6.7	6.3	2.8	81.6
	Scenario 4: Response rates decreasing, toxicity rates constant								
Response rate P_i	0.3	0.15	0.15	0.05	0.2	0.10	0.10	0.05	
Toxicity rate Q_i	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	
Pr(closed), toxicity	0	0.01	0.01	0.06	0.01	0.04	0.03	0.06	
Pr(selected), efficacy	0.74	0.04	0.04	0	0.12	0.01	0.01	0	
Sample size used (n_i)	32.8	9.1	9.2	4.8	13.3	6.5	6.6	4.5	86.9
	Scenario 5: Response rates constant, toxicity rates increasing								
Response rate P_i	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	
Toxicity rate Q_i	0.05	0.10	0.15	0.20	0.10	0.15	0.20	0.25	
Pr(closed), toxicity	0	0.04	0.08	0.36	0.04	0.17	0.23	0.60	
Pr(selected), efficacy	0	0.01	0	0.01	0.01	0.01	0.01	0.01	
Sample size used (n_i)	8.3	8.3	8.3	5.6	8.1	6.8	6.7	3.5	55.6
	Scenario 6: Response rates constant, toxicity rates constant								
Response rate P_i	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	
Toxicity rate Q_i	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	
Pr(closed), toxicity	0	0	0	0.05	0.01	0.03	0.03	0.07	
Pr(selected), efficacy	0	0.01	0.01	0.01	0.01	0.02	0.01	0.01	
Sample size used (n_i)	8.1	8.4	8.0	7.3	8.2	7.5	7.3	6.8	61.6

combination with increasing and constant toxicity rates (see Table 1). We perform 1000 simulations for each scenario, using scenarios 1 to 4 to estimate power, and scenarios 5 and 6 to assess the type I error rates. The results are presented in Table 1. Proportions that are less than 0.005 are reported as 0. We choose the boundary values for the monitoring rules (such as 0.8 in the toxicity rule) by calibration, such that the design has the desired properties. The features of the design are summarized below.

First, the rules monitoring dose escalation and toxicity in our proposed design are very conservative. For example, doses with a toxicity rate of 25% have a chance of about 60% to be closed by toxicity (scenarios 1, 3, and 5).

Second, the design uses adaptive randomization to assign more patients to arms with higher efficacy, while closing arms with high toxicity rates. These two effects determine the sample size for each arm. To see their net effects, consider scenarios with either constant toxicity rates or constant response

rates. In scenario 4, all arms have a 5% toxicity rate, and the response rates for the eight arms range from 5% to 30%. The corresponding average sample sizes used by them range from 4.5 to 32.8. The design successfully assigns many more patients to the most effective arm. In scenario 5, all arms have a 5% response rate. The toxicity rates for the eight arms range from 5% to 25%, and the corresponding average sample sizes used by them decrease from 8.3 to 3.5. So the design assigns fewer patients to more toxic arms.

Third, it is clear that this design can reduce the overall sample size. Consider scenarios 2 and 4, where the toxicity rates for all arms are at a constant rate of 5%. The highest response rates are 40% and 30%, and the average total sample sizes are 70.5 and 86.9 in scenarios 2 and 4, respectively. The design reduces the sample sizes by stopping early for efficacy. All treatment arms in scenario 6 have a 5% response rate and a 5% toxicity rate, and the average total sample size is 61.6 due to early stopping for futility.

Fourth, the design controls the type I error rate, and has good power to select the arms with the highest response rates. Type I error rates can be seen from scenarios 5 and 6. In scenario 6, all arms have a null response rate of 5% and a low toxicity rate of 5%. The chance for each of these arms to be selected is always less than 2%, and the overall probability of designating some null arm as effective is 8%. The power of the design can be seen from scenarios 1 to 4. In scenario 4, for example, arm 1 has the highest response rate (30%), and is selected with probability 0.74. Arm 5 has the second highest response rate (20%) and is selected with probability 0.12. In sum, the probability of selecting one of these two best performing arms is 86%. The best two doses are indicated by bold numbers in Table 1. Although scenarios 1 to 4 are designed to estimate power, they also provide estimates for type I error rates because they all have arms with 5% response rates. In all of these four scenarios, the estimated type I error rate for each null arm is zero in Table 1. Comparing with scenarios 5 and 6, the type I error rates are even lower in scenarios 1 to 4. This is due to the competition for power from those efficacious arms. This phenomenon shows the advantage of building comparisons between arms into the design, as opposed to separate evaluation for each arm. An example of the latter is provided below.

