Published in final edited form as:

Stat Med. 2012 October 30; 31(24): 2872–2881. doi:10.1002/sim.5365.

Two-stage designs for Phase 2 dose-finding trials

Anastasia Ivanova^{1,*†}, Changfu Xiao¹, and Yevgen Tymofyeyev²

¹Department of Biostatistics, University of North Carolina, Chapel Hill, NC 27599-7420, U.S.A.

²Janssen R&D, 1125 Trenton-Harbourton Road, Titusville, NJ 08560, U.S.A.

SUMMARY

We propose a Bayesian adaptive two-stage design for the efficient estimation of the maximum dose or the minimum effective dose in a dose-finding trial. The new design allocates subjects in stage two according to the posterior distribution of the target dose location. Simulations show that the proposed two-stage design is superior to equal allocation and to a two-stage strategy where only one dose is left in the second stage.

Keywords

Dose ranging; Minimum effective dose; Maximum dose; Phase 2 trials; two-stage designs

1. INTRODUCTION

The objective of a non-oncology Phase 2 trial is to select a dose or a small subset of promising doses that are further investigated in a Phase 3 trial. Several Phase 2 trials are usually conducted, including Phase 2A and Phase 2B trials. The primary endpoint of a Phase 2A trial is an efficacy endpoint often measured by a continuous biomarker. A typical Phase 2B trial is larger than a Phase 2A trial. It enrolls a few hundred patients, investigates several doses, and the treatment time is commonly longer than in a Phase 2A trial. Adverse events monitored in a Phase 2 trial comprise a secondary endpoint. Phase 2A and Phase 2B trials can have various objectives such as finding the minimum effective dose or finding a dose with the best adverse event – efficacy trade off. The minimum effective dose (MED) is the smallest dose with a discernible useful effect (ICH E4 Guideline, 1994) [1]. The MED can be defined as the dose where the mean efficacy outcome is equal to a certain target, with the placebo used as a reference. Mean efficacy is usually assumed to be non-decreasing with dose. Both efficacy and safety endpoints are often taken into consideration when selecting a Phase 3 dose, as increasing the dose can result in both higher efficacy and increased adverse event rates. A common approach is to quantify efficacy and adverse event rate trade-off through a utility function. Such a function incorporates both efficacy and safety into a measure of overall clinical “utility” [2–4]. Utility is often defined as a linear combination of the efficacy measure and adverse event rates, with no or little negative weight given to

*Correspondence to: Anastasia Ivanova, Department of Biostatistics, CB #7420, The University of North Carolina at Chapel Hill, Chapel Hill, NC 27599-7420, U.S.A..

†aivanova@bios.unc.edu

moderate adverse events, and very heavy negative weight given to severe adverse events. Utility function typically has an “umbrella” or “inverse U” shape, and the objective of a trial is to maximize overall clinical “utility” of the drug. We will refer to the dose that maximizes the utility function as the optimal dose. When unrelated treatments are studied, no order is assumed among mean responses and the goal is simply to find the treatment with the maximum mean response. An additional objective in Phase 2 can be to test efficacy and adverse event rates at the estimated MED or the optimal dose against placebo and/or an active control. Therefore a recommended assignment strategy for a Phase 2 study will provide good quality of estimation of the target dose and increased sample size at the estimated target dose to yield better power of treatment comparisons.

A number of fully sequential adaptive designs have been proposed for dose-finding Phase 2 trials in non-oncology setting (Berry *et al.* [3]; Smith *et al.* [5]; Ivanova *et al.* [6]; Ivanova *et al.* [4]). The logistics of implementing a fully sequential dose finding study can be daunting. On the other hand, a parallel group design with equal allocation to all doses is the most common design in Phase 2 non-oncology trials. A two-stage design is a reasonable compromise between multistage and single-stage approaches. Miller *et al.* [7] investigated a two-stage strategy for a dose-ranging study that is optimal across several parametric models. They concluded that the proposed two-stage strategy offers minor benefit compared to a single-stage design in terms of the efficiency of estimation of the target dose. Dragalin *et al.* [8] investigated optimal two-stage designs for two correlated binary endpoints that follow a bivariate probit model and concluded that two-stage strategy is superior to equal allocation.

In this paper, we propose a Bayesian two-stage design for dose-finding Phase 2 non-oncology trials under the following common set-ups: 1) estimating the MED under the assumption of non-decreasing dose-response curve, 2) estimating the MED under isotonic matrix order, a set-up arising when several different administration schedules are investigated, 3) estimating the dose with the highest response, 4) estimating the dose with the highest response under the umbrella order assumption.

