When is an adaptive design useful in clinical dose-finding trials?

Frank Miller[1]

Abstract

During the development process for new drugs, dose-finding trials have to be conducted and the choice of their design is an important issue. Traditionally, the standard design is a balanced design where equally large groups of patients are treated with different doses of the new drug or with a control. However, it has been identified that other innovative designs might be more efficient: Optimal designs which use non-balanced allocation to dose, and adaptive designs where the allocation to the doses can be changed during the study based on results collected earlier in the study. In a simulation study we will compare efficiencies of balanced non-adaptive, optimal non-adaptive, adaptive two-stage and fully sequential adaptive designs. In all situations considered one can gain from applying optimal design theory. However, when moving from the optimal non-adaptive design to an adaptive design, there are situations where the design is improved and other situations where there is only a minor or no gain. Based on our considered situations, we generalize our observations to answer when an adaptive design is useful.

Key words: Adaptive design; Clinical trial; Dose-finding; Efficiency; Fully sequential design; Interim analysis; Optimal design; Two-stage design.

1 Introduction

The development process of a new drug is divided into several phases: In Phase I, the tolerability of the drug is investigated in clinical trials and the aim is to identify a maximal tolerated dose. As the Phase I investigations often are conducted with healthy volunteers, no or limited information on the

[1]Department of Statistics, Stockholm University, Sweden
e-mail: frank.miller@stat.su.se, Phone: +46 8 162 976

effect of the new drug is collected. In Phase II, patients are investigated. The objective in Phase II is to show that the drug has effect, and to obtain information about the dose-response profile. Often, a Phase II trial for dose-finding includes some hundred patients. In Phase III, trials are conducted comparing one chosen dose with placebo or with a control treatment (in some cases also two or rarely three doses of the new drug are investigated in Phase III). In Phase III trials, patients are usually treated during a longer time (if a chronic disease is to be treated) and patients come from a broader patient population than the Phase II trials. The dose-response profile estimated in Phase II is important such that the dose(s) for the expensive Phase III trials are chosen in the most adequate way.

It is therefore especially desirable to choose an efficient design for the Phase II dose-finding trial. With a more efficient design, there may be the opportunity to learn better about the new drug without increasing the number of patients needed (sample size). Traditionally, a common design for a dose-finding study allocates patients in a balanced way to placebo (or other control) and to 3-5 doses of the new drug. The same number of patients are therefore treated with each treatment out of 4-6 possible options. In recent years, the discussion about adaptive designs became popular. The trial may be started with a balanced design or with another design. After obtaining a certain number of results, the data is analysed (interim analysis) and the design for the remaining patients is chosen based on the results from the first patients. For example, it might turn out that the doses investigated in the first part (Stage 1) gave all similar and good effect such that the interest would be to investigate smaller doses after the interim analysis (in Stage 2).

In the first years when adaptive design ideas were discussed in this context, the usefulness of adaptivity had been overestimated. One reason was that no comprehensive simulation-based or analytical comparisons were available. Some small simulation studies which compared a non-adaptive design with an optimized adaptive design in a specific situation often showed great benefit of the adaptive design. However, a major shortcoming was that the adaptive design was optimized - but the non-adaptive design was suboptimal. This puts the non-adaptive design at a disadvantage but the interpretation was anyway that the large gain was due to the feature of adaptivity alone. A further shortcoming was the investigation of specific situations (sometimes assuming variances for observations much lower than usual in clinical trials) which turned out to be beneficial for the adaptation.

The organisation PhRMA (Pharmaceutical Research and Manufacturers of America) founded a working group on adaptive dose-ranging studies which

more comprehensively investigated these designs. They published their investigations in form of two white papers, see Bornkamp et al. (2007); Dragalin et al. (2010). They investigate several non-adaptive and adaptive designs for a variety of dose-response scenarios and compare different ways of analysing the data. Their work shows the potential of innovative methods applied to dose-finding. In some publications, situations were identified where the gain of adaptations was limited, see Miller et al. (2007); Dragalin et al. (2010); Jones et al. (2011); Dette et al. (2013); McCallum and Bornkamp (2015).