3.4 Comparison with a Conventional Design

We run a second set of 1000 simulations for each of the six scenarios using a conventional design consisting of separate phase I and phase II trials. For easy comparison, we use total sample sizes of 100 for the combined phase I and phase II trials, which correspond to the sample sizes used in the simulations of the proposed design. In the phase I trial, we use the same modified “3 + 3” method in Section 2 for dose-escalation decisions. We then use the admissible doses in the subsequent phase II trial. After simulating the phase I trial,

we divide the remaining patients into groups of equal size, to receive one of the admissible doses. If the division cannot be done evenly, we let the arms of the lower dose levels have one more patient than the other arms. Initially we assume an independent Beta(0.05, 0.95) prior for each arm, and then update this distribution using all the data from the phase I and II trials. There is no logistic model, interim monitoring, or adaptive randomization. We use the same final decision rules that we followed in the example applying the proposed design, that is, the arm with the highest $\Pr(P_i > 0.05 | \text{data})$ is selected if this posterior probability is greater than 0.95. If no arm has this posterior probability greater than 0.95, then the trial results are negative.

We compare the results obtained using our proposed design with those obtained using the conventional design. The comparison results for scenarios 1 to 4 are illustrated in Figure 3. The conventional design always uses 100 patients, while the proposed design requires fewer patients, ranging from 70.5 to 86.9 in the four scenarios. Even using data from fewer patients, the proposed design has better power (Figure 3c and 3d). For example, in scenario 4, the proposed design has 74% power to select the best performing arm (best in the sense of having the highest response rate), whereas the conventional design has only 57% power to select the best arm.

While the conventional design assigns equal numbers of patients to each dose level, the proposed design assigns more patients to more effective doses. Consequently, the proposed design, compared to the conventional design, results in higher proportions of patients achieving response during the trial (Figure 3b). Using scenario 4 again as an example, we find that 20% of patients achieve response under the proposed design, compared to only 14% of patients under the conventional design.

Other designs, such as Simon’s optimal two-stage phase II design (Simon, 1989), also save patient resource by doing

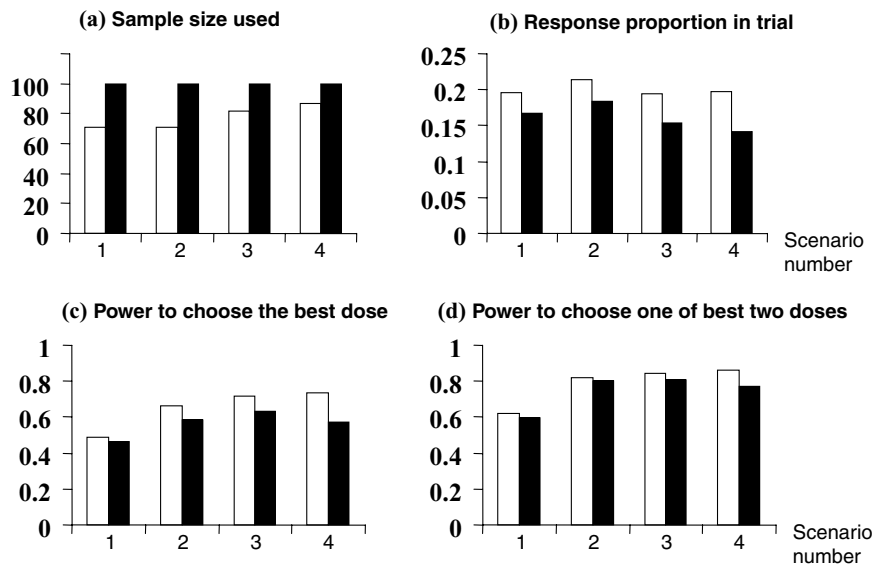


Figure 3. Comparisons between the proposed design (white bars) and the conventional design (black bars).

interim monitoring. However, it is not clear how Simon's design can be extended to a trial involving multiple doses in combination. Using Simon's design for each dose would cost too many patients but still not have direct comparisons between different doses, and thus have difficulty in identifying the optimal dose.