2. THE MODEL

Let $\{d_1, \dots, d_K\}$ be the set of ordered dose levels selected for a trial with d_1 denoting placebo. The methodology we propose can be used with continuous or binary outcomes. If treatment response is normally distributed, ignoring the monotonicity, a conjugate prior density (Gelman *et al.* [9], p. 78) can be specified as

$$\mu_j | \sigma^2 \sim N(\mu_{0j}, \sigma^2 / k_{0j}), j=1, 2, \dots, K, \text{ and } \sigma^2 \sim IG(v_0, \sigma_0^2),$$

where IG denotes inverse gamma distribution. Let n_j be the number of subjects assigned to d_j , $N = n_1 + \dots + n_K$. Subjects' response at d_j , $\mathbf{y}_j$, $j = 1, 2, \dots, K$ is a vector of n_j i.i.d. $N(\mu_j, \sigma^2)$ random variables. The posterior of $\boldsymbol{\mu} = (\mu_1, \dots, \mu_K)'$ conditional on σ^2 and $\mathbf{y}$ is

$$\mu_j | \sigma^2, \mathbf{y} \sim N(M_j; V_j), j=1, 2, \dots, K, \text{ and } \sigma^2 | \mathbf{y} \sim IG(V_n, \sigma_n), \quad (1)$$

where $M_j = (k_{0j}\mu_{0j} + n_j\bar{y}_j)/(k_{0j} + n_j)$, $V_j = \sigma^2 / (k_{0j} + n_j)$, $v_n = v_0 + N/2$, and

$$\sigma_n = \sigma_0 + \frac{1}{2} \sum_{j=1}^K \left\{ (n_j - 1) s_j^2 + \frac{k_{0j} n_j}{k_{0j} + n_j} (\bar{y}_j - \mu_{0j})^2 \right\},$$

with $(\bar{y}_1, \dots, \bar{y}_K)'$ denoting the unrestricted maximum likelihood estimates, and s_j^2 denoting the empirical variance of y_j .

If treatment response is binary (yes/no) with mean response vector $\boldsymbol{\mu} = (\mu_1, \dots, \mu_K)'$, subjects' response at d_j , y_j , $j = 1, 2, \dots, K$ is a vector of n_j i.i.d. *Bernoulli*(μ_j) random variables. Responses at d_j can be summarized as $m_j = y_{j1} + \dots + y_{jn_j}$. Assuming *Beta*(α, β) prior on μ_j , $j = 1, 2, \dots, K$, the posterior distribution of μ_j conditional on y_j is

$$\mu_j | y_j \sim \text{Beta}(\alpha + m_j, \beta + n_j - m_j), j = 1, 2, K.$$

We follow the isotonic transformation approach of Dunson and Neelon [10] and Gunn and Dunson [11] and map unconstrained mean vector $\boldsymbol{\mu}$ from $R^K \rightarrow \Omega$ to obtain the posterior distribution for the restricted means. Here $\Omega \subset R^K$ is defined by a set of inequalities on the elements of $\boldsymbol{\mu}$. Since the posterior distribution (1) of unconstrained parameter vector $\boldsymbol{\mu}$ follows a simple conjugate form, we can easily obtain the draws via Gibbs sampling algorithm, and transform draws to the constrained draws from the posterior density for the constrained parameter vector, $\boldsymbol{\mu}^*$, using the isotonic transformation approach. In the following sections, we consider three types of constraints that define Ω : non-decreasing, umbrella and matrix order.

3. TWO-STAGE DESIGN TO FIND THE MINIMUM EFFECTIVE DOSE

3.1 Estimating the MED

In this section dose-response is assumed to be non-decreasing with dose, $\mu_1 \leq \dots \leq \mu_K$, and the goal is to find the minimum effective dose, MED, defined as the dose with the mean response of $\mu_1 + \eta$, where $\eta > 0$ is the minimum clinically important difference specified before the trial. Under the assumption of non-decreasing dose-response $\mu_1 \leq \dots \leq \mu_K$, in non-Bayesian set-up, the restricted maximum likelihood estimates for components of $\boldsymbol{\mu}$,

$\hat{\boldsymbol{\mu}}^* = (\hat{\mu}_1^*, \dots, \hat{\mu}_K^*)'$, can be obtained from the unrestricted maximum likelihood estimates, $(\bar{y}_1, \dots, \bar{y}_K)'$ [12] as:

$$\hat{\mu}_j^* = \min_{t \in U_j} \max_{s \in L_j} \left(\frac{\sum_{h=s}^t n_h \bar{y}_h}{\sum_{h=s}^t n_h} \right), \quad (2)$$

for $j = 1, 2, \dots, K$. Here $L_j = \{1, \dots, j\}$ and $U_j = \{j, \dots, K\}$. Transformation (2) is a least-squares projection from R^K to the restricted space Ω . In the Bayesian setting, following [10] we project draws from the unconstrained posterior density (1) onto Ω using a minimal distance mapping. We then work with transformed draws.

3.2 Two stage design

The two-stage strategy we propose is described below. Let N_1 and N_2 be the total sample sizes in two stages respectively. The issue of the optimal split of the total sample size, $N_1 + N_2$, between the two stages is considered in Section 6.

Step 1. In stage 1, assign N_1/K subjects to each dose.

Step 2. Update the prior using stage 1 data to obtain unconstrained posterior density of μ . Transform each of D draws from the unconstrained posterior density of μ to follow non-decreasing order as described in Section 3.1. For each draw, the location of the MED is determined as the dose with the value closest to $\mu_1 + \eta$. These locations are summarized as the posterior distribution for the location of the MED $\pi = (\pi_1, \dots, \pi_K)$.

