



HAL
open science

Penalized optimal designs for dose-finding

Luc Pronzato

► **To cite this version:**

Luc Pronzato. Penalized optimal designs for dose-finding. *Journal of Statistical Planning and Inference*, 2010, 140 (1), pp.283-296. 10.1016/j.jspi.2009.07.012 . hal-00416014

HAL Id: hal-00416014

<https://hal.science/hal-00416014>

Submitted on 11 Sep 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Penalized optimal designs for dose-finding[☆]

Luc Pronzato

*Laboratoire I3S, CNRS/Université de Nice-Sophia Antipolis
Bât. Euclide, Les Algorithmes, 2000 route des lucioles, BP 121
06903 Sophia Antipolis cedex, France*

Abstract

Optimal design under a cost constraint is considered, with a scalar coefficient setting the compromise between information and cost. It is shown that for suitable cost functions, by increasing the value of the coefficient one can force the support points of an optimal design measure to concentrate around points of minimum cost. An example of adaptive design in a dose-finding problem with a bivariate binary model is presented, showing the effectiveness of the approach.

Key words:

optimal experimental design, penalized experimental design, compound optimal design, constrained optimal design, dose finding, bivariate binary model

2000 MSC: 62K05, 62L05, 62P10

1. Introduction and motivation

In the papers (Dragalin and Fedorov, 2006; Dragalin et al., 2008) the authors make a stimulating step towards a clear formalization of the necessary compromise between individual and collective ethics in experimental design for dose-finding studies. The idea is to use a penalty function that accounts for poor efficacy and for toxicity, and to construct a penalized D -optimal design (or equivalently a cost-constrained D -optimal design, the cost being defined through the penalty function, see (Fedorov and Hackl, 1997, Chap. 4)). Using a parametric model for the dose/efficacy-toxicity responses, a Gumbel or Cox model in (Dragalin and Fedorov, 2006) or a bivariate probit model in (Dragalin et al., 2008), the Fisher information matrix can be calculated and locally optimal designs can be obtained through (now standard) algorithmic methods. This shows a neat advantage over more standard approaches. Indeed, D -optimal design alone (and its extensions to Bayesian, minimax and adaptive variants) favors collective ethics in the sense that future patients will benefit from the information collected in the trials, but it neglects individual ethics

[☆]This work was partly accomplished while the author was invited at the Isaac Newton Institute for Mathematical Sciences, Cambridge, UK, in July 2008. The support of the Newton Institute and of CNRS are gratefully acknowledged.

Email address: pronzato@i3s.unice.fr (Luc Pronzato)

in the sense that the patients enrolled in the trials may receive doses with low efficacy or high toxicity. In contrast, dose-finding approaches based on up-and-down (Kpamegan and Flournoy, 2001) or bandit methods (Hardwick and Stout, 2001) focus on individual ethics by determining the best dose associated with some predefined criterion (e.g., the probability of efficacy and no toxicity) but may result in a poor learning from the trial and thus in poorly adapted dosage regimen for future individuals. By a suitable tuning of the penalty function, the approach used in (Dragalin and Fedorov, 2006; Dragalin et al., 2008) makes a clear compromise between the efficient treatment of individuals in the trial (by preventing the use of doses with low efficacy or high toxicity) and the precise estimation of the model parameters (accompanied with measures of statistical accuracy), to be used for making efficient decisions for future treatments. As such, it has a great potential in combining traditional phase I and phase II clinical trials into a single one, thereby accelerating the drug development process.

The aim of the present paper is to introduce flexibility in setting the compromise between the information gained (in terms of precision of parameter estimation) and the cost of the experiment (in terms of poor success for the patients enrolled in the trial). We show in particular that, for suitable penalty functions, by increasing the weight set on the penalty one guarantees that all doses in the experiment have a small cost (and concentrate around the optimal dose when this one is unique). This permits the avoidance of extreme doses generally suggested by optimal design for parameter estimation.

Penalized optimal design is introduced in Sect. 2, where the two equivalent formulations of the constrained design problem considered in (Dragalin and Fedorov, 2006; Dragalin et al., 2008) are presented. The maximization of information under a cost constraint is considered in Sect. 3. An example with a bivariate binary model is presented in Sect. 4, illustrating the effectiveness of the approach.

2. Penalized D -optimal design

The setting is rather standard for optimal design of experiments. The admissible domain for the experimental variables x (design points) is a compact subset of \mathbb{R}^d denoted by \mathcal{X} , and we denote by $\theta \in \mathbb{R}^p$ the (p -dimensional) vector of parameter of interest in a parametric model generating the log-likelihood $l(Y, x; \theta)$ for an observation Y at the design point x . We suppose that $\theta \in \Theta$, a compact subset of \mathbb{R}^p . For N independent observations $\mathbf{Y} = (Y_1, \dots, Y_N)$ at non-random design points $\mathbf{X} = (x_1, \dots, x_N)$, the log-likelihood at θ is $l(\mathbf{Y}, \mathbf{X}; \theta) = \sum_{i=1}^N l(Y_i, x_i; \theta)$. Let $\mathbf{M}(X, \theta)$ denote the corresponding Fisher information matrix, $\mathbf{M}(X, \theta) = -\mathbb{E}_\theta \{ \partial^2 l(\mathbf{Y}, \mathbf{X}; \theta) / (\partial \theta \partial \theta^\top) \}$. Because of the independence of the Y_i 's, $\mathbf{M}(X, \theta)$ can be written as $\mathbf{M}(X, \theta) = \sum_{i=1}^N \mu(x_i, \theta)$. When $N(x_i)$ denotes the number of observations made at $x = x_i$, we get the following normalized information matrix per observation $\mathbf{M}(\xi, \theta) = (1/N) \mathbf{M}(X, \theta) = \sum_{i=1}^K [N(x_i)/N] \mu(x_i, \theta)$, where K is the number of distinct design points and ξ denotes the design measure (a probability measure on \mathcal{X}) that puts mass $N(x_i)/N$ at x_i . Following the usual approximate design approach, we shall relax the constraints on design measures and consider ξ as any element of Ξ , the set of probability measures on \mathcal{X} , so that $\mathbf{M}(\xi, \theta) = \int_{\mathcal{X}} \mu(x, \theta) \xi(dx)$.

In a regression model with independent and homoscedastic observations satisfying $\mathbb{E}_\theta(Y|x, \theta) = \eta(x, \theta)$, with $\eta(x, \theta)$ differentiable with respect to θ for any x , we have

$$\mu(x, \theta) = \mathcal{I} \frac{\partial \eta(x, \theta)}{\partial \theta} \frac{\partial \eta(x, \theta)}{\partial \theta^\top} \quad (1)$$

with $\mathcal{I} = \int [\varphi'(t)/\varphi(t)]^2 \varphi(t) dt$ the Fisher information for location, where $\varphi(\cdot)$ is the probability density function of the observation errors and $\varphi'(\cdot)$ its derivative.

In a dose-response problem with single response $Y \in \{0, 1\}$ (efficacy or toxicity response at the dose x for instance) and $\text{Prob}\{Y = 1|x, \theta\} = \pi(x, \theta)$, we have $l(Y, x; \theta) = Y \log[\pi(x, \theta)] + (1 - Y) \log[1 - \pi(x, \theta)]$. Therefore, assuming that $\pi(x, \theta)$ is differentiable with respect to θ for any x ,

$$\mu(x, \theta) = \frac{\partial \pi(x, \theta)}{\partial \theta} \frac{\partial \pi(x, \theta)}{\partial \theta^\top} \frac{1}{\pi(x, \theta)[1 - \pi(x, \theta)]}.$$

Bivariate extensions, where both efficacy and toxicity responses are observed at a dose x , are considered in (Dragalin and Fedorov, 2006) with Gumbel and Cox models and in (Dragalin et al., 2008) with a bivariate probit model. See also the example in Sect. 4. Besides a few additional technical difficulties, the main difference with the single response case is the fact that $\mu(x, \theta)$ may have rank larger than one, so that less than p support points in ξ may suffice to estimate θ consistently. The same situation occurs for regression models when $\dim(\eta) > 1$ so that (1) may have rank larger than one. We shall always assume that $\mu(x, \theta)$ is bounded on \mathcal{X} .

