Bayesian Adaptive Randomization for A Phase 2 Dose-Ranging Clinical Trial

Morgan de Ferrante[1], Ming-Dauh Wang[2]
[1]Edwards Lifesciences
[2]Regeneron Pharmaceuticals

Abstract
Bayesian Adaptive designs allow for key study design specifications to be modified as information is collected from observed data during a clinical trial. When carried out in the proper circumstances, adaptive designs can lead to more efficient allocation of resources and shorter study spans. Bayesian Adaptive Randomization, specifically, establishes a well-defined framework for updating allocation ratios based on dose performance, allowing to increase the probability of treatment allocation to more promising doses of a study drug, while also preserving the benefits of randomization. To study potential advantages of adaptive randomization and its statistical properties, we simulated data for a planned Phase II dose-ranging study of an experimental treatment for pain due to osteoarthritis. Bayesian interim analysis was conducted to implement adaptive randomization using the following three efficacy dose-response models for comparison: the Emax model, normal dynamic linear model (NDLM), and analysis of variance (ANOVA). Adaptive randomization was conducted based on predictive treatment efficacy and safety, and with the intention of avoiding exposure to unnecessary risk when high dose safety is uncertain. Through our simulation, we showed increased efficiency by our proposed approach to adaptive randomization that enables optimization of patient allocation balancing between efficacy and safety, as well as decreased enrollment of patients to unnecessary high doses.

Key Words: Bayesian statistics, adaptive design, adaptive randomization, interim analysis, predictive probability, pain due to osteoarthritis

1. Introduction

Adaptive randomization in a clinical trial considers mid-course modification of randomization ratios for doses studied in the trial by rewarding higher randomization probabilities to doses that are more likely to be successful at the end of the trial. For efficacy it balances resource assignment toward doses more promising to be efficacious, thus enhancing efficiency for cost-effectiveness. For safety it avoids exposing patients to doses with higher dose related adverse reaction liabilities. Seeing the potential benefits of adaptive randomization, we proposed a Bayesian adaptive randomization design to a planned Phase 2 dose-ranging clinical trial of an experimental treatment for pain due to osteoarthritis. The Bayesian approach allows utilization of all cumulative data for interim randomization ratio adjustment, and it better facilitates utilization of prior information to inform the adjustment. We differentiate our proposed approach by considering penalty for doses that are predicted to provide limited additional efficacy but are liable to increased safety concerns. We studied the performance of the proposed approach by simulation under various scenarios, for which a more advanced clinical program of another drug candidate with a different mechanism for the same disease indication provided benchmark information.

We delineate the methods of our proposed Bayesian adaptive randomization approach in the next section. Results of our simulations for the study of the performance of the approach are summarized in Section 3. Section 4 presents impact of variations of design parameters on the design performance. Conclusions and discussion of our research are given in Section 5.

2. Methods

2.1 Simulations

Simulations were developed to incorporate both safety and efficacy information into the adaptive randomization design. Trials were replicated (N = 500) with 5 different cohorts of 'patients' entering the study with prespecified sample sizes in each cohort. Each cohort would have 80 patients, leaving a total of 400 patients in each trial. The primary efficacy outcome of interest was change from baseline in the Western Ontario and McMaster Universities Osteoarthritis Index (WOMAC), with which a higher reduction is expected as indicating a greater decrease in pain. Additionally, there was a dose limiting adverse event of particular concern.

Efficacy data were simulated from an Emax design model (see Table 1), using information from a previous Phase II clinical trial for the same indication of pain due to osteoarthritis, with doses of 0 units (placebo), 1 unit, 3 units, 6 units, and 9 units. The Emax dose response curve was based on the following values: ED50 = 1, E_{max} = - 1.4, E_0 = - 2.25 and σ = 2.4. E_{max} refers to the maximum effect a dose can have, and ED50 is the dose that achieves 50% of that maximum effect. E_0 is the zero dose, or placebo, effect. For efficacy, we define an active dose to be efficacious if the effect size (ES = (μ_0 - μ_i) / σ) is greater than 0.4, where μ_0 is the response for those receiving placebo and μ_i is the response for doses i = 1,2,3,4. Thus, we have the following utility measure for efficacy of a dose:

$$Probability\ of\ Success\ for\ Efficacy = \Pr(ES > 0.4)$$

Safety data were simulated under the assumption of a binomial model (see Table 1), using similar probabilities of experiencing the adverse event of interest observed from a prior Phase 2 study of another clinical program. That is, for doses of 0 units, 1 unit, 3 units, 6 units and 9 units, we used p = 0.01, 0.025, 0.05, 0.075, and 0.125, respectively, to generate simulated data. The utility measure for safety of a dose is thus defined as:

$$Probability\ of\ Success\ for\ Safety = 1 - \Pr(adverse\ event) = 1 - p$$

The probability of success (POS) of a dose is defined below as a function of both the probability of success for efficacy and the probability of success for safety:

$$POS = (Probability\ of\ Success\ for\ Efficacy) * (Probability\ of\ Success\ for\ Safety)^4$$

Dose allocation ratios were then based on standardized POS's across the doses. Probability of success for safety in the POS was initially raised to the first power, but we found it would not allow safety data to have enough impact on allocation ratios. After more simulations, we found raising it to the fourth power was ideal for our case in that it sufficiently decreased allocation to "unsafe" doses, while also not over-diminishing their standardized POS values.

Table 1. Design Models

Efficacy Design Model	$\mu_i = E_0 + \frac{E_{max} * dose_i}{ED50 + dose_i}$, $\varepsilon_i \sim N(0, \sigma)$, $Response_i = \mu_i + \varepsilon_i$
Safety Design Model	$X_i \sim \text{Bin}(n_i, p_i)$ p_i = *probability of experiencing an adverse event at dose i*; n_i = *number of subjects at dose i*

2.2 Analysis

For efficacy, three Bayesian analysis models were constructed to analyze the simulated data in the context of an adaptive design. Our first analysis model is the Emax model, which is non-linear and commonly used in analysis of pharmacodynamic data. With four parameters in the model (see Table 2), the Emax model is ideal for situations with more than four comparative doses. One major benefit of the Emax model is that the parameters are easy to interpret and understand, but we must note an assumption of this model is that the dose-response relationship is monotonic.[2] It will be the smoothest of our three models. The next model is the normal dynamic linear model (NLDM), which is more complex in specifications. NDLM assumes the change in mean from one dose to another follows a linear model, and the slope (ξ) is updated by a random walk. NDLM borrows information from neighboring doses to estimate mean responses, and the treatment response at one dosage level is related to the treatment response of the previous dose level.[3] This model does not require monotonicity of the response curve and still allows for smoothing.[4] NDLM is more flexible but is more difficult to interpret and less commonly used. The final model is analysis of variance (ANOVA), which does not require monotonicity of the dose-response curve and is the simplest of our three models. ANOVA does not allow for smoothing of the dose-response curve. For safety, simulated data were analyzed by logistic regression using Bayesian analysis. Although the safety data were analyzed using only one model, cumulative sample sizes across doses differed for the three efficacy analysis models through adaptation, and thus different posterior distributions for the safety model parameter were generated depending on the choice of the efficacy analysis model. Efficacy and safety analysis models are outlined in Table 2.

For each individual trial, data for 400 individuals were simulated. The first cohort of 80 subjects used allocation to each dose with equal probability (n = 16 to each dose). Efficacy and safety responses for this cohort were analyzed using the three efficacy analysis models and the safety analysis model, for which Markov Chain Monte Carlo (MCMC) simulation using a Gibbs sampler generated posterior distributions for the efficacy effect sizes and adverse event rates of the doses. Due to convergence issues, the Emax analysis model was built using semi-informative prior distributions as well as non-informative priors for the model parameters. This often occurred in cases of lower ED50 in the design model (see "Adjustments on Design Model Parameters" on Page 8) likely due to lack of information below ED50. All remaining efficacy and the safety analysis models had non-informative prior distributions. POS for each dose was calculated using the resulting posterior distributions for the effect size and adverse event rate of the dose, and the probability of treatment allocation to the dose for the next cohort of 80 subjects is proportional to the POS of the dose as compared to those of other doses. For the third cohort, POS for each dose

was calculated using the responses from all 160 subjects cumulative in the trial, and the updated POS's for all doses using both efficacy and safety data determined the third cohort's allocation ratios. This process was repeated for the fourth and final cohort.

