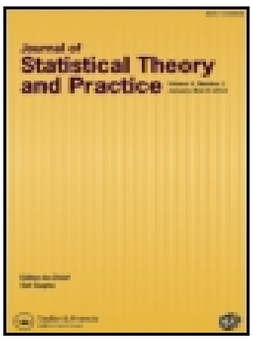

Journal of Statistical Theory and Practice

ISSN: 1559-8608 (Print) 1559-8616 (Online) Journal homepage: http://www.tandfonline.com/loi/ujsp20

Adaptive designs for quantal dose-response experiments with false answers

Norbert Benda, Paul-Christian Bürkner, Fritjof Freise, Heinz Holling & Rainer Schwabe

To cite this article: Norbert Benda, Paul-Christian Bürkner, Fritjof Freise, Heinz Holling & Rainer Schwabe (2016): Adaptive designs for quantal dose-response experiments with false answers, Journal of Statistical Theory and Practice, DOI: 10.1080/15598608.2016.1213209

To link to this article: http://dx.doi.org/10.1080/15598608.2016.1213209

Accepted author version posted online: 15 Jul 2016.
Published online: 15 Jul 2016.

JOURNAL OF STATISTICAL THEORY AND PRACTICE
http://dx.doi.org/10.1080/15598608.2016.1213209

Adaptive designs for quantal dose-response experiments with false answers

Norbert Benda[a], Paul-Christian Bürkner[b], Fritjof Freise[c], Heinz Holling[b], and Rainer Schwabe[c]

[a]Federal Institute for Drugs and Medical Devices, Bonn, Germany; [b]Institute for Psychology, University of Münster, Münster, Germany; [c]Institute for Mathematical Stochastics, University of Magdeburg, Magdeburg, Germany

ABSTRACT

In psychophysical experiments with quantal dose-response a common problem is the occurrence of lapses due to inattention or, as in forced choice experiments, the occurrence of incorrect guesses—or both. For these situations an optimized sequential design is proposed based on the Fisher information evaluated at the maximum likelihood estimate. This sequential design is compared to the classical Robbins–Monro method of stochastic approximation and modifications thereof in the original nonparametric approach, as well as adapted to the current model with false answers and to Wu's (1985) original method based on the maximum likelihood estimate by means of a simulation study. In these simulations the optimized sequential design turns out to perform substantially better than its nonparametric competitors, in particular, in the situation of starting points (initial guesses of the parameter of interest) that are far from the true value. Overall, the optimized version also outperforms the original maximum likelihood based method by Wu (1985).

ARTICLE HISTORY
Received 11 April 2016
Accepted 12 July 2016

KEYWORDS
Psychophysics; quantal response; false answers; forced choice experiments; sequential design; stochastic approximation

AMS SUBJECT CLASSIFICATION
62L05; 62L20; 62K05; 91E30

1. Introduction

The purpose of a sensitivity experiment is the estimation of a dose-response curve that relates a stimulus level to the probability of a (positive) response. Often the main parameter of interest is the dose for which the probability of a response is 50%, that is, the median of the response curve. In a biological assay experiment this refers to the median effective dose (ED50), the dose where half of a population of animals treated with a certain agent dies. In psychophysical research the dose-response curve is known as the psychometric function and represents the probability of detecting a stimulus given a fixed intensity level. Here, the median of the response curve is often considered as a threshold, since the early days of psychophysics by Gustav Theodor Fechner (1860), who described the psychophysical experiment still by a deterministic model, where the threshold intensity leads to an abrupt change in response. This article refers mainly to psychophysics. Nevertheless, the results are also applicable to biological assay problems.

In section 2 we specify the model for binary response with false answers, and in section 3 we start with the derivation of the locally optimal design for estimating the ED50 as a benchmark for competing designs. However, the locally optimal design cannot be used in

CONTACT Fritjof Freise fritjof.freise@ovgu.de University of Magdeburg, Institute for Mathematical Statistics, PF 4120, 39016 Magdeburg, Germany.

practice because it requires the prior knowledge of the parameters of the response curve to be estimated. Hence, adaptive procedures have to be taken into account that aim at estimating the unknown parameters on-line and incorporating those estimates into the generation of a suitable design. A commonly used method in adaptive, that is, sequentially generated designs for quantal response problems is the stochastic approximation procedure of Robbins and Monro (1951), which is nonparametric in the sense that it does not make use of the exact shape of the underlying response curve. For some modifications of that procedure see Wetherill and Glazebrook (1986). Section 3 reviews briefly this method, as well as Wu's related maximum likelihood-based method (Wu 1985). Adapting the sequential optimization based on the Fisher information to the model with false answers, we present an optimized design. Section 4 gives some considerations about the existence of the maximum likelihood estimator needed for the construction of the designs and the final estimate.

In section 5 we review some modifications of the Robbins–Monro procedure that are based on averaging (Polyak 1990; Ruppert 1988; Bather 1989) and that are more generally applicable. Simulations given in section 6 compare the optimized design with the other procedures. The article concludes with a discussion of the results in section 7.

2. Model specification

The most common models to describe the probability of a response $\tilde{\pi}(x)$ at a stimulus level are the probit and the logit model. Typically, the stimulus level x is measured on a logarithmic scale. Both models can be motivated by the assumption of an individual threshold at each observation modeled by a latent random variable Y^* that follows either a normal distribution in the probit model or a logistic distribution in the logit model. For a given stimulus x there should be a response $\tilde{Y} = 1$ when the stimulus exceeds (or equals) the individual threshold Y^* and $\tilde{Y} = 0$ otherwise. Then the probability $\tilde{\pi}(x)$ for an ideal response at stimulus x is given by $\tilde{\pi}(x) = F(x)$, where $F(x) = P(Y^* \leq x)$ denotes the distribution function of Y^*.

It is well known that the probit and the logit model agree very closely in a sufficiently large neighborhood of the ED50, and therefore the logit model is often preferred because of its greater computational ease. In this model we have

$$\tilde{\pi}(x) = \frac{1}{1 + \exp(-(x - \mu)/\sigma)}, \tag{1}$$

where μ and σ are the location and the scale parameter of the logistic distribution. Since $\tilde{\pi}(\mu) = \frac{1}{2}$, the parameter μ represents the median of the response curve, $\mu = \mathrm{ED50}$.