We will here in this simulation-based investigation quantify the gain of several innovative design features, namely we first improve the traditional balanced design as much as possible without using adaptations. When we then add on adaptivity, we want to figure out what the adaptivity itself contributes to the good properties of an adaptive design. Moreover, we consider two different adaptive designs: a two-stage design with a single interim analysis (as mentioned before) and a fully sequential design where each new patient is assigned to a treatment determined based on all available data from the ongoing study.

Consequently, we consider in this article the following four different types of designs:

1. Balanced (non-adaptive) design with k treatment arms (placebo and $k-1$ doses). The number of patients in each treatment arm is equal.

2. Optimal non-adaptive design. Optimal dose allocation is done according to knowledge prior to the study. Optimality is measured by a certain criterion, which will be specified in Section 2.

3. Adaptive two-stage design (with optimal allocation). Allocation ratios are updated once during the study. Two different optimal designs are used before and after the interim analysis.

4. Fully sequential adaptive design (with optimal dynamic allocation). Optimal allocation for each patient based on all information which is collected until the inclusion of this patient.

These four designs will be described with more details in Section 3.

In this article, we use a similar setting as considered by Miller et al. (2007); Fackle-Fornius et al. (2015) but investigate other, varying scenarios of prior knowledge. By this, we get a feeling in which situations an adaptive design is useful.

2 Model assumptions and objective of the trial

In our considered dose-finding trial we assume the possibility to use k treatment arms: $x_0 = 0$ (placebo) and $k-1$ doses of the new drug $0 < x_1 < \cdots < x_{k-1} \leq x_{\max}$. We use in this article five doses $x_1 = 20, x_2 = 40, x_3 = 60, x_4 = 80, x_5 = 100$ mg. A main aim of phase II in drug development is to obtain knowledge about the dose-efficacy and dose-safety profile of the drug. In this article, we focus on the dose-efficacy-profile, only. We assume that the primary outcome measuring the drug effect of patient i in dose group d is $Y_{di}, d = 0, \ldots, k-1$ following the $E_{\max}$-sigmoid model,

$$Y_{di} \sim N(f(x_d, \vartheta), \sigma^2)$$

with

$$\vartheta = (E_0, E_{\max}, ED_{50}, \alpha)^\top, \quad f(x, \vartheta) = E_0 + \frac{E_{\max} x^\alpha}{ED_{50}^\alpha + x^\alpha}.$$

For modelling of both efficacy and safety, we refer to Magnusdottir and Nyquist (2015) who consider a bivariate $E_{\max}$ model for simultaneous inference.

Here, we want to estimate the dose-efficacy in relation to placebo,

$$f(x, \vartheta) - f(0, \vartheta).$$

The placebo effect $f(0, \vartheta)$ is treated here as a nuisance parameter. However, not all parts of the dose-response curve are of equal importance. Especially, we do not need precision in estimates for the part of the curve with low effects below some threshold of clinical importance, δ. Therefore, Miller et al. (2007) used the following objective: if there exist doses within the dose range up to $x_{\max}$ with an effect of at least δ compared to placebo, we want to estimate

$$f(x, \vartheta) - f(0, \vartheta), \quad x \in [x_\delta, x_{\max}],$$

where x_δ is the dose with effect $= \delta$, i.e. $f(x_\delta, \vartheta) - f(0, \vartheta) = \delta$ or $x_\delta = (\delta/(E_{\max} - \delta))^{(1/\alpha)} ED_{50}$. If we have a drug without clinical relevant effect up to dose $x_{\max}$, we want to estimate the effect at the highest dose

$$f(x_{\max}, \vartheta) - f(0, \vartheta).$$

This objective is illustrated in Figure 1. We call this objective "estimation of the interesting part of the dose-response curve".