Step 3. Let N_2 be the number of subjects available for stage 2 and $\pi_m = \max_{j=2, \dots, K}(\pi_j)$. In stage 2, $\pi_j / (1 - \pi_1 + \pi_m) N_2$ of subjects are assigned to dose d_j , $j = 2, \dots, K$ and $\pi_m / (1 - \pi_1 + \pi_m) N_2$ subjects are assigned to placebo. That is, subjects are allocated proportionally to the posterior of the MED location except for placebo, where the number of subjects is set equal to the number allocated to the most likely target dose.

Step 4. After stage 2, the data from both stages are combined. The estimated MED is the dose d_j , $j = 2, \dots, K$, such that the posterior mean of $\mu_j - \mu_1$ is the closest to η .

3.3 Comparing the MED with placebo

It is often of interest to compare the target dose with placebo and/or an active control. This comparison should account for both the multiplicity of treatments in stage 2 and the selection processes (i.e. interim analysis). The classical Dunnett's test adjusts for the original number of hypotheses but does not take into account selection process. Our simulations show that when used with our two-stage design, the Dunnett's test is conservative in terms of controlling of the family-wise type I error rate for comparing the estimated MED to placebo. These conclusions are similar to those in Koenig *et al.* [13]. Another approach is to use a combination test with the weighted inverse normal combination function applied together with the closed testing principle [14]. In the combination test, t -test p-values are calculated for each dose and each stage. The closed testing principle with Simes' test of intersection hypotheses is then used within each stage. The overall p-value for each dose is calculated by applying a weighted inverse normal combination function to the two adjusted p-values, p and q ,

$$C(p, q) = 1 - \Phi[\sqrt{w}\Phi^{-1}(1 - p) + \sqrt{1 - w}\Phi^{-1}(1 - q)].$$

. Here w , $w \geq 0$, is a pre-defined weight, and Φ is the cumulative distribution function of the standard normal distribution. The adaptive combination test performs well when one treatment comparison is made in stage 2. However, due to the closed testing principle, the adaptive combination test becomes more conservative if more arms are selected after the first stage. Because our proposed two-stage design can have any number of treatment arms in stage 2, both testing procedures will be conservative. Instead, we propose to simulate the

distribution of the Dunnett's p-values under the null hypothesis following the proposed two-stage design, then the critical value for, say, a one-sided 0.025 level test can be obtained as 0.025 percentile of the Dunnett's p-values distribution. The critical value might depend on the number of doses K , as well as the variance of the outcome. Our simulations show that the critical value increases as the variance of outcome increases. Therefore, computing critical value using the variance value that is equal to or lower than the true variance will yield a conservative choice that preserves type I error rate. Such a strategy will be more powerful, even for very conservative choices of the guesstimate, compared to the Dunnett's test that always preserves the type I error rate.

4. TWO-STAGE DESIGN TO FIND THE MED WHEN TWO ADMINISTRATION SCHEDULES ARE INVESTIGATED

Several administration schedules can be studied in a phase 2 trial. Often schedules can be ordered based on intensity, for example, twice-a-day administration is more intense than once a day, with twice-a-day yielding higher or same mean response compared to once-a-day. This leads to two-dimensional monotonicity assumption: 1) the mean response is non-decreasing with dose given the schedule, and 2) the mean response is non-decreasing with schedule given the dose. Let $\mu_{11}, \dots, \mu_{1K}$ be the vector of mean responses for once-a-day schedule and $\mu_{21}, \dots, \mu_{2K}$ for twice a day. We have $\mu_{11} \leq \dots \leq \mu_{1K}$, $\mu_{21} \leq \dots \leq \mu_{2K}$, and $\mu_{1i} \leq \mu_{2i}$ for any $i = 2, \dots, K$. Since the first dose is placebo, we additionally have $\mu_{11} = \mu_{21}$. This order is often referred to as matrix order [15]. The maximum likelihood estimates under matrix order restriction, μ^* , can be computed using the Dykstra *et al.* algorithm [15] that can be found in the Appendix.

The goal of the trial with two administration schedules can be to estimate one MED or to estimate two MEDs, one for each administration schedule. The two-stage design we propose will work with either goal. The two-stage design is similar to the one in Section 3. In the final analysis, depending on the objective, either one MED or two MEDs, one for each administration schedule are selected. When comparing the estimated MED with placebo as in Section 3, we generated the critical value from the distribution of Dunnett's p-values obtained under the null hypothesis. As in the case of a single administration schedule, the critical value depends on the number of doses as well as on σ^2 .

