

Assessing the similarity of dose response and target doses in two non-overlapping subgroups

Frank Bretz¹, Kathrin Möllenhoff², Holger Dette²,
Wei Liu³, Matthias Trampisch⁴

July 12, 2016

1 Novartis Pharma AG, CH-4002 Basel, Switzerland

2 Department of Mathematics, Ruhr-Universität Bochum, Germany

3 S3RI and School of Mathematics, University of Southampton, SO17 1TB, UK

4 Boehringer Ingelheim Pharma GmbH & Co. KG, Biostatistics + Data Sciences / BDS, Germany

Abstract

We consider two problems that are attracting increasing attention in clinical dose finding studies. First, we assess the similarity of two non-linear regression models for two non-overlapping subgroups of patients over a restricted covariate space. To this end, we derive a confidence interval for the maximum difference between the two given models. If this confidence interval excludes the equivalence margins, similarity of dose response can be claimed. Second, we address the problem of demonstrating the similarity of two target doses for two non-overlapping subgroups, using again a confidence interval based approach. We illustrate the proposed methods with a real case study and investigate their operating characteristics (coverage probabilities, Type I error rates, power) via simulation.

Keywords and Phrases: equivalence testing, multiregional trial, target dose estimation, subgroup analyses

1 Introduction


Establishing dose response and selecting optimal dosing regimens is a fundamental step in the investigation of any new compound, be it a medicinal drug, an herbicide or fertilizer, a molecular entity, an environmental toxin, or an industrial chemical (1). This has been recognized for many years, especially in the drug development area, where patients are exposed to a medicinal drug once it has been released on the market (2). An indication of the importance of properly conducted dose response studies is the early publication of the tripartite ICH E4 guideline, which gives recommendations on the design and conduct of studies to assess the relationship between doses, blood levels and clinical response throughout the clinical development of a new drug (3).

Very often clinical trials are analyzed beyond the primary study objectives by assessing efficacy and safety profiles in clinically relevant subgroups, such as different gender, age classes, grades of disease severity, etc. A rising area of particular importance are global clinical trials, which are run in different countries and potentially serve different submissions. For example, many pharmaceutical companies focus on running global clinical trials that include a major Japanese subpopulation for later regulatory submission in Japan. A natural question is then whether the dose response results for the Japanese and the non-Japanese populations are consistent (4; 5).

To illustrate the general problem, assume that we are interested in assessing similarity for (a) two dose response curves or for (b) two same target doses, say for male/female or Japanese/non-Japanese patients. For question (a) we thus want to show that the maximum difference in response between two (potentially different) non-linear parametric regression models is smaller than a pre-specified margin. Figure 1a displays an example, where the two dose response curves follow an Emax and a logistic model. The maximum response difference over the dose range is indicated by the arrow. For question (b) we want to show that two same target doses do not differ relevantly. Figure 1b displays the minimum effective dose (*MED*) derived from the two previous dose response models. Here, the *MED* is defined as the smallest dose which demonstrates a clinically relevant benefit over placebo, as indicated by the horizontal line in Figure 1b. If we succeed in demonstrating either (a) or (b), evidence is provided that the difference in response over the entire dose range under investigation or the two target doses differ at most marginally. In practice, such a result may provide sufficient evidence that the same dose can be administered in both subgroups (i.e., the doses for males/females or Japanese/non-Japanese patients are the same).

Demonstrating similarity of target doses or dose response curves in each of several subgroups has not been addressed in much detail so far in the literature. One exception is (6), who proposed a non-standard bootstrap approach for question (a) which addresses the specific form of the interval hypotheses. In particular, data has to be generated under the null

(a)



(b)




Figure 1: Assessing similarity for (a) two dose response curves and (b) two same target doses.

hypothesis using constrained least squares estimates. In this paper we consider different methods to address both questions (a) and (b). In Section 2 we address problem (a) using the results from (7) and derive a confidence interval for the maximum difference between the two given non-linear regression models over the entire covariate space of interest. If this confidence interval excludes the equivalence margins, similarity of dose response can be claimed. In Section 3, we consider asymptotic methods to derive confidence intervals for the difference between two same target doses to address problem (b). Again, if such a confidence interval excludes a pre-specified relevance margin, similarity in dose can be claimed. In Section 4 we provide some concluding remarks. Technical details are left for the Appendix.

2 Assessing similarity of two dose response curves

We consider the non-linear regression models

$$Y_{\ell,i,j} = m_{\ell}(\vartheta_{\ell}, d_{\ell,i}) + \varepsilon_{\ell,i,j}, \quad j = 1, \dots, n_{\ell,i}, \quad i = 1, \dots, k_{\ell}, \quad \ell = 1, 2, \quad d_{\ell,i} \in \mathcal{D}, \quad (1)$$

where $Y_{\ell,i,j}$ denotes the j th observed response at the i th dose level $d_{\ell,i}$ under the ℓ th dose response model m_{ℓ} . The error terms $\varepsilon_{\ell,i,j}$ are assumed to be independent and identically distributed with expectation 0 and variance σ_{ℓ}^2 . Further, $n_{\ell} = \sum_{i=1}^{k_{\ell}} n_{\ell,i}$ denotes the sample size in group ℓ where we assume $n_{\ell,i}$ observations in the i th dose level ($i = 1, \dots, k_{\ell}$, $\ell = 1, 2$). We further assume that for both regression models the different dose levels are attained

on the same (restricted) covariate region \mathcal{D} . For the purpose of this paper, we assume \mathcal{D} to be the dose range under investigation, although the results in this section can be generalized to include other covariates. The functions m_1 and m_2 in (1) denote the (non-linear) regression models with fixed but unknown p_1 - and p_2 -dimensional parameter vectors ϑ_1 and ϑ_2 , respectively. Note that both the regression models m_1 and m_2 and the parameters ϑ_1 and ϑ_2 may be different. In particular, the design matrices for the two regression models may be unequal. This implies that we do not assume the same doses to be investigated for $\ell = 1, 2$ and that the sample sizes n_ℓ can be unequal. We refer to (8) for an overview of several linear and non-linear regression models commonly employed in clinical studies.

2.1 Methodology

Using results from (7), we derive in the following a confidence interval for the maximum absolute difference between the two given non-linear regression models m_1 and m_2 over the entire covariate space \mathcal{D} . We use this confidence interval in order to derive a test demonstrating similarity of the two dose response curves.

Let $U(Y_1, Y_2, d)$ denote a $1 - \alpha$ pointwise upper confidence bound on the difference curve $m_2(\vartheta_2, d) - m_1(\vartheta_1, d)$, i.e. $P\{m_2(\vartheta_2, d) - m_1(\vartheta_1, d) \leq U(Y_1, Y_2, d)\} \geq 1 - \alpha$ for all $d \in \mathcal{D}$, where α denotes the pre-specified significance level and Y_ℓ the vector of observations from group $\ell = 1, 2$. Similarly, let $L(Y_1, Y_2, d)$ denote a $1 - \alpha$ pointwise lower confidence bound on $m_2(\vartheta_2, d) - m_1(\vartheta_1, d)$. Using these pointwise confidence bounds we can deduce a confidence interval for the maximum absolute difference between the two models $\max_{d \in \mathcal{D}} |m_2(\vartheta_2, d) - m_1(\vartheta_1, d)|$ over the region \mathcal{D} , that is

$$P \left\{ \max_{d \in \mathcal{D}} |m_2(\vartheta_2, d) - m_1(\vartheta_1, d)| \leq \max \left\{ \max_{d \in \mathcal{D}} U(Y_1, Y_2, d), -\min_{d \in \mathcal{D}} L(Y_1, Y_2, d) \right\} \right\} \geq 1 - \alpha. \quad (2)$$