Given the described objective, we want to search for a good, or "optimal", design. To do this, we need to further formalize the objective and can then

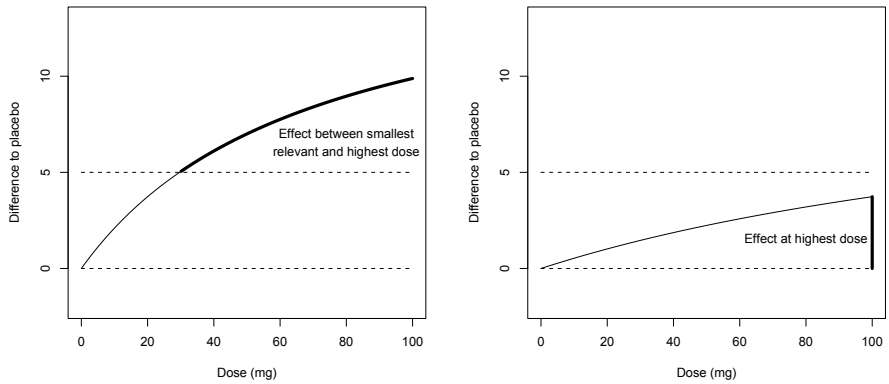

Figure 1: Objective of the trial: estimation of the placebo-adjusted effect between smallest relevant and highest dose (left picture) or – if no such relevant dose exists in the considered dose range – estimation of the effect at highest dose (right picture). The smallest relevant dose is the dose with effect $\delta = 5$.

apply optimal design theory. For a general background on optimal design of experiments, we refer to Silvey (1980); Atkinson et al. (2007).

Let us consider a non-adaptive design which is characterised by allocation ratios $w_d \geq 0, \sum_{j=0}^{k-1} w_j = 1$, for the dose $x_d, d = 0, \ldots, k-1$. We estimate the dose-response curve with the least square estimation in our assumed $E_{\max}$-sigmoid model. The variance of the estimated difference in the effect between a dose x and placebo (dose 0) is approximately proportional to

$$d(x, \xi, \vartheta) = (g(x, \vartheta) - g(0, \vartheta))^\top M^{-1}(\xi, \vartheta) (g(x, \vartheta) - g(0, \vartheta))$$

where

$$M(\xi, \vartheta) = \sum_{j=1}^{k} w_j g(x_j, \vartheta) g^\top(x_j, \vartheta)$$

is the information matrix, and

$$\begin{aligned} g(x, \vartheta) &= \left(\frac{\partial f(x, \vartheta)}{\partial E_0}, \frac{\partial f(x, \vartheta)}{\partial E_{\max}}, \frac{\partial f(x, \vartheta)}{\partial ED_{50}}, \frac{\partial f(x, \vartheta)}{\partial \alpha} \right)^\top \\ &= \left(1, \frac{x^\alpha}{ED_{50}^\alpha + x^\alpha}, \frac{-E_{\max} \alpha ED_{50}^{\alpha-1} x^\alpha}{(ED_{50}^\alpha + x^\alpha)^2}, \frac{E_{\max} ED_{50}^\alpha x^\alpha (\log x - \log ED_{50})}{(ED_{50}^\alpha + x^\alpha)^2} \right)^\top \end{aligned}$$

is the gradient of the dose-efficacy-function with respect to ϑ, see also Miller et al. (2007); Fackle-Fornius et al. (2015).

If there exists no dose within the dose range up to $x_{\max}$ with an effect of at least δ compared to placebo, we just have to maximise $1/d(x_{\max}, \xi, \vartheta)$. If there exist doses within this dose range with the required effect, we want to minimise the average variance of the estimates for $f(x, \vartheta) - f(0, \vartheta), x \in [x_\delta, x_{\max}]$, and use the I_L-criterion with $L = 1$, see Fedorov (1972); Dette and O'Brien (1999). Therefore we search a design ξ with large value of the following criterion function:

$$\Phi(\xi, \vartheta) = \begin{cases} \left\{ \frac{1}{x_{\max} - x_\delta} \int_{x_\delta}^{x_{\max}} d(x, \xi, \vartheta)\ dx \right\}^{-1}, & \text{if } f(x_{\max}, \vartheta) - f(0, \vartheta) > \delta, \\ 1/d(x_{\max}, \xi, \vartheta), & \text{if } f(x_{\max}, \vartheta) - f(0, \vartheta) \leq \delta. \end{cases} \tag{1}$$

We define the relative efficiency of an arbitrary design ξ with respect to the balanced design ξ_0 (used as reference design) by

$$\text{Eff}(\xi, \vartheta) = \Phi(\xi, \vartheta)/\Phi(\xi_0, \vartheta).$$

For example, an efficiency of 1.25 means that the balanced design would need 25% more patients than the design under consideration to obtain estimates with approximately the same precision.