5. TWO-STAGE DESIGN TO FIND THE OPTIMAL DOSE

In this section, we consider a problem of finding the maximum of a utility function. We assume the umbrella order $\mu_1 \leq \dots \leq \mu_{h-1} \leq \mu_h \leq \mu_{h+1} \leq \dots \leq \mu_K$, where the location of the peak, h , is unknown. First, assuming a known peak location k , the restricted estimates can be obtained as follows:

$$\mu_j^{*k} = \min_{t \in U_k^k} \max_{s \in L_j^k} \left(\frac{\sum_{h=s}^t n_h \bar{y}_h}{\sum_{h=s}^t n_h} \right), \quad (3)$$

for $j = 1, 2, \dots, K$. Here U_j^k and L_j^k denote subsets of $\{1, \dots, K\}$ such that the ordering $\mu_{j'} \leq \mu_j$ is known for all $j' \in L_j^k$ and the ordering $\mu_{j'} \geq \mu_j$ is known for all $j' \in U_j^k$. To allow for a peak at an unknown location, k , we choose μ^* by minimizing the distance across different choices of peak:

$$\mu^* = \min_{k \in \{1, \dots, K\}} \{(\mu^{*k} - \mu) \Sigma_{\mu}^{-1} (\mu^{*k} - \mu)'\}, \quad (4)$$

where $\Sigma_{\mu} = \text{diag}(V_1, \dots, V_K)$.

As in Gunn and Dunson [11], we transform the unrestricted draws using formulae (3) and (4) and then consider the draws of μ^* to be draws from a Bayesian posterior.

The two-stage design for estimating the optimal dose is as follows.

Step 1. In stage 1, assign N_1/K subjects to each dose.

Step 2. Update the prior using stage 1 data to obtain unconstrained posterior density of μ . Transform each of D draws from the unconstrained posterior density of μ to follow umbrella order to obtain the posterior distribution for the maximum of the umbrella $\pi = (\pi_1, \dots, \pi_K)$.

Step 3. Let N_2 be the number of subjects available for stage 2 and $\pi_m = \max_{j=2, \dots, K} (\pi_j)$. In stage 2, $\pi_j / (1 - \pi_1 + \pi_m) N_2$ of subjects are assigned to dose d_j , $j = 2, \dots, K$ and $\pi_m / (1 - \pi_1 + \pi_m) N_2$ subjects are assigned to placebo. That is, subjects are allocated proportional to the posterior of the optimal dose location except for placebo, where the number of subjects is set equal to the number allocated to the most likely target dose.

Step 4. The optimal dose is estimated from combined stage 1 and 2 data, as the mode of the posterior distribution for the optimal dose location.

Methods similar to that in Section 3.3 are used to compare the estimated optimal dose with placebo. The critical value for the test is obtained as 0.025 percentile of the distribution of Dunnett's p-values under the null following a two-stage design to estimate the optimal dose. Simulations show that the critical value depends on the number of doses K , however, unlike the MED case, it does not depend on unknown variance σ^2 .

If no order among treatment means is assumed, the two-stage design described above can be used except the posterior distribution for the maximum mean $\pi = (\pi_1, \dots, \pi_K)$ is obtained from untransformed rather than transformed draws. That is, allocation to a dose is proportional to the posterior probability of the dose having the highest response. A similar approach was used to randomize patients continuously in a dose-finding trial [16], and stems from the work of Thompson [17].

6. SIMULATION STUDY

We conducted a simulation study to compare the performance of the proposed two-stage design with a two-stage design where only one treatment arm is selected for the second

stage, and with a single stage design with equal allocation. In both of the comparator designs, we select the mode of the posterior distribution of the target dose location as the dose investigated in stage 2 and the estimated target dose after the trial. In the select-one-dose two-stage design where a single arm is carried in stage 2, the estimation is based on stage 1 data; in a single stage design, the estimations is based on all data. Results are based on 10000 simulation runs.

In simulations with a normally distributed outcome, the total number of subjects in a trial was 180 to estimate the MED, 252 to estimate the MED with two administration schedules and 100 to estimate the optimal dose. This sample sizes were selected via simulations to yield at least 80% power for target dose – placebo comparison for a set of plausible scenarios. Scenarios 1–5 in Table I were used for the MED simulations and scenarios 6–10 for the optimal dose simulations. All dose-response shapes are from [18]. Table II displays scenarios with two administration schedules. The dose-response curves are from [18] with doses (0.05, 0.15, 0.30) and (0.40, 0.70, 1.0) for scenario 1, and doses (0.05, 0.45, 0.85) and (0.30, 0.70, 1.0) for scenario 2. For a scenario with mean vector μ outcomes at d_j follow a normal distribution $N(\mu_j, \sigma^2)$ with $\sigma = 0.65$. The conjugate prior for μ_j follows the conditional distribution $\mu_j | \sigma^2 = N(\mu_{0j}, \sigma^2/k_{0j})$, $j = 1, 2, \dots, K$, and $\sigma^2 \sim IG(a_0, b_0)$, with $\mu_{0j} = 0$, $\sigma^2 = 1$, $k_{0j} = 0.001$, and $a_0 = b_0 = 0.001/2$. We obtain the draws from the posterior density (1) via Gibbs sampling algorithm, and transform draws to the constrained draws from the posterior density of the constrained parameter, μ^* , as described in Section 3.1, 4 or 5. The above process is repeated 1500 times discarding the first 500 iterations as a burn-in. For testing against placebo, we used the combination test with $w = 0.5$ with Simes' method for the two-stage design with a single arm carried in stage 2. We used Dunnett's test in a single stage design.