The proof is given in Appendix A. For moderate sample sizes the pointwise confidence bounds $U(Y_1, Y_2, d)$ and $L(Y_1, Y_2, d)$ can be derived from the delta method (9). Let $u_{1-\alpha}$ denote the $1 - \alpha$ quantile of the standard normal distribution. Then,

$$U(Y_1, Y_2, d) = m_2(\hat{\vartheta}_2, d) - m_1(\hat{\vartheta}_1, d) + u_{1-\alpha} \hat{\rho}(d)$$

and

$$L(Y_1, Y_2, d) = m_2(\hat{\vartheta}_2, d) - m_1(\hat{\vartheta}_1, d) - u_{1-\alpha} \hat{\rho}(d)$$

are the desired $1 - \alpha$ asymptotic pointwise upper and lower confidence bounds, respectively, for $m_2(\vartheta_2, d) - m_1(\vartheta_1, d)$. Here, $\hat{\vartheta}_\ell$ denotes the least squares estimate of ϑ_ℓ and

$$\hat{\rho}^2(d) = \frac{\hat{\sigma}_1^2}{n_1} \left(\frac{\partial}{\partial \vartheta_1} m_1(\hat{\vartheta}_1, d) \right)^T \hat{\Sigma}_1^{-1} \left(\frac{\partial}{\partial \vartheta_1} m_1(\hat{\vartheta}_1, d) \right) + \frac{\hat{\sigma}_2^2}{n_2} \left(\frac{\partial}{\partial \vartheta_2} m_2(\hat{\vartheta}_2, d) \right)^T \hat{\Sigma}_2^{-1} \left(\frac{\partial}{\partial \vartheta_2} m_2(\hat{\vartheta}_2, d) \right) \quad (3)$$

is an estimate of the variance of $m_2(\hat{\vartheta}_2, d) - m_1(\hat{\vartheta}_1, d)$. In (3) $\hat{\sigma}_\ell^2$, is the common variance estimate in the ℓ th group ($\ell = 1, 2$) and $\hat{\Sigma}_\ell = \sum_{i=1}^{k_\ell} \frac{n_{\ell,i}}{n_\ell} \frac{\partial}{\partial \vartheta_\ell} m_\ell(x_{\ell,i}, \hat{\vartheta}_\ell) \left(\frac{\partial}{\partial \vartheta_\ell} m_\ell(x_{\ell,i}, \hat{\vartheta}_\ell) \right)^T$. Note that the matrix $\frac{\hat{\sigma}_\ell^2}{n_\ell} \hat{\Sigma}_\ell^{-1}$ is a consistent estimator of the covariance matrix of $\hat{\vartheta}_\ell$ ($\ell = 1, 2$). Next we are interested in demonstrating that the maximum absolute difference in response between the two regression models in (1) over the covariate space \mathcal{D} is not larger than a pre-specified margin $\delta > 0$. Formally, we test the null hypothesis

$$H : \max_{d \in \mathcal{D}} |m_2(\vartheta_2, d) - m_1(\vartheta_1, d)| \geq \delta \quad (4)$$

against the alternative hypothesis

$$K : \max_{d \in \mathcal{D}} |m_2(\vartheta_2, d) - m_1(\vartheta_1, d)| < \delta. \quad (5)$$

Consequently, using the confidence interval (2), equivalence is claimed if

$$\max \left\{ \max_{d \in \mathcal{D}} U(Y_1, Y_2, d), -\min_{d \in \mathcal{D}} L(Y_1, Y_2, d) \right\} < \delta.$$

Thus, we reject the null hypothesis H at level α and assume similarity of m_1 and m_2 if

$$-\delta < \min_{d \in \mathcal{D}} L(Y_1, Y_2, d) \quad \text{and} \quad \max_{d \in \mathcal{D}} U(Y_1, Y_2, d) < \delta. \quad (6)$$

2.2 Case study

To illustrate the methodology described in Section 2.1, we consider a dose finding trial for a weight loss drug given to patients suffering from overweight or obesity. This trial aims at comparing the dose response relationship for two regimens, namely a once-daily (o.d.) and a twice-daily (b.i.d.) application of the drug. The primary objective in this trial is not to apply a joint model that includes both regimen, but rather treat both regimen separately and assess the similarity of dose response. Because this study has not been completed yet, we simulate data based on the assumptions made at the trial design stage. For confidentiality reasons, we use blinded dose levels and all chosen dose levels denote the total daily dose. These limitations do not change the utility of the calculations below.


In this trial, the dose levels for the o.d. and b.i.d. regimens are given by 0.033, 0.1, 1 and 0.067, 0.3, 1, respectively. Patients are thus randomized to receive either placebo or one of the six active treatments. In total, we assume that 350 patients are allocated equally across the seven arms, resulting in a sample size of 50 patients per treatment group. The primary endpoint of the study was the percentage of weight loss after a treatment duration of 20 weeks, with smaller values corresponding to a better treatment effect.

We used the `nls` function in R (10) to compute the non-linear least squares estimates $\hat{\vartheta}_\ell$ of ϑ_ℓ and the standard errors necessary for calculating $U(Y_1, Y_2, d)$ and $L(Y_1, Y_2, d)$ from Section

2.1. The R code for this example and all other calculations in this paper is available from the authors upon request.

For this example, we fitted two Emax models: $m_1(\vartheta_1, d) = \vartheta_{1,1} + \vartheta_{1,2} \frac{d}{\vartheta_{1,3} + d}$ for the o.d. regimen and $m_2(\vartheta_2, d) = \vartheta_{2,1} + \vartheta_{2,2} \frac{d}{\vartheta_{2,3} + d}$ for the b.i.d. regimen, where $\vartheta_1 = (\vartheta_{1,1}, \vartheta_{1,2}, \vartheta_{1,3})$ and $\vartheta_2 = (\vartheta_{2,1}, \vartheta_{2,2}, \vartheta_{2,3})$. For the data set at hand, $\hat{\vartheta}_1 = (0.03, -5.17, 7.94)$ and $\hat{\vartheta}_2 = (-0.09, -6.56, 31.24)$. Figure 2a displays the fitted dose response models $m_1(\hat{\vartheta}_1, d)$ and $m_2(\hat{\vartheta}_2, d)$, $d \in [0, 1]$, together with the individual observations, where the y -axis is truncated to $[-7, 1]$ for better readability. Figure 2b displays the difference $m_2(\hat{\vartheta}_2, d) - m_1(\hat{\vartheta}_1, d)$ together with the associated 90% pointwise confidence intervals for each dose $d \in [0, 1]$. The maximum upper confidence bound for $\alpha = 0.1$ is $\max_{d \in \mathcal{D}} U(Y_1, Y_2, d) = 2.137$ at dose $d = 0.08$ and the minimum lower confidence bound is $\min_{d \in \mathcal{D}} L(Y_1, Y_2, d) = -1.848$ at the maximum dose $d = 1$. That is, the maximum difference in response between the two regimens over the dose range $\mathcal{D} = [0, 1]$ lies between -1.848 and 2.137 . Therefore, similarity of the dose response curves can be claimed at level $\alpha = 0.1$ as long as δ is larger than 2.137 , according to (6).

(a)



(b)




Figure 2: Plots for the weight loss case study. (a) The fitted Emax model m_1 (m_2) for the o.d. (b.i.d.) regimen is given by the solid (dashed) line with observations marked by “x” (“o”). (b) Mean difference curve with associated pointwise 90% confidence bounds. Bold dots denote the maximum upper and minimum lower confidence bound over $\mathcal{D} = [0, 1]$.

2.3 Simulations

We conducted a simulation study to investigate the operating characteristics of the method described in Section 2.1. We investigated coverage probabilities of the confidence intervals as well as Type I error rates and power of the test (6) for different scenarios. To simplify the simulations, we assumed balanced designs and that dose is the only covariate. For all simulations below, we generated data as follows:

Step 1: Specify the models m_1, m_2 , their parameters ϑ_1, ϑ_2 , a common variance σ^2 and the actual dose levels $d_{\ell,i}$.