Here as well as in the probit model we have

$$\lim_{x \to \infty} \tilde{\pi}(x) = 1 \text{ and } \lim_{x \to -\infty} \tilde{\pi}(x) = 0. \tag{2}$$

This corresponds to the idea that a very high dose must always deliver a positive response, whereas a very low dose always leads to a negative response. A common problem in psychophysics is, however, the occurrence of false answers. In some ophthalmologic experiments, for instance, in perimetry, the person to be examined is asked whether he or she has recognized a stimulus with a given luminance. Inattention leads quite often to some "lapses." Even at levels far higher than the threshold, subjects tend to give sometimes a "no" answer,

and they tend to some "yes" answers at levels far lower than the threshold. This behavior can be described by the introduction of the conditional probability α_0 of a false negative answer given $Y^* \leq x$, when there should be a response, and the conditional probability α_1 of a false positive answer given $Y^* > x$, when there should be no response.

The observed response Y is again binary with $Y = 1$ for a positive and $Y = 0$ for a negative answer, and the probability $\pi(x) = P(Y = 1)$ of a positive answer can be obtained by the formula of total probability,

$$P(Y = 1) = P(Y = 1|Y^* \leq x)P(Y^* \leq x) + P(Y = 1|Y^* > x)P(Y^* > x) \tag{3}$$

$$= (1 - \alpha_0)F(x) + \alpha_1(1 - F(x)). \tag{4}$$

Hence, the dose-response curve changes to

$$\pi(x) = \alpha_1 + (1 - \alpha_0 - \alpha_1)\tilde{\pi}(x), \tag{5}$$

leading to

$$\lim_{x\to\infty} \pi(x) = 1 - \alpha_0 \text{ and } \lim_{x\to-\infty} \pi(x) = \alpha_1. \tag{6}$$

In a "forced choice" experiment, where the experimental subject has to choose between n possibilities, guessing leads to a correct response with a probability of $1/n$. Such an experiment can be described by Eq. (5), setting α_1 equal to the probability of a correct guess $1/n$ and α_0 to 0. An ophthalmologic experiment to measure visual acuity, for example, uses spatial discrimination tasks where the examined person has to decide between two patterns, leading to a correct guess with a probability of 0.5.

In a bioassay experiment the parameter α_1 can be regarded as a natural mortality whereas α_0 can be interpreted as a natural immunity. The case $\alpha_1 = \lambda$ and $\alpha_0 = 0$ is known as Abbott's formula (Abbott 1925).

Even after incorporating false answers or a guessing probability into our model we are still interested in the median μ of the original dose-response curve $\tilde{\pi}$, because $\tilde{\pi}$ is the probability of a (correct) response for all questions whose outcome is not due to a lapse. Hence, we still call μ the threshold, leading now to a probability of

$$\pi(\mu) = \frac{1 - \alpha_0 + \alpha_1}{2} \tag{7}$$

for a response.

Since in a psychophysical experiment the condition and the vigilance of the experimental person or the patient limit the duration of a session, the experiment should be designed to render a good estimate of the threshold using few stimulus presentations. A classical approach for this problem is sequential analysis. In bioassay experiments, where the need for a good design arises especially for ethical reasons, sequential procedures are often not applicable, since these experiments take normally several weeks to be carried out. In psychophysics, however, we have almost immediate answers and a good design should be adaptive, taking previous answers into account.

3. Adaptive designs

In many psychophysical experiments the slope parameter σ changes little between different experimental sessions. A good estimate can be obtained from previous experiments. Thus, it can often be appropriate to consider the scale parameter σ to be known. This is assumed throughout the entire article, in addition to the knowledge of both probabilities α_0 and α_1 of false answers.

Let x_n be the stimulus intensity in the nth experiment and let $y_n \in \{0, 1\}$ be the corresponding result with $y_n = 1$ for a response and $y_n = 0$ for a nonresponse. For the sake of convenience we use always the notation "response/nonresponse," although in psychophysics a response can be a "yes" answer and a nonresponse a "no" answer. In a forced choice experiment the response $y = 1$ refers to a "correct choice" and $y = 0$ to an "incorrect choice."

Note that $\alpha_1 = \lim_{x\to-\infty} P(Y = 1|x)$ and $\alpha_0 = \lim_{x\to\infty} P(Y = 0|x)$.

The asymptotic variance of the maximum likelihood estimator $\hat{\mu}_{ML}$ of the threshold μ is essentially given by the inverse $\left(I_{(x_1,\ldots,x_n)}(\mu)\right)^{-1}$ of the Fisher information,

$$I_{(x_1,\ldots,x_n)}(\mu) = \sum_{i=1}^{n} \frac{(\pi'(x_i))^2}{\pi(x_i)(1-\pi(x_i))} \tag{8}$$

in model (5) with σ known when n experiments are performed at the stimulus levels $x_1, \ldots, x_n$, respectively. This variance is minimized or, equivalently, the Fisher information is maximized when all stimulus levels $x_1, \ldots, x_n$ are chosen equal to the optimal level x^* that maximizes the information

$$I_x(\mu) = \frac{(\pi'(x))^2}{\pi(x)(1-\pi(x))} \tag{9}$$

for a single experiment performed at x. Applying the special form of model (5) and setting $z := \exp((x-\mu)/\sigma)$ we get

$$I_x(\mu) = \left(\frac{1-\alpha_0-\alpha_1}{\sigma}\right)^2 \frac{z^2}{(1+z)^2} \frac{1}{(1-\alpha_1+\alpha_0 z)(\alpha_1+(1-\alpha_0)z)} \tag{10}$$

and it is sufficient to maximize $I_x(\mu)$ as a function of z.

