Bayesian Adaptive Design for Delayed Binary Response Dose-Finding Studies

Xiaobi Huang* Haoda Fu†

Abstract

Bayesian adaptive design is a popular concept in recent dose-finding studies. The idea of adaptive design is to use accrued data to make adaptation or modification to an ongoing trial to improve the efficiency of the trial. While performing the interim analysis, most current methods only use data from patients who have completed the study. However, in certain therapeutic areas subjects are usually studied for months to observe a treatment effect. Thus, a large proportion of them have not completed the study at the interim analysis. Fu and Manner (2010) proposed a Bayesian integrated two-component prediction model to incorporate subjects who have not yet completed the study at the time of interim analysis. In this paper, we extend this method to accommodate delayed binary response and illustrate the Bayesian adaptive design through a simulation example.

Key Words: Adaptive Design; Bayesian method; Dose-finding; Delayed responses; BITP model; Emax model

1. Introduction

One major goal of Phase II dose-ranging studies is to identify the right dose for Phase III registration studies. It is one of the most challenging steps in drug development. According to Powell (2005), poor dose selection is one of the key reasons accounting for 50% failure rate in Phase III clinical studies. There are often several doses tested in Phase II studies to establish a dose-response curve. The chosen dose will be carried to the more expensive and pivotal Phase III registration studies. Traditional Phase II dose-ranging studies usually assign equal number of patients to different parallel dose arms and estimate the dose-response relationship at the study completion. Chow and Chang (2012) has shown that the traditional fixed designs are not efficient. It is also well agreed that adaptive designs can efficiently improve the learning of the dose-ranging studies (Berry et al. 2010; Krams et al. 2007). In an adaptive design, investigators can prespecify one or more interim time to analyze the data early to make adjustments to the rest of the study, such as dropping dose arms, changing allocation ratio or stopping early for futility.

There are many published references describing various adaptive designs (Jennison and Turnbull 2000; Berry et al. 2010; Chow and Chang 2012). Most of these methods only use the end of study values (endpoints) from patients who have completed the study. Such designs can be very useful when patients can complete the study quickly, or when the treatment effects can show up quickly and early responses are highly predictive of the final responses. In the later case, some traditional imputation methods such as last observation carried forward (LOCF) can be applied to impute the final responses (Weir et al. 2007). However, it is very common that clinical studies can take several months to even years to complete and the treatment effects can take a long time to show. In these studies, most patients have not completed the study at the interim time. It is very inefficient to only use patients with observed endpoints in the interim analysis. How to utilize intermediate data for patients who have not completed the study at the interim time postulates a common challenge for many adaptive designs.

*Biostatistics and Programming, Sanofi Genzyme, Cambridge, MA, USA

†Eli Lilly and Company, Indianapolis, IN, USA

Fu and Manner (2010) proposed an integrated two-component prediction (ITP) model to utilize the intermediate outcomes to predict the endpoints in the clinical trials with delayed responses. The ITP model was nicely integrated in a Bayesian adaptive design to estimate the dose-response curve. However, their method can only handel continuous outcomes. In this manuscript, we extend the method in Fu and Manner (2010) to binary outcomes and propose a binary integrated two-component prediction (BITP) model.

In an ideal hypothetical world, the sponsor of study owns the data and can retrieve and look at the data and adjust the trial design at any time. However, in the realistic world, the number of interim analyses are limited because there are many operational aspects to consider. Nowadays, many clinical trials are multi-regional with many different sites and even countries involved. It often takes weeks for the sponsor to receive the interim data after the interim data cut-off date. Then the data need to be cleaned to pass the quality control process and it will take another several weeks for data analyses and decision making. The whole process takes tremendous amount of resource. In addition, regulatory agencies often raise concerns of introduced bias to studies with too many interims. Thus, although theoretically a clinical study can have multiple interim analyses, the one-interim design is still the most popular and practical in the pharmaceutical research.

The goal of this manuscript is to propose a Bayesian adaptive design for dose-finding studies with delayed binary outcomes. Section 2 describes the BITP model and the utility based decision making. The method is evaluated through simulation and the results are presented in section 3. In section 4, we conclude this manuscript with a discussion.

2. Method

In order to handle delayed responses, we need to develop a method to utilize the information at interim to predict future outcomes at study completion. Then we will evaluate different decisions by analyzing the predictions from the interim analysis. In section 2.1, we will briefly introduce the integrated two-component prediction (ITP) model with continuous responses. Then we will generalize the ITP model to accommodate binary responses (BITP model) in section 2.2. Section 2.3 will illustrate how to make decisions to drop arms by optimizing a utility function.