In its now traditional form, local optimal design consists in determining a measure ξ^* that maximizes a concave function $\Psi(\cdot)$ of the Fisher information matrix $\mathbf{M}(\xi, \theta)$ for a given value of θ . We assume that $\Psi(\cdot)$ is monotone for the Loewner ordering (therefore, $\Psi(a\mathbf{M})$ is a non-decreasing function of $a \in \mathbb{R}^+$ for any non-negative definite matrix \mathbf{M}) and shall pay special attention to local D -optimal design, for which $\Psi(\mathbf{M}) = \log \det \mathbf{M}$. We shall denote $\Lambda_{\min}(\mathbf{M})$ and $\Lambda_{\max}(\mathbf{M})$ the minimum and maximum eigenvalues of the matrix \mathbf{M} .

In many circumstances, besides the optimality criterion $\psi(\xi) = \Psi[\mathbf{M}(\xi, \theta)]$, it is desirable to introduce a constraint of the form $\Phi(\xi, \theta) \leq C$ for the design measure, accounting for the cost of the experiment. In dose-finding problems, the introduction of such a constraint allows one to take individual ethical concerns into account. For instance, when both the efficacy and toxicity responses are observed, one can relate $\Phi(\xi, \theta)$ to the probability of success (efficacy and no toxicity) for a given dose, as done in (Dragalin and Fedorov, 2006; Dragalin et al., 2008). See also Sect. 4. We suppose that the cost function $\Phi(\xi, \theta)$ is linear in ξ , that is

$$\Phi(\xi, \theta) = \int_{\mathcal{X}} \phi(x, \theta) \xi(dx).$$

It simply expresses that the total cost of an experiment with observations at $\mathbf{X} = (x_1, \dots, x_N)$ is the sum of the costs induced by single observations at the x_i 's. The

extension to nonlinear constraints is considered, e.g., in (Cook and Fedorov, 1995) and (Fedorov and Hackl, 1997, Chap. 4). Also, we shall restrict our attention to the case where a single (scalar) constraint is present, some of the issues caused by the presence of several constraints are addressed in the same references. We shall always assume that $\phi(x, \theta)$ is bounded on \mathcal{X} .

The fact that a single cost function is present permits to consider the problem of maximizing information per cost-unit, which can be formulated as a design problem without constraint. (We shall see in Sect. 3 that this is different, however, from maximizing information per observation under a cost constraint of the form $\Phi(\xi, \theta) \leq C$.) Suppose that $\phi(x, \theta) > 0$ for all $x \in \mathcal{X}$. The approach used in (Dragalin and Fedorov, 2006; Dragalin et al., 2008) formulates the problem as follows.

Problem $P_1(\theta)$: maximize the total information for N observations, that is, maximize $\Psi[N\mathbf{M}(\xi, \theta)]$ with respect to N and $\xi \in \Xi$ satisfying the total cost constraint $N\Phi(\xi, \theta) = N \int_{\mathcal{X}} \phi(x, \theta) \xi(dx) \leq C$.

By considering this problem more closely, one can show that it is equivalent to another design problem *without constraint*. Indeed, for any ξ , the optimal value of N is $N^*(\xi) = C/\Phi(\xi, \theta)$, so that an optimal measure $\xi^* \in \Xi$ for $P_1(\theta)$ maximizes $\Psi[C\mathbf{M}(\xi, \theta)/\Phi(\xi, \theta)]$. Denote now $\nu = N\xi$, which is a measure on \mathcal{X} not normalized to 1; we have $\int_{\mathcal{X}} \nu(dx) = N$ which becomes a free variable. $P_1(\theta)$ is then equivalent to the maximization of $\Psi[\mathbf{M}(\nu, \theta)]$ under the constraint $\Phi(\nu, \theta) \leq C$. The constraint is saturated at the optimum, i.e. $\Phi(\nu^*, \theta) = C$, which we can thus set as an active constraint. Imposing $\Phi(\nu, \theta) = C$ and defining $\xi'(dx) = \nu(dx)\phi(x, \theta)/C$ we obtain $\int_{\mathcal{X}} \xi'(dx) = 1$ and $\mathbf{M}(\nu, \theta) = \int_{\mathcal{X}} \mu(x, \theta) \nu(dx) = \int_{\mathcal{X}} C [\mu(x, \theta)/\phi(x, \theta)] \xi'(dx) = \mathbf{M}'(\xi', \theta)$. The constraint design problem $P_1(\theta)$ is thus equivalent to a standard unconstrained one, with $\mu(x, \theta)$ simply replaced by $C\mu(x, \theta)/\phi(x, \theta)$. Call $P_2(\theta)$ this problem.

Problem $P_2(\theta)$: maximize $\Psi[\mathbf{M}'(\xi, \theta)]$ with respect to $\xi \in \Xi$.

The equivalence between $P_1(\theta)$ and $P_2(\theta)$ is further evidenced by considering the necessary and sufficient conditions for optimality expressed by the Equivalence Theorem, see Kiefer and Wolfowitz (1960) for D -optimality and, e.g., Pukelsheim (1993) for a general formulation. For $P_1(\theta)$ with $\psi(\xi) = \log \det[\mathbf{M}(\xi, \theta)/\Phi(\xi, \theta)]$, the measure ξ^* is optimal if and only if

$$\forall x \in \mathcal{X}, \text{ trace}[\mu(x, \theta)\mathbf{M}^{-1}(\xi^*, \theta)] \leq p \frac{\phi(x, \theta)}{\Phi(\xi^*, \theta)} \quad (2)$$

(note that the condition does not depend on the normalization constant $\int_{\mathcal{X}} \xi^*(dx)$). For $P_2(\theta)$ with $\psi(\xi') = \log \det \mathbf{M}'(\xi', \theta)$, ξ'^* is optimal in Ξ if and only if

$$\forall x \in \mathcal{X}, C \text{ trace} \left[\frac{\mu(x, \theta)}{\phi(x, \theta)} \mathbf{M}'^{-1}(\xi'^*, \theta) \right] \leq p \quad (3)$$

(note that \mathbf{M}' is proportional to C which thus cancels out). The two conditions are equivalent: (3) can be written as (2) by setting $\xi^*(dx) = C \xi'^*(dx)/\phi(x, \theta)$.

One should note that the value of C plays no role in the definition of optimal designs for $P_1(\theta)$ and $P_2(\theta)$. For the dose-response problem considered in (Dragalin and Fedorov, 2006; Dragalin et al., 2008) this has the important consequence that the prohibition of excessively low (with poor efficacy) or high (with high toxicity) doses can only be obtained by an ad-hoc modification of the penalty function $\phi(x, \theta)$. Indeed, this is the only way to modify the optimal design and hopefully to change its support. This can be contrasted with the solution of the constrained design problem that we present in the next section and then consider in the rest of the paper.

3. Maximizing information per observation under cost constraint

3.1. Problem statement

A direct formulation of the optimal design problem under constraint is as follows.

Problem $P_3(\theta)$: maximize $\Psi[\mathbf{M}(\xi, \theta)]$ with respect to $\xi \in \Xi$ under the constraint $\Phi(\xi, \theta) \leq C$.

We say that a design measure $\xi \in \Xi$ is θ -admissible if $\Phi(\xi, \theta) \leq C$. We suppose that a strictly θ -admissible measure exists in Ξ : $\Phi(\xi, \theta) < C$ for some $\xi \in \Xi$. Then, a necessary and sufficient condition for the optimality of a θ -admissible $\xi^* \in \Xi$ is the existence of $\lambda^* \geq 0$ such that $\lambda^* [C - \Phi(\xi^*, \theta)] = 0$ with $\xi^* = \xi^*(\lambda^*)$ maximizing the design criterion $\mathcal{L}_\theta(\xi, \lambda) = \Psi[\mathbf{M}(\xi, \theta)] + \lambda^* [C - \Phi(\xi, \theta)]$ (the Lagrangian of the problem) with respect to $\xi \in \Xi$. Moreover, the Lagrange coefficient λ^* minimizes $\mathcal{L}_\theta[\xi^*(\lambda), \lambda]$ with respect to $\lambda \in \mathbb{R}^+$. When $\Psi(\cdot) = \log \det(\cdot)$, the necessary and sufficient condition for the optimality of a θ -admissible $\xi^* \in \Xi$ becomes

$$\exists \lambda^* \geq 0 \text{ such that } \begin{cases} \lambda^* [C - \Phi(\xi^*, \theta)] = 0 \\ \text{and} \\ \forall x \in \mathcal{X}, \text{ trace}[\mu(x, \theta)\mathbf{M}^{-1}(\xi^*, \theta)] \leq p + \lambda^* [\phi(x, \theta) - \Phi(\xi^*, \theta)]. \end{cases} \quad (4)$$

In practice, ξ^* can be determined by maximizing

$$H_\theta(\xi, \lambda) = \Psi[\mathbf{M}(\xi, \theta)] - \lambda \Phi(\xi, \theta) \quad (5)$$

for an increasing sequence (λ_i) of Lagrange coefficients λ , starting at $\lambda_0 = 0$ and stopping at the first λ_i such that the associated optimal design ξ^* satisfies $\Phi(\xi^*, \theta) \leq C$, see, e.g., Mikulecká (1983) (for C large enough, the unconstrained optimal design for $\Psi(\cdot)$ is optimal for the constrained problem). The parameter λ can thus be used to set the tradeoff between the maximization of $\Psi[\mathbf{M}(\xi, \theta)]$ (gaining information) and minimization of $\Phi(\xi, \theta)$ (reducing cost). It can be considered as a penalty coefficient, and $\phi(\cdot, \theta)$ as a penalty function that penalizes a design for the cost it induces. An optimal design for $H_\theta(\xi, \lambda)$ is thus called a *penalized optimal design*.