Solving $\frac{\partial}{\partial z} I_{x_n}(\mu) = 0$ we can maximize $I_{x_n}(\mu)$ by the unique positive solution z^* of

$$\alpha_0(1-\alpha_0)z^3 + \frac{1}{2}(\alpha_0\alpha_1 + (1-\alpha_0)(1-\alpha_1))(z^2 - z) - \alpha_1(1-\alpha_1) = 0\,. \tag{11}$$

Consequently, $I_x(\mu)$ is maximized by

$$x^* = \mu + \sigma \log(z^*)\,. \tag{12}$$

In a symmetric situation, $\alpha_0 = \alpha_1$, $I_x(\mu)$ is maximized by $z^* = 1$ and, hence, by $x^* = \mu$, as could be expected. Nevertheless, there are a number of situations where one error is much higher than the other. Especially for $\alpha_0 = 0$ we get

$$z^* = \frac{\sqrt{1+8\alpha_1}+1}{2} \tag{13}$$

which applies in particular to forced choice experiments and is in accordance with findings of, for example, Lord (1968). In this situation we have $1 \leq z^* < 2$. Similarly, we get for $\alpha_1 = 0$:

$$z^* = \frac{2}{\sqrt{1 + 8\alpha_0} + 1}, \tag{14}$$

where $1/2^* \leq 1$. Note that, in general,

$$\left.\begin{array}{l} x^* > \mu \\ x^* = \mu \\ x^* < \mu \end{array}\right\} \Leftrightarrow \left\{\begin{array}{l} z^* > 1 \\ z^* = 1 \\ z^* < 1 \end{array}\right\} \Leftrightarrow \left\{\begin{array}{l} \alpha_0 < \alpha_1 \\ \alpha_0 = \alpha_1 \\ \alpha_0 > \alpha_1 \end{array}\right. \tag{15}$$

Hence, the optimal level x^* is larger than the threshold μ if $\alpha_0 < \alpha_1$ and smaller if $\alpha_1 < \alpha_0$. Note that if α_0 is fixed, z^* and hence x^* are increasing functions of α_1. For fixed α_1 the optimal level x^* is decreasing in α_0.

An important application of Eq. (12) is a forced choice experiment with two possible answers, where $\alpha_0 = 0$ and $\alpha_1 = 0.5$ and $\pi(\mu) = 0.75$. Here we get the famous golden section number $z^* = (\sqrt{5} + 1)/2$, leading to a probability of a correct response of 0.81.

However, the locally optimal Fisher information obtained from repetitions of the optimal level $x^* = \mu + \sigma \log(z^*)$ can only serve as a benchmark, as such a design cannot be realized due to the fact that μ is unknown. Therefore, one must look for competing designs which aim at approximating the locally optimal Fisher information by an adaptive selection of the levels $x_1, \ldots, x_n$ based on the previous observations.

For estimating the $100p$th percentile ED$100p$, that is, the dose that leads to a response with a probability p, Robbins and Monro (1951) proposed to choose

$$x_{n+1} = x_n - c_n(y_n - p) \tag{16}$$

as the stimulus level for the $(n+1)$st experiment. After N trials ED$100p$ is estimated by x_{N+1}. According to Chung (1954) and Sacks (1958) the estimate x_{N+1} is asymptotically normal and its asymptotic variance is minimized within this class of estimators by setting $c_n = 1/(\lambda n)$ where $\lambda = \pi'(\mathrm{ED}\,100p)$ is the slope of the response curve π at the percentile ED$100p$ of interest.

In particular, for estimating the threshold μ in model (5) with false answers the probability p has to be replaced by $\pi(\mu)$ in recursion (9) and we have $\lambda = (1 - \alpha_0 - \alpha_1)/(4\sigma)$ in that situation. Hence, we get the optimal Robbins–Monro procedure in the model with false answers as

$$x_{n+1} = x_n - \frac{4\sigma}{(1 - \alpha_0 - \alpha_1)n}\left(y_n - \frac{1 - \alpha_0 + \alpha_1}{2}\right). \tag{17}$$

Note that for $\alpha_0 < \alpha_1$ the Robbins–Monro procedure (17) increases the dose by larger steps proportional to $(1 - \alpha_0 + \alpha_1)/2$ in the case of a nonresponse than it decreases the dose, $(1 + \alpha_0 - \alpha_1)/2$, in the case of a response. As the Robbins–Monro procedure is, finally, oscillating around the threshold μ, this observation is in agreement with the fact that slightly larger levels should be preferred according to Eq. (15).

Wu (1985) suggested an adaptive design that uses for the stimulus level x_{n+1} the maximum-likelihood estimate $\hat{\mu}_{ML}$ of μ corresponding to the observations made at $(x_1, \ldots, x_n)$. The asymptotic properties of this procedure have been established by Ying

and Wu (1997). For the logistic regression model of Eq. (1) Wu's proposition meets the idea of McLeish and Tosh (1990), who used the Fisher information to obtain a good adaptive design.

Recall that after n experiments the Fisher information is given by

$$I_{(x_1,\ldots,x_n)}(\mu) = \sum_{i=1}^{n} \frac{(\pi'(x_i))^2}{\pi(x_i)(1-\pi(x_i))}\,. \tag{18}$$

Since the threshold μ is unknown, the Fisher information has to be estimated, which can be done by plugging in the maximum likelihood estimate $\hat{\mu}_{ML}$ for μ. Hence, for maximizing the estimated information $I_{(x_1,\ldots,x_n,x_{n+1})}(\hat{\mu}_{ML})$ of a subsequent experiment, we are led to choose the $(n+1)$st stimulus level by maximizing

$$I_{x_{n+1}}(\hat{\mu}_{ML}) = \frac{(\pi'(x_{n+1}))^2}{\pi(x_{n+1})(1-\pi(x_{n+1}))} \tag{19}$$

which, in turn, is the locally optimal level $x^* = \hat{\mu}_{ML} + \sigma\log(z^*)$ when $\hat{\mu}_{ML}$ would be the true threshold. Hence, for our optimized design we choose the $(n+1)$st stimulus level as

$$x_{n+1} = \hat{\mu}_{ML}(x_1,\ldots,x_n;\sigma) + \sigma\log(z^*)\,, \tag{20}$$

where $\hat{\mu}_{ML}(x_1,\ldots,x_n;\sigma)$ is the maximum likelihood estimate of μ associated with the first n observations when σ is given.

In order to construct a specific design applying Eq. (20) at each stage of the experiment, we make use of the information we obtained so far by employing the current maximum likelihood estimate of the threshold. The difference from Wu's (1985) design of taking the maximum likelihood estimate as the next design point,

$$x_{n+1} = \hat{\mu}_{ML}(x_1,\ldots,x_n;\sigma)\,, \tag{21}$$

lies in the additional term $\sigma\log(z^*)$.