2.1 Integrated Two-Component Prediction (ITP) Model for Continuous Responses

During the interim analysis of dose ranging trial, the goal is to utilize the information collected to make recommendations for the next step. The main challenge with adaptive design for delayed responses is that a large portion of subjects may have not completed the study at the time of interim analysis. Furthermore, the time course of the drug effect is usually a nonlinear curve with a diminishing return shape and the variance of the responses also changes over time. Thus, Fu and Manner (2010) designed an Integrated Two-component Prediction (ITP) model for continuous data to extrapolate the information from partially completed subjects to predict their results at the completion of the study:

$$Y_{ijl} = (\theta_i + s_{ij} + \epsilon_{ijl})\frac{1 - e^{k_i t_{ijl}}}{1 - e^{dk_i}}$$

where Y_{ijl} is the observed change from baseline from dose level i, subject j at time l; t_{ijl} is the time covariate for dose level i, subject j at time l and d is the duration of the study. This model consists of two components. The first component $(\theta_i + s_{ij} + \epsilon_{ijl})$ describes the drug effect at study completion, where θ_i represents the predicted drug effect at dose level i at study completion ($t_{ijl} = d$). This component is a mixed-effect analysis of variance

(ANOVA) model. Parameter s_{ij} is the random effect for dose i subject j, which describes the between subject variability and parameter ϵ_{ijl} is the residual error, which describes the within subject variability. The second component $(1 - e^{k_i t_{ijl}}) \setminus (1 - e^{dk_i})$ describes the time course of the treatment effect, where k_i is a shape parameter to control how quickly the drug is taking effect over the treatment time. The integration of the time effect (second component) to the drug effect (first component) enables the model to predict drug effect at study completion while subjects have not yet completed the study or have dropped out, as well as to capture the increase of variability as treatment time increases.

2.2 Binary Integrated Two-Component Prediction (BITP) Model

While Fu and Manner (2010) focused in continuous responses, in some situations the major outcomes are binary. Thus it is imperative to extend the method to handle delayed binary responses. While the delayed binary responses share similar characteristics as continuous responses: (1) many patients will not be able to complete the study at the time of interim analysis, (2) the drug effect follows a nonlinear curve with diminishing return shape over treatment time, (3) the variance of responses changes as time changes, there are still some differences compared to continuous responses. The distribution of the binary responses is specified as Bernoulli distribution or binomial distribution, in which the probability is modeled through a link function (e.g. logit link) to linear predictors, thus there is no level-1 residual error ϵ_{ijl} and the probability is determined completely by the covariates. The variability associated with the binary responses comes from the Bernoulli or binomial distribution with the probability determined by the model. Thus the within subject variability is associated with the probability and is not modeled directly. Secondly, while we usually model change from baseline if the data is continuous, there is no change from baseline when the data is binary. Thus, we designed the following Binary Integrated Two-Component Prediction (BITP) model to handle delayed binary responses:

$$logit(P_{ijl}) = a_{0i} + (\theta_i + s_{ij}) \frac{1 - e^{k_i t_{ijl}}}{1 - e^{dk_i}}$$
$$Y_{ijl} \sim Bernoulli(P_{ijl})$$

where Y_{ijl} is the observed binary outcome for dose level i, subject j at time l and P_{ijl} is the corresponding probability of event occurring; t_{ijl} is the time covariate for dose level i, subject j at time l and d is the duration of the study. Parameter a_{0i} are the logit of the baselines for each subject, which is invariant over time. Similar to the continuous ITP model, the BITP model also consists of two components. The first component is drug effect, which is a generalized mixed effect ANOVA model. The parameter θ_i represents the drug effect of does level i at study completion and s_{ij} is the random effect for subject j at dose level i. There is no individual residual error in this component. The second component is the time effect, which is similar to the time effect in the continuous ITP model. Figure 1(a) describes how the second component, the time effect, extrapolates the drug effect during a 3-month study via a diminishing return curve. The larger the absolute value of the shape parameter k, the quicker the drug takes effect and reaching a steadier state. Figure 1(b) illustrates the drug effect (the logit of the interested probability) over time. Parameter theta decides effect size, while the time effect component extrapolates it across the treatment time.