Table 2. Analysis Models

Efficacy Analysis Models	**Emax Model**	$Y_{ij} \sim N(\mu_i, \sigma^2)$ $\mu_i = E_0 + \frac{E_{max} * dose}{ED50 + dose}$
	Normal Dynamic Linear Model (NDLM)	$Y_{ij} \mid dose_i \sim N(\theta_i, \sigma^2)$ $\theta_i \sim N(\theta_{i-1} + \xi_{i-1}, \tau_i^2)$ $\xi_i = \xi_i + \tau_i$ $\tau_i = \tau * (dose_i - dose_{i-1})$
	Analysis of Variance (ANOVA)	$Y_{ij} = \mu + \tau_j + \varepsilon_{ij}$ $\varepsilon_{ij} \sim N(0, \sigma^2)$
Safety Analysis Model	**Logistic Regression**	$logit(p_i) = \beta_0 + \beta_1 * dose_i$

2.3 Incorporating a High-Dose Penalty

When safety information is difficult to collect or estimate, due to lack of historical data or long latent periods before adverse event onset, we may need to rely on other methods to reduce the risk of adverse events for patients. We propose incorporating a high-dose penalty to address this issue. For doses that have already reached a high threshold of efficacy, for example doses at or above ED75 (which corresponds to the dose that achieves 75% of the maximum dose effect, i.e. 75% of E_{max}), we may decrease allocation to those that have similar efficacy and are unnecessarily high (to avoid regulatory and potential unknown safety concerns). As an example, in our simulation case doses of 3, 6 and 9 units are all above ED75, so we can decrease the probability of allocation to doses of 6 and 9 through adaptation (see Figure 2). In Figure 1 showing the method steps, E75 is the efficacy response value that corresponds to ED75. ψ is a parameter that determines the impact of the high-dose penalty. A higher value of ψ will correspond to a greater reduction in allocation to doses affected by the high-dose penalty. In our simulation, we used $\psi = 100$.

Figure 1. Method Steps for High-Dose Penalty

Method Steps

1. For each simulated trial, identify response values above E75 -> $r_1, r_2, \ldots$
2. Identify r_{min} as the smallest of $r_1, r_2, \ldots$, i.e. $r_{min} = \min(r_1, r_2, \ldots)$
3. Calculate the high dose penalty for doses corresponding to r1, r2, … using the following formula:

$$W_i = \frac{1}{\psi} \frac{(r_{min} - r_i)^2}{(E75 - E_{max})^2}$$

4. Multiply the POS calculated from the adaptive method by either 1 or W_i for each active dose (1 when no penalty is necessary and W_i when a penalty is to be added)
5. Calculate the allocation probabilities

In Figure 2, r_{min} is the response value for the 3 units dose, and r_2 and r_3 correspond to 6 and 9 units. W_i will penalize doses with response values closer to r_{min} much more heavily, i.e. the closer the dose response value is to r_{min}, the more penalty we apply for randomizing to the corresponding dose.