Step 2: Generate $n_{\ell,i}$ values $m_{\ell}(\vartheta_{\ell}, d_{\ell,i})$ at each dose $d_{\ell,i}$.

Step 3: Generate normally distributed residual errors $\varepsilon_{\ell,i,j} \sim N(0, \sigma^2)$ and use the final response data

$$Y_{\ell,i,j} = m_{\ell}(\vartheta_{\ell}, d_{\ell,i}) + \varepsilon_{\ell,i,j}, \quad j = 1, \dots, n_{\ell,i}, \quad i = 1, \dots, k_{\ell}, \quad \ell = 1, 2. \quad (7)$$

This procedure is repeated using 10,000 simulation runs.


2.3.1 Coverage probabilities

In the following we report the coverage probabilities of the confidence intervals for the maximum absolute difference derived in (2) under two different scenarios.

Scenario 1 We start with the comparison of a linear and a quadratic model. More specifically, we chose the linear model $m_1(d) = d$ and the quadratic model $m_2(d) = 3\delta_1 + (1 - 4\delta_1)d + \delta_1 d^2$, $d \in [1, 3]$; see Figure 3a for $\delta_1 = 1$. We assumed identical dose levels $d_{\ell,i} = i$, $i = 1, 2, 3$ for both regression models $\ell = 1, 2$. Consequently, the two curves coincide at the two boundary doses $d = 1, 3$, and the maximum difference δ_1 occurs at dose $d = 2$. For each configuration of $\sigma^2 = 1, 2, 3$ and $\delta_1 = 1, 2, 3$ we used (7) to simulate $n_{\ell,i} = 10(50)$ observations at each dose level $d_{\ell,i}$, resulting in $n_{\ell} = 30(150)$, $\ell = 1, 2$.

The left side of Table 1 displays the coverage probabilities for $\alpha = 0.05, 0.1$. We observe that the nominal level of $1 - \alpha$ is reached in all cases under consideration, which confirms (2). The confidence intervals are more accurate for larger sample sizes and smaller variances, because we used the asymptotic quantiles from the normal distribution. If, instead, we select the quantiles from the t distribution, the simulated coverage probabilities are closer to the nominal $1 - \alpha$ level (results not shown here). Note that the confidence bounds perform better for larger values of δ_1 . This effect can be explained by a careful look at the proof given in Appendix A and the particular example under consideration. First note that the maximum absolute difference δ_1 between the two curves is attained at a single point, say

(a) Scenario 1



(b) Scenario 2




Figure 3: Graphical illustration of the two scenarios used for the simulations. Open dots in the left panel indicate the actual dose levels. In the right panel they indicate the doses where the maximum distance to the reference curve m_1 (dashed line) is observed.

d_0 ; see Figure 3a. If this difference is large then either $\max_{d \in \mathcal{D}} U(Y_1, Y_2, d) = U(Y_1, Y_2, d_0)$ or $-\min_{d \in \mathcal{D}} L(Y_1, Y_2, d) = L(Y_1, Y_2, d_0)$ with high probability and consequently there is equality either in (16) or (17) in Appendix A. The same effect appears for increasing sample sizes and smaller values of δ_1 as in this case the parameter estimates and approximation of the coverage probability of the confidence interval are more precise.

Scenario 2 We now consider the comparison of two different Emax models, where the maximum distances with respect to the same reference model are 0.25, 0.5, 1, 1.5 and 2. More specifically, we compared the reference Emax model $m_1(d) = 1 + \frac{9.70d}{6.70+d}$ with

$$m_2^1(d) = 1 + \frac{6.88d}{3.60+d}, \quad m_2^2(d) = 1 + \frac{5.66d}{2.25+d}, \quad m_2^3(d) = 1 + \frac{4.52d}{1+d}, \quad m_2^4(d) = 1 + \frac{4.05d}{0.48+d}, \quad m_2^5(d) = 1 + \frac{3.82d}{0.22+d}, \quad (8)$$

where the dose range is given by $\mathcal{D} = [0, 4]$. Note that the placebo response at $d = 0$ is 1 and the response at the highest dose $d = 4$ is 4.62 for all five models; see Figure 3b. The difference curve is given by $m_2^h(\vartheta_2^h, d) - m_1(\vartheta_1, d)$ for $h = 1, 2, 3, 4, 5$. Note that the dose which produces the maximum difference is different for each h . More precisely, these doses are given by 1.4, 1.28, 1.04, 0.82 and 0.61 for $h = 1, \dots, 5$; see again Figure 3b. The maximum absolute distance attained at each of these doses is denoted by $\delta_\infty = \max_{d \in \mathcal{D}} |m_2^h(\vartheta_2^h, d) - m_1(\vartheta_1, d)|$.

δ_1 σ^2		Coverage probabilities				Type I error rates			
		$\alpha = 0.05$		$\alpha = 0.1$		$\alpha = 0.05$		$\alpha = 0.1$	
		$n_\ell = 30$	$n_\ell = 150$	$n_\ell = 30$	$n_\ell = 150$	$n_\ell = 30$	$n_\ell = 150$	$n_\ell = 30$	$n_\ell = 150$
1	1	0.987	0.950	0.953	0.906	0.012	0.050	0.046	0.095
1	2	0.999	0.956	0.991	0.906	0.001	0.042	0.009	0.088
1	3	1.000	0.971	0.999	0.923	0.000	0.031	0.001	0.077
2	1	0.949	0.952	0.901	0.907	0.047	0.049	0.096	0.105
2	2	0.960	0.951	0.913	0.901	0.039	0.048	0.079	0.095
2	3	0.977	0.950	0.936	0.902	0.025	0.047	0.065	0.097
3	1	0.951	0.954	0.906	0.908	0.053	0.048	0.102	0.100
3	2	0.952	0.954	0.905	0.907	0.048	0.047	0.094	0.099
3	3	0.949	0.952	0.900	0.903	0.052	0.049	0.098	0.099

Table 1: Simulated coverage probabilities and Type I error rates for different configurations of δ_1 , σ^2 , α , and n_ℓ under Scenario 1.

We assumed identical dose levels $d_{\ell,i} = i - 1$, $i = 1, 2, 3, 4, 5$ for both regression models $\ell = 1, 2$. For each configuration of $\sigma^2 = 1, 2, 3$ and $\delta_\infty = 0.25, 0.5, 1, 1.5, 2$, we used (7) to simulate $n_{\ell,i} = 30$ observations at each dose level $d_{\ell,i}$, resulting in $n_\ell = 150$, $\ell = 1, 2$.

The left side of Table 2 displays the coverage probabilities for $\alpha = 0.05, 0.1$. As already seen under Scenario 1, the confidence intervals are more accurate for smaller variances (and larger sample sizes, results not shown here) and for increasing values of δ_∞ . As before, asymptotically the coverage probability is at least $1 - \alpha$ under all configurations investigated here.

2.3.2 Type 1 error rates

For the Type I error rate simulations we investigated the two scenarios from Figure 3 for each configuration of $\alpha = 0.05, 0.1$ and $\sigma^2 = 1, 2, 3$. Further, we set $\delta = \delta_\infty$ in (4). For a fixed configuration, we generated data according to (7), fit both models, performed the hypothesis test (6) and counted the proportion of rejecting the null hypothesis H . Note that due to the choice of δ both Scenarios 1 and 2 belong to the null hypothesis H defined in (4). Thus, rejecting H would be a Type I error, i.e. we would erroneously claim similarity of the two dose response curves.