We choose to use Bayesian method to fit the BITP model. There are many ways to

choose priors. If history data is limited, we can always use the noninformative priors as:

$$\theta_i \propto 1$$
$$k_i \propto 1$$
$$a_{0i} \propto 1$$
$$s_{ij} \sim N(0, \tau^2)$$
$$log(\tau^2) \propto 1$$
$$log(\sigma^2) \propto 1$$

A Markov Chain Monte Carlo (MCMC) algorithm was used to sample the posterior parameters.

2.3 Adaptive Design Decision Making

Based on the interim data I, we summarize the interim information from the posterior distribution $P(\Theta|I)$ of the BITP model, where Θ represents all parameters of interests. Among a set of all possible decisions $\{\delta|\delta = 1, 2, , n\}$ such as dropping dose arms or change allocation ratio, we can evaluate each decision delta by obtaining the posterior predictive distribution and compare different decisions by optimizing a utility function. We illustrate how the optimal decision is chosen in the following steps:

1. Fit BITP model by using the interim data to obtain the posterior distribution $p(\Theta|I)$ and generate M sets of parameters from the posterior distribution, each generated at an iteration of the MCMC sampler.

2. For each decision delta, use the M sets of posterior parameters to generate M sets of posterior predictive data sets under the posterior predictive distribution:

$$p(Y_\delta|I) = \int p(Y_\delta|\Theta)p(\Theta|I)d\Theta$$

   Thus, we have M sets of predictive trial observations for each decision δ, denoted as $\{Y_\delta^j\}_{j=1}^M$.

3. A utility function $g(.)$ is then constructed and calculated under each predictive data $\{Y_R^j\}$.

4. Estimate $E(g(Y_\delta)|I)$ by taking the average of the $\{g(Y_\delta^j)\}_{j=1}^M$, to minimize the Monte Carlo error.

5. Estimate $E(g(Y_\delta)|I)$ for each decision delta and choose the optimal decision which maximize the utility function:

$$\hat{\delta} = argmax_{\{1,2,\ldots,n\}} E(g(Y_\delta)|I).$$

There are many choices for the utility function depending on the questions to answer by the interim analysis. In dose ranging trials, the most important question is which decision will be the most efficient in estimating the dose response curve. Often the primary goal is to maximize our learning on ED_{90}, which is the minimum dose to achieve 90% of the maximum drug efficacy. ED_{90} can be estimated from an Emax model (Rowland and Tozer, 1995). By denoting ED_{90} as a function of predictive observations Y_δ as $E\hat{D}_{90} = h(Y_\delta)$, we can define our utility function as $g(Y_\delta) = -\{h(Y_\delta) - E[h(Y_\delta)|I]\}^2$, and $E\{g(Y_\delta)|I\} = -var\{h(Y_\delta)|I\}$. The expectation of this utility increases as the precision of estimated ED90 increases. More details will be illustrated via the simulation example in Section 3.

3. Simulation and Results

We illustrate how to use the BITP model and the utility based decision rule to conduct dose-finding clinical studies with delayed binary responses through a simulation example.

3.1 Study Design

Consider a randomized, double blinded, parallel-group placebo controlled Phase II dose-finding clinical study. The primary objective is to learn the dose-response relationship to find a clinical dose for Phase III development. The planned sample size is 324 patients with 36 subjects allocated to the placebo arm and 288 patients allocated to active treatment arms. Initially, there are 8 treatment arms and 1 placebo arm with equal number of subjects allocated. The patients are enrolled and screened during prestudy period and randomized to the placebo arm or one of the 8 treatment arms in a staggered manner. Once the patients started the study, they will be administered with the allocated treatment from week 0 to week 12 and will have biweekly visits during the treatment period. Thus there will be 7 measurements taken for each patient including baseline. An interim analysis is planned at some time during the study and after which 4 less informative active treatment arms will be dropped and the incoming patients will be randomized to the 4 more informative arms and the placebo arm. We assume the enrollment rate is approximately 18 patients per month, so there are roughly 21 months from the first group of patients started the treatments until the last group of patients finished the treatments.

3.2 Simulation Setups and Implementations

Most pharmacodynamics variables exhibit a nonlinear dose-response nature which is often described by the Emax model (Rowland and Tozer 1995; Shargel et al. 2012). Because our response is binary, we assume the logit of the probability of interested responses by the end of the study follows a true dose-response curve of an Emax shape:

$$logit(P) = E_0 + \frac{E_{max} dose^h}{ED_{50}^h + dose^h}$$

with $E_0 = -2, E_{max} = 4$ translated to $\sim 12\%$ probability at baseline and a maximum of $\sim 88\%$ probability at maximum efficacy. We assume the true ED_{50}, the minimum dose to achieve 50% of the efficacy, is 40 mg, and the hill parameter $h = 1$. Figure 2 displays the true dose response curve under this model.