4. Existence of the maximum likelihood estimate

In order to apply the proposed optimized design as well as Wu's MLE design, we have to calculate the maximum likelihood estimate for μ. After n experiments let $x_1^{(1)},\ldots,x_r^{(1)}$ be the stimulus levels associated with a response $(y=1)$ and $x_1^{(0)},\ldots,x_s^{(0)}$ the stimulus levels associated with a nonresponse $(y=0)$, where $n=r+s$.

In the usual logistic model with known σ and without guessing or lapses, that is, $\alpha_0=\alpha_1=0$, the log-likelihood is given by

$$l(\mu) = \sum_{i=1}^{r}\log\frac{1}{1+\exp(-(x_i^{(1)}-\mu)/\sigma)} + \sum_{j=1}^{s}\log\frac{1}{1+\exp((x_j^{(0)}-\mu)/\sigma)}\,. \tag{22}$$

A finite estimate for μ exists as soon as there is at least one response and one nonresponse. This follows from the continuity of $\tilde{\pi}$, since $\lim_{\mu\to\infty} l(\mu) = \lim_{\mu\to-\infty} l(\mu) = -\infty$. The estimate is in fact unique due to the convexity of $-\log\tilde{\pi}$ and $-\log(1-\tilde{\pi})$.

In the situation including guessing or the possibility for lapses, existence and uniqueness do not follow this easily. Even assumptions that would guarantee the existence

of the maximum likelihood estimate of the parameter vector (μ, σ) in the two parameter logistic model with unknown σ are not sufficient, as the following example shows.

According to Silvapulle (1981) the simple logistic regression model of Eq. (1) with the parameterization commonly used in generalized linear models $\text{logit}(\tilde{\pi}) = \beta_1 + \beta_2 x$ leads to a well-defined (finite and unique) maximum likelihood estimation of $(\beta_1, \beta_2) = (-\mu/\sigma, 1/\sigma)$ if and only if $r, s > 0$ and there exist $x_i^{(0)} < x_j^{(1)} < x_k^{(0)}$ or $x_i^{(1)} < x_j^{(0)} < x_k^{(1)}$, that is, the responses and nonresponses are not separated.

Now consider, for example, the situation of a forced choice experiment with two answers, that is, $\alpha_0 = 0$, $\alpha_1 = 0.5$, and $r < s$, $x_1^{(0)} > \ldots > x_s^{(0)}$, and $x_i^{(1)} = x_i^{(0)}, 1 \leq i \leq r$. It is easy to see that the supremum of the likelihood function is $(1/2)^n$, leading to an infinite maximum likelihood estimate. The problem that occurs in this example is that there are too few responses compared to what has to be expected by pure guessing.

Therefore, we have to derive conditions that are feasible in the situation of false answers for a general setting of parameters. First note that the log-likelihood for μ is given by

$$l(\mu) = \sum_{i=1}^{r} \log \frac{1 - \alpha_0 + \alpha_1 \exp(-(x_i^{(1)} - \mu)/\sigma)}{1 + \exp(-(x_i^{(1)} - \mu)/\sigma)} + \sum_{j=1}^{s} \log \frac{\alpha_0 + (1 - \alpha_1) \exp(-(x_j^{(0)} - \mu)/\sigma)}{1 + \exp(-(x_j^{(0)} - \mu)/\sigma)}. \tag{23}$$

If either $r = 0$ or $s = 0$, then the log-likelihood is monotonic in μ. Hence, as in the case of the usual logit model, at least one response and one nonresponse are necessary for the existence of a finite estimate. Let $r \geq 1$ and $s \geq 1$.

If a finite maximum likelihood estimate $\hat{\mu}_{ML}$ exists then $l(\hat{\mu}_{ML})$ is at least as large as the maximum of the limits $\lim_{\mu \to -\infty} l(\mu) = r \log(1 - \alpha_0) + s \log(\alpha_0)$ and $\lim_{\mu \to \infty} l(\mu) = r \log(\alpha_1) + s \log(1 - \alpha_1)$. If, on the other hand, some μ^* exists such that $l(\mu^*) \geq \max\{\lim_{\mu \to \infty} l(\mu), \lim_{\mu \to -\infty} l(\mu)\}$ then there exists a finite estimate $\hat{\mu}_{ML}$ since the log-likelihood function is continuous and $l(\mu^*) \leq l(\hat{\mu}_{ML})$. Thus, a finite estimate $\hat{\mu}_{ML}$ exists if and only if there exist a μ^* such that

$$l(\mu^*) \geq \max\{r \log(\alpha_1) + s \log(1 - \alpha_1), r \log(1 - \alpha_0) + s \log(\alpha_0)\} \tag{24}$$

in the case $\alpha_1 > 0$ and $\alpha_0 > 0$, when both false positive and false negative answers may occur. In the case when there are only false positive answers, that is, $\alpha_1 > 0$ and $\alpha_0 = 0$, we have $\lim_{\mu \to -\infty} l(\mu) = -\infty$ and consequently

$$l(\mu^*) \geq r \log(\alpha_1) + s \log(1 - \alpha_1). \tag{25}$$

In the case of $\alpha_0 > 0$ and $\alpha_1 = 0$, when there are only false negative answers, this becomes

$$l(\mu^*) \geq r \log(1 - \alpha_0) + s \log(\alpha_0). \tag{26}$$

The following theorem gives sufficient conditions for the existence of the maximum likelihood estimate, which secure that inequality (24) holds. Similar to the results of Silvapulle (1981), the conditions are based on the location of the design points where responses and nonresponses occur.