Optimal designs for $E_{\max}$ dose-finding models based on other optimality criteria have been derived in the literature. Burman (2015) presents D-optimal designs for the $E_{\max}$-sigmoid model. Magnusdottir (2013) derives c-optimal designs for the bivariate $E_{\max}$ model. Dette et al. (2008) consider MED-optimal designs for estimation of the minimum effective dose.

Usually, there exists prior knowledge about possible dose-response scenarios before the study starts. Miller et al. (2007) have investigated an example based on an AstraZeneca study with seven possible dose-response scenarios. In this paper, we use as example three possible scenarios: an optimistic scenario, a pessimistic scenario, and a scenario with good effects only at high doses, see Table 1 and Figure 2. In Table 1, the parameter E_0 is not included, since we treat the placebo effect as nuisance parameter and our investigation does not depend on the value of E_0. Further, a standard deviation of $\sigma = 10$ is assumed leading to reasonable signal-to-noise ratios in clinical studies. With the knowledge before the study, we assume that the prior probabilities for the three scenarios are 0.35, 0.35, and 0.30, respectively.

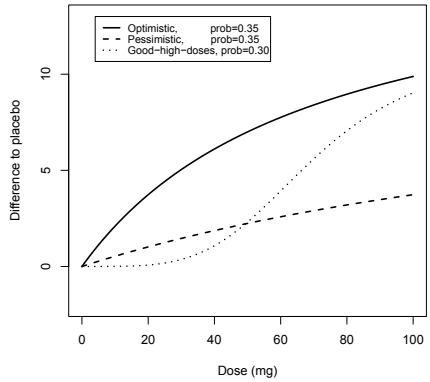

Figure 2: Prior knowledge about dose-efficacy scenarios (Example 1)

Table 1: Parameters for the dose-efficacy scenarios (Example 1)

Scenario	$E_{\max}$	ED_{50}	α	prior prob. π_j
Optimistic	16.8	70	1	0.35
Pessimistic	11.2	200	1	0.35
Good-high-doses	11.2	70	4	0.30

3 Description of the considered designs

Once we have described the assumed model and the objective of the trial, we can describe the four considered designs in more detail. In the simulations in the next sections, we consider a total sample size of $n = 300$ patients.

3.1 Balanced design

Since we have $n = 300$ patients and $k = 6$ treatment arms, we have $n/k = 50$ patients per treatment arm.

Table 2: Optimal non-adaptive design

Dose x_d	0	20	40	60	80	100
Weight w_d	0.385	0.038	0.062	0.096	0.119	0.300
n_d	115	11	19	29	36	90

3.2 Optimal non-adaptive design

As mentioned before, a non-adaptive design is characterised by the allocation ratios $w_0, w_1, \ldots, w_{k-1}$ for the treatment arms 0 (placebo) and $1, 2, \ldots, k-1$ (active doses). In common notation for experimental design, a design characterised by the observational points (here doses) and their allocation ratios (weights) is often called ξ and can be written as

$$\xi = \begin{pmatrix} x_0 & x_1 & \ldots & x_{k-1} \\ w_0 & w_1 & \ldots & w_{k-1} \end{pmatrix}.$$

Our prior knowledge consists of three different scenarios which we can call ϑ_j with $j = 1$ for the optimistic, $j = 2$ for the pessimistic and $j = 3$ for the good-high-doses scenario. We are interested in the design ξ (i.e. the weights w_d) which maximises the average efficiencies for the scenarios,

$$\sum_{j=1}^{3} \pi_j \mathrm{Eff}(\xi, \vartheta_j), \tag{2}$$

using the prior probabilities $\pi_1 = \pi_2 = 0.35, \pi_3 = 0.30$ (see Table 1). We call this design optimal non-adaptive design. We have calculated this optimal design by a numerical method using a first order exchange algorithm, see e.g. Atkinson et al. (2007), see Table 2. The patient numbers n_d are obtained by rounding $300w_d$ (while ensuring a total sum of 300).