In reality, our therapeutic window is limited due to safety and other concerns. Suppose we are going to dose as high as the highest safe dose identified from Phase I studies and evaluate the adaptive design under three dosing scenarios in our simulation, where the highest safe dose is 400 mg under the low dose range scenario, 600 mg under the median dose range scenario and 1000 mg under the high dose range scenario:

Low dose range: 5, 10, 20, 40, 60, 100, 200, 400 mg;

Median dose range: 5, 15, 40, 80, 120, 200, 400, 600 mg;

High dose range: 20, 40, 80, 160, 300, 500, 800, 1000 mg.

Figure 3 illustrate how doses are distributed on the true response curve under these scenarios. Under low dose range, the drug effect just starts to approaching plateau and only a portion of the dose-response curve is covered, while under the high dose range, the drug effect has almost achieved the plateau and a larger portion of the dose response curve is covered and the medium dose range lies in between.

In addition, we assume the shape parameter for time effect $k_i = -0.75$. The random effect $s_{ij} \sim N(0, \tau^2)$, where $\tau = 0.45$. The missing data mechanism in this study is missing at random (MAR) with 22.5% dropout rate. After a patient is dropped out, all of his future data are missing. We chose to perform the interim analysis at month 10, where more than half of the patients have enrolled in the study and there are a good portion of patients who have completed the study. Table 1 illustrates how subjects are allocated at month 10.

For one simulated data, with noninformative priors, parameters of the BITP model are estimated using the Markov Chain Monte Carlo (MCMC) algorithm at the time of interim analysis. To illustrate in more details, after burn-in, we obtain 1000 posterior samples of parameters of the BITP model $\{a_{0i}, k_i, \theta_i, \tau_i, i = 1, 2, ..., 9\}$. To select 4 treatment arms to keep from 8 treatment arms, there are 70 possible combinations ($C_8^4 = 70$). For each combination of treatment arms, we assign the rest of the subjects in those kept arms and obtain 1000 posterior predictive trial data based on the 1000 sets of posterior parameters. For each predictive trial data, an Emax model is fitted to the last observation of each patient to estimate ED_{50}, thus for each combination of treatment arms, we have 1000 estimated ED_{50}. In this example, we consider ED_{50} to be our parameter of interest, because for a fixed h, ED_x is proportional to ED_{50}. We can derive ED_X from $E_{max}ED_X^h/(ED_{50}^h + ED_X^h) = X\%E_{max}$ and $ED_X = ED_{50}(1-X\%)^{-(1/h)}$. Therefore, the precision of ED_{50} estimate directly indicates the precision of ED_{90} or ED_{80} estimates. Thus, we choose the treatment combination with the most precise estimate of ED_{50}:

$$argmin_{\delta \in \{1,2,...,70\}} var\{h(Y_\delta)|I\}$$

where $\hat{ED_{50}} = h(Y_\delta)$. For each given interim time and dosing scenario, we repeat the above simulation for 500 times and evaluate the loss by the loss function:

$$Loss = \frac{1}{500}\sum_{i=1}^{500}(\hat{ED_{50,i}} - ED_{50})^2$$

For comparison purpose, we also evaluate the same loss function for the traditional parallel design with 1 placebo arm and 8 treatment arms (N = 36 / arm).

3.3 Simulation Results

Figure 4 displays the simulation results of interim analysis for different dosing scenarios. The box plots present differences of the estimated and true ED_{50}. The figure suggests that the adaptive design is more efficient than the traditional design. Table 2 summarizes the numerical comparison between adaptive design versus traditional design under the three dosing scenarios. Because the low dose range scenario covers the smallest portion of the dose response curve among the three dosing scenarios simulated, the learning of the curve is less efficient than the other two scenarios and we observed a smaller saving compared to the other two scenarios. Also, one can see that the loss is the smallest under the high dose range scenario and the largest under low dose range scenario, suggesting having doses approaching dose-response plateau will enhance the learning accuracy. Compared to traditional design with equal number of subjects (N=324). The adaptive design reduced the loss by approximately 5%, 21% and 19% under the low, medium and high dose range scenarios. When comparing the adaptive design (N=324) to the traditional design with more number of subjects (N = 360), adaptive design still produces comparable losses.