Remark 1.

One may notice that maximizing $H_\theta(\xi, \lambda)$ for $\lambda \geq 0$ is equivalent to maximizing $(1 -$

$\gamma) \Psi[\mathbf{M}(\xi, \theta)] + \gamma [-\Phi(\xi, \theta)]$ with $\gamma = \lambda/(1 + \lambda) \in [0, 1)$ (one may refer to Cook and Wong (1994) for the equivalence between constrained and compound optimal designs).

Similarly to the case of unconstrained optimal design with a strictly concave criterion, the optimal matrix $\mathbf{M}(\xi^*, \theta)$ is unique when the function $\Psi(\cdot)$ is strictly concave on the set of positive definite matrices (but the optimal design measure ξ^* is not necessarily unique).

3.2. Properties

Let $\xi^*(\lambda)$ denote an optimal design for $H_\theta(\xi, \lambda)$ given by (5). One can easily check that both $\Psi\{\mathbf{M}[\xi^*(\lambda), \theta]\}$ and $\Phi[\xi^*(\lambda), \theta]$ are non-increasing functions of λ , see Cook and Wong (1994) for examples. Two questions naturally arise concerning the tradeoff between maximization of $\Psi[\mathbf{M}(\xi, \theta)]$ and minimization of $\Phi(\xi, \theta)$.

- (i) How fast does the cost $\Phi(\xi, \theta)$ decrease when λ increases?
- (ii) How big is the loss of information (decrease of $\Psi[\mathbf{M}(\xi, \theta)]$) when $\Phi(\xi, \theta)$ decreases?

Suppose that $\mu(x, \theta)$ and $\phi(x, \theta)$ are continuous in $x \in \mathcal{X}$ (\mathcal{X} is a compact subset of \mathbb{R}^d). For any $\xi \in \Xi$, we define

$$\Delta_\theta(\xi) = \Phi(\xi, \theta) - \phi_\theta^*,$$

where

$$\phi_\theta^* = \phi(x^*, \theta) \quad \text{with} \quad x^* = x^*(\theta) = \arg \min_{x \in \mathcal{X}} \phi(x, \theta). \quad (6)$$

Note that we do not assume here that x^* is unique.

We focuss our attention on D -optimality and denote by ξ_D^* a D -optimal design that maximizes $\log \det \mathbf{M}(\xi, \theta)$ with respect to $\xi \in \Xi$. We assume that $\Delta_\theta(\xi_D^*) > 0$ (otherwise ξ_D^* maximizes $\log \det \mathbf{M}(\xi, \theta) - \lambda \Phi(\xi, \theta)$ for any $\lambda \geq 0$) and that $\log \det \mathbf{M}(\xi_D^*, \theta) > 0$ (otherwise $\mathbf{M}(\xi, \theta)$ is singular for any $\xi \in \Xi$). We then have the following results concerning the two questions above.

Proposition 1. Let $\xi^* = \xi^*(\lambda)$ be an optimal design that maximises $H_\theta(\xi, \lambda)$ given by (5) with respect to $\xi \in \Xi$, with $\Psi(\mathbf{M}) = \log \det \mathbf{M}$. It satisfies

- (i) $\Delta_\theta(\xi^*) \leq p/\lambda$;
- (ii) for any ξ such that $\Delta_\theta(\xi) > 0$, any $a > 0$ and any $\lambda \geq a/\Delta_\theta(\xi)$,

$$\log \det \mathbf{M}(\xi^*, \theta) \geq \log \det \mathbf{M}(\xi, \theta) + p \log\{a/[\lambda \Delta_\theta(\xi)]\} - a, \quad (7)$$

moreover, $\Lambda_{\min}[\mathbf{M}(\xi^*, \theta)] > \delta/(p + \lambda[\bar{\phi}_\theta - \phi_\theta^*])$ for some positive constant δ , where $\bar{\phi}_\theta = \max_{x \in \mathcal{X}} \phi(x, \theta)$.

The proof is given in Appendix.

Property (i) shows the guaranteed cost-reduction obtained when λ is increased and (ii) shows that $\Lambda_{\min}[\mathbf{M}(\xi^*, \theta)]$ decreases not faster than $1/\lambda$. Notice that taking $\xi = \xi_D^*$ in (7)

ensures $\det \mathbf{M}(\xi^*, \theta) \geq \det[\mathbf{M}(\xi_D^*, \theta) \exp(-a/p)]$ for $\lambda = a/\Delta_\theta(\xi_D^*)$ and $\log \det \mathbf{M}(\xi^*, \theta) \geq \log \det \mathbf{M}(\xi_D^*, \theta) + p \log[\Delta_\theta(\xi^*)/\Delta_\theta(\xi_D^*)]$ for any $\lambda \geq 0$ (take $a = \lambda\Delta_\theta(\xi^*)$ in (16)).

By increasing λ in (5) one reduces the total cost of the experiment, but this does not necessarily imply that the cost $\phi(x, \theta)$ will be small for every support point of the experiment. However, the property below shows that for suitable penalty functions the additional flexibility offered by the tuning parameter λ in (5) compared to problems $P_1(\theta)$ and $P_2(\theta)$ can be used to obtain a small cost value $\phi(\hat{x}_i, \theta)$ at every support point \hat{x}_i of an optimal design $\xi^*(\lambda)$.

Proposition 2. Let $\xi^* = \xi^*(\lambda)$ be an optimal design that maximises $H_\theta(\xi, \lambda)$ given by (5) with respect to $\xi \in \Xi$, with $\Psi(\mathbf{M}) = \log \det \mathbf{M}$. Take any design measure $\xi_\lambda \in \Xi$ such that $\Delta_\theta(\xi_\lambda) \geq p/\lambda$. The cost at any support point \hat{x} of ξ^* satisfies

$$\phi(\hat{x}, \theta) \leq \phi_\theta^* + 2\Delta_\theta(\xi_\lambda) \text{trace}[\mu(\hat{x}, \theta)\mathbf{M}^{-1}(\xi_\lambda, \theta)]. \quad (8)$$

The proof is given in Appendix.

Remark 2. When x^* is unique, Proposition 2 implies that if there exist designs $\xi_\lambda \in \Xi$ such that

$$\Delta_\theta(\xi_\lambda) \geq p/\lambda \text{ and } \forall \epsilon > 0, \limsup_{\lambda \rightarrow \infty} \sup_{\|x-x^*\| > \epsilon} \frac{2\Delta_\theta(\xi_\lambda) \text{trace}[\mu(x, \theta)\mathbf{M}^{-1}(\xi_\lambda, \theta)]}{\phi(x, \theta) - \phi_\theta^*} < 1, \quad (9)$$

then the supporting points of ξ^* converge to x^* as $\lambda \rightarrow \infty$. The choice of suitable designs ξ_λ is central for testing if (9) is satisfied, designs with support points in the neighborhood of x^* being good candidates, see Example 1 and Sect. 4. Notice that, when $\text{rank}[\mu(x, \theta)] < p$, for (9) to be satisfied the ξ_λ 's must necessarily have support points that approach x^* as $\lambda \rightarrow \infty$. Indeed, suppose that it is not the case. It means that there exists $\gamma > 0$ such that, for all λ larger than some λ_0 , the support points $x_\lambda^{(i)}$ of ξ_λ satisfy $\|x_\lambda^{(i)} - x^*\| > \gamma$. Replace \mathcal{X} by $\mathcal{X}' = \mathcal{X} \setminus \mathcal{B}(x^*, \gamma) \cup \{x^*\}$, that is, remove the ball $\mathcal{B}(x^*, \gamma) = \{x : \|x - x^*\| \leq \gamma\}$ from \mathcal{X} but keep x^* . Then, ξ_λ is a design measure on \mathcal{X}' for $\lambda > \lambda_0$ and (9) would indicate that the optimal design ξ^* on \mathcal{X}' is the delta measure δ_{x^*} , which is impossible since the optimal information matrix must have full rank. The same is true if the designs ξ_λ have one support point at x^* and the others outside the ball $\mathcal{B}(x^*, \gamma)$ for λ larger than some λ_0 . Finally, note that the support points of $\xi^*(\lambda')$ for $\lambda' > \lambda$ must also satisfy (8) for the same ξ_λ (since $\Delta_\theta(\xi_\lambda) > p/\lambda'$).