The right side of Table 1 displays the simulated Type I error rates under Scenario 1. We observe that the simulated Type I error rate is bounded by the nominal significance level α for all configurations investigated here, indicating that the hypothesis test (6) is indeed a level- α test, even under total sample sizes as small as 30. Note also that the significance level is actually well exhausted under many configurations. For small sample sizes and small values of δ the test becomes conservative, matching the observed performance of the

			Coverage probabilities		Type I error rates	
(m_1, m_2)	δ_∞	σ^2	$\alpha = 0.05$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.1$
(m_1, m_2^1)	0.25	1	1.000	1.000	0.000	0.000
		2	1.000	1.000	0.000	0.000
		3	1.000	1.000	0.000	0.000
(m_1, m_2^2)	0.5	1	0.994	0.960	0.006	0.040
		2	1.000	0.993	0.000	0.007
		3	1.000	1.000	0.000	0.000
(m_1, m_2^3)	1	1	0.954	0.893	0.036	0.107
		2	0.963	0.903	0.047	0.097
		3	0.983	0.942	0.015	0.058
(m_1, m_2^4)	1.5	1	0.952	0.899	0.048	0.101
		2	0.962	0.913	0.038	0.087
		3	0.949	0.897	0.051	0.103
(m_1, m_2^5)	2	1	0.945	0.902	0.055	0.098
		2	0.942	0.889	0.068	0.118
		3	0.941	0.896	0.065	0.116

Table 2: Simulated coverage probabilities and Type I error rates for different model choices and configurations of σ^2 and α under Scenario 2, for $n_\ell = 150$, $\ell = 1, 2$.

confidence bounds shown in the left side of Table 1. Again, this conservatism disappears for large sample sizes.

The right side of Table 2 displays the simulated Type I error rates under Scenario 2. As before, the simulated Type I error rate is bounded by the nominal significance level α under all configurations. However, we observe that the test is very conservative for small values of δ_∞ , as already expected from the previously reported results on the coverage probabilities.

2.3.3 Power

We now consider testing the null hypothesis H in (4) for $\delta = 1$, where in fact the maximum difference is smaller than 1. We start with the comparison of the models from Scenario 1 for different values of δ_1 under the alternative; see Figure 4. The dose levels remain the same as under Scenario 1. For each configuration of $\sigma^2 = 1, 2, 3$ and $\delta_1 = 0, 0.25, 0.5, 0.75, 0.9$, we used (7) to simulate $n = 10(30, 50)$ observations under m_1 and m_2 at each dose level $d_{\ell,i}$, resulting in $n_\ell = 30(90, 150)$, $\ell = 1, 2$. Table 3 summarizes the results for $\alpha = 0.05, 0.1$. The power increases with decreasing values of δ_1 . For large values of σ^2 the power remains small, even for $\delta_1 = 0$. In these cases we need larger sample sizes n_ℓ in order to achieve reliable

results, as otherwise, due to the large variances, the confidence intervals in (2) become too wide and hence the test very conservative.




Figure 4: Graphical illustration of Scenario 1 used for the power simulations. Open dots indicate the actual dose levels.

Regarding Scenario 2, we tested the null hypothesis H in (4) using $\delta = 1$ and generating data under the models m_1 , m_2^1 and m_2^2 defined in (8). Hence we simulated the performance of the test under the alternative K in (5) for different choices of σ and α . For the sake of brevity we restrict ourselves again to a fixed total sample size of $n_\ell = 150$, $\ell = 1, 2$. Table 4 displays the simulated power. We observe that the test achieves high power, even for larger variances. However, the power decreases for an increasing true maximum distance between the models and for increasing variances.

2.4 Placebo-adjusted modeling

So far we assessed the similarity of two dose response models in terms of the maximum difference over the dose range under investigation. Sometimes one might be interested in adjusting for the placebo response, that is, the treatment effect relative to the placebo response, before comparing two dose response curves. In this case one has to modify the results from Section 2.1 as follows. Different to model (1), we consider the placebo-adjusted responses

$$Y_{\ell,i,j} = m_\ell(\vartheta_\ell, d_{\ell,i}) - m_\ell(\vartheta_\ell, 0) + \varepsilon_{\ell,i,j}, \quad j = 1, \dots, n_{\ell,i}, \quad i = 1, \dots, k_\ell, \quad \ell = 1, 2, \quad d_{\ell,i} \in \mathcal{D}.$$

δ_1	σ^2	$\alpha = 0.05$			$\alpha = 0.1$		
		$n_\ell = 30$	$n_\ell = 90$	$n_\ell = 150$	$n_\ell = 30$	$n_\ell = 90$	$n_\ell = 150$
0.00	1	0.211	0.966	0.999	0.426	0.988	0.999
0.25	1	0.170	0.939	0.997	0.377	0.974	0.999
0.50	1	0.102	0.731	0.917	0.268	0.843	0.958
0.75	1	0.046	0.306	0.444	0.143	0.433	0.583
0.90	1	0.023	0.111	0.144	0.074	0.195	0.245
0.00	2	0.002	0.544	0.911	0.046	0.749	0.967
0.25	2	0.001	0.479	0.867	0.045	0.692	0.941
0.50	2	0.001	0.302	0.628	0.030	0.500	0.770
0.75	2	0.000	0.119	0.247	0.012	0.245	0.391
0.90	2	0.000	0.050	0.098	0.011	0.128	0.181
0.00	3	0.000	0.196	0.651	0.007	0.434	0.822
0.25	3	0.000	0.162	0.576	0.005	0.382	0.758
0.50	3	0.000	0.098	0.365	0.004	0.263	0.558
0.75	3	0.000	0.040	0.142	0.002	0.128	0.276
0.90	3	0.000	0.021	0.050	0.001	0.072	0.126

Table 3: Simulated power for $\delta = 1$ and different configurations of δ_1 , σ^2 , α , and n_ℓ in Scenario 1.

(m_1, m_2)	δ_∞	σ^2	$\alpha = 0.05$	$\alpha = 0.1$
(m_1, m_1)	0	1	0.986	0.996
(m_1, m_2^1)	0.25	1	0.980	0.992
(m_1, m_2^2)	0.5	1	0.871	0.938
(m_1, m_1)	0	2	0.719	0.873
(m_1, m_2^1)	0.25	2	0.657	0.833
(m_1, m_2^2)	0.5	2	0.442	0.655
(m_1, m_1)	0	3	0.350	0.622
(m_1, m_2^1)	0.25	3	0.286	0.553
(m_1, m_2^2)	0.5	3	0.183	0.400

Table 4: Simulated power for different model choices and configurations of σ^2 and α under Scenario 2, for $\delta = 1$ and $n_\ell = 150$, $l = 1, 2$.

The confidence interval for the maximum absolute difference between the placebo-adjusted curves is then given by

$$P \left\{ \max_{d \in \mathcal{D}} |(m_2(\vartheta_2, d) - m_2(\vartheta_2, 0)) - (m_1(\vartheta_1, d) - m_1(\vartheta_1, 0))| \leq \max \left\{ \max_{d \in \mathcal{D}} U'(Y_1, Y_2, d), -\min_{d \in \mathcal{D}} L'(Y_1, Y_2, d) \right\} \right\} \geq 1 - \alpha,$$

where U' (Y_1, Y_2, d) and L' (Y_1, Y_2, d) denote the pointwise confidence bounds for the placebo-adjusted differences derived by the delta method. For example,

$$U' (Y_1, Y_2, d) = (m_2(\hat{\vartheta}_2, d) - m_2(\hat{\vartheta}_2, 0)) - (m_1(\hat{\vartheta}_1, d) - m_1(\hat{\vartheta}_1, 0)) + u_{1-\alpha}\hat{\rho}'(d),$$

where $\hat{\rho}'(d)$ is calculated for the difference of two placebo-adjusted dose response curves. Proceeding, the null hypothesis of interest becomes

$$H' : \max_{d \in \mathcal{D}} |(m_2(\vartheta_2, d) - m_2(\vartheta_2, 0)) - (m_1(\vartheta_1, d) - m_1(\vartheta_1, 0))| \geq \delta$$

and following (6) we reject H' if

$$-\delta < \min_{d \in \mathcal{D}} L' (Y_1, Y_2, d) \quad \text{and} \quad \max_{d \in \mathcal{D}} U' (Y_1, Y_2, d) < \delta. \quad (9)$$