The approach to maximize average efficiencies (2) is called "optimal-on-average approach" in contrast to a "maximin approach" where the minimal efficiency over the scenarios is maximised, see Fackle-Fornius et al. (2015).

3.3 Adaptive two-stage design

For this design, we start as above with the optimal non-adaptive design but only for the first part of the study with 100 patients. Given the weights in Table 2, the patient numbers are 38, 4, 6, 10, 12, 30 for doses $0, 20, \ldots, 100$,

respectively (rounding $100w_d$ and ensuring a total sum of 100). Based on the results Y_{di} of 100 patients, we calculate posterior probabilities for the three scenarios according to the Bayes formula. With these posterior probabilities, we calculate a new optimal non-adaptive design for the whole 300 patients. We do this in the same way as we did before with the only exception that we optimise restricted to the fact that we have already a certain number of patients treated with each dose. In practice, it takes some time to collect and analyse the data from the first part of the study while recruitment of new patients is ongoing. Therefore we assume that even patient number 101 to 140 are included according to the starting design. Then, for patient number 141 to 300, the new optimal non-adaptive design is applied.

3.4 Fully sequential adaptive design

As before we assume that we include 140 patients according to the optimal non-adaptive design. When patient 141 enters the study, the results of the first 100 patients are analysed, posterior probabilities calculated and the treatment of this patient is chosen in order to maximise our optimality criterion. We continue in this way, updating the posterior probabilities on an ongoing basis for determining the treatment of the next patient. We assume throughout the study that we have a lag of 40 patients, i.e. when including patient number $i+1$, the results of $i-40$ patients are available.

4 Efficiency for designs

For each scenario and each design, we performed 5000 simulations. Based on these simulations, we calculated relative efficiencies between the designs.

We need simulations as available asymptotic formulae for the efficiency is to crude for finite sample sizes when adaptive designs are considered. Therefore, we compute mean squared errors (MSE) of the estimates in the simulations. More precisely, for a certain design and a certain simulation scenario, we obtain the MSE at dose x for estimation of $f(x) - f(0)$ by

$$\mathrm{MSE}(x, \xi, \vartheta) = \frac{1}{s} \sum_{l=1}^{s} \left[\{\hat{f}_l(x) - \hat{f}_l(0)\} - \{f(x) - f(0)\} \right]^2,$$

where s is the total number of simulations and $\hat{f}_l$ denotes the estimated function in the lth run of the simulation. We replace then in equation (1) the function $d(x, \xi, \vartheta)$ by $\mathrm{MSE}(x, \xi, \vartheta)$ and calculate Φ; if the scenario is such

Table 3: Efficiency gain (Example 1)

Scenario	Efficiency gain		
	from balanced to optimal non-adaptive	from optimal non-adaptive to adaptive two-stage	from adaptive two-stage to fully sequential
Optimistic	+12%	+10%	+3%
Pessimistic	+62%	+1%	±0%
Good-high-doses	-4%	+7%	+8%
Overall	+24%	+6%	+3%

that the first part of the Φ-formula applies we use numerical integration over $[x_\delta, x_{\max}]$ where $\delta = 5$ and x_δ depends on the scenario. Then $\Phi(\xi_b, \vartheta)/\Phi(\xi_a, \vartheta)$ gives us the relative efficiency of Design ξ_b relative to Design ξ_a.

We went step by step and calculated the efficiency gain from the balanced design to the optimal non-adaptive design, from the optimal non-adaptive design to the adaptive two-stage design and finally from the adaptive two-stage to the fully sequential design. Results are summarized in Table 3.