For dose-response problems, the property (9) has the important consequence that excessively high or low doses can be prohibited by choosing C small enough in $P_3(\theta)$ or, equivalently, λ large enough in (5). Its effectiveness very much depends on the choice of the penalty function $\phi(\cdot, \theta)$, and in particular on its local behavior around x^* . Contrary to what intuition might suggest, it requires $\phi(\cdot, \theta)$ to be sufficiently flat around x^* . Indeed, in that case a design ξ supported in the neighborhood of x^* can at the same time have a small cost $\Phi(\xi, \theta)$ and be dispersed enough to carry significant information through $\log \det \mathbf{M}(\xi, \theta)$. This is illustrated by the example below.

Example 1.

The example, simple enough to make the optimal designs calculable explicitly, illustrates the influence of λ on the support points of an optimal design for (5) (or, equivalently, the influence of C in problem $P_3(\theta)$). We consider a linear regression model $\eta(x, \theta) = \mathbf{f}^\top(x)\theta$ with homoscedastic errors and $\mathbf{f}(x) = (1 \ x \ x^2)^\top$, so that

$$\mu(x) = \begin{pmatrix} 1 \\ x \\ x^2 \end{pmatrix} (1 \ x \ x^2) = \begin{pmatrix} 1 & x & x^2 \\ x & x^2 & x^3 \\ x^2 & x^3 & x^4 \end{pmatrix}.$$

Note that $\mu(\cdot)$ does not depend on θ so that we write μ instead of μ_θ ; the same will be true for the penalty functions considered in the example, and we shall write $\phi(x)$ instead of $\phi(x, \theta)$.

We take $\mathcal{X} = [-1, 1]$, $\Psi(\cdot) = \log \det(\cdot)$ and consider several penalty functions $\phi(x)$, symmetric with respect to $x = 0$ where they reach their minimum value: $\phi(x) = 1 + x^{2q}$, $q = 1, 2, 3, 4$; $\phi(x) = 1 + x^2 + x^4$; $\phi(x) = 1/(1 - x^2)$; the minimum cost ϕ^* equals one for each of them. For any λ , the optimal designs ξ^* for (5) are symmetric with respect to 0 and take the form

$$\xi^* = \left\{ \begin{array}{ccc} -z & 0 & z \\ \frac{1-\alpha}{2} & \alpha & \frac{1-\alpha}{2} \end{array} \right\}, \quad (10)$$

where the first row gives the support points and the second their respective weights. This gives $\det \mathbf{M}(\xi^*) = \alpha(1 - \alpha)^2 z^6$. The D -optimal design ξ_D^* corresponds to $z = 1$ and $\alpha = 1/3$.

For $\phi(x) = 1 + x^{2q}$, q integer, the optimal designs for problems P_1 and P_2 correspond to $\alpha = \min\{q/[3(q-1)], 1/2\}$ and $z = \min\{[3/(2q-3)]^{1/2q}, 1\}$ (note that $z < 1$ only for $q \geq 4$). The penalty functions $\phi(x) = 1 + x^2 + x^4$ and $\phi(x) = 1/(1 - x^2)$ respectively give the optimal designs defined by $(\alpha = 3/5, z = 1)$ and $(\alpha = 5/9, z = \sqrt{3/5})$.

The optimal designs ξ^* for P_3 obtained for the different $\phi(\cdot)$ mentioned above are given in Table 1, together with the optimal value $\lambda^*(C)$ of the Lagrange coefficient associated with C . When there is no solution, it means that $\lambda^*(C) = \infty$. When there is one, then $\Phi(\xi^*) = C$.

In order to check if the support points of ξ^* concentrate around $x^* = 0$ when λ increases (without computing ξ^*), we use (8) with the test designs

$$\xi_\lambda = \xi_\lambda(\gamma) = \left\{ \begin{array}{ccc} -\gamma & 0 & \gamma \\ 1/3 & 1/3 & 1/3 \end{array} \right\}.$$

For $\phi(x) = 1 + x^{2q}$ we get $\Delta(\xi_\lambda) = 2\gamma^{2q}/3$ and the condition (8) then gives $\hat{x}^{2q} \leq [4\gamma^{2q}/3] \text{trace}[\mu(\hat{x})\mathbf{M}^{-1}(\xi_\lambda)]$. Noticing that $\text{trace}[\mu(\gamma t)\mathbf{M}^{-1}(\xi_\lambda)] = P(t) = 3(1 - 3/2t^2 + 3/2t^4)$ independently of γ (a property of D -optimal design for polynomial regression), we obtain that a support point \hat{x} of ξ^* must satisfy $t^{2q} \leq 4P(t)/3$, with $t = |\hat{x}|/\gamma$. For $q = 3$ we obtain $t \leq [1 + (1 + \beta^{1/3})^2]^{1/2}/\beta^{1/6}$ with $\beta = 4 + 2\sqrt{2}$, that is $t \lesssim 2.2252$. For $q = 4$, we get $t \leq \sqrt{2}$. Since we need to have $\Delta(\xi_\lambda) \geq p/\lambda = 3/\lambda$, we take $\gamma = [9/(2\lambda)]^{1/(2q)}$ (which

corresponds to $\alpha = 1$ and $\tilde{\xi}_\lambda = \xi_\lambda$ in the proof of Proposition 1-(ii)). It gives $|\hat{x}| \leq \hat{x}_{\max}$ with $\hat{x}_{\max} \simeq 2.860 \lambda^{-1/6}$ for $q = 3$ and $\hat{x}_{\max} = \sqrt{2}[9/(2\lambda)]^{1/8} \simeq 1.707 \lambda^{-1/8}$ for $q = 4$. This is consistent with the rate of decrease of z in Table 1 when C tends to 1 (and $\lambda^*(C)$ tends to infinity). When $q = 1, 2$ all t are admissible and we cannot obtain a bound on $|\hat{x}|$; the same situation occurs for the penalty function $\phi(x) = 1 + x^2 + x^4$.

The case $\phi(x) = 1/(1 - x^2)$ illustrates that it is the local behavior of $\phi(x)$ around the minimum x^* that influences the support of ξ^* when λ tends to infinity. Indeed, $1/(1 - x^2)$ tends to infinity for x tending to ± 1 but equals $1 + x^2 + x^4 + \mathcal{O}(x^6)$ around $x^* = 0$. Condition (8) then becomes $\phi(\gamma t) - \phi(0) = \gamma^2 t^2 / (1 - \gamma^2 t^2) \leq 2 \Delta(\xi_\lambda) P(t) = (4/3) \gamma^2 P(t) / (1 - \gamma^2)$. The bound obtained for t now depends on γ ; the best bound (minimum) for $|\hat{x}|$ is $\hat{x}_{\max} \simeq 0.9649$, obtained at $\gamma \simeq 0.7385$, and $\Delta(\xi_\lambda) \geq p/\lambda$ imposes $\lambda \gtrsim 3.7516$. Therefore, we only learn from (8) that the support of ξ^* is included in $[-0.9649, 0.9649]$ for λ large enough. This is consistent with the behavior of the support points $-z, z$ of ξ^* as λ tends to infinity, which do not converge to zero ($\lim_{\lambda \rightarrow \infty} z = \lim_{C \rightarrow 1^+} z = 1/\sqrt{3}$, see Table 1). \square

3.3. Nonlinear cost-constrained design

In a nonlinear situation, like in the example presented in Sect. 4, where $\mathbf{M}(\xi, \theta)$ and $\Phi(\xi, \theta)$ depend on θ , robustness with respect to misspecifications of θ can be achieved by considering average-optimal design. Problem $P_3(\theta)$ is then transformed into: maximize $\mathbb{E}_\theta\{\Psi[\mathbf{M}(\xi, \theta)]\}$ with respect to $\xi \in \Xi$ under the constraint $\mathbb{E}_\theta\{\Phi(\xi, \theta)\} \leq C$, where the expectation \mathbb{E}_θ is calculated for some prior probability measure ν for θ . For D -optimality, the optimality condition (4) becomes

$$\exists \lambda^* \geq 0 \text{ such that } \begin{cases} \lambda^* [C - \mathbb{E}_\theta\{\Phi(\xi^*, \theta)\}] = 0 \\ \text{and} \\ \forall x \in \mathcal{X}, \mathbb{E}_\theta\{\text{trace}[\mu(x, \theta) \mathbf{M}^{-1}(\xi^*, \theta)]\} \leq \\ p + \lambda^* \mathbb{E}_\theta\{\phi(x, \theta) - \Phi(\xi^*, \theta)\}. \end{cases}$$