4. Discussion

In order to utilize as much information as possible to better estimate the dose-response curve in Phase II delayed binary response dose-finding trials, we have proposed a BITP model to predict the clinical trial data using interim information. This model can be considered as an extension of the ITP model by adapting it to incorporate the binary data. We then apply a utility based decision rule to decide the adjustment of the future trial design. In addition, we have discussed the importance of timing of the interim analysis.

The BITP model is able to deal with generalized nonlinear longitudinal models with binary responses, as well as to incorporate subjects who have not completed the study at the time of interim analysis, which account for a large portion in the delayed response study. Subjects enrolled after interim analysis are able to be allocated to more informative arms, thus to improve the accuracy of estimates. Simulation results suggest that the adaptive design improves the learning accuracy of the dose-response under three dosing scenarios when the interim adaption is conducted at an appropriate time. The accuracy of estimating the dose-response also relies on the dosing scenario. Generally, the more the dose-response curve is covered, the better we can estimate the dose response curve. However, in reality our therapeutic window is often limited by safety and other concerns. Our simulation suggests that the Bayesian adaptive design has the largest saving compared to the traditional design when the therapeutic window is limited.

One concern of the adaptive design is the timing of the interim analysis. The optimal timing is an optimal allocation of information collected before and after the interim in order to optimize our learning of the dose-response curve. It is very important to guarantee sufficient information at the interim time so that the BITP model can give a more accurate prediction of the trial data to help us making correct interim decision. Binary responses are associated with more uncertainty compared to continuous responses, thus when the responses are binary, the interim analysis should be scheduled later to accumulate sufficient information.In our example, we chose to perform the interim analysis at month 10, where more than half of the planned patients have enrolled in the study and a good portion of them (135 out of 324, 42%) have completed the study. This will also depend on the enrollment rate and other characteristics of the study and the next step of this research is to conduct simulation according to the actual trial characteristics to determine the interim time.

REFERENCES

Berry,S. M., Carlin, B. P., Lee, J. and Muller, P. (2010), *Bayesian Adaptive Methods for Clinical Trials* (1st ed.), Boca Raton, FL: Chapman & Hall/CRC.

Chow, S. C. and Chang, M. (2012), *Adaptive Design Methods in Clinical Trials* (2nd ed.), Boca Raton, FL: CRC Press.

Fu,H. and Manner, D. (2010), "Bayesian Adaptive Dose-Finding Studies with Delayed Responses," *Journal of Biopharmaceutical Statistics*, 20, 1055–1070.

Jennison, C. and Turnbull, B. W. (2000), *Group Sequential Methods with Applications to Clinical Trials* (1st ed.), Boca Raton, FL: Chapman & Hall/CRC.

Krams, M., Burman, C. F., Dragalin, V., Gaydos, B., Grieve, A., Pinheiro, J. and Maurer, W. (2007), "Adaptive Designs in Clinical Drug Development: Opportunities, Challenges, and Scope Reflections following PhRMAs November 2006 Workshop," *Journal of Biopharmaceutical Statistics*, 17(6), 957–964.

Powell, B. (2005), "FDA Experience with End of Phase IIa meetings: An Attempt to Improve Development Decisions," http:www.fda.govohrmsdocketsac05slides20054194S1_08_Powell.ppt.

Rowland, M. and Tozer, T. N. (1995), *Clinical Pharmacokinetics: Concepts and Applications* (3rd ed.), Baltimore, MD: Williams & Wikins.

Shargel, L., Yu, A. and Wu-Pong, S. (2012), *Applied Biopharmaceutics & Pharmacokinetics* (6th ed.), New York: McGraw-Hill Medical.

Weir, C. J., Spiegelhalter, D. J. and Grieve, A. P. (2007), "Flexible Design and Efficient Implementation of Adaptive Dose-Finding Studies," *Journal of Biopharmaceutical Statistics*, 17(6), 1033–1050.

Table 1: Allocation of subject with different timing of the interim analysis.

Interim time (months)	No. subjects enrolled at interim time	No. subjects in each droppedl treatment arm	No. subjects in each kept treatment arm
10	180	20	52
Subjects are allocated to the placebo arm in the same time, the total number of subjects receiving placebo does not exceed 36 subjects.			

Table 2: Numerical Results of Interim Analysis Under Various Dosing Scenarios.