To illustrate this methodology, we revisit the weight loss case study from Section 2.2. The individual model fits remain the same, i.e. $m_1(\vartheta_1, d) = 0.03 - 5.17 \frac{d}{7.94+d}$ for the o.d. regimen and $m_2(\vartheta_2, d) = -0.09 - 6.56 \frac{d}{31.24+d}$ for the b.i.d. regimen. Figure 5a displays the placebo-adjusted model fits $m_1(\hat{\vartheta}_1, d) - m_1(\hat{\vartheta}_1, 0)$ and $m_2(\hat{\vartheta}_2, d) - m_2(\hat{\vartheta}_2, 0)$, $d \in [0, 1]$, together with the individual observations, where only the range $[-7, 1]$ is displayed on the vertical axis for better readability. Figure 5b displays the difference $(m_2(\hat{\vartheta}_2, d) - m_2(\hat{\vartheta}_2, 0)) - (m_1(\hat{\vartheta}_1, d) - m_1(\hat{\vartheta}_1, 0))$ together with the associated 90% pointwise confidence intervals for each dose $d \in [0, 1]$. In this example, the estimated placebo effects are very similar compared with the original fits. Thus, the placebo-adjusted difference curve and its confidence bounds do not change much compared with the previous results of Section 2.2; see Figure 2. The maximum upper confidence bound for $\alpha = 0.1$ is $\max_{d \in \mathcal{D}} U' (Y_1, Y_2, d) = 2.255$, again observed at dose $d = 0.08$, and the minimum lower confidence bound is $\min_{d \in \mathcal{D}} L' (Y_1, Y_2, d) = -1.730$ at dose $d = 1$. That is, the maximum placebo-adjusted difference between the two regimens over the dose range $\mathcal{D} = [0, 1]$ lies between -1.730 and 2.255 . Therefore, similarity of the placebo-adjusted dose response curves can be claimed according to (9) as long as δ is larger than 2.255 .

3 Assessing the similarity of two target doses

This section focuses on assessing the similarity of two target doses. We consider the difference between the minimum effective doses (*MEDs*) of two dose response curves from two non-overlapping subgroups. We derive confidence intervals and statistical tests to decide at a given level α whether the absolute difference of two *MEDs* is smaller than a prespecified margin η . Furthermore, we illustrate the proposed methodology by revisiting the case study from 2.2 and investigate its operating characteristics.

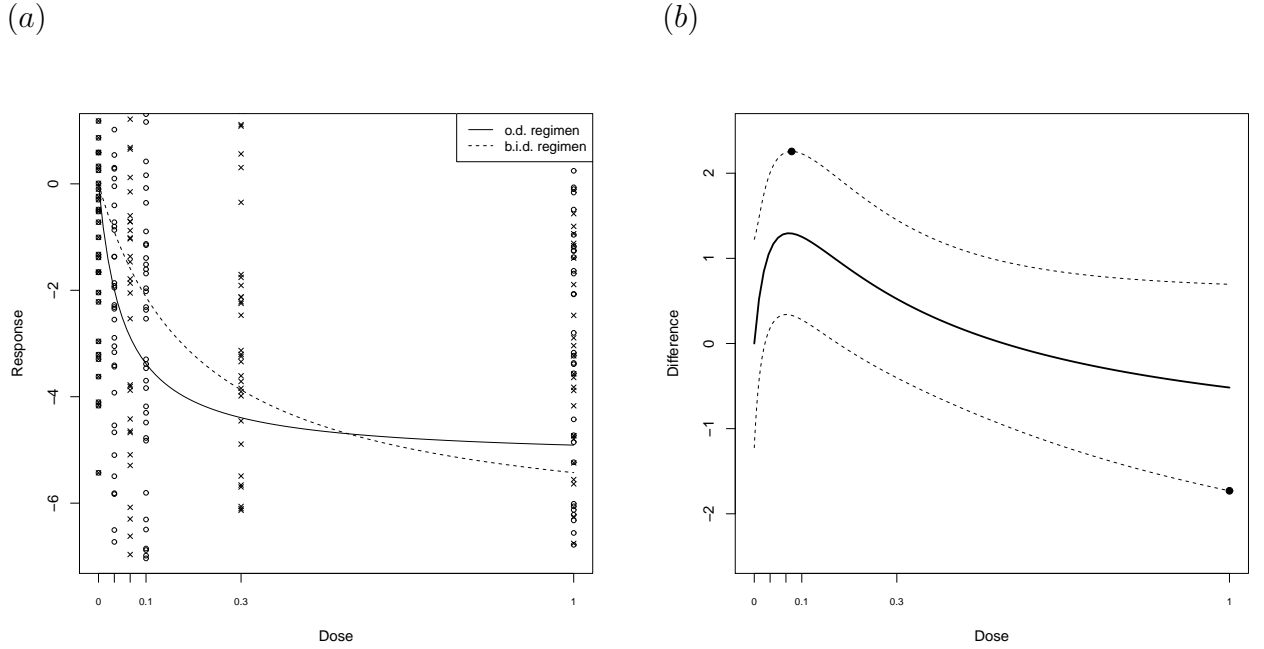


Figure 5: Placebo-adjusted plots for the weight loss case study. (a) The placebo-adjusted Emax model fit m_1 (m_2) for the o.d. (b.i.d.) regimen is given by the solid (dashed) line with observations marked by “x” (“o”). (b) Mean difference curve with associated pointwise 90% confidence bounds. Bold dots denote the maximum upper and minimum lower confidence bound over $\mathcal{D} = [0, 1]$.

3.1 Methodology

Following (1), the MED is defined as the smallest dose that produces a clinically relevant response Δ on top of the placebo effect (i.e. at dose $d = 0$). That is,

$$MED_\ell = MED_\ell(\vartheta_\ell) = \inf_{d \in \mathcal{D}} \{m_\ell(\vartheta_\ell, 0) < m_\ell(\vartheta_\ell, d) - \Delta\}, \ell = 1, 2. \quad (10)$$

From now on we assume strict monotonicity of the dose response curves m_ℓ such that (10) becomes

$$MED_\ell = MED_\ell(\vartheta_\ell) = m_\ell^{-1}(\vartheta_\ell, m_\ell(\vartheta_\ell, 0) + \Delta), \ell = 1, 2,$$

where the inverse is calculated with respect to d for fixed model parameters ϑ_1 and ϑ_2 . Estimates for the MED are then given by

$$\widehat{MED}_\ell = m_\ell^{-1}(\hat{\vartheta}_\ell, m_\ell(\hat{\vartheta}_\ell, 0) + \Delta), \ell = 1, 2,$$

where $\hat{\vartheta}_1$ and $\hat{\vartheta}_2$ are the non-linear least squares estimators for the true parameters. Due to the asymptotic normality of the estimates $\hat{\vartheta}_1$ and $\hat{\vartheta}_2$, the estimated difference of the MED s

is approximately normal distributed (11). To be more precise, the delta method (12) gives

$$\widehat{MED}_1 - \widehat{MED}_2 - (MED_1 - MED_2) \approx \mathcal{N}(0, \tau^2), \quad (11)$$

for

$$\tau^2 = \left(\frac{\partial}{\partial \vartheta_1} m_1^{-1}(\vartheta_1, \Delta_1) \right)^T \frac{\sigma_1^2 \Sigma_1^{-1}}{n_1} \frac{\partial}{\partial \vartheta_1} m_1^{-1}(\vartheta_1, \Delta_1) + \left(\frac{\partial}{\partial \vartheta_2} m_2^{-1}(\vartheta_2, \Delta_2) \right)^T \frac{\sigma_2^2 \Sigma_2^{-1}}{n_2} \frac{\partial}{\partial \vartheta_2} m_2^{-1}(\vartheta_2, \Delta_2)$$