First, we investigated what is the optimal way to split the total sample size between the two stages. The best proportion for the optimal dose problem was selected based on average power for placebo – target dose comparison. Figure 1 presents power averaged over corresponding scenarios for the three set-ups. The probability of a correct selection follows a similar pattern and similar results were observed for other values of σ^2 . Allocating 0.58 of the sample size in stage 1 gives the best average power when the MED is estimated under non-decreasing or matrix order, while allocating 0.42 of the sample is the best to estimate the optimal dose. In the simulation study, we used proportion 0.5 in all set-ups allocating equal number of patients in stage 1 and stage 2.

Table III reports simulation results for the MED estimation, including the probability of selecting each dose as the target dose and the probability of correctly rejecting the null hypothesis of equality of placebo mean response with the estimated MED. Adaptive design yields the same probability of correct selection compared to the equal allocation. However, it assigns more subjects to the target dose on average compared to equal allocation which leads to much better power. The number of subjects assigned to the estimated MED is equal to 36 for equal allocation, compared to the median number of 46 for the adaptive two-stage strategy with 36 and 52 being the 25th and 75th percentiles. The new adaptive strategy yields much higher probability of correctly selecting the MED and higher power for comparing with placebo than the two-stage strategy where one dose is left in stage 2. The

critical value for the adaptive two-stage design was obtained by simulating 40,000 trials under the null hypothesis using true $\sigma^2 = 0.65$. Though the critical value depends on σ^2 , the critical values obtained for σ^2 in interval $[0.4, 0.9]$ were almost the same as for $\sigma^2 = 0.65$. Since the critical value increases as variance increases, our recommendation is to use a low bound as an estimate for the variance to obtain the critical value.

We also compared the new two-stage strategy to estimate the MED with the multi-stage t -statistic design from [6]. A total 180 subjects were assigned in 20 cohorts of size 9. In the first four cohorts, 5 subjects in each cohort were assigned to placebo to provide a good estimate of placebo response early in the trial. After that, 3 subjects in each cohort received placebo. Non-placebo assignments throughout the trial were determined according to design in [6]. The total number of subjects assigned to placebo was 68 subjects. This multi-stage strategy yielded slightly better probability of selecting the correct MED 0.71, 0.73, 0.47, 0.78 and 0.81 for scenarios 1–5 correspondingly, compared to 0.67, 0.67, 0.47, 0.79 and 0.76 for the Bayesian two-stage design. The number of subjects assigned to the estimated MED was significantly higher: median (25th; 75th percentiles) were 72 (52; 88) for the t -statistic design compared to 46 (36; 52) for the two-stage strategy, yielding much higher power of the MED - placebo comparison.

Table IV shows simulation results for selecting the MED in case of two administration schedules. We simulated trials where the goal was to identify one MED. The adaptive strategy yields better estimation and significantly larger power compared to the other two designs. The number of subjects assigned to the estimated target dose is 36 for equal allocation, compared to the median number of 47 for the adaptive two-stage strategy with 37 and 56 being the 25th and 75th percentiles. Comparing Tables III and IV we conclude that the more assumptions are utilized, the more benefit the adaptive two-stage strategy offers compared to a single-stage design.

Table V displays results for the optimal dose estimation when assuming an umbrella order. Our conclusions are very similar to the ones for the MED, except this time both two-stage strategies yield much higher power than equal allocation. The number of subjects assigned to the estimated optimal dose is 20 for equal allocation, compared to the median number of 29 for the adaptive two-stage strategy with 23 and 32 being the 25th and 75th percentiles. We repeated simulations in case no order is assumed among the mean responses. The results for the scenarios and sample size used in the simulation study are very similar, with umbrella estimation yielding slightly higher power for some scenarios. Our other simulations show that utilizing the umbrella assumption is more beneficial when the standard error of the estimates is large or sample size is small.

We compared the two-stage strategy to estimate the optimal dose with the multistage design from [4]. As recommended in [4], we assigned 40% of the total sample size of 100 in stage 1 allocating 40 subjects equally among doses. Subsequently, subjects were assigned in cohorts of 6. One or two subjects in each cohort received placebo, and the rest received the active treatment. The number of placebo assignments in each cohort was varied in order to keep the total number of placebo assignments approximately equal to the number of assignments at the best dose. This was done to ensure good power of optimal dose – placebo comparison

at the end of the trial. Mean responses were estimated after each step assuming umbrella order. The multi-stage strategy did not improve the likelihood of selecting the optimal dose: the probabilities of correctly selecting the optimal dose were 0.53, 0.82, 0.96, 0.86 and 0.84 for scenarios 6–10 compared to 0.51, 0.80, 0.95, 0.85 and 0.83 in Bayesian two-stage design. The number of subjects assigned to the estimated optimal dose was also similar: median (25th; 75th percentiles) were 29 (28; 30) for the multi-stage design compared to 29 (23; 32) for the two-stage strategy. Multi-stage design is more efficient compared to the two-stage design when more doses are studied, larger sample size is used or the variability of the outcome, σ^2 , is smaller.