Apart from additional numerical cost (which remains reasonable when ν is a discrete measure with a limited number of support points), the introduction of a prior probability for θ does not raise any special difficulty. This is used by Haines et al. (2003), with $\phi(x, \theta) = \mathbb{I}_{[Q_R(\theta), \infty)}(x)$, where $\mathbb{I}_{\mathcal{A}}(x)$ is the indicator function of the set \mathcal{A} (1 if $x \in \mathcal{A}$, 0 otherwise) and $Q_R(\theta)$ is a quantile of the probability of toxicity, parameterized by θ , defining the maximum acceptable probability of toxicity (note that $\mathbb{E}_\theta\{\phi(x, \theta)\} = \nu\{Q_R(\theta) \leq x\}$, the prior probability that x exceeds Q_R).

Another common approach to overcome the issue of dependence of the optimum design in θ consists in designing the experiment sequentially. In adaptive D -optimal design for instance, the design point after N observations is taken as

$$x_{N+1} = \arg \max_{x \in \mathcal{X}} \text{trace}[\mu(x, \hat{\theta}^N) \mathbf{M}^{-1}(\xi_N, \hat{\theta}^N)], \quad (11)$$

where $\hat{\theta}^N$ is the current estimated value of θ . By alternating between estimation based on previous observations and determination of the next design point where to observe,

$\phi(x)$	C	$\lambda^*(C)$	α	z
$1 + x^2$	$5/3 \leq C$	0	1/3	1
	$1 < C \leq 5/3$	$\frac{5-3C}{(C-1)(2-C)}$	$2 - C$	1
	$C \leq 1$	∞	— no solution —	
$1 + x^4$	$5/3 \leq C$	0	1/3	1
	$4/3 \leq C \leq 5/3$	$\frac{5-3C}{(C-1)(2-C)}$	$2 - C$	1
	$1 < C \leq 4/3$	$3/[2(C-1)]$	2/3	$[3(C-1)]^{1/4}$
	$C \leq 1$	∞	— no solution —	
$1 + x^6$	$5/3 \leq C$	0	1/3	1
	$3/2 \leq C \leq 5/3$	$\frac{5-3C}{(C-1)(2-C)}$	$2 - C$	1
	$1 < C \leq 3/2$	$1/(C-1)$	1/2	$[2(C-1)]^{1/6}$
	$C \leq 1$	∞	— no solution —	
$1 + x^8$	$5/3 \leq C$	0	1/3	1
	$14/9 \leq C \leq 5/3$	$\frac{5-3C}{(C-1)(2-C)}$	$2 - C$	1
	$1 < C \leq 14/9$	$3/[4(C-1)]$	4/9	$[9(C-1)/5]^{1/8}$
	$C \leq 1$	∞	— no solution —	
$1 + x^2 + x^4$	$7/3 \leq C$	0	1/3	1
	$1 < C \leq 7/3$	$\frac{7-3C}{(C-1)(3-C)}$	$(3-C)/2$	1
	$C \leq 1$	∞	— no solution —	
$1/(1-x^2)$	$1 < C$	$\frac{2}{C(C-1)}$	$\frac{C}{3C-2}$	$\frac{(3C-2)^{1/2}}{(3C)^{1/2}}$
	$C \leq 1$	∞	— no solution —	

Table 1: Optimal designs ξ^* for problem P_3 in Example 1, see (10). The optimal support points $\pm z$ concentrate around 0 when $\lambda^*(C) \rightarrow \infty$ for the penalty functions $1 + x^4$, $1 + x^6$ and $1 + x^8$.

one forces the empirical design measure to progressively adapt to the correct (true) value of the model parameters. Adaptive design is considered in (Dragalin and Fedorov, 2006; Dragalin et al., 2008), but the convergence of the procedure (strong consistency of the parameter estimator and convergence of the empirical design measure to the optimal non-sequential design for the true value of the model parameters) is left as an open issue.

The difficulty of proving the consistency of the estimator when design variables are sequentially determined is usually overcome by considering an initial experiment (non adaptive) that grows in size when the total number of observations increases, see, e.g., Chaudhuri and Mykland (1993). Although this number is often severely limited in practise, especially for clinical trials, we think that it is reassuring to know that, *for a given initial experiment*, adaptive design guarantees suitable asymptotic properties under reasonable conditions. Using simple arguments, one can show that this is indeed the case when *the design space is finite*, which forms a rather natural assumption in the context of clinical trials. The case of adaptive D -optimal design is considered in (Pronzato, 2009b) (notice that it also covers the situation considered by Dragalin and Fedorov (2006); Dragalin et al. (2008), which can be formulated as a standard D -optimal design problem).

In the case of adaptive penalized D -optimal design, the design point after N observations is taken as

$$x_{N+1} = \arg \max_{x \in \mathcal{X}} \left\{ \text{trace}[\mu(x, \hat{\theta}^N) \mathbf{M}^{-1}(\xi_N, \hat{\theta}^N)] - \lambda_N \phi(x, \hat{\theta}^N) \right\}. \quad (12)$$

Following an approach similar to that in (Pronzato, 2009b), one can show that when \mathcal{X} is finite, λ_N is the optimal Lagrange coefficient for problem $P_3(\hat{\theta}^N)$, and under rather standard regularity assumptions, this procedure is asymptotically “optimal” in the sense that the estimated value of the parameters (by least-squares in a nonlinear regression model or by maximum-likelihood in Bernoulli trials) converges a.s. to its true value $\bar{\theta}$ and the information matrix tends a.s. to the penalized D -optimal matrix at $\bar{\theta}$ as $N \rightarrow \infty$, see Pronzato (2009a). (Note that the true optimal design for sequential dependent observations is extremely difficult to construct, see, e.g., Gautier and Pronzato (1998, 2000) for suboptimal attempts.) Moreover, the estimator is asymptotically normal, with variance-covariance matrix given by the inverse of the usual information matrix, similarly to the non-adaptive case, see Pronzato (2009a). The strong consistency of $\hat{\theta}^N$ is preserved when λ_N is taken as a control parameter that tends to infinity not too fast (more slowly than $N/(\log \log N)$). As in (Pronzato, 2000), by letting λ_N tend to infinity one focusses more and more on cost minimization and thus obtain design measures that converge weakly to the delta measure at $x^* = \arg \min_{x \in \mathcal{X}} \phi(x, \bar{\theta})$ (and when, moreover, the property (9) is satisfied, all design points tend to concentrate around x^*).

As an illustration, an example with a nonlinear model with binary responses is considered in the next section, first for local penalized optimal design and then for adaptive penalized optimal design through simulations.

4. Example: Cox model for efficacy-toxicity response

The example is taken from (Dragalin and Fedorov, 2006) and concerns a problem with bivariate binary responses. Let Y (respectively Z) denote the binary outcome indicating efficacy (resp. toxicity) for a trial at dose x . The set of available doses consists of 11 points equally spaced in the interval $[-3, 3]$, $\mathcal{X} = \{x^{(1)}, \dots, x^{(11)}\}$, $x^{(i)} < x^{(i+1)}$, $i = 1, \dots, 10$. We write $\text{Prob}\{Y = y, Z = z|x, \theta\} = \pi_{yz}(x, \theta)$, $Y, y, Z, z \in \{0, 1\}$. The following six-parameter model is used in (Dragalin and Fedorov, 2006) (we refer to that paper for motivations and justifications):

$$\begin{aligned}\pi_{11}(x, \theta) &= \frac{e^{a_{11}+b_{11}x}}{1 + e^{a_{01}+b_{01}x} + e^{a_{10}+b_{10}x} + e^{a_{11}+b_{11}x}} \\ \pi_{10}(x, \theta) &= \frac{e^{a_{10}+b_{10}x}}{1 + e^{a_{01}+b_{01}x} + e^{a_{10}+b_{10}x} + e^{a_{11}+b_{11}x}} \\ \pi_{01}(x, \theta) &= \frac{e^{a_{01}+b_{01}x}}{1 + e^{a_{01}+b_{01}x} + e^{a_{10}+b_{10}x} + e^{a_{11}+b_{11}x}} \\ \pi_{00}(x, \theta) &= (1 + e^{a_{01}+b_{01}x} + e^{a_{10}+b_{10}x} + e^{a_{11}+b_{11}x})^{-1}\end{aligned}$$