and $\Delta_\ell = m_\ell(\vartheta_\ell, 0) + \Delta$, $\ell = 1, 2$. The variance τ^2 can be estimated by replacing ϑ_ℓ and Σ_ℓ by their estimates $\hat{\vartheta}_\ell$ and $\hat{\Sigma}_\ell$, $\ell = 1, 2$; see Section 2.1. The corresponding estimator is denoted by $\hat{\tau}^2$. It then follows from (11) that

$$P \left\{ MED_1 - MED_2 \in \left[\widehat{MED}_1 - \widehat{MED}_2 - u_{1-\alpha/2} \hat{\tau}, \widehat{MED}_1 - \widehat{MED}_2 + u_{1-\alpha/2} \hat{\tau} \right] \right\} \xrightarrow{n_1, n_2 \rightarrow \infty} 1 - \alpha, \quad (12)$$

and an asymptotic $(1 - \alpha)$ -confidence interval for the difference of the MED s is given by

$$\left[\widehat{MED}_1 - \widehat{MED}_2 - u_{1-\alpha/2} \hat{\tau}, \widehat{MED}_1 - \widehat{MED}_2 + u_{1-\alpha/2} \hat{\tau} \right].$$

In order to derive a test for similarity of two target doses we consider the problem of testing

$$H'' : |MED_1 - MED_2| \geq \eta \quad \text{against} \quad K'' : |MED_1 - MED_2| < \eta. \quad (13)$$

In Appendix B we show that rejecting H'' if

$$|\widehat{MED}_1 - \widehat{MED}_2| < c, \quad (14)$$

gives an asymptotic (uniformly most powerful) level α test, where c is the unique solution of the equation

$$\alpha = \Phi \left(\frac{c - \eta}{\hat{\tau}} \right) - \Phi \left(\frac{-c - \eta}{\hat{\tau}} \right). \quad (15)$$

Note that (15) can easily be solved by using Newton's algorithm (13).

3.2 Case study revisited

To illustrate the methodology in the previous subsection, we revisit the weight loss case study from Section 2.2. Recall the individual model fits $m_1(\hat{\vartheta}_1, d) = 0.03 - 5.17 \frac{d}{7.94+d}$ for the o.d. regimen and $m_2(\hat{\vartheta}_2, d) = -0.09 - 6.56 \frac{d}{31.24+d}$ for the b.i.d. regimen. We chose a clinically relevant difference of $\Delta = -3$. That is, a weight loss of 3% compared to the placebo response is assumed to be a clinically relevant effect on top of the placebo response at dose $d = 0$. Therefore, $\widehat{MED}_1 = m_1^{-1}(\hat{\vartheta}_1, 0.03 - 3) = 0.073$, $\widehat{MED}_2 = m_2^{-1}(\hat{\vartheta}_2, -0.09 - 3) = 0.176$ and $\widehat{MED}_1 - \widehat{MED}_2 = -0.103$. Figure 6(a) displays the model fits $m_\ell(\hat{\vartheta}_\ell, d)$, together with the estimates \widehat{MED}_ℓ , $\ell = 1, 2$.

The $1 - \alpha$ confidence interval for the true difference $MED_1 - MED_2$ is then given by $[-0.103 - u_{1-\alpha/2} 0.119, -0.103 + u_{1-\alpha/2} 0.119]$. For example,

$MED_1 - MED_2 \in [-0.338, 0.133]$ for $\alpha = 0.05$ and $MED_1 - MED_2 \in [-0.300, 0.094]$ for $\alpha = 0.1$. Applying the test in (14) for $\alpha = 0.05$ allows us to claim similarity of the two MED s whenever $\eta > 0.3$ as we have for the unique solution of (15) $c > 0.103 = |\widehat{MED}_1 - \widehat{MED}_2|$ in this case. Figure 6(b) displays the value of c as a function of η . For $\alpha = 0.1$ we obtain by similar calculations that η has to be larger than 0.255 in order to claim similarity.




Figure 6: Plots for the revisited weight loss case study. (a) The fitted Emax model m_1 (m_2) for the o.d. (b.i.d.) regimen is given by the solid (dashed) line, together with the estimated MED s for $\Delta = -3$. (b) Plot of the unique solution c of equation (15) as a function of η . The dashed lines indicate the absolute difference of the MED estimates and the minimum choice of η in order to claim similarity for $\alpha = 0.05$.

3.3 Simulations

We now report the results of a simulation study to investigate the operating characteristics of the method described in Section 3.1. Adapting the data generation algorithm from Section 2.3, we investigated the coverage probabilities of the confidence intervals in (12) as well as the Type I error rates and power of the test (14) for different scenarios. All results were obtained using 10,000 simulation runs.




Figure 7: Graphical illustration of Scenarios 3 and 4 used for the simulations. (a) displays the shifted Emax models with $\delta_1 = 2$. (b) displays the curves for Scenario 4, together with the MED s corresponding to $\Delta = 1.6$.

3.3.1 Coverage probabilities

Scenario 3 We start with the comparison of two shifted Emax models $m_1(d, \vartheta_1) = \delta_1 + 5d/(1+d)$ and $m_2(d, \vartheta_2) = 5d/(1+d)$ over $\mathcal{D} = [0, 4]$, with identical dose levels $d_{\ell,i} = i - 1, i = 1, \dots, 5$ for both regression models $\ell = 1, 2$; see Figure 7a. Because the models are shifted by the constant δ_1 , the true difference $MED_1 - MED_2 = 0$ regardless of the value for Δ . For each configuration of $\sigma^2 = 1, 2$ and $\delta_1 = 1, 2, 3$ we used (7) to simulate $n_{\ell,i} = 6(30)$ observations at each dose level $d_{\ell,i}$, resulting in $n_\ell = 30(150)$, $\ell = 1, 2$.

The left side of Table 5 displays the coverage probabilities for $\alpha = 0.05, 0.1$. We observe that the coverage probability is at least $1 - \alpha$ under all configurations. The confidence intervals are more accurate for larger sample sizes and smaller variances, which confirms the asymptotic result from (12). Furthermore, the simulated differences between the MED estimates are very close to the true difference under all configurations (results not shown here).

Scenario 4 We now consider the comparison of the Emax model $m_1(d, \vartheta_1) = 1 + 4d/(2+d)$ with the linear model $m_2(d, \vartheta_2) = 1 + 0.8d$ for the same set of doses as in Scenario 3. Note that the responses at doses $d = 0$ and $d = 3$ are the same in both models; see Figure 7b. For each configuration of $\sigma^2 = 1, 2, 3$ and $\Delta = 0.8, 1.6, 2.4$, we used again (7) to simulate

δ_1 σ^2		Coverage probabilities				Type I error rates			
		$\alpha = 0.05$		$\alpha = 0.1$		$\alpha = 0.05$		$\alpha = 0.1$	
		$n_\ell = 30$	$n_\ell = 150$	$n_\ell = 30$	$n_\ell = 150$	$n_\ell = 30$	$n_\ell = 150$	$n_\ell = 30$	$n_\ell = 150$
1	1	0.979	0.959	0.941	0.907	0.050	0.050	0.103	0.103
2	1	0.982	0.958	0.945	0.909	0.053	0.048	0.105	0.105
3	1	0.980	0.961	0.946	0.908	0.053	0.052	0.099	0.105
1	2	0.996	0.967	0.977	0.917	0.049	0.051	0.104	0.101
2	2	0.996	0.968	0.978	0.922	0.052	0.049	0.103	0.101
3	2	0.995	0.966	0.976	0.916	0.045	0.049	0.100	0.099

Table 5: Simulated coverage probabilities and Type I error rates for different configurations of δ_1 , σ^2 , α , and n_ℓ under Scenario 3.