For binary outcomes, we present simulation results for MED estimation (Table VI). The total sample size was 90 subjects per trial. Beta prior with parameters $\alpha = \beta = 0.01$ was used. We used Fisher's exact test to test the MED against placebo at the end of the trial. Simulations show that the type I error rate was preserved for all the designs with the empirical type I error rate being the highest for equal allocation and the smallest for the select-one-dose two-stage design. Therefore we did not use any adjustments for multiple comparisons. The conclusions are very similar to the ones for trials with continuous outcome: the new adaptive design has significantly better power than equal allocation and allows for better estimation of the target dose compared to the select one dose two-stage design. The select-one-dose two-stage design often has similar or even better power than the two-stage adaptive design for MED estimation. This is in part due to the fact that a dose higher than the MED is selected after stage 1. Comparing power values for the MED is only meaningful when designs yield comparable percentages of selecting each dose as the MED. Conclusions regarding the optimal dose estimation are very similar to those for continuous outcome (simulation results are available from the authors). As in the case of continuous outcome, methods that utilize umbrella assumption performed similar with minor improvement in quality of estimation of the optimal dose. Advantage of using umbrella assumption is more pronounced when the optimal dose is estimated based on less data as in the select-one-dose two-stage design where the dose is estimated based on stage 1 data only.

We investigated three-stage strategies similar to the proposed Bayesian two-stage design. Our conclusion was that adding a stage to a two-stage design does not improve power and selection probability, for the scenarios and sample sizes considered, enough to justify additional complexity.

7. DISCUSSION

An adaptive two-stage design is a reasonable alternative to equal allocation or multistage strategies. Compared to a single stage design with equal allocation, it yields larger sample size at the estimated target dose and hence provides better power for treatment comparison. The logistics of a two-stage trial are more complex compared to a single stage design but easier than a multi-stage approach. The two-stage approach allows for an interim analysis after stage 1 to stop the trial for futility or efficacy.

Often, there is a set of covariates believed to be associated with response. When the MED is defined using placebo as reference and the mean response is modeled with identity link

function using a linear model with covariates, the target doses for different levels of covariate coincide. Therefore the proposed two-stage adaptive strategy can be easily extended to the case when adjustment with respect to covariate is needed. Another possible extension is finding the optimal dose when several administration schedules are considered.

Acknowledgments

Ivanova's work was supported in part by the National Institutes of Health RO1 CA120082-01A1. The authors thank an associate editor and anonymous reviewers for helpful comments.

References

1. International conference on harmonisation of technical requirements for registration of pharmaceuticals for human use. ICH harmonised tripartite guideline: E4 dose-response information to support drug registration. Current step 4 version. 1994 Mar 10. Available at: <http://www.ich.org/LOB/media/MEDIA480.pdf>
2. Dragalin V, Fedorov V. Adaptive designs for dose-finding based on efficacy-toxicity response. *Journal of Statistical Planning and Inference*. 2006; 136:1800–1823.
3. Berry, DA.; Müller, P.; Grieve, AP.; Smith, M.; Parke, T.; Blazek, R.; Mitchard, N.; Krams, M. Adaptive Bayesian designs for dose-ranging drug trials. In: Gatsonis, C.; Kass, RE.; Carlin, B.; Carriquiry, A.; Gelman, A.; Verdine, I.; West, M., editors. *Case Studies in Bayesian Statistics V*. New York: Springer-Verlag; 2001. p. 99–181.
4. Ivanova A, Liu K, Snyder E, Snavely D. An adaptive design for identifying the dose with the best efficacy/tolerability profile with application to a crossover dose finding study. *Statistics in Medicine*. 2009; 28:2941–2951. [PubMed: 19731265]
5. Smith MK, Jones I, Morris MF, Grieve AP, Tan K. Implementation of a Bayesian adaptive design in proof of concept study. *Pharmaceutical Statistics*. 2006; 5:39–50. [PubMed: 17080927]
6. Ivanova A, Bolognese J, Perevozskaya I. Adaptive design based on t -statistic for dose-response trials. *Statistics in Medicine*. 2008; 27:1581–1592. [PubMed: 18241082]
7. Miller F, Guilbaud O, Dette H. Optimal Designs for Estimating the Interesting Part of a Dose-Effect Curve. *Journal of Biopharmaceutical Statistics*. 2007; 17:1097–1115. [PubMed: 18027219]
8. Dragalin V, Fedorov V, Wu Y. Adaptive designs for selecting drug combinations based on efficacy-toxicity response. *Journal of Statistical Planning and Inference*. 2008; 138:352–373.
9. Gelman, A.; Carlin, JB.; Stern, HS.; Rubin, DB. *Bayesian Data Analysis*. London: Chapman and Hall; 1995.
10. Dunson D, Neelon B. Bayesian inference on order-constrained parameters in generalized linear models. *Biometrics*. 2003; 59:286–295. [PubMed: 12926713]
11. Gunn L, Dunson D. A transformation approach for incorporating monotone or unimodal constraints. *Biostatistics*. 2005; 6(3):434–449. [PubMed: 15831579]
12. Robertson, T.; Wright, FT.; Dykstra, RL. *Ordered Restricted Statistical Inference*. New York: Wiley; 1988.
13. Koenig F, Brannath W, Bretz F, Posch M. Adaptive Dunnett tests for treatment selection. *Statistics in Medicine*. 2008; 27:1612–1625. [PubMed: 17876763]
14. Posch M, Koenig F, Branson M, Brannath W, Dunger-Baldauf C, Bauer P. Testing and estimation in flexible group sequential designs with adaptive treatment selection. *Statistics in Medicine*. 2005; 24:3697–3714. [PubMed: 16320264]
15. Dykstra RL, Robertson T. An Algorithm for isotonic regression for two or more independent variables. *Annals of Statistics*. 1982; 10:708–716.
16. Thall PF, Inoue LYT, Martin TG. Adaptive decision making in a lymphocyte infusion trial. *Biometrics*. 2002; 58:560–568. [PubMed: 12229990]
17. Thompson WR. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*. 1933; 25:285–294.