4. Summary and Discussion

4.1 Summary

Combination therapies in cancer treatment are becoming commonplace. Even though the toxicity profile of each agent may be known, the toxicity profile of the agents used in combination must be established. Due to the difficulty in identifying the "highest" dose in a two-dimensional dose plane, a dose-selection procedure is necessary for combination therapies. This contrasts with the case of single-agent trials. We have proposed a design to fulfill these requirements for trials of combination therapies. By our design, efficacy and toxicity data are accumulated throughout the trial and adaptive randomization is applied such that sample size can be reduced, and patients can be assigned to more effective and less toxic treatments. The proposed design also eliminates combination doses with lower efficacy or intolerable toxicity, and stops the trial whenever the posterior probability for safety, efficacy, or futility crosses a prespecified boundary. Berry and Eick (1995) discussed the advantages of adaptive assignment versus balanced randomization in clinical trials.

Bayesian approaches for clinical research are elaborated by Berry (2003, 2004), and are reviewed in a recent article in the journal *Nature* (Berry, 2006). It is intuitively appealing to use Bayesian posterior probabilities in interim monitoring and final decision making. However, multiple interim analyses inflate the final type I error rate. Currently, simulation studies are used to choose the boundary values to achieve desired frequentist operating characteristics. This process involves a search over potential boundary values, and can be computationally expensive. It would be desirable to find out formulas for the relationship between boundary values on posterior probabilities and the operating characteristics. These formulas should involve as input the maximum sample size, number of combination doses, number of interim analyses, and the prior distributions of the parameters. These results can be developed into software. Once the software is ready, it will be posted on <http://biostatistics.mdanderson.org/SoftwareDownload/>. Currently, this website has more than 30 software packages for statistical computation, data analysis, and clinical trial design (free download).

Inoue, Thall, and Berry (2002) proposed a seamless randomized phase II/III trial design. Although potentially very useful, implementation of the seamless randomized phase II/III trial design may be limited by regulatory guidelines. Our parallel phase I/II design seamlessly expands phase I into phase II. This design should be less constrained by regulatory guidelines and readily applied in practice.

4.2 Modification of the Design in Various Situations

For simplicity, a modified "3 + 3" design is used in the dose-escalation phase. We may replace it by another design to take advantage of the single-agent toxicity information from

previous trials. The toxicity probability model by Wang and Ivanova (2005) can be used to do that. Their model can be viewed as a natural extension of the continual reassessment method (O'Quigley, Pepe, and Fisher, 1990) to a two-dimensional plane. We may also use the models by Korn and Simon (1993) to incorporate information on the numbers of each type of single-agent toxicity. These extensions warrant further research.

We have used a logistic regression model to monitor the response rates of all the combination dose levels. Even though the logistic regression model is only approximately true, the operating characteristics of the design are sufficiently good for practical use. Interaction terms between covariates can be induced in the model to further improve the performance of the design. In future research, we may consider dynamic modeling, that is, including interaction terms when they achieve a certain level of significance and eliminating them when they do not.

In Section 2, we have used a search over three dose levels for a each agent as an easy illustration of the design. In practice, three doses for each agent may not be enough. We may use two-fold increments for each agent, starting at a very low dose level. We may also allow the possibility to open a new dose level that was not preplanned.

The number of combination doses is often large, and thus it requires a large sample size to do dose escalation and selection. For rare diseases such as many types of cancer, this can be a serious problem. To reduce sample size, one may use mathematical models such as that by Korn and Simon (1993) as guidelines to select a small number of combination dose levels. One may also use informative prior distributions in the design. These prior distributions can be derived from historical data on each single agent. Certainly meta-analyses of results from several small trials will also be useful. Results on these small trials may stimulate large multicenter trials.

Patients with fatal diseases may be willing to accept greater levels of toxicity to achieve higher rates of response. For trials involving such patients, the proposed design can be adapted to apply a trade-off between efficacy and toxicity for dose selection. In this case, we would use two logistic regression models: one for response rates and the other for toxicity rates. Details of this design, such as the quantification of the trade-off, require further consideration.