$n_{\ell,i} = 6(30)$ observations at each dose level $d_{\ell,i}$, resulting in $n_\ell = 30(150)$, $\ell = 1, 2$.

The left side of Table 6 displays the coverage probabilities for $\alpha = 0.05, 0.1$. As before, asymptotically the coverage probability is at least $1 - \alpha$ under all configurations investigated here, except for small sample sizes and $\Delta = 2.4$ (in which case the *MEDs* coincide). This is a direct consequence of the definition of the *MED*. Inverting an Emax model $m(\vartheta, d) = y = \vartheta_1 + \vartheta_2 d / (\vartheta_3 + d)$ gives $m_1^{-1}(\vartheta, y) = \vartheta_3(y - \vartheta_1) / (\vartheta_1 + \vartheta_3 - y)$. Therefore higher values of Δ result in being closer to the pole of m^{-1} , which is at $\vartheta_1 + \vartheta_2 = 5$ in this case. However, further simulations show that the results get better for larger sample sizes and the coverage probabilities converge quickly to their nominal values. Finally, the simulated differences between the *MED* estimates are very close to the true difference under all configurations, except in the case where the *MEDs* coincide (i.e. $\Delta = 2.4$; results not shown here).

Δ σ^2		Coverage probabilities				Type I error rates			
		$\alpha = 0.05$		$\alpha = 0.1$		$\alpha = 0.05$		$\alpha = 0.1$	
		$n_\ell = 30$	$n_\ell = 150$	$n_\ell = 30$	$n_\ell = 150$	$n_\ell = 30$	$n_\ell = 150$	$n_\ell = 30$	$n_\ell = 150$
0.8	1	0.964	0.965	0.929	0.908	0.036	0.025	0.077	0.068
1.6	1	0.953	0.946	0.922	0.903	0.051	0.040	0.098	0.087
2.4	1	0.920	0.949	0.877	0.916	0.069	0.057	0.137	0.116
0.8	2	0.989	0.968	0.960	0.928	0.050	0.026	0.101	0.061
1.6	2	0.967	0.956	0.936	0.913	0.053	0.045	0.111	0.088
2.4	2	0.918	0.932	0.870	0.901	0.069	0.060	0.143	0.124

Table 6: Simulated coverage probabilities and Type I error rates for different configurations of Δ , σ^2 , α , and n_ℓ under Scenario 4.

3.3.2 Type 1 error rates

For the Type I error rate simulations we investigated the two scenarios from Figure 7. We start with Scenario 3. Because $|MED_1 - MED_2| = 0$ for all values of Δ , we chose $\eta = 0$. For a fixed configuration of parameters, we generated data according to (7), fit both models, performed the hypothesis test (14) and counted the proportion of rejecting the null hypothesis H'' . The right side of Table 5 displays the simulated Type I error rates under Scenario 3. We observe that the simulated Type I error rate is well exhausted at the nominal significance level α for all configurations investigated here, indicating that the hypothesis test (14) is indeed a level- α test, even under total sample sizes as small as 30.

The right side of Table 6 displays the simulated Type I error rates under Scenario 4. As before, the simulated Type I error rate is bounded by the nominal significance level α under almost all configurations. The test can be liberal for small sample sizes and large values of Δ , matching the observed performance of the confidence bounds shown in the left side of Table 6. Again, this size inflation disappears for large sample sizes.

3.3.3 Power

For the power simulations we again considered the two scenarios from Figure 7 and start with Scenario 3. Because $|MED_1 - MED_2| = 0$ for all values of Δ , the power of the test depends only on the given threshold η . For the concrete simulations, we set $\Delta = 1$ and used $\delta_1 = 1$ for convenience. For each configuration of $\sigma^2 = 1, 2, 3$ and $\eta = 0.1, 0.2, 0.5, 1$, we used (7) to simulate $n = 10(30, 50)$ observations under m_1 and m_2 at each dose level $d_{\ell,i}$, resulting in $n_{\ell} = 30(90, 150)$, $\ell = 1, 2$. All configurations belong to the alternative in (13). Table 7 summarizes the results for $\alpha = 0.05, 0.1$. The power increases with increasing values of η . The power decreases for larger values of σ^2 , especially for small values of η . In these cases we need larger sample sizes n_{ℓ} in order to achieve reliable results.

For the final set of simulations, we revisit Scenario 4 and investigate the power for different values of σ^2 and Δ . We set $\eta = 0.8$ and $n_{\ell} = 30, 150$ for $\ell = 1, 2$ and summarize the results in Table 8. In alignment with all former results, the performance of the test is worse in case of $\Delta = 2.4$ due to the already mentioned numerical problems when calculating the $MEDs$. In general, the power increases with increasing sample sizes and decreasing variances under all observed configurations. The power converges to 1 for $n_1, n_2 \rightarrow \infty$.

4 Conclusions

The choice of the equivalence margins δ and η in (4) and (13), respectively, is a delicate problem. This choice depends on the particular application and has to be made by clinical experts, possibly with input from statisticians and other quantitative scientists. Regulatory

η	σ^2	$\alpha = 0.05$			$\alpha = 0.1$		
		$n_\ell = 30$	$n_\ell = 90$	$n_\ell = 150$	$n_\ell = 30$	$n_\ell = 90$	$n_\ell = 150$
1	1	0.979	1.000	1.000	0.989	1.000	1.000
0.5	1	0.679	0.988	0.999	0.783	0.995	1.000
0.2	1	0.116	0.364	0.641	0.226	0.543	0.784
0.1	1	0.061	0.086	0.123	0.123	0.179	0.232
1	2	0.823	0.997	1.000	0.893	0.999	1.000
0.5	2	0.400	0.853	0.975	0.524	0.915	0.987
0.2	2	0.078	0.167	0.283	0.163	0.288	0.462
0.1	2	0.055	0.066	0.077	0.109	0.132	0.156

Table 7: Simulated power for different configurations of η , σ^2 , α , and n_ℓ in Scenario 3.

Δ	σ^2	$\alpha = 0.05$		$\alpha = 0.1$	
		$n_\ell = 30$	$n_\ell = 150$	$n_\ell = 30$	$n_\ell = 150$
0.4	1	0.914	1.000	0.958	1.000
0.8	1	0.116	0.625	0.261	0.778
1.6	1	0.057	0.080	0.118	0.152
2.4	1	0.090	0.118	0.163	0.233
0.4	2	0.668	0.999	0.806	0.999
0.8	2	0.089	0.324	0.183	0.523
1.6	2	0.058	0.060	0.116	0.122
2.4	2	0.081	0.093	0.165	0.189

Table 8: Simulated power for different configurations of Δ , σ^2 , α , and n_ℓ in Scenario 4.

guidance documents are available in specific settings, such as for the problem of demonstrating bioequivalence. For example, (14) discusses how the thresholds for bioequivalence hypotheses of the form considered in this paper can be defined in various settings. For the comparison of curves as considered in this paper we refer to Appendix 1 of (14), with emphasis on dissolution profiles on the basis of specific measures.

In this paper we investigated the problem of assessing similarity of dose response curves or target doses for two non-overlapping subgroups of patients for normally distributed responses. We leave several extensions of this basic problem for further research. For example, in certain applications it may be necessary to demonstrate similarity of dose response curves or target doses for more than two non-overlapping subgroups (such as more than two geographic regions or age classes). Finally, the research of this paper was motivated by the

need of comparing the dose response information from males with female or Japanese with non-Japanese patients. In practice, one may equally be interested in comparing males or Japanese with the overall population rather than the complementary subgroup. These cases are more difficult to handle because the data for the specific subgroup of interest is also included in the overall population, thus introducing correlations through nested data structure that need special attention. Again, we leave this topic for future research.