Dosing scenarios	Interim time (months)	No. subjects in interim analysis	No. subjects in each dropped arm	No. subjects in each kept arm [1]	No. subjects completed the study at interim	Average loss for adaptive design (N = 324)	Average loss for traditional design (N=324)	Average loss for traditional design (N=360)
Low dose range	10	180	20	52	135	1201.07	1263.19	1114.45
Medium dose range	10	180	20	52	135	702.35	891.09	678.76
High dose range	10	180	20	52	135	545.40	671.34	567.73

[1]. Number of subjects in each dropped arm is also the number of subjects per arm at interim.

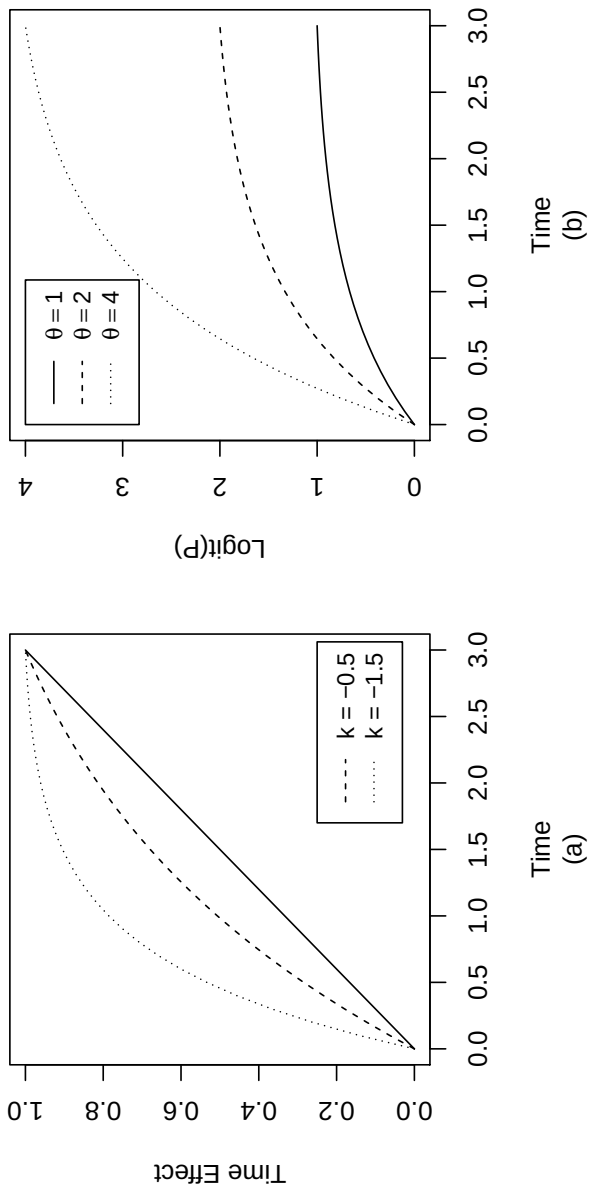

Figure 1: Illustration of time effect and drug effect: shape parameter k governs how how quickly the drug takes effect; parameter θ is the drug effect after 3 months of dosing. Panel (a) illustrates the time effect and how it changes as k changes, $k = -0.5, -1.5$. Panel (b) illustrates the drug effect over time. K is fixed at -1 and $\theta = 1, 2, 4$.