We see a quite large efficiency gain of +24% from the balanced design to the optimal design, which is mainly due to an efficiency gain if the underlying scenario is the pessimistic one (+62% efficiency gain; note that the increased allocation to the highest dose is especially important in this scenario). Surprisingly, we cannot improve it much further with the considered adaptive designs: if we use the adaptive two-stage design, we gain additionally +6%, and if we go even further and adapt after each patient, we can gain +3% more efficiency. With these moderate efficiency gains for the adaptive design it is in most situations hard to justify the additional complexity of an adaptive trial.

Why is there not more gain from adaptive dosing? Let us consider for example the optimistic scenario and the comparison between the optimal non-adaptive design and the adaptive two-stage design with average efficiency gain of 10% (see Table 3). The cumulative distribution function for the MSE from the simulations of the optimal non-adaptive design and the adaptive two-stage design are shown in Figure 3 (a). We can see that the cumulative distribution function for the adaptive design is mostly above the function for the non-adaptive optimal design which reflects smaller MSE. However, if we look closer into the tail region of the distribution (Figure 3 (b)), we see that this is reversed in the part above 0.97 (i.e. for the largest 3% of the MSE). In the part with large MSE, the adaptive design is even worse than the

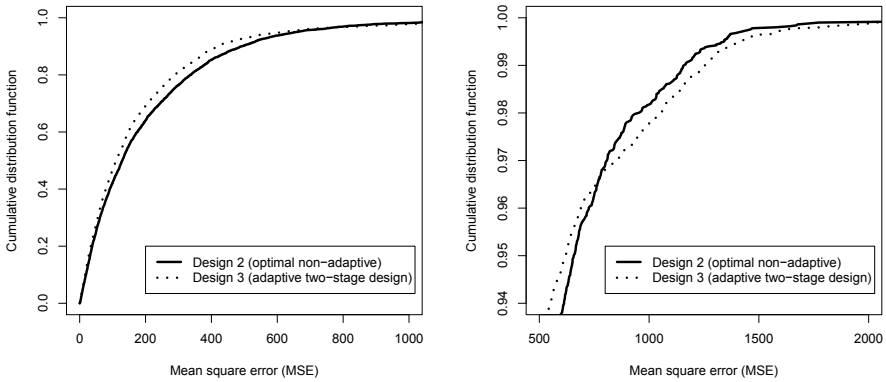

Figure 3: Cumulative distribution function of the MSE for the optimal non-adaptive design and the adaptive two-stage design (optimistic scenario). (a) Left figure: whole cumulative distribution, (b) Right figure: tail of the cumulative distribution

non-adaptive. Since the large MSE have also a high impact on the average efficiency of a design, these approximately 3% of the simulations contribute that the efficiency of the adaptive two-stage design becomes not too good compared to the optimal non-adaptive design. How can the high MSEs be interpreted? Due to the assumed variability, some of the simulations have interim data suggesting a totally different dose-response-shape compared to the true scenario. In these cases, the problem for the adaptive design is that an inferior design is chosen for the part after the interim analysis based on the wrong interim estimate. This has then an additional negative impact on the precision of the final estimate.

5 When is an adaptive design useful?

We modify now the prior assumptions in order to investigate in which cases adaptive designs are useful compared to the non-adaptive optimal design. In our first modification (Example 2), we change the optimistic scenario to a "realistic" ($E_{\max} = 11.2, ED_{50} = 70, \alpha = 1$) which is closer to the other two scenarios, see Figure 4. Otherwise, we change nothing: we keep the other two

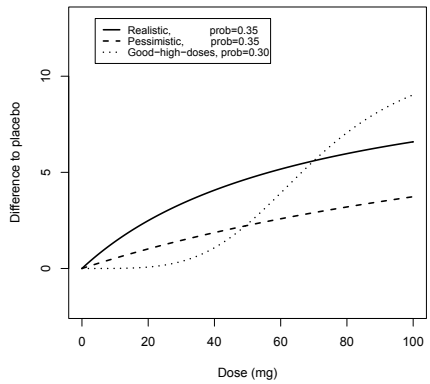

Figure 4: Prior knowledge about dose-efficacy scenarios (Example 2)

scenarios and prior probabilities.