Figure 2. High-Dose Penalty with Emax Model

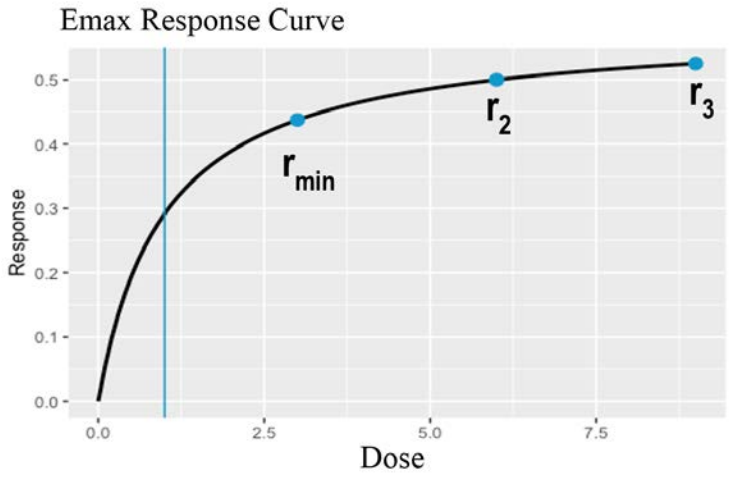

3. Simulation Results

3.1 Results from Adaptation Considering Efficacy and Safety

The density plots in Figure 3 show the posterior distributions of effect sizes for each dose at the end of one simulated trial through adaptive randomization. The vertical line corresponds to an effect size of 0.4, making the area under the curve to the right of the vertical line as the probability of success for efficacy for each dose. For example, we can see that by the Emax model and NDLM, a dose of 1 unit was estimated to have around

10% probability of success at the end of the trial, whereas by ANOVA, we see about 20% probability of success. Density curves for doses are more spread out by the Emax model, showing greater disparities in effect size between doses with this model than with NDLM and ANOVA.

Figure 3. Posterior distributions for efficacy for one simulated trial

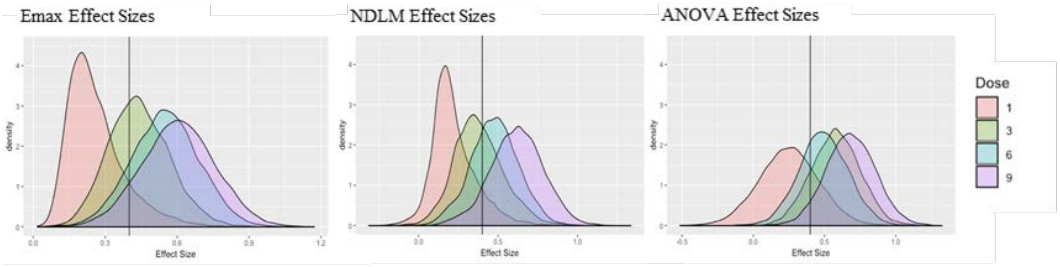

In Figure 4, we can see the average effect sizes for each dose across 500 simulated trials, allowing us to compare the effect sizes calculated in each model to the "true" dose-response curve. Likely due to lack of information below the true ED50 by the Emax model and model smoothing by NDLM, there was underestimation of the dose response curve. Overall Emax model and ANOVA were less biased in terms of effect sizes.

Figure 4. Average effect sizes across doses as compared to the "Truth"

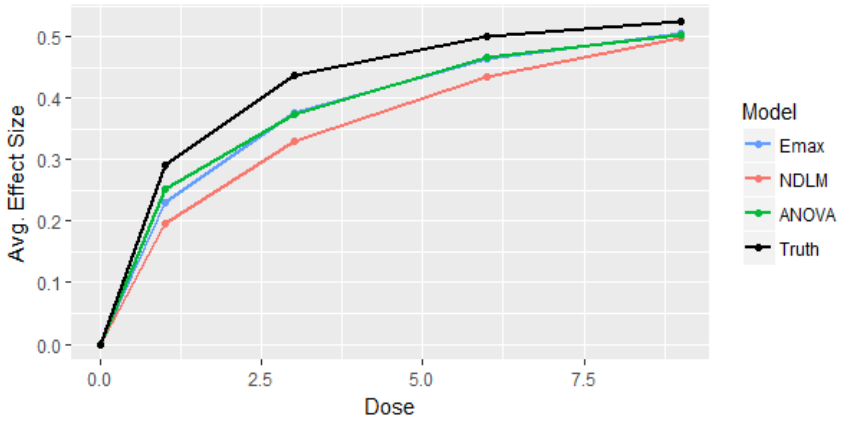

After calculating an average probability of success for each dose across 500 simulated trials, we can observe the cumulative subject totals after 5 cohorts of 80 subjects (see Figure 5). The probability of subject allocation for placebo was fixed at 0.2, i.e., we always had 16 subjects allocated to placebo for each cohort. By the Emax model, after the initial randomization only a few subjects were allocated to the 1-unit dose in each subsequent cohort. The Emax model and NDLM resulted in similar randomization probabilities, and thus had similar cumulative subject totals in each dose. Adaptive randomization by the ANOVA model exhibited higher robustness against adaptation and tendency toward preserving equal randomization.