Theorem 4.1. Let $r \geq 1$ and $s \geq 1$. Let $x_i^{(1)}, i = 1, \ldots, r$, and $x_j^{(0)}, j = 1, \ldots, s$, be defined as at the beginning of this section. Then either of the following statements (i) and (ii) is sufficient for the existence of a finite maximum likelihood estimate of μ in model (5) when σ is known:

(i)
$$\alpha_0^s(1-\alpha_0)^r \geq \alpha_1^r(1-\alpha_1)^s \quad \text{and}$$
$$\alpha_0 \sum_{i=1}^{r} \exp(-x_i^{(1)}/\sigma) < (1-\alpha_0) \sum_{j=1}^{s} \exp(-x_j^{(0)}/\sigma)\,. \tag{27}$$

(ii)
$$\alpha_0^s(1-\alpha_0)^r \leq \alpha_1^r(1-\alpha_1)^s \quad \text{and}$$
$$(1-\alpha_1) \sum_{i=1}^{r} \exp(x_i^{(1)}/\sigma) > \alpha_1 \sum_{j=1}^{s} \exp(x_j^{(0)}/\sigma)\,. \tag{28}$$

Proof. (i) Let $0 \leq \theta < \infty$ and define the function

$$\tilde{l}(\theta) = \sum_{i=1}^{r} \log \frac{(1-\alpha_0)\exp(x_i^{(1)}/\sigma) + \alpha_1\theta}{\exp(x_i^{(1)}/\sigma) + \theta} + \sum_{j=1}^{s} \log \frac{\alpha_0 \exp(x_j^{(0)}/\sigma) + (1-\alpha_1)\theta}{\exp(x_j^{(0)}/\sigma) + \theta}\,. \tag{29}$$

Note that, for $0 < \theta < \infty$, Eq. (29) shows the log-likelihood for our model after reparameterization using $\theta = \exp(\mu/\sigma)$. Including $\theta = 0$ it is a continuously differentiable extension of the log-likelihood. Also note that we can restrict out attention to the finiteness of the maximum likelihood estimate for $0 < \theta < \infty$, because in this case $\hat{\mu}_{ML} = \sigma \log(\hat{\theta}_{ML})$.

Since $\alpha_0^s(1-\alpha_0)^r \geq \alpha_1^r(1-\alpha_1)^s$ we have $\lim_{\theta\to 0} \tilde{l}(\theta) = \tilde{l}(0) \geq \lim_{\theta\to\infty} \tilde{l}(\theta)$. Moreover, we get

$$\frac{\partial}{\partial\theta}\tilde{l}(\theta)\bigg|_{\theta=0} = (1-\alpha_0-\alpha_1)\left(\sum_{j=1}^{s} \frac{1}{\alpha_0 \exp(x_j^{(0)}/\sigma)} - \sum_{i=1}^{r} \frac{1}{(1-\alpha_0)\exp(x_i^{(1)}/\sigma)}\right) > 0\,, \tag{30}$$

which follows from the second condition. Hence by the continuity of $\tilde{l}$ and its derivative there exist $\theta^* > 0$ in a neighborhood of 0 such that $\tilde{l}(\theta^*) \geq \tilde{l}(0)$. The arguments stated prior to the inequality in Eq. (24) yield the finiteness of the maximum likelihood estimate $\hat{\theta}_{ML}$ and hence of $\hat{\mu}_{ML}$.

The second part follows by symmetry.

It seems that the complications concerning the existence are due to "boundary conditions" imposed by α_0 and α_1. The curve fitted to the data describes the relative frequencies of the responses. Hence, no finite maximum likelihood estimate will exist if the relative frequency of responses is smaller than the guessing probability α_1 or the relative frequency of nonresponses is smaller than the probability α_0 of lapses. The value that is chosen for σ influences the existence in these cases, too. In examples a large value of σ, that is, a smaller slope of the curve, appeared to be beneficial.

This is, as we have seen after Eq. (22), in contrast to the logistic model without guessing or lapses. Here the curve describing the probability can be shifted to fit arbitrary relative frequencies, independent of σ. If in this model, where $\alpha_0 = \alpha_1 = 0$, the estimate exists for a given sample of size n, it exists also if new observations are added. With guessing or lapses in the model, this does not hold.

5. Modifications of the Robbins–Monro procedure

As already stated in section 3, the minimal asymptotic variance for the Robbins–Monro procedure is achieved for $c_n = 1/(\lambda n)$, where λ denotes the slope of the response curve π at the percentile of interest. In more general situations, where the scale parameter σ and hence the slope λ is not known, it is not possible to choose the step length c_n optimally. One way to avoid this problem is to substitute λ by a strongly consistent estimator for the slope. This adaptive algorithm using an estimate for λ has the minimal asymptotic variance property (Venter 1968; Lai and Robbins 1979; Schwabe 1986; see also Morgan 1992). A disadvantage of this approach is that the slope estimates, which are based on the observations from the recursion, are likely to be very unstable for small and moderate sample sizes. Additionally, they are sensitive to bad starting values.

Polyak (1990) and, independently, Ruppert (1988) proposed using a Robbins–Monro stochastic approximation scheme

$$x_{n+1} = x_n - c_n(y_n - \pi(\mu)) \tag{31}$$

with larger step sizes $c_n = c/n^\gamma$, where $\frac{1}{2} < \gamma < 1$, and estimate the threshold μ by the mean $\bar{x}_n = \frac{1}{n}\sum_{i=1}^{n} x_i$ of the stimulus levels $x_1, \ldots, x_n$. For every c this procedure shares the optimality properties of the Robbins–Monro procedure with optimal step size. Hence, no prior knowledge of σ is requested for that procedure. In our simulation study we chose $\gamma = 3/4$ and $c = 1/\lambda$, that is, equal to the optimal value for the Robbins–Monro case.

Alternatively, Bather (1989) proposed a recursion

$$x_{n+1} = \bar{x}_n - nc_n(\bar{y}_n - \pi(\mu)) \tag{32}$$

based on the mean $\bar{x}_n$ of the levels and the proportion $\bar{y}_n$ of the responses. Also, here the threshold μ is estimated by the mean $\bar{x}_n$. Bather motivated his proposition by maximum likelihood arguments. Based on simulations he argued for larger step sizes $c_n = c/n^\gamma$, $\frac{1}{2} < \gamma < 1$, as for the Polyak–Ruppert scheme. Schwabe and Walk (1996) established that also Bather's procedure shares the properties of the Robbins–Monro procedure with optimal step size. Moreover, Bather's procedure shows a better performance in handling bad starting points. For the simulations the same values for c and γ as in the Polyak and Ruppert case were used.