The efficiency gains based on 5000 simulations per scenario and design are summarised in Table 4. We have an even larger gain from the balanced to the optimal non-adaptive design. We explain this as follows: The scenarios are more similar compared to the example before. Therefore, a non-adaptive design can be better optimised for all scenarios simultaneously. But when we go further to adaptive designs we have even an efficiency loss. When there is a good understanding prior to the trial with a few possible scenarios which are not too different, then there is no need to introduce interim analyses for design modification. In these cases, interim analyses could rather lead to an inferior design due to the variability in the data used for interim decisions.

In our second modification (Example 3), we use exactly the same scenarios and prior probabilities as in the main example (Example 1, see Figure 2). We change only the underlying variability. Instead of the assumption $\sigma = 10$ in Example 1 and 2, we use $\sigma = 6$. Again, we performed 5000 simulations per scenario and design with results shown in Table 5. As before, we have a good gain with the optimal non-adaptive design. In contrast to the examples before, we see also a good gain with the adaptive two-stage design (+16%) which is similar for all three scenarios (+18, +13, +16%, respectively). An additional gain of 5% is made when using the fully sequential design. In this example, we have seen that the adaptive dosing has value which could justify the more

complicated logistics of an adaptive trial. Here, the possible scenarios are sufficiently different in relation to the underlying variability. This makes it possible to obtain good information from the interim analysis to improve the prior probabilities to reliable posterior probabilities for the scenarios.

6 Discussion

We have seen that there is a large gain in efficiency when optimal design theory is applied and an optimal non-adaptive design is chosen instead of a traditional balanced design. Using interim data to change allocation ratios (adaptive dosing) is an attractive concept but it depends on the situation whether it can lead to a further gain in efficiency or not. This is in line with the observation from a recent simulation study, McCallum and Bornkamp (2015), concluding that "the benefit of including an interim analysis has not been shown to universally to improve the performance of a dose-finding study".

If the possible scenarios are similar or the variance is large, decisions based on interim data could lead into the wrong direction. In these cases, an optimal non-adaptive design might be the better choice. If differences between the possible scenarios are large (in relation to the variability of data in interim analysis), there is a clear gain from adaptive dosing.

Most investigations in this context are forced to build on simulation studies as available asymptotic formulae for the efficiency of adaptive designs are to crude for finite sample sizes. However, in the context of a simplified dose-response model (one-parameter model), Dette et al. (2013) successfully derived explicit expressions for the asymptotic efficiency of the adaptive design which are precise and suited for comparison. Based on these expressions, they

Table 4: Efficiency gain (Example 2)

Scenario	Efficiency gain		
	from balanced to optimal non-adaptive	from optimal non-adaptive to adaptive two-stage	from adaptive two-stage to fully sequential
Realistic	+48%	-4%	-1%
Pessimistic	+64%	+2%	-6%
Good-high-doses	+8%	-1%	+2%
Overall	+42%	-1%	-2%

Table 5: Efficiency gain in case of small variance $\sigma^2 = 6^2$ (Example 3)

Scenario	Efficiency gain		
	from balanced to optimal non-adaptive	from optimal non-adaptive to adaptive two-stage	from adaptive two-stage to fully sequential
Optimistic	+13%	+18%	+5%
Pessimistic	+79%	+13%	+4%
Good-high-doses	-3%	+16%	+6%
Overall	+31%	+16%	+5%

were able to compare the efficiency of non-adaptive and adaptive design in this setting. The combined evidence from simulation studies and algebraic investigations for simplified models gives a good picture about when adaptive designs are useful.

Once it is concluded that an adaptive design is useful, it remains the question whether it is feasible for the study to be planned. Logistical issues need to be resolved before such a design can be applied. Miller et al. (2014) discuss an example where a simplified adaptive dose-finding design was conducted. The use of an interim analysis offers the important benefit that the study can be stopped early (so called futility stopping) when results in the interim suggest that no dose of the drug will be useful.