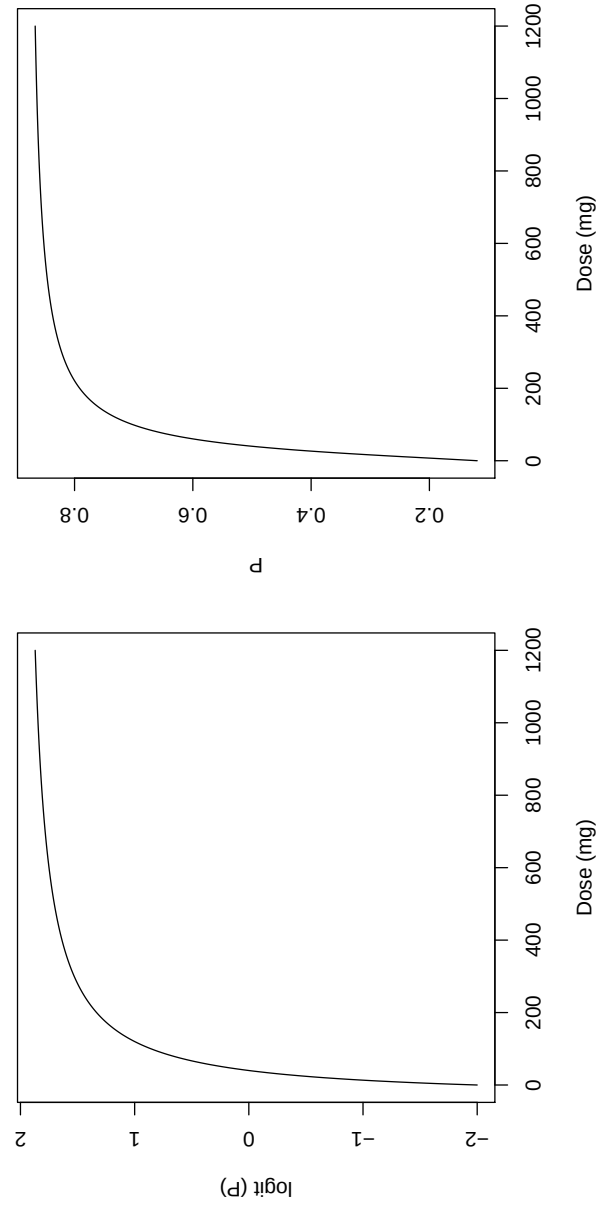

Figure 2: True dose response curve follows an Emax model with $E_0 = -2$, $E_{max} = 4$, $ED_{50} = 40$ mg and $h = 1$. Panel (a) is the dose-response relationship between $logit(P)$ and dose; Panel (b) is the dose-response relationship between the interested probability P and dose.