In analogy to the sequentially optimized design, where the $(n+1)$st stimulus level is an adjusted maximum likelihood estimate, the three stochastic approximation procedures can be modified, such that they converge to $\mu + \sigma\log(z^*)$ instead of μ. The adjusted recursions are then

$$x_{n+1} = x_n - c_n(y_n - \pi(\mu + \sigma\log(z^*))) \tag{33}$$

for Robbins–Monro and Polyak–Ruppert and

$$x_{n+1} = \bar{x}_n - nc_n(\bar{y}_n - \pi(\mu + \sigma\log(z^*)))\,, \tag{34}$$

for Bather. Now an estimate of μ can be calculated by $x_{n+1} - \sigma\log(z^*)$ and $\bar{x}_{n+1} - \sigma\log(z^*)$, respectively.

6. Simulations

In our simulations we compared Wu's maximum likelihood based design from Eq. (21), the sequentially optimized design from Eq. (20), the Robbins–Monro procedures of Eqs. (16) and (17), and the averaging based methods and their corresponding adjusted methods from section 5. As in the previous sections, it is assumed that α_0, α_1 and σ are known.

To quantify their performance we estimated the root mean squared error (RMSE) from 10,000 simulation runs each. For Wu's design in Eq. (21) and the sequentially optimized design (20) we used steps of the Robbins–Monro procedure and its adjusted version, respectively, whenever the maximum likelihood estimate did not exist according to the criterion given in the theorem of section 4.

Since in a symmetric situation Eq. (20) coincides with Eq. (21), we chose different error rates. For the simulation studies we considered the cases ($\alpha_0 = 0$, $\alpha_1 = 0.2$) and ($\alpha_0 = 0$, $\alpha_1 = 0.5$), which refer (for example) to a forced choice experiment with five and with two choices, respectively. For a moderate sample size the starting value should have an important influence on the behavior of the design. Therefore, we chose x_1 to $\mu + i\sigma$, $i = -10, -9, \ldots, 10$. For a more detailed inspection in the neighborhood of the threshold a finer grid $i = -5, -4.9, -4.8, \ldots, 5$ was used. Note that the doses can be rescaled such that $\sigma = 1$, since σ is known. Hence all results are stated for $\sigma = 1$. For the sample sizes we chose $N = 40$, $N = 60$, and $N = 100$.

The estimated RMSE as a function of the starting value is shown in Figures 1 and 2. The left column displays the whole interval of initial values, while the right column shows the interval from -5 to 5 for a more detailed view. The estimated bias for $N = 60$ is displayed in Figure 3. The curves for the RMSE and the bias were fitted by using a cubic smoothing spline, where the smoothing parameter was chosen by cross validation.

From the figures it can be observed that for all designs the optimal starting value x_1 is larger than the threshold μ, in the sense of smallest RMSE.

As could be expected, the optimized design behaves better than Wu's design for most initial values. Asymptotically this gain seems rather small, since using Eq. (10) the asymptotic relative efficiency of Wu's design with respect to the optimized design, which is given by $I_\mu(\mu)/I_{x^*}(\mu)$, is 0.98 for $\alpha_1 = 0.2$ and 0.92 for $\alpha_1 = 0.5$. Only for very large initial values larger than 6 or 7, the optimized design behaves worse.

Nevertheless, the main advantage lies in the robustness against a poor starting value, which affects the optimized design much less than the others. For $\alpha_1 = 0.2$ the curve of the RMSE is almost constant over the central interval from -5 to 5. The interval increases with the sample size. Since the variability is higher if $\alpha_1 = 0.5$, this effect is visible only for larger sample sizes, namely, $N = 100$. This is also visible in the graphs of the estimated bias.

For $\alpha_1 = 0.5$ the optimized design is uniformly better than the Robbins–Monro procedure. In the case $\alpha_1 = 0.2$ the results of the Robbins–Monro procedure are comparable to those of the MLE-based methods, if the initial value is close enough to the threshold μ. For small sample sizes the Robbins–Monro procedure is superior to the other stochastic approximation methods, if a good starting value was chosen. This gain diminishes for increasing sample size.

Bather's method shows a better handling of bad starting points for larger sample sizes, as already mentioned. This property can be seen in the results from the Polyak–Ruppert