Figure 5. Average effect sizes across doses as compared to the "Truth"

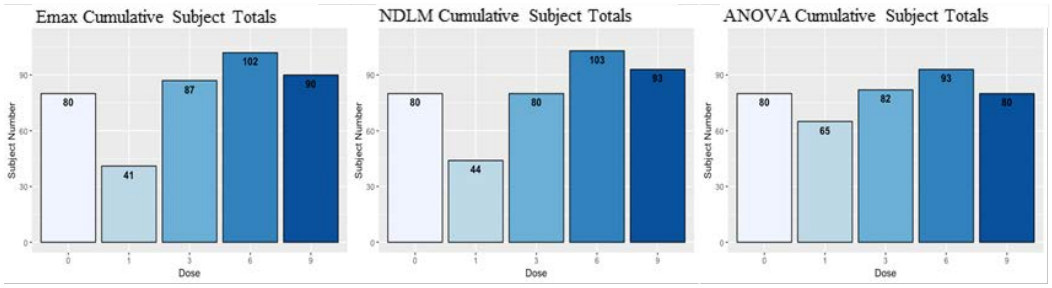

3.2 Results with Consideration of High-Dose Penalty

As we can see from Figure 6, imposing a high-dose penalty can reduce the probability of allocation to higher doses, and in our scenario, it mimics the results of the design incorporating safety data. The penalty affects allocation probabilities for doses above 3 units in a dose-increasing manner. Figure 7 demonstrates the effect of these allocation probabilities on the cumulative subject totals of a simulated trial. Adding a high-dose penalty resulted in about a 19% reduction in the number of subjects allocated to the highest dose, as compared with the design incorporating safety data having a 14% reduction in allocation to the highest dose (Figures 5 & 7). Note the proposed high-dose penalty is only applied to the Emax analysis model, since NDLM and ANOVA do not have Emax and ED50 as built-in parameters, and other forms of penalty might be proposed for the NDLM and ANOVA analysis models.

Figure 6. Average allocation probabilities for different adaptive designs across doses

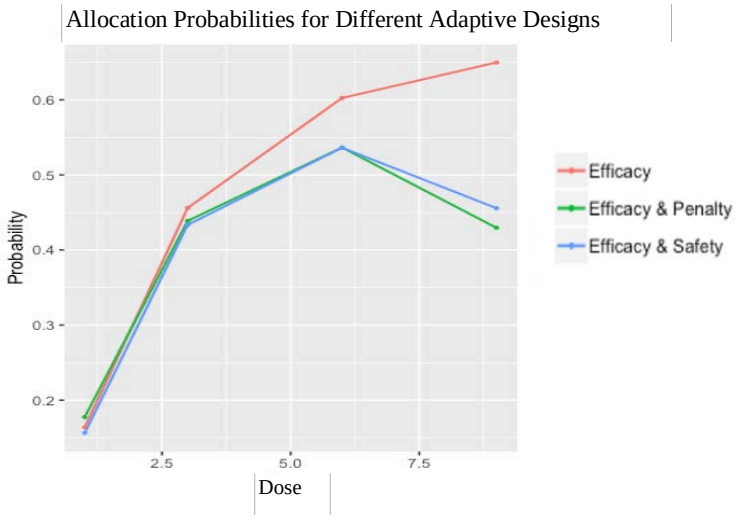

Figure 7. Cumulative subject totals of a simulated trial with and without a high-dose penalty