18. Bretz F, Pinheiro JC, Branson M. Combining multiple comparisons and modeling techniques in dose-response studies. *Biometrics*. 2005; 61:738–748. [PubMed: 16135025]

Appendix

The algorithm for computing maximum likelihood estimates under matrix order [15]

Step 1. Let $\hat{\mu}^{(1)} = (\hat{\mu}_{ij}^{(1)})$ denote the isotonic regression of μ over rows, i.e. $\mu^{(1)}$ minimizes $\sum_{i=1}^2 \sum_{j=1}^K (\mu_{ij} - f_{ij})^2 n_{ij} + (\mu_{10} - f_{10})^2 n_{10}$ subject to $f_{10} \leq f_{1j} \leq f_{2j}$ for $j = 1, \dots, K$. Let $R^{(1)} = (r_{ij}^{(1)}) = (\hat{\mu}_{ij}^{(1)} - \mu_{ij})$ be the first set of ‘row increments’.

Step 2. Let $\tilde{\mu}^{(1)} = (\tilde{\mu}_{ij}^{(1)})$ denote the isotonic regression over columns of $\mu + R^{(1)}$, i.e. $\mu^{(1)}$ minimizes $\sum_{i=1}^2 \sum_{j=1}^K (\mu_{ij} + r_{ij}^{(1)} - f_{ij})^2 n_{ij} + (\mu_{10} + r_{10}^{(1)} - f_{10})^2 n_{10}$ subject to $f_{10} \leq f_{i1} \leq \dots \leq f_{iK}$ for $i = 1, 2$. Call $C^{(1)} = \mu^{(1)} - (\mu + R^{(1)})$ the first set of ‘column increments’. Note that $\mu^{(1)} = \mu + R^{(1)} + C^{(1)}$.

Step 3. At the beginning of the m th cycle, $\mu^{\hat{m}}$ is obtained by isotonizing $\mu + C^{(m-1)}$ over rows. The m th set of row increments is defined by $R^{(m)} = \mu^{\hat{m}} - (\mu + C^{(m-1)})$, so that $\mu^{\hat{m}} = \mu + C^{(m-1)} + R^{(m)}$. Next obtain $\mu^{\tilde{m}}$ by isotonizing $\mu + R^{(m)}$ over columns. The m th set of column increments is given by $C^{(m)} = \mu^{\tilde{m}} - (\mu + R^{(m)})$ or, equivalently, $\mu^{\tilde{m}} = \mu + R^{(m)} + C^{(m)}$.

Figure 1.

Power averaged over all scenarios plotted against the proportion allocated in stage 1 with proportion of 1.0 corresponding to a single stage design. Solid line corresponds to the optimal dose estimation, dashed line to the MED estimation and dotted line to the MED estimation with two administration schedules.

Table I

Dose-response scenarios. Scenarios 1–5 are to illustrate the MED estimation, scenarios 6–10 the optimal dose estimation. The MED and the optimal dose are shown in bold. Placebo is dose one.

Scenario	Model	Mean Response
1	Emax	(0.20,0.34, 0.68 ,0.76,0.78)
2	linear in log-dose	(0.20,0.27, 0.59 ,0.74,0.80)
3	Linear	(0.20,0.23,0.47, 0.68 ,0.80)
4	Truncated-logistic	(0.20,0.20,0.22, 0.54 ,0.80)
5	Logistic	(0.20,0.21, 0.58 ,0.79,0.80)
6	Quadratic	(0.20,0.60, 0.79 ,0.75,0.50)
7	Double-logistic	(0.20,0.37, 0.79 ,0.59,0.50)
8	Exponential	(0.20,0.22,0.29,0.43, 0.80)
9	Step 1	(0.20,0.50,0.50, 0.80 ,0.50)
10	Step 2	(0.20,0.40, 0.80 ,0.60,0.40)

Table II

Scenarios for the two administration schedules. The MEDs are in bold. Placebo is dose

Scenario	Model	Group	Mean Response
1	Linear	A	(0.2,0.23,0.29,0.38)
		B	(0.2,0.44, 0.62 ,0.80)
2	Logistic	A	(0.2,0.21, 0.58 ,0.80)
		B	(0.2,0.34,0.78,0.80)

Estimating the MED when outcome is continuous. The probability of selecting each dose as the MED, and the probability of correctly rejecting the null hypothesis that response at the estimated MED is equal to placebo response (Power). The best results are in bold.