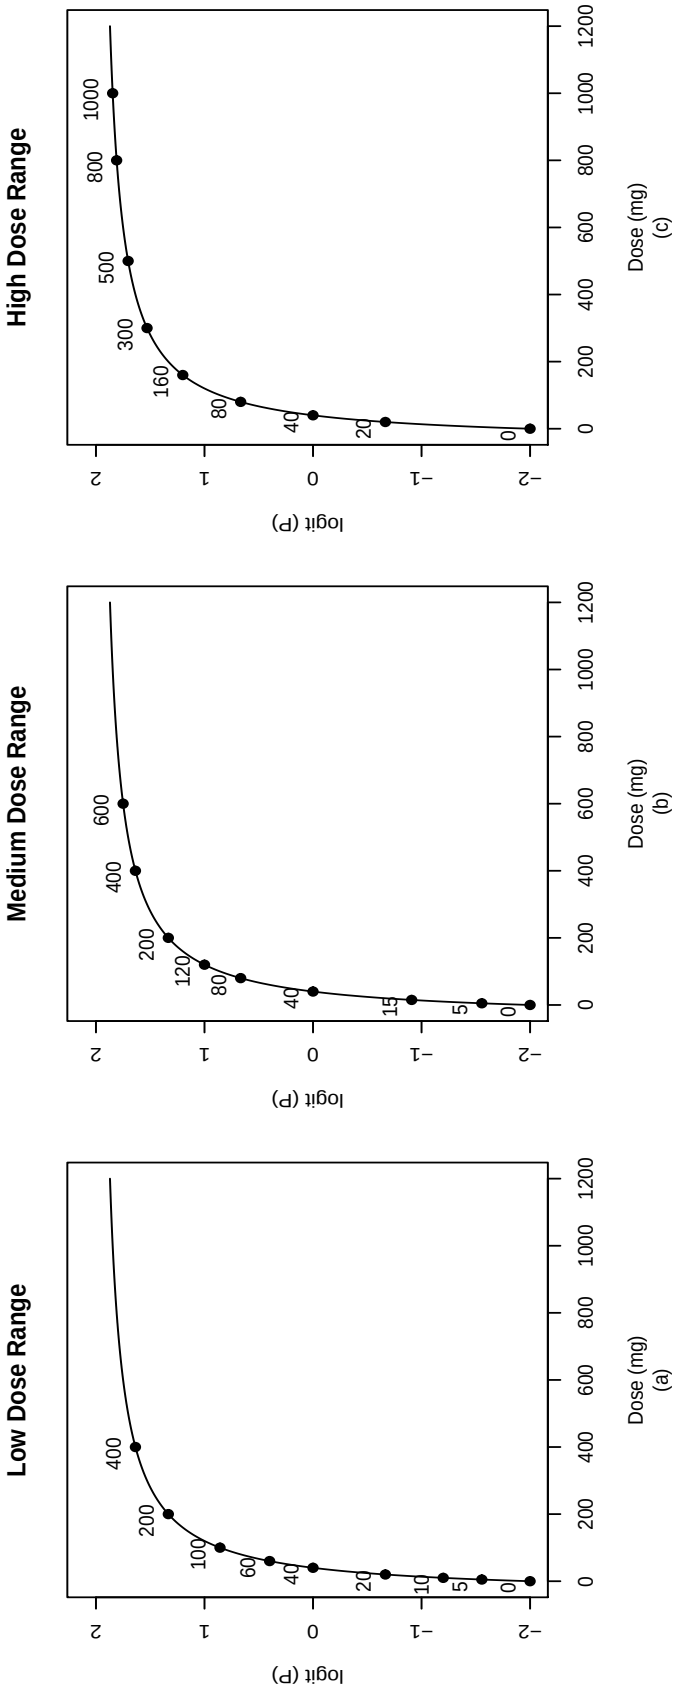

Figure 3: Distribution of doses on the true dose-response curve under different dosing scenarios.

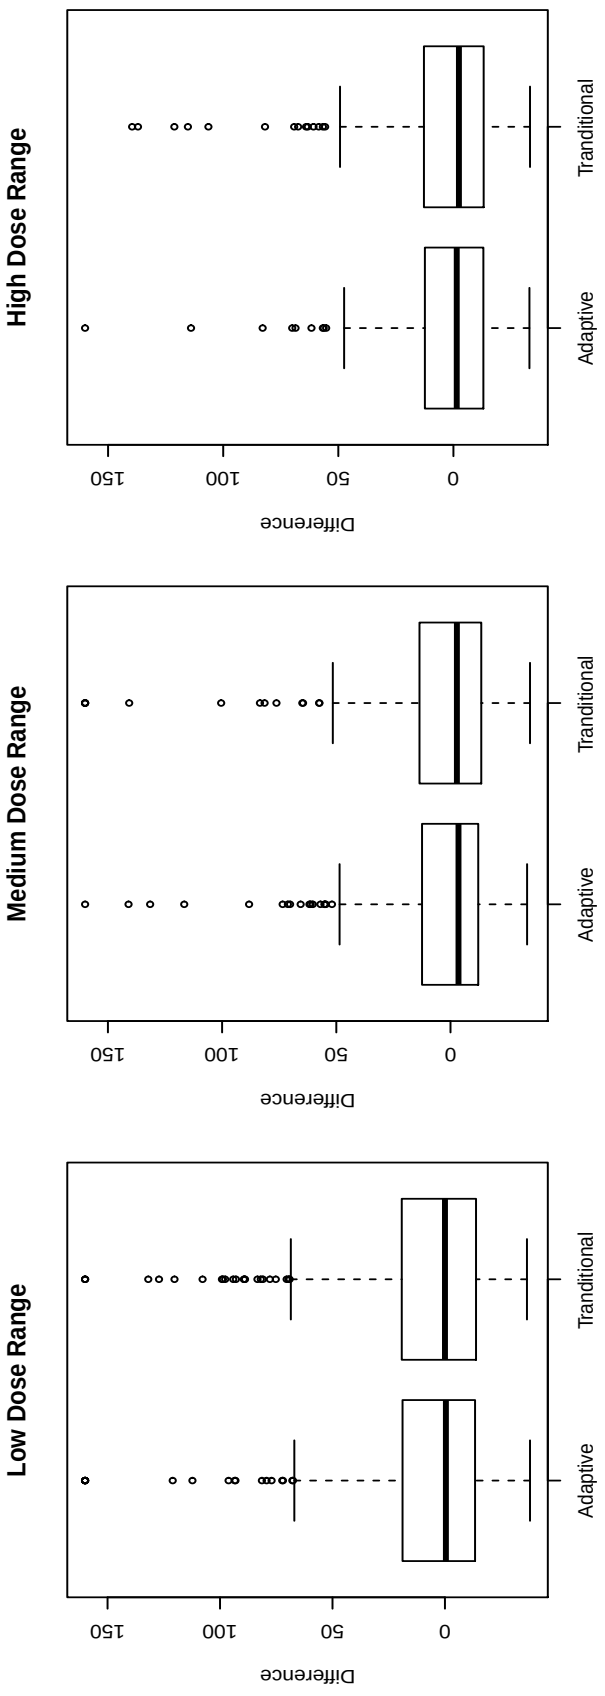

Figure 4: Performance of adaptive design of interim analysis and comparison to traditional design. The interim analysis is conducted at month 10. The panels are the boxplots of differences between the estimated and true ED_{50}. Panel (a) presents results from the low dose range scenario. Panels (b) presents results from the medium dose range scenario. Panels 3 presents results from the high dose range scenario.