A further important point for adaptive designs is that the statistical inference at the end of the trial need to take the adaptivity into account. Note that in the context of clinical studies significance tests and their frequentist properties (type I error and power) are of importance. For example Miller (2010) has derived a trend-test which controls the type I error for an adaptive two stage dose-finding trial.

References

Atkinson, A. C., Donev, A. N., and Tobias, R. D. *Optimum experimental design, with SAS*. Oxford University Press, Oxford, 2007.

Bornkamp, B., Bretz, F., Dmitrienko, A., Enas, G., Gaydos, B., Hsu, C. H., König, F., Krams, M., Liu, Q., Neuenschwander, B., Parke, T., Pinheiro, J., Roy, A., Sax, R., and Shen, F. Innovative approaches for designing

and analyzing adaptive dose-ranging trials. *Journal of Biopharmaceutical Statistics*, 17:965–995, 2007.

Burman, C. F. From optimal design theory to optimizing designs of clinical trials. In Fackle-Fornius, E., editor, *Festschrift in Honor of Hans Nyquist on the Occasion of his 65th Birthday*. Department of Statistics, Stockholm University, 2015.

Dette, H. and O'Brien, T. E. Optimality criteria for regression models based on predicted variance. *Biometrika*, 86:93–106, 1999.

Dette, H., Bretz, F., Pepelyshev, A., and Pinheiro, J. Optimal designs for dose finding studies. *Journal of the American Statistical Association*, 103: 1225–1237, 2008.

Dette, H., Bornkamp, B., and Bretz, F. On the efficiency of two–stage response–adaptive designs. *Statistics in Medicine*, 32:1646–1660, 2013.

Dragalin, V., Bornkamp, B., Bretz, F., Miller, F., Padmanabhan, S. K., Patel, N., Perevozskaya, I., Pinheiro, J., and Smith, J. R. A simulation study to compare new adaptive dose-ranging designs. *Statistics in Biopharmaceutical Research*, 2:487–512, 2010.

Fackle-Fornius, E., Miller, F., and Nyquist, H. Implementation of maximin efficient designs in dose-finding studies. *Pharmaceutical Statistics*, 14:63–73, 2015.

Fedorov, V. V. *Theory of Optimal Experiments*. Academic Press, New York, 1972.

Jones, B., Layton, G., Richardson, H., and Thomas, N. Model-based Bayesian adaptive dose-finding designs for a Phase II trial. *Statistics in Biopharmaceutical Research*, 3:276–287, 2011.

Magnusdottir, B. T. c-optimal designs for the bivariate Emax model. In Ucinski, D., Atkinson, A., and Patan, M., editors, *mODa 10 - Advances in Model-Oriented Design and Analysis: Proceedings of the 10th International Workshop in Model-Oriented Design and Analysis*, pages 153–162. Springer International Publishing: Switzerland, 2013.

Magnusdottir, B. T. and Nyquist, H. Simultaneous estimation of parameters in the bivariate Emax model. *Statistics in Medicine*, 2015. doi: 10.1002/sim.6585.

McCallum, E. and Bornkamp, B. Accounting for parameter uncertainty in two-stage designs for phase II dose-response studies. In Sverdlov, O., editor, *Modern Adaptive Randomized Clinical Trials. Statistical and Practical Aspects.*, Boca Raton, 2015. Chapman & Hall/CRC.

Miller, F. Adaptive dose-finding: Proof of concept with type I error control. *Biometrical Journal*, 52:577–589, 2010.

Miller, F., Guilbaud, O., and Dette, H. Optimal designs for estimating the interesting part of a dose-effect curve. *Journal of Biopharmaceutical Statistics*, 17:1097–1115, 2007.

Miller, F., Björnsson, M., Svensson, O., and Karlsten, R. Experiences with an adaptive design for a dose-finding study in patients with osteoarthritis. *Contemporary Clinical Trials*, 37:189–199, 2014.

Silvey, S. D. *Optimal design.* Chapman and Hall, London, 1980.