Table III

Scenario	Design	d_2	d_3	d_4	d_5	Power %
Emax	Two-stage adaptive	0.19	0.67	0.10	0.04	89
	Two stage, select-one-dose	0.29	0.46	0.14	0.11	84
	Equal allocation	0.23	0.60	0.11	0.06	83
Linear in log-dose	Two-stage adaptive	0.07	0.67	0.22	0.04	90
	Two stage, select-one-dose	0.18	0.51	0.21	0.10	84
	Equal allocation	0.11	0.62	0.22	0.05	84
Linear	Two-stage adaptive	0.02	0.44	0.47	0.07	87
	Two stage, select-one-dose	0.09	0.43	0.36	0.12	83
	Equal allocation	0.03	0.45	0.45	0.07	80
Truncated-logistic	Two-stage adaptive	0.00	0.04	0.79	0.17	86
	Two stage, select-one-dose	0.02	0.15	0.64	0.19	72
	Equal allocation	0.00	0.05	0.79	0.16	76
Logistic	Two-stage adaptive	0.04	0.76	0.18	0.02	90
	Two stage, select-one-dose	0.12	0.63	0.19	0.06	84
	Equal allocation	0.06	0.76	0.17	0.02	82

Estimating the MED when two administration schedules are considered and outcome is continuous. The probability of selecting each dose as the MED, and the probability of correctly rejecting the null hypothesis that response at the estimated MED is equal to placebo response (Power). The best results are in bold.

Table IV

Scenario	Design	Group	d_2	d_3	d_4	Power %
1	Two-stage adaptive	A	0.00	0.01	0.15	89
		B	0.25	0.53	0.06	
	Two stage, select-one-dose	A	0.02	0.05	0.22	73
		B	0.24	0.35	0.12	
2	Equal allocation	A	0.00	0.02	0.17	78
		B	0.24	0.47	0.10	
	Two-stage adaptive	A	0.00	0.64	0.11	91
		B	0.14	0.11	0.00	
	Two stage, select-one-dose	A	0.03	0.44	0.13	81
		B	0.26	0.12	0.02	
	Equal allocation	A	0.01	0.60	0.10	82
		B	0.17	0.11	0.01	

Table V

Estimating the optimal dose when outcome is continuous. The probability of selecting each dose as the optimal dose, and the probability of correctly rejecting at least one null hypothesis (Power). The best results are in bold.

Scenario	Design	d_2	d_3	d_4	d_5	Power %
Quadratic	Two-stage adaptive	0.10	0.51	0.38	0.03	89
	Two stage, select-one-dose	0.13	0.45	0.36	0.06	90
	Equal allocation	0.10	0.50	0.38	0.02	80
Double-logistic	Two-stage adaptive	0.01	0.80	0.14	0.05	83
	Two stage, select-one-dose	0.04	0.66	0.20	0.10	83
	Equal allocation	0.01	0.80	0.14	0.05	73
Exponential	Two-stage adaptive	0.00	0.01	0.04	0.95	87
	Two stage, select-one-dose	0.02	0.03	0.09	0.86	83
	Equal allocation	0.00	0.01	0.04	0.95	77
Step 1	Two-stage adaptive	0.06	0.06	0.82	0.06	82
	Two stage, select-one-dose	0.11	0.10	0.69	0.10	83
	Equal allocation	0.05	0.05	0.85	0.05	71
Step 2	Two-stage adaptive	0.01	0.83	0.15	0.01	85
	Two stage, select-one-dose	0.05	0.69	0.21	0.05	84
	Equal allocation	0.01	0.82	0.16	0.01	73

Estimating the MED when outcome is binary. The probability of selecting each dose as the MED, and the probability of correctly rejecting the null hypothesis that response at the estimated MED is equal to placebo response (Power). The best results are in bold.

Table VI

Scenario	Design	d_2	d_3	d_4	d_5	Power %
Emax	Two-stage adaptive	0.17	0.70	0.10	0.04	83
	Two stage, select-one-dose	0.25	0.41	0.22	0.12	82
	Equal allocation	0.21	0.64	0.10	0.05	72
Linear in log-dose	Two-stage adaptive	0.05	0.71	0.21	0.04	82
	Two stage, select-one-dose	0.13	0.47	0.27	0.13	83
	Equal allocation	0.07	0.67	0.21	0.04	69
Linear	Two-stage adaptive	0.01	0.46	0.48	0.06	77
	Two stage, select-one-dose	0.07	0.40	0.37	0.15	79
	Equal allocation	0.01	0.45	0.47	0.06	62
Truncated-logistic	Two-stage adaptive	0.00	0.03	0.82	0.15	74
	Two stage, select-one-dose	0.02	0.10	0.63	0.25	74
	Equal allocation	0.00	0.04	0.80	0.16	55
Logistic	Two-stage adaptive	0.02	0.81	0.15	0.02	83
	Two stage, select-one-dose	0.08	0.58	0.23	0.10	84
	Equal allocation	0.03	0.80	0.15	0.02	64