with $\theta = (a_{11}, b_{11}, a_{10}, b_{10}, a_{01}, b_{01})^\top$. The log-likelihood function of a single observation (Y, Z) at dose x is then $l(Y, Z, x; \theta) = YZ \log \pi_{11}(x, \theta) + Y(1 - Z) \log \pi_{10}(x, \theta) + (1 - Y)Z \log \pi_{01}(x, \theta) + (1 - Y)(1 - Z) \log \pi_{00}(x, \theta)$ and elementary calculations show that the contribution to the Fisher information matrix is

$$\mu(x, \theta) = \frac{\partial \mathbf{p}^\top(x, \theta)}{\partial \theta} (\mathbf{P}^{-1}(x, \theta) + [1 - \pi_{11}(x, \theta) - \pi_{10}(x, \theta) - \pi_{01}(x, \theta)]^{-1} \mathbf{1} \mathbf{1}^\top) \frac{\partial \mathbf{p}(x, \theta)}{\partial \theta^\top}$$

where $\mathbf{p}(x, \theta) = [\pi_{11}(x, \theta), \pi_{10}(x, \theta), \pi_{01}(x, \theta)]^\top$, $\mathbf{P}(x, \theta) = \text{diag}\{\mathbf{p}(x, \theta)\}$ and $\mathbf{1} = (1, 1, 1)^\top$. Note that $\mu(x, \theta)$ is generally of rank 3. As in (Dragalin and Fedorov, 2006), we take $\theta = (3, 3, 4, 2, 0, 1)^\top$. The D -optimal design is then supported on $x^{(1)}, x^{(4)}, x^{(5)}$ and $x^{(10)}$ with respective weights 0.3318, 0.3721, 0.1259 and 0.1701.

Next section illustrates how a locally optimal design for (5) depends on the choice of λ and ϕ .

4.1. Locally optimal design

We first choose a penalty function given by the inverse of the probability $\pi_{10}(x, \theta)$ of efficacy and no toxicity (probability of success) and take

$$\phi_1(x, \theta) = \pi_{10}^{-1}(x, \theta).$$

The Optimal Safe Dose (OSD), maximizing $\pi_{10}(x, \theta)$, is $x^{(5)} = -0.6$. Figure 1 presents the optimal designs $\xi^*(\lambda)$ for λ varying between 0 and 100 along the horizontal axis. The weight associated with each $x^{(i)}$ on the vertical axis is proportional to the thickness of the plot. The design measure tends to give more and more weight to the OSD $x^{(5)}$ as λ increases, but some support points remain far from $x^{(5)}$.




Figure 1: Optimal designs $\xi^*(\lambda)$ as function of $\lambda \in [0, 100]$ for the cost function $\pi_{10}^{-1}(x, \theta)$: each horizontal dotted line corresponds to a point in \mathcal{X} , the thickness of the plot indicates the associated weight.

Consider now the cost function

$$\phi_2(x, \theta) = \{\pi_{10}^{-1}(x, \theta) - [\max_x \pi_{10}(x, \theta)]^{-1}\}^2, \quad (13)$$

also related to the probability of success, but more flat than $\pi_{10}^{-1}(x, \theta)$ around its minimum (at the OSD). Using the condition (8) with the test designs $\xi_{\lambda,1}$, giving weights $(1-\alpha)/2, \alpha$ and $(1-\alpha)/2$ at $x^{(4)}, x^{(5)}$ and $x^{(6)}$ respectively, and $\xi_{\lambda,2}$, giving weights $1-\beta$ and β at $x^{(4)}$ and $x^{(5)}$ respectively, we obtain that the support of $\xi^*(\lambda)$ is included in $\{x^{(4)}, x^{(5)}, x^{(6)}\}$ for $\alpha \gtrsim 0.9993$ and $\beta \gtrsim 0.4508$, showing that the optimal designs concentrate on three doses around the optimal one when λ is large enough. Figure 2 presents the optimal designs $\xi^*(\lambda)$ for λ varying between 0 and 100 along the horizontal axis. It shows that for $\lambda \gtrsim 75$ the optimum designs are supported on $x^{(4)}$ and $x^{(6)}$ only, with weights approximately 1/2 each, that is, all patients in a trial defined by $\xi^*(\lambda)$ receive a dose close to the optimal one, $x^{(5)}$. Note, however, that none receives the OSD (compare with Figure 2). The situations changes for larger values of λ , see Figure 3 where the optimal design is supported on $\{x^{(4)}, x^{(5)}, x^{(6)}\}$ for $\lambda \gtrsim 160$ and the weight of the optimal dose $x^{(5)}$ increases with λ .

Finally, one may also consider a penalty function that puts more stress on toxicity avoidance, for instance

$$\phi_3(x, \theta) = \pi_{10}^{-1}(x, \theta)[1 - \pi_{\cdot 1}(x, \theta)]^{-1}, \quad (14)$$

which is used in (Dragalin and Fedorov, 2006), with $\pi_{\cdot 1}(x, \theta) = \pi_{01}(x, \theta) + \pi_{11}(x, \theta)$ denoting the marginal probability of toxicity (the dose minimizing $\phi_3(x, \theta)$ is then $x^{(4)}$). For λ large enough the optimal designs then tend to avoid large doses, compare Figure 4 with Figures 1 and 2.

4.2. Adaptive design

As an illustration of the behavior of the adaptive scheme (12), we present now some simulation results (using the value $\theta = (3, 3, 4, 2, 0, 2)^\top$). For comparison, we use the




Figure 2: Same as Figure 1, but for the cost-function (13).




Figure 3: Same as Figure 2, but for $\lambda \in [0, 1000]$.

up-and-down rule of Ivanova (2003) (which is also considered by Dragalin and Fedorov (2006), see also (Kpamegan and Flournoy, 2001, p. 221)), defined by

$$x_{N+1} = \begin{cases} \max\{x^{(i_N-1)}, x^{(1)}\} & \text{if } Z_N = 1, \\ x^{(i_N)} & \text{if } Y_N = 1 \text{ and } Z_N = 0, \\ \min\{x^{(i_N+1)}, x^{(11)}\} & \text{if } Y_N = 0 \text{ and } Z_N = 0, \end{cases} \quad (15)$$

where the index $i_N \in \{1, \dots, 11\}$ is defined by $x^{(i_N)} = x_N$ and (Y_N, Z_N) denotes the observation for x_N . The stationary allocation distribution $\xi_{u\&d}$ is log-concave and is approximately given by

$$\xi_{u\&d}(\theta) \simeq \left\{ \begin{array}{cccccccc} x^{(1)} & x^{(2)} & x^{(3)} & x^{(4)} & x^{(5)} & x^{(6)} & x^{(7)} & x^{(8)} \\ 1.70 \cdot 10^{-3} & 2.12 \cdot 10^{-2} & 0.146 & 0.426 & 0.345 & 5.88 \cdot 10^{-2} & 1.90 \cdot 10^{-3} & 1.13 \cdot 10^{-5} \end{array} \right\}$$




Figure 4: Same as Figure 1, but for the cost-function (14).

(the total weight on $x^{(9)}, x^{(10)}, x^{(11)}$ is less than 10^{-7}). Note that the mode is at $x^{(4)}$, one dose below the OSD $x^{(5)}$. See Durham and Flournoy (1994); Giovagnoli and Pintacuda (1998); Ivanova (2003) for analytical results.

We consider trials on 36 patients, organized in a similar way as in (Dragalin and Fedorov, 2006): the allocation for the first 10 patients uses the up-and-down rule above, starting with the lowest dose $x^{(1)}$; after the 10th patient, the up-and-down rule is still used until the first observed toxicity ($Z_N = 1$); we then switch to the adaptive design rule (12), with the restriction that we do not allow allocation at a dose one step higher than the maximum level tested so far (following recommendations for practical implementation, see Dragalin and Fedorov (2006)). The parameters are estimated by maximum likelihood (the log-likelihood $\sum_i l(Y_i, Z_i, x_i; \theta)$ being regularized by the addition of the term $0.01 \|\theta\|^2$, which is equivalent to maximum *a posteriori* estimation with the normal prior $\mathcal{N}(\mathbf{0}, 50 \mathbf{I})$, with \mathbf{I} the 6-dimensional identity matrix).