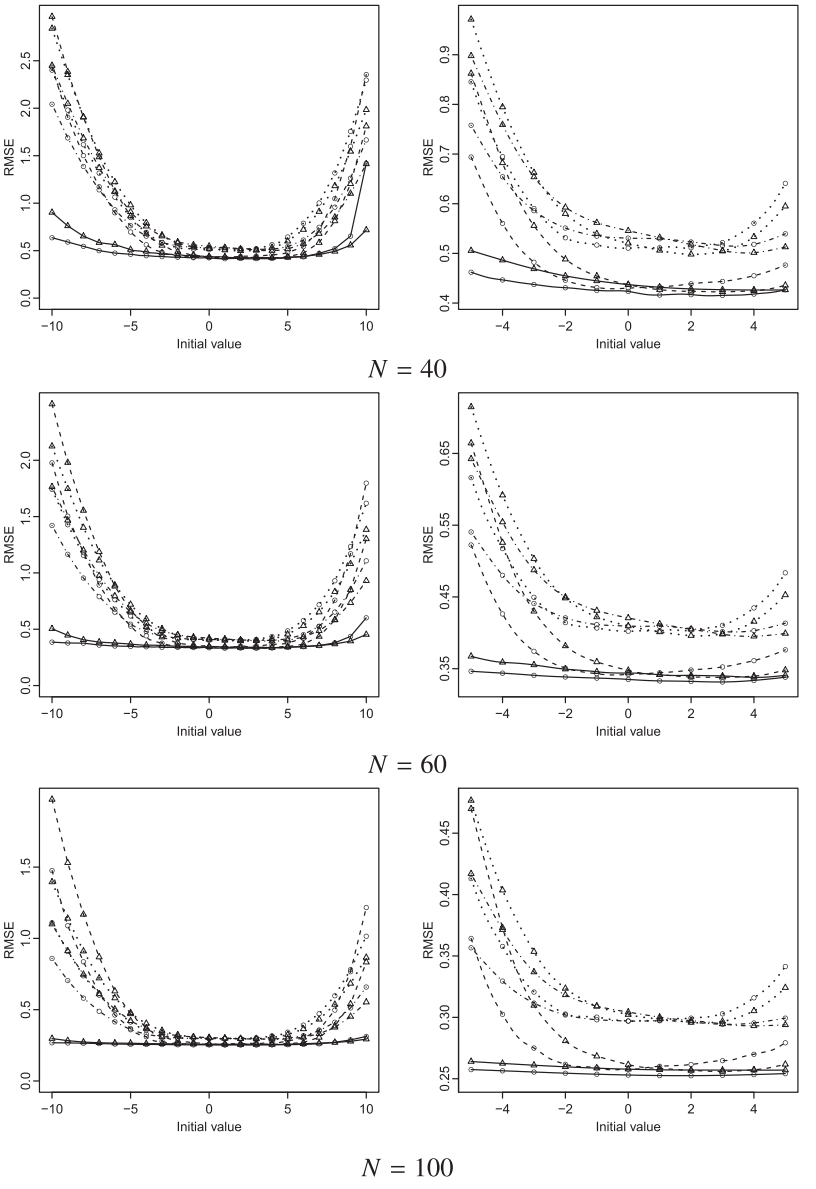

Figure 1. Estimated RMSE from the simulations for $\alpha_0 = 0$ and $\alpha_1 = 0.2$. Methods are displayed in different line styles: maximum likelihood based methods (solid), Robbins–Monro (dashed), Polyak–Ruppert (dotted), and Bather (dash-dotted). Versions are marked by different symbols: original (triangles) and adapted (circles).

scheme as well, but an even larger sample size or a more extreme (bad) starting point is needed to make use of it.

For a poor starting value, which is lower than μ, the adjusted schemes perform better than the originals. In the case of $\alpha_1 = 0.5$ and small to moderate sample sizes their behavior is comparable to $(N = 60)$ or even better than $(N = 40)$ Wu's MLE design.

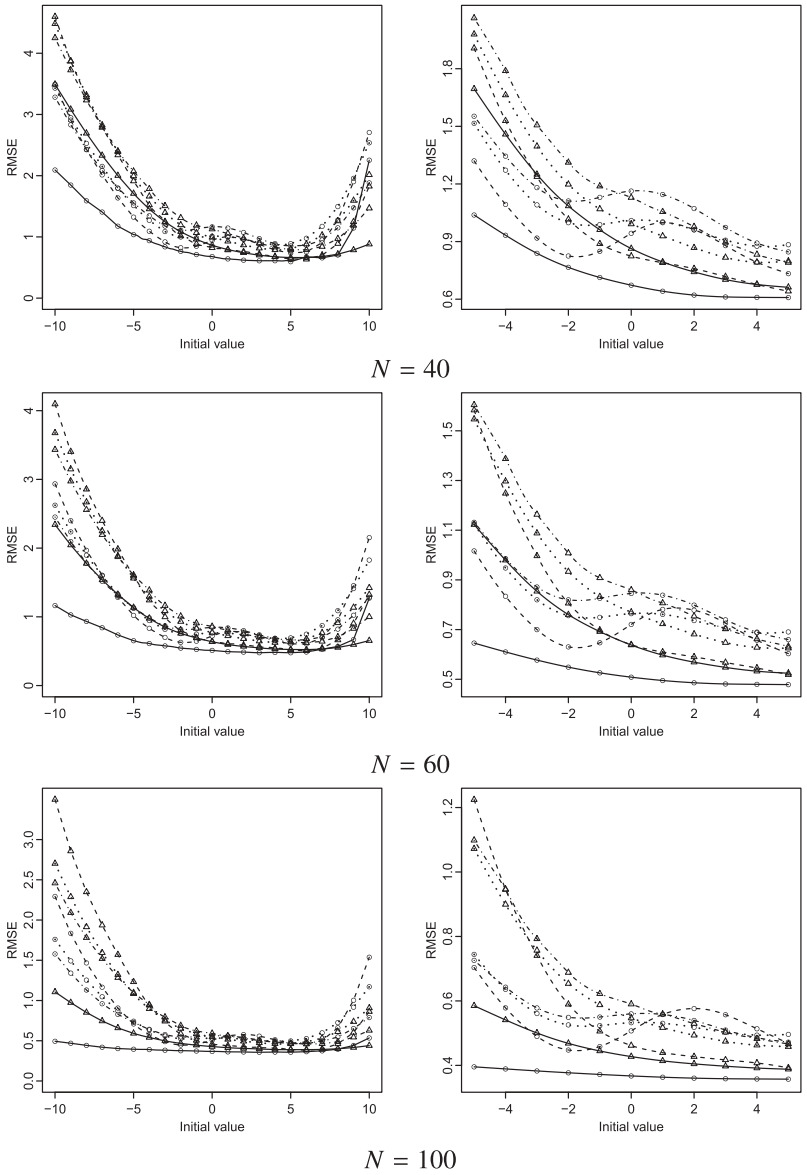

Figure 2. Estimated RMSE from the simulations for $\alpha_0 = 0$ and $\alpha_1 = 0.5$. Methods are displayed in different line styles: maximum likelihood-based methods (solid), Robbins–Monro (dashed), Polyak–Ruppert (dotted) and Bather (dash-dotted). Versions are marked by different symbols: original (triangles) and adapted (circles).

These roles are changing around the threshold, depending on α_1 and the sample size. While for $\alpha_1 = 0.2$ this point is above μ, it is below the threshold for $\alpha_1 = 0.5$.

For $\alpha_1 = 0.5$ there is a plainly visible local maximum around the threshold for the adjusted stochastic approximation methods. This bump is not as pronounced for $\alpha_1 = 0.2$. As Figure 3 illustrates, this behavior is apparently caused by a substantial positive bias. A possible explanation of this phenomenon may be related to the discussion of the Robbins–