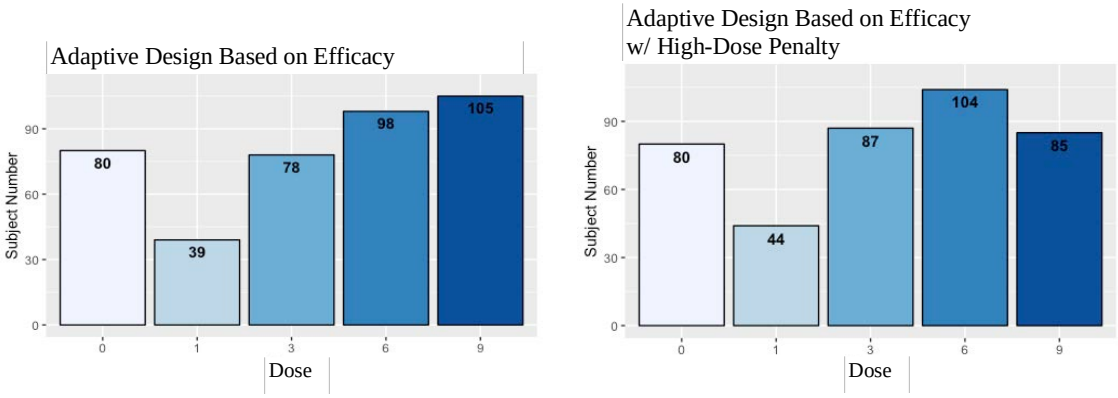

4. Adjustments in Design Model Parameters

In addition to comparing results by the three analysis models in the context of our Phase II osteoarthritis mimicked data, we wanted to see how slight variations in the design model parameters would affect our results. The plot in Figure 8 shows the assumed "truth" for each design model with varying ED50 (which greatly affects the slope and shape of the dose-response curve).

Figure 8. Effect sizes of the true Emax curve for varying ED50 values

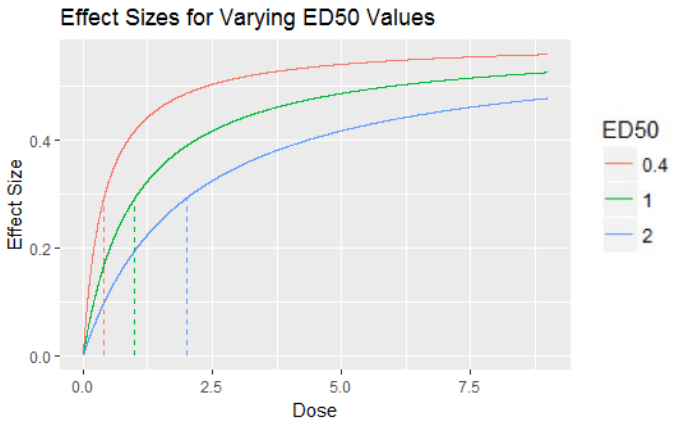

In looking at average effect sizes across all 500 trials (see Figure 9), when ED50 in our "true" model was reduced to 0.4, the ANOVA analysis model became our least biased model in terms of average effect sizes across doses. When ED50 in the design model was increased to 2, the Emax model had the least biased average effect sizes. Across all the models, ED50 = 0.4 leads to the highest unstandardized overall probability of success. In looking at cumulative subject totals in the three design model scenarios, we discovered that higher ED50 in the design model would lead to increases in allocation to higher doses and a much more drastic effect on allocation probabilities (when compared to traditional equal randomization designs). For example, by the Emax model with ED50 of 0.4, we have a resulting 52 subjects allocated to the 1-unit dose. When the design model was constructed

with ED50 of 2, there were only 29 subjects allocated to the 1-unit dose. ANOVA and NDLM were less sensitive to adjustments in ED50 of the design model, telling us the slope/shape of the "true" dose-response curve will have a lesser impact on these models.