ACKNOWLEDGEMENTS

This research was supported by National Institute of Health grant CA100632. We thank the associate editor and reviewer for their constructive comments.

REFERENCES

- Berry, D. A. (2003). Statistical innovations in cancer research. In *Cancer Medicine*, D. Kufe, R. Pollock, R. Weichselbaum, R. Bast, T. Gansler, J. Holland, and E. Frei (eds), 465–478. London: BC Decker.
- Berry, D. A. (2004). Bayesian statistics and the efficiency and ethics of clinical trials. *Statistical Science* **19**, 175–187.
- Berry, D. A. (2006). Bayesian clinical trials. *Nature Reviews Drug Discovery* **5**, 27–36.

- Berry, D. A. and Eick, S. G. (1995). Adaptive assignment versus balanced randomization in clinical trials: A decision analysis. *Statistics in Medicine* **14**, 231–246.
- Conaway, M. R., Dunbar, S., and Peddada, S. D. (2004). Design for single- or multiple-agent phase I trials. *Biometrics* **60**, 661–669.
- Gehen, E. A. (1961). Determination of number of patients required in a preliminary and a follow-up trial of a new chemotherapeutic agent. *Journal of Chronic Diseases* **13**, 346–353.
- Inoue, L. Y. T., Thall, P. F., and Berry, D. A. (2002). Seamlessly expanding a randomized phase II trial to phase III. *Biometrics* **58**, 823–831.
- Issa, J. P., Garcia-Manero, G., Giles, F. J., Mannari, R., Thomas, D., Faderl, S., Bayar, E., Lyons, J., Rosenfeld, C. S., Cortes, J., and Kantarjian, H. M. (2004). Phase 1 study of low-dose prolonged exposure schedules of the hypomethylating agent 5-aza-2'-deoxycytidine (decitabine) in hematopoietic malignancies. *Blood* **103**, 1635–1640.
- Kantarjian, H. M., O'Brien, S., Cortes, J., Giles, F. J., Faderl, S., Issa, J. P., Garcia-Manero, G., Rios, M. B., Shan, J., Andreeff, M., Keating, M., and Talpaz, M. (2003). Results of decitabine (5-aza-2'-deoxycytidine) therapy in 130 patients with chronic myelogenous leukemia. *Cancer* **98**, 522–528.
- Korn, E. and Simon, R. (1993). Using the tolerable-dose diagram in the design of phase I combination chemotherapy trials. *Journal of Clinical Oncology* **11**, 794–801.
- Koshy, M., Dorn, L., Bressler, L., Molokie, R., Lavelle, D., Talischy, N., Hoffman, R., van Overveld, W., and DeSimone, J. (2000). 2-deoxy 5-azacytidine and fetal hemoglobin induction in sickle cell anemia. *Blood* **96**, 2379–2384.
- 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.
- Simon, R. (1989). Optimal 2-stage designs for phase-II clinical trials. *Controlled Clinical Trials* **10**, 1–10.
- Simon, R. and Korn, E. L. (1990). Selecting drug combination based on total equivalence dose (dose intensity). *Journal of the National Cancer Institute* **82**, 1469–1476.
- Simon, R. and Korn, E. L. (1991). Selecting combinations of chemotherapeutic drugs to maximize dose intensity. *Journal of Pharmaceutical Statistics* **1**, 247–258.
- Storer, B. E. (1989). Design and analysis of phase I clinical trials. *Biometrics* **45**, 925–937.
- Thall, P. F., Millikan, R. E., Mueller, P., and Lee, S. J. (2003). Dose-finding with two agents in phase I oncology trials. *Biometrics* **59**, 487–496.
- Wang, K. and Ivanova, A. (2005). Two-dimensional dose finding in discrete dose spaces. *Biometrics* **61**, 217–222.
- Wei, L. J. and Durham, S. (1978). The randomized play-the-winner rule in medical trials. *Journal of the American Statistical Association* **73**, 840–843.
- Zelen, M. (1969). Play the winner rule and the controlled clinical trial. *Journal of the American Statistical Association* **64**, 131–146.

Received February 2006. Revised July 2006.

Accepted July 2006.