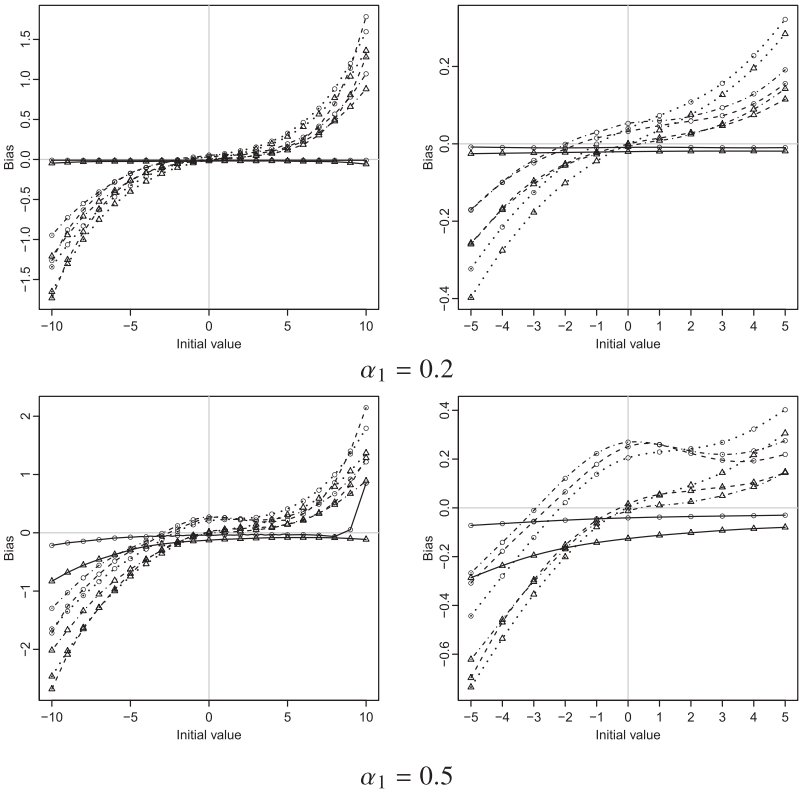

Figure 3. Estimated bias from the simulations for $N = 60$. Methods are displayed in different line styles: maximum likelihood based methods (solid), Robbins–Monro (dashed), Polyak–Ruppert (dotted), and Bather (dash-dotted). Versions are marked by different symbols: original (triangles) and adapted (circles).

Monro procedure following Eq. (17). There it was mentioned that after a nonresponse the dose is increased by a larger amount than it would be decreased in the case of a response. Hence, observations at doses larger than the threshold and nonresponses smaller than the threshold, but close to it, most likely lead to new estimates and doses that are larger than the threshold. This effect is also visible in the bias of the original version of the Robbins–Monro procedure. Explicit calculations of the bias up to $N = 10$ also exhibit this effect. Since the averaging methods are similar to the Robbins–Monro procedure, the behavior of their bias and RMSE should be similar, too. The maximum may be flattened by a smaller step size, but then the influence of bad starting values would increase.

7. Discussion

In a logistic regression model with false answers or guessing probability, a sequentially optimized design is compared with Wu's MLE design and some derivatives of the Robbins–Monro procedure adapted to the model with false answers, when the variance of the underlying logistic distribution is fixed. Although Wu's MLE design leads almost always to a greater root mean squared error than the optimized design, the comparison

between the Robbins–Monro procedure and the optimized design depends on the starting value and the sample size. For a small sample size the Robbins–Monro procedure should be preferred only if a good prior guess of the threshold can be expected, leading to a good starting value. If there is no or small prior information and the threshold should not be overestimated, the adjusted version is a reasonable choice. The optimized design is much more robust against a bad choice of the initial stimulus level and uniformly better than the Robbins–Monro procedure.

Even for smaller sample sizes, the optimized design can be recommended, since a slight loss in efficiency for a good starting value is compensated by a considerable gain in robustness.

References

Abbott, W. S. 1925. A method of computing the effectiveness of an insecticide. *Journal of Economic Entomology* 18 (2):265–67.

Bather, J. A. 1989. Stochastic approximation: A generalisation of the Robbins–Monro procedure. In *Proceedings of the Fourth Prague Symposium on Asymptotic Statistics*, ed. P. Mandl and M. Hušková, 13–27. Prague, Czech Republic: Charles University.

Chung, K. L. 1954. On a stochastic approximation method. *Annals of Mathematical Statistics* 25 (3):463–83.

Fechner, G. T. 1860. *Elemente der Psychophysik*. English translation: Howes, D. H., E. C. Boring, (eds.) and H. E. Adler (transl.). 1966. *Elements of psychophysics*. New York, NY: Holt Rinehart & Winston.

Lai, T. L., and H. Robbins. 1979. Adaptive design and stochastic approximation. *Annals of Statistics* 7 (6):1196–221.

Lord, F. M. 1968. An analysis of the verbal scholastic aptitude test using Birnbaum's three-parameter logistic model. *Educational and Psychological Measurement* 28 (4):989–1020.

McLeish, D. L., and D. H. Tosh. 1990. Sequential designs in bioassay. *Biometrics* 46 (1):103–16.

Morgan, B. J. T. 1992. *Analysis of quantal response data*. London, UK: Chapman & Hall.

Polyak, B. T. 1990. A new method of stochastic approximation type. *Automation and Remote Control* 51 (7):937–46.

Robbins, H., and S. Monro. 1951. A stochastic approximation method. *Annals of Mathematical Statistics* 22 (3):400–7.

Ruppert, D. 1988. Efficient estimators from a slowly convergent Robbins–Monro process. Technical report 781. School of Operations Research and Industrial Engineering, Cornell University, Ithaca, NY.

Sacks, J. 1958. Asymptotic distribution of stochastic approximation procedures. *Annals of Mathematical Statistics* 29 (2):373–405.

Schwabe, R. 1986. Strong representation of an adaptive stochastic approximation procedure. *Stochastic Processes and their Applications* 23 (1):115–30.

Schwabe, R., and H. Walk. 1996. On a stochastic approximation procedure based on averaging. *Metrika* 44 (2):165–80.

Silvapulle, M. J. 1981. On the existence of maximum likelihood estimators for the binomial response models. *Journal of the Royal Statistical Society B* 43 (3):310–13.

Venter, J. H. 1968. An extension of the Robbins–Monro procedure. *Annals of Mathematical Statistics* 38 (1):181–90.

Wetherill, G. B., and K. D. Glazebrook. 1986. *Sequential methods in statistics*, 3rd ed. London, UK: Chapman & Hall.

Wu, C. F. J. 1985. Efficient sequential designs with binary data. *Journal of the American Statistical Association* 80 (392):974–84.

Ying, Z., and C. F. J. Wu. 1997. An asymptotic theory of sequential designs based on maximum likelihood recursions. *Statistica Sinica* 7 (1):75–91.