We first use the penalty function $\phi_1(x, \theta) = \pi_{10}^{-1}(x, \theta)$, with minimum value at the OSD $x^{(5)}$, $\phi_1(x^{(5)}, \theta) \simeq 1.2961$. Figure 5 shows the progress of a typical trial with $\lambda_N \equiv 2$. The symbols indicate the values of the observations at the given points: Δ for $(Y = 0, Z = 0)$, \triangleright for $(Y = 1, Z = 0)$, \diamond for $(Y = 1, Z = 1)$ and ∇ for $(Y = 0, Z = 1)$. The up-and-down rule is used until $N = 15$ where toxicity is observed. The next dose should have been $x^{(4)}$ but the adaptive design rule (12) selects $x^{(6)}$ instead.

As noticed by a referee, in a practical implementation one would certainly be more cautious than what is shown on Figure 5: although toxicity is observed for subjects 17 and 19, the design recommends escalation to a higher dose. As shown on Figure 6, the substitution of the penalty function $\phi_3(x, \theta)$ given by (14) for $\phi_1(x, \theta)$ permits to avoid this type of behavior (compare with Figure 5).

We now perform 1,000 independent repetitions of similar trials, using three different adaptive rules: (i) the up-and-down rule (15) used along the whole trial, for the 36 patients; (ii) the up-and-down rule followed by adaptive D -optimal design (11); and




Figure 5: Graphical presentation of a trial with $\phi_1(x, \theta) = \pi_{10}^{-1}(x, \theta)$ and $\lambda_N \equiv 2$ in (12); \triangle is for $(Y = 0, Z = 0)$, \triangleright for $(Y = 1, Z = 0)$, \diamond for $(Y = 1, Z = 1)$ and ∇ for $(Y = 0, Z = 1)$.

(iii) the up-and-down rule followed by adaptive penalized D -optimal design (12), with $\phi_1(x, \theta) = \pi_{10}^{-1}(x, \theta)$ and $\lambda_N \equiv 2$. Since the choice $\lambda_N \equiv 2$ may seem arbitrary, we also considered the situation (iv) where λ_N is adapted to the estimated value of θ . To reduce the computational cost, we only adapt λ_N once in the trial, at the value N_s when we abandon the up-and-down rule. The value $\lambda_N = \lambda_{N_s}^*$, $N = N_s, \dots, 36$, is chosen as the solution for λ of $\Phi[\xi^*(\lambda), \hat{\theta}_{ML}^{N_s}] = C_\gamma = (1 + \gamma) \min_{x \in \mathcal{X}} \phi(x, \hat{\theta}_{ML}^{N_s})$; we take $\gamma = 0.52$ because it yields $\lambda = 2$ when θ is substituted for $\hat{\theta}_{ML}^{N_s}$, that is, $\Phi[\xi^*(2), \theta] / \phi_\theta^* \simeq 1.52$ (it corresponds to allowing an average reduction of about 34% for the probability of success compared to $\max_{x \in \mathcal{X}} \pi_{10}(x, \hat{\theta}_{ML}^{N_s})$). The solution for λ is easily obtained by dichotomy, since we know that the solution satisfies $0 < \lambda \leq (1 + 1/\gamma) p / C_\gamma$, see Proposition 1-(i), and $\Phi[\xi^*(\lambda), \theta]$ decreases when λ increases (note that this is more economical in terms of computations than the Lagrangian approach proposed in Sect. 2.3 of Cook and Fedorov (1995) in a more general situation). Table 2 summarizes the results in terms of the following performance measures: $\Phi_1(\xi_{36}, \theta) = (1/36) \sum_{i=1}^{36} \phi_1(x_i, \theta)$, the total cost of the experiment; $J(\xi_{36}, \theta) = \det^{-1/6}[\mathbf{M}(\xi_{36}, \theta)]$, which indicates the precision of the estimation of θ ; $\hat{x}_{\{t=i\}}^*$, the number of times the estimated OSD at the end of the trial, that is, $\arg \max_{x \in \mathcal{X}} \pi_{10}(x, \hat{\theta}_{ML}^{36})$, coincided with $x^{(i)}$, for $i = 4, 5, 6$, and $\hat{x}_{\{t<4\}}^*$ (resp. $\hat{x}_{\{t>6\}}^*$), the number of times the estimated OSD was smaller than $x^{(4)}$ (resp. larger than $x^{(6)}$); finally $\#x^{(11)}$, the percentage of patients that received the highest dose $x^{(11)}$. The values of $\Phi_1(\xi, \theta)$ and $J(\xi, \theta)$ for the designs $\xi_{u\&d}$, ξ_D^* and $\xi^*(\lambda = 2)$ computed at the true value of θ are also indicated.

Table 2 reveals that the up-and-down rule (15) is very cautious: its associated cost $\Phi_1(\xi_{36}, \theta)$ is low, the extreme dose $x^{(11)}$ has never been used over the 1,000 repetitions. On the other hand, it fails at providing a precise estimation of the model parameters, and the OSD is estimated at values higher than $x^{(6)}$ in almost 14% of the cases. Parameter estimation is naturally more precise with adaptive D -optimal design, and the true OSD




Figure 6: Graphical presentation of a trial, in the same conditions as in Figure 5 but with the penalty $\phi_3(x, \theta)$ given by (14).

design rule	$\Phi_1(\xi, \theta)$	$J(\xi, \theta)$	$\widehat{x}_{\{t < 4\}}^*$	$\widehat{x}_{\{t = 4\}}^*$	$\widehat{x}_{\{t = 5\}}^*$	$\widehat{x}_{\{t = 6\}}^*$	$\widehat{x}_{\{t > 6\}}^*$	$\#x^{(11)}$
(i): (15)	1.87	28.02	20	386	369	86	139	0
$\xi_{u\&d}(\theta)$	1.47	29.4						
(ii): (15)–(11)	3.16	17.23	0	198	705	78	19	5%
$\xi_D^*(\theta)$	4.45	14.99						
(iii): (15)–(12), $\phi_1, \lambda_N \equiv 2$	2.25	19.22	3	231	693	59	14	1.6%
$\xi^*(\lambda = 2, \theta), \phi_1$	1.97	17.00						
(iv): (15)–(12), ϕ_1, λ_N^*	2.38	18.78	0	223	682	70	25	2.3%
(v): (15)–(12), $\phi_3, \lambda_N \equiv 2$	2.09	21.08	4	330	575	61	23	0.5%

Table 2: Performance measures of different adaptive designs for 36 patients (sample mean over 1,000 repetitions for $\Phi_1(\xi) = \int_{\mathcal{X}} \phi_1(x, \theta) \xi(dx)$ and $J(\xi, \theta) = \det^{-1/6}[\mathbf{M}(\xi, \theta)]$; $\widehat{x}_{\mathcal{A}}$ is the number of times $\arg \max_{x \in \mathcal{X}} \pi_{10}(x, \hat{\theta}_{ML}^{36}) \in \mathcal{A}$ and $\#x^{(11)}$ is the percentage of patients that received the highest dose $x^{(11)}$ over the 1,000 repetitions). The values of $\Phi_1(\xi, \theta)$ and $J(\xi, \theta)$ for $\xi_{u\&d}$, ξ_D^* and $\xi^*(\lambda = 2)$ are also indicated. In (iii) and (iv), the adaptive rule (12) uses the penalty function $\phi_1(x, \theta) = \pi_{10}^{-1}(x, \theta)$, (v) uses the penalty $\phi_3(x, \theta)$ given by (14).

$x^{(5)}$ is recognized in more than 70% of the cases. However, this successful behavior in terms of collective ethics is obtained at the price of having about 5% of patients receiving a dose as high as $x^{(11)}$; also, the associated value of $\Phi_1(\xi_{36}, \theta)$ is rather high. The adaptive penalized D -optimal design appears to make a good compromise between the two strategies: the value of $\Phi_1(\xi_{36}, \theta)$ is close to that of the up-and-down rule, the value of $\det^{-1/6}[\mathbf{M}(\xi_{36}, \theta)]$ is close to that obtained for adaptive D -optimal design. It recognized $x^{(5)}$ as the OSD in about 70% of the cases and only 1.6% of the patients received the dose $x^{(11)}$ when $\lambda_N \equiv 2$ (2.3% when λ_N is adapted). Of course, other choices of λ would set other compromises. Also, other penalty functions can be used to define design rules more reluctant at allocating high doses: the performance of the adaptive penalized D -optimal design with the penalty (14) is indicated on line (v) of the table; comparison with line (iii) shows that the less precise estimation of the OSD is balanced by a more cautious strategy (only 0.5% of the patients receive the highest dose $x^{(11)}$).