Figure 9. Average of mean effect sizes across doses for varying ED50 values

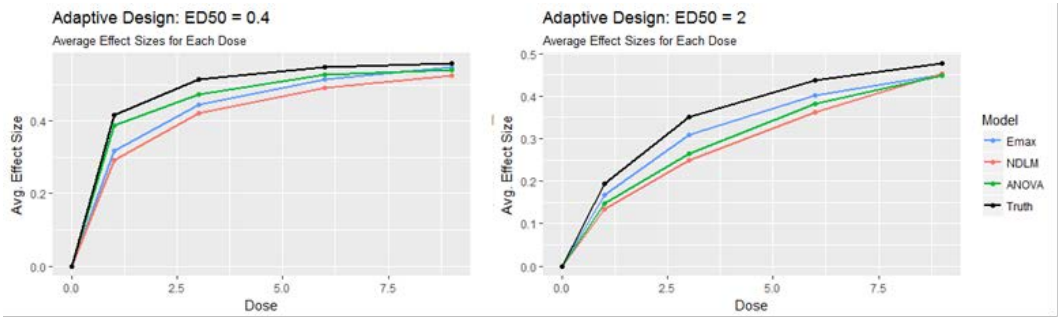

5. Conclusions and Discussion

Our simulation study demonstrated that by using adaptive randomization in clinical trials, we can increase the number of patients allocated to more efficacious doses, as well as decrease the number of patients exposed to doses with potential safety concerns. In the presented scenario and other scenarios we studied, Bayesian analysis by the Emax analysis model resulted in the most drastic effect in adaptive dose allocation ratios, yet Emax and NDLM analysis models had similar dose allocation. ANOVA more closely resembled a traditional, equal randomization design. The Emax analysis model was more sensitive to changes in prior distributions and adjustments in design model specifications. Non-informative prior distributions for the Emax analysis model in many cases may lead to non-convergence of the MCMC simulations. For this reason, the Emax analysis model should be more carefully examined when it is used for adaptive randomization. NDLM and ANOVA were relatively robust to changes in prior distributions in terms of MCMC convergence, and ANOVA was less sensitive to changes in the slope of the design model.

Adaptive methods have the potential to improve the efficiency and flexibility of clinical trials and increase the success rate of new treatments.[5] As FDA Commissioner Scott Gottlieb, MD. stated, "by enriching the enrollment in the trial for patients with characteristics that are likely to predict clinical success, it has the potential to make the development process more efficient." It should be noted that these benefits of adaptation are only realized in certain circumstances. Adaptive randomization is not ideal in situations where there are extended latent periods between when a patient is given a drug and the response in terms of efficacy and safety. Adaptive designs often require a considerable amount of preparation to account for logistical concerns and the avoidance of operational biases. For reference, "Practical Considerations and Strategies for Executing Adaptive Clinical Trials" provides an in-depth account of some of these logistical concerns.[6] Overall, adaptive randomization can provide a very useful framework for increasing clinical trial efficiency and improving multi-arm comparisons of treatment effects, but implementation should be carefully considered in terms of design and analysis before use.

Acknowledgements

The authors are appreciative of Meilin Huang of Genentech and Bret Musser of Regeneron Pharmaceuticals for their insightful input to this research.

References

1. Pallmann, P., Bedding, A. W., Choodari-Oskooei, B., Dimairo, M., Flight, L., Hampson, L.V., … Jaki, T. (2018). Adaptive designs in clinical trials: why use them, and how to run and report them. *BMC medicine*, *16*(1), 29. doi:10.1186/s12916-018-1017-7.
2. Macdougall, J. (2006). Analysis of dose–response studies—E_{max} model. In: Ting N. (eds) Dose Finding in Drug Development. Statistics for Biology and Health. Springer, New York, NY.
3. Holm Hansen, C., Warner, P., Parker, R.A., Walker, B.R., Critchley, H O., & Weir, C. J. (2015). Development of a Bayesian response-adaptive trial design for the Dexamethasone for Excessive Menstruation study. Statistical methods in medical research, 26(6), 2681–2699. doi:10.1177/0962280215606155
4. Smith, M.K., Jones, I.G., Morris, M.F., Grieve, A.P., & Tan, K.K. (2006). Implementation of a Bayesian adaptive design in a proof of concept study. *Pharmaceutical statistics, 5*, 39–50.
5. Lin J., Lin L.A., Sankoh, S. (2016) A general overview of adaptive randomization design for clinical trials. J Biometric and Biostatistics 7:294. doi:10.4172/2155-6180.1000294
6. He, W., Kuznetsova, O.M., Harmer, M., Leahy, C., Anderson, K., Dossin, N., … Schindler, J. (2012). Practical considerations and strategies for executing adaptive clinical trials. Drug Information Journal, 46(2), 160–174. https://doi.org/10.1177/0092861512436580