Acknowledgements This work has been supported in part by the Collaborative Research Center "Statistical modeling of nonlinear dynamic processes" (SFB 823, Project C1) of the German Research Foundation (DFG). Kathrin Möllenhoff's research has received funding from the European Union Seventh Framework Programme [FP7 20072013] under grant agreement no 602552 (IDEAL - Integrated Design and Analysis of small population group trials). The authors would like to thank Georgina Bermann for many helpful discussions and bringing this case study to our attention.

References

- [1] Ruberg, S.J. (1995). Dose response studies I. Some design considerations. *Journal of Biopharmaceutical Statistics*, **5**: 1–14.
- [2] Bretz F., Hsu J., Pinheiro J., Liu Y. (2008). Dose Finding – A Challenge in Statistics. *Biometrical Journal*, 50(4): 480–504.
- [3] ICH Harmonized Tripartite Guideline. (1994). Topic E4: Dose-response information to support drug registration. Available at <http://www.ich.org/LOB/media/MEDIA480.pdf>
- [4] Malinowski, H.J., Westelinck, A., Sato, J., Ong, T. (2008). Same Drug, Different Dosing: Differences in Dosing for Drugs Approved in the United States, Europe, and Japan. *J. Clin. Pharmacol.*, 48, 900–908.
- [5] Ueasaka, H. (2009). Statistical Issues in global drug development with a close look at Japan and the Asian region. Invited presentation at the 57th Session of the International Statistical Institute. Available at <http://www.statssa.gov.za/isi2009/ScientificProgramme/IPMS/0248.pdf>
- [6] Dette, H., Moellenhoff, K., Volgushev, S., Bretz, F. (2016) Equivalence of regression curves. Available at <http://arxiv.org/pdf/1505.05266v3.pdf>

- [7] Liu, W., Bretz, F., Hayter, A.J., Wynn, H.P. (2009). Assessing Nonsuperiority, Noninferiority, or Equivalence When Comparing Two Regression Models Over a Restricted Covariate Region. *Biometrics*, **65**(4):1279–1287.
- [8] Pinheiro, J., Bretz, F., Branson, M. (2006). Analysis of dose response studies: Modeling approaches. Pp. 146–171 in Ting, N. (ed.) *Dose finding in drug development*, Springer, New York.
- [9] Serfling, R. (1980). *Approximation Theorems of Mathematical Statistics*. Wiley.
- [10] R Core Team (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL: <https://www.R-project.org/>
- [11] Dette, H., Bretz, F., Pepelyshev, A., Pinheiro, J. (2008). Optimal designs for dose-finding studies. *Journal of the American Statistical Association*, 103(483), 1225-1237
- [12] Serfling, R. J. (1980). *Approximation Theorems of Mathematical Statistics*. New York: Wiley.
- [13] Deuffhard, P. (2001). *Newton methods for nonlinear problems: affine invariance and adaptive algorithms*. Springer.
- [14] European Medical Agency (2010). *Guideline on the investigation of bioequivalence*. Available at http://www.ema.europa.eu/docs/en_GB/document_library/Scientific_guideline/2010/01/WC500070039.pdf
- [15] Lehmann, E.L., Casella, G. (1986). *Testing statistical hypotheses*. Springer.

A. Coverage probability of the confidence interval for the maximum absolute difference

In the following we prove equation (2) from Section 2.1. To this end, let $d_0 \in \mathcal{D}$ such that

$$\max_{d \in \mathcal{D}} |m_2(\vartheta_2, d) - m_1(\vartheta_1, d)| = |m_2(\vartheta_2, d_0) - m_1(\vartheta_1, d_0)|.$$

Hence

$$\begin{aligned} \mathcal{P} &= P \left\{ \max_{d \in \mathcal{D}} |m_2(\vartheta_2, d) - m_1(\vartheta_1, d)| \leq \max \left\{ \max_{d \in \mathcal{D}} U(Y_1, Y_2, d), -\min_{d \in \mathcal{D}} L(Y_1, Y_2, d) \right\} \right\} \\ &= P \left\{ |m_2(\vartheta_2, d_0) - m_1(\vartheta_1, d_0)| \leq \max \left\{ \max_{d \in \mathcal{D}} U(Y_1, Y_2, d), -\min_{d \in \mathcal{D}} L(Y_1, Y_2, d) \right\} \right\} \\ &\geq P \left\{ |m_2(\vartheta_2, d_0) - m_1(\vartheta_1, d_0)| \leq \max \left\{ U(Y_1, Y_2, d_0), -L(Y_1, Y_2, d_0) \right\} \right\}. \end{aligned}$$

Now we distinguish two cases. If $m_2(\vartheta_2, d_0) - m_1(\vartheta_1, d_0) \geq 0$ we have

$$\mathcal{P} \geq P\left\{m_2(\vartheta_2, d_0) - m_1(\vartheta_1, d_0) \leq U(Y_1, Y_2, d_0)\right\} \xrightarrow{n_1, n_2 \rightarrow \infty} 1 - \alpha, \quad (16)$$

as $U(Y_1, Y_2, d)$ is a $1 - \alpha$ pointwise upper confidence bound on $m_2(\vartheta_2, d) - m_1(\vartheta_1, d)$. Otherwise, $m_2(\vartheta_2, d_0) - m_1(\vartheta_1, d_0) \leq 0$ and the same argument applies to $L(Y_1, Y_2, d)$, yielding

$$\mathcal{P} \geq P\left\{m_2(\vartheta_2, d_0) - m_1(\vartheta_1, d_0) \geq L(Y_1, Y_2, d_0)\right\} \xrightarrow{n_1, n_2 \rightarrow \infty} 1 - \alpha. \quad (17)$$

B. Asymptotic level of the test for similarity of two target doses

We show that the test (14) defined in Section 3.1 has asymptotic level α , that is

$$\lim_{n_1, n_2 \rightarrow \infty} P\left(|\widehat{MED}_1 - \widehat{MED}_2| \leq c\right) \leq \alpha \quad (18)$$

under the null hypothesis. First note that the solution of equation (15) is unique as the function $c \rightarrow \Phi\left(\frac{c-\eta}{\hat{\tau}}\right) - \Phi\left(\frac{-c-\eta}{\hat{\tau}}\right)$ is strictly increasing with limits -1 and 1 as $c \rightarrow -\infty$ and ∞ , respectively. Next, let $t = MED_1 - MED_2$, $\hat{t} = \widehat{MED}_1 - \widehat{MED}_2$ and denote the power function of the test by

$$G_{n_1, n_2}(\theta) = P(|\hat{t}| < c).$$

The assertion (18) is then equivalent to

$$\lim_{n_1, n_2 \rightarrow \infty} G_{n_1, n_2}(\theta) \leq \alpha \quad \text{for all } |t| \geq \eta. \quad (19)$$

A standard calculation shows that

$$\begin{aligned} G_{n_1, n_2}(t) &= P(|\hat{t}| \leq c) = P(-c \leq \hat{t} \leq c) = P\left(\frac{-c-t}{\hat{\tau}} \leq \frac{\hat{t}-t}{\hat{\tau}} \leq \frac{c-t}{\hat{\tau}}\right) \\ &\xrightarrow{n_1, n_2 \rightarrow \infty} \tilde{G}(t) := \Phi\left(\frac{c-t}{\tau}\right) - \Phi\left(\frac{-c-t}{\tau}\right) \end{aligned}$$

Now consider the problem of testing the hypotheses $H: |t| \geq \eta$ against $K: |t| < \eta$ for normally distributed data $X \sim \mathcal{N}(t, \tau^2)$. A simple calculation shows that the (asymptotic) power function \tilde{G} coincides with the power of the test, which rejects the null hypothesis $H: |t| \geq \eta$ whenever $|X| \leq c$. Considering the discussion in Lehmann et al. (15, p. 81), it follows that this test is uniformly most powerful and unbiased of size α . This implies $\tilde{G}(t) \leq \tilde{G}(\eta) = \alpha$ for all $|t| \geq \eta$ and proves (19).