In order to limit more severely the number of patients that receive very high doses, we finally consider a compromise strategy that implements a smoother (and less arbitrary) transition between up-and-down and adaptive penalized D -optimal design.

Letting λ_N increase with N in (12) might be controversial in the context of clinical trials (as mentioned by a referee, it implies a loss of randomization feature in the trial and brings potential for operational bias, not to mention the difficulties induced concerning the statistical analysis of the data accumulated during the trial due to the nonhomogeneous treatments at the beginning and the end of the trial). Thus, letting λ_N depend on N might not be a realistic practical situation in this context. It is instructive, however, to investigate through simulations the performance achieved by such “non-stationary” designs. We thus consider much longer trials, with $N_T = 240$ patients enrolled. Define $x^{**}(\theta) = \arg \min_{x \in \mathbb{R}} \phi(x, \theta)$ and $h(x, \theta) = \partial \phi(x, \theta) / \partial x$. From the implicit function theorem,

$$\nabla_{\theta} x^{**}(\theta) = \frac{dx^{**}(\theta)}{d\theta} = - \left[\frac{\partial h(x, \theta)}{\partial x} \Big|_{x=x^{**}(\theta)} \right]^{-1} \frac{\partial h(x, \theta)}{\partial \theta} \Big|_{x=x^{**}(\theta)}$$

and, when using the up-and-down rule (15) the estimator $\hat{\theta}_{ML}^N$ asymptotically satisfies

$$\sqrt{N} V_N^{-1/2} [x^{**}(\hat{\theta}_{ML}^N) - x^{**}(\bar{\theta})] \xrightarrow{d} z \sim \mathcal{N}(0, 1), \quad N \rightarrow \infty,$$

where $V_N = [\nabla_{\theta} x^{**}(\hat{\theta}_{ML}^N)]^{\top} \mathbf{M}^{-1}(\xi_N, \hat{\theta}_{ML}^N) [\nabla_{\theta} x^{**}(\hat{\theta}_{ML}^N)]$. Based on that, we decide to switch from the up-and-down rule to the adaptive one when $\sqrt{V_N/N} < x^{(2)} - x^{(1)}$, the interval between two consecutive doses. If N_s is the index of the patient for which the rule changes, we take λ_{N_s} as the solution for λ of $\Phi[\xi^*(\lambda), \hat{\theta}_{ML}^{N_s}] = C_{\gamma} = (1+\gamma) \min_{x \in \mathcal{X}} \phi(x, \hat{\theta}_{ML}^{N_s})$ with $\gamma = 0.5$ (thus targeting 33% of decrease with respect to the maximum of $\pi_{10}(x, \hat{\theta}_{ML}^{N_s})$). The value of λ_{N_T} at the end of the trial is chosen as the solution for λ of the same equation with $\gamma = 0.1$ (allowing only 9% of decrease with respect to the maximum of $\pi_{10}(x, \hat{\theta}_{ML}^{N_s})$). In between λ_N increases at a logarithmic rate, that is, $\lambda_N = \lambda_{N_s} [1 + a \log(N/N_s)]$, $N = N_s, \dots, N_T$, with $a = (\lambda_{N_T}/\lambda_{N_s} - 1) / \log(N_T/N_s)$. When uncertainty on the OSD is large, that is when $\sqrt{V_N/N} > [x^{(2)} - x^{(1)}] / 2$, we also restrict the allocations at high doses by adapting the design space, taken as $\mathcal{X}_N = \{x^{(1)}, \dots, x^{(i_N)}\}$ at step N : the maximum dose $x^{(i_N)}$ allowed in (12) is never more than one step higher than previous dose and is smaller than previous dose if toxicity was observed.

The results obtained for 150 repetitions of the experiment are summarized in Table 3 (the results obtained when the up-and-down rule (15) is used for the 240 patients are also indicated). One may notice the precise estimation of the OSD for the adaptive penalized design compared to the up-and-down rule (it even does slightly better than the up-and-down rule both in terms of $\Phi_1(\xi, \theta)$ and $\det^{-1/6}[\mathbf{M}(\xi, \theta)]$). At the same time, only about 0.11% of the patients received the maximal dose $x^{(11)}$.

5. Conclusions

We have shown that constrained optimal design can be formulated in a way that allows a clear balance between gaining information and minimizing a cost. A dose-finding

	$\Phi_1(\xi, \theta)$	$J(\xi, \theta)$	$\widehat{x^*}_{\{t < 4\}}$	$\widehat{x^*}_{\{t = 4\}}$	$\widehat{x^*}_{\{t = 5\}}$	$\widehat{x^*}_{\{t = 6\}}$	$\widehat{x^*}_{\{t > 6\}}$	$\#x^{(11)}$
(15)	1.54	29.04	0	21	116	11	2	0
(15)–(12), $\phi_1, \lambda_N \nearrow$	1.52	27.87	0	13	135	1	1	0.11%

Table 3: Performance measures of adaptive design (12) with increasing λ_N for 240 patients (sample mean over 150 repetitions for $\Phi_1(\xi) = \int_{\mathcal{X}} \phi_1(x, \theta) \xi(dx)$ and $J(\xi, \theta) = \det^{-1/6}[\mathbf{M}(\xi, \theta)]$; $\widehat{x^*}_{\mathcal{A}}$ is the number of times $\arg \max_{x \in \mathcal{X}} \pi_{10}(x, \hat{\theta}_{ML}^{200}) \in \mathcal{A}$ and $\#x^{(11)}$ is the percentage of patients that received the highest dose $x^{(11)}$ over the 150 repetitions). The adaptive rule (12) uses the penalty function $\phi_1(x, \theta) = \pi_{10}^{-1}(x, \theta)$.

example with bivariate binary responses has illustrated the potential of adaptive penalized D -optimal design to set compromises between individual ethics (cost of the experiment, related to the probability of success through the penalty function) and collective ethics (information gained from the trial, to be used for future patients). Further developments and numerical studies are required to define suitable rules for selecting cost functions and for choosing the value (or the sequence of values) for the penalty coefficients λ_N .

Appendix

Proof of Proposition 1.

(i) Since ξ^* is optimal, we have for all $x \in \mathcal{X}$: $\text{trace}[\mu(x, \theta)\mathbf{M}^{-1}(\xi^*, \theta)] \leq p + \lambda[\phi(x, \theta) - \Phi(\xi^*, \theta)]$, see (4). This is true in particular at a x^* defined by (6) and $\text{trace}[\mu(x^*, \theta)\mathbf{M}^{-1}(\xi^*, \theta)] \geq 0$ gives the result.

(ii) For any $a > 0$, take $\lambda \geq a/\Delta_\theta(\xi)$ and define $\tilde{\xi} = (1 - \alpha)\xi + \alpha\delta_{x^*}$ with δ_{x^*} the delta measure at a point x^* satisfying (6) and $\alpha = 1 - a/[\lambda\Delta_\theta(\xi)]$. This gives $\Phi(\tilde{\xi}, \theta) - \phi_\theta^* = a/\lambda$ and $\log \det \mathbf{M}(\tilde{\xi}, \theta) \geq p \log(1 - \alpha) + \log \det \mathbf{M}(\xi, \theta)$. Therefore,

$$\begin{aligned} \log \det \mathbf{M}(\xi^*, \theta) - \lambda[\Phi(\xi^*, \theta) - \phi_\theta^*] &\geq \log \det \mathbf{M}(\tilde{\xi}, \theta) - \lambda[\Phi(\tilde{\xi}, \theta) - \phi_\theta^*] \\ &\geq p \log a - a - p \log \Delta_\theta(\xi) + \log \det \mathbf{M}(\xi, \theta) - p \log \lambda. \end{aligned} \quad (16)$$

Since $\Phi(\xi^*, \theta) \geq \phi_\theta^*$, the result follows.

When $\phi(x, \theta)$ is bounded by $\bar{\phi}_\theta$, the optimality of ξ^* implies that for all $x \in \mathcal{X}$, $\text{trace}[\mu(x, \theta)\mathbf{M}^{-1}(\xi^*, \theta)] \leq B = p + \lambda(\bar{\phi}_\theta - \phi_\theta^*)$. Write $\mu(x, \theta) = \mathbf{F}_\theta^\top(x)\mathbf{F}_\theta(x)$ with $\mathbf{F}_\theta^\top(x) = [\mathbf{f}_{1,\theta}(x), \dots, \mathbf{f}_{m,\theta}(x)]$ and $\mathbf{f}_{i,\theta}(x)$ a p -dimensional vector, $i = 1, \dots, m$. From the inequality $\text{trace}[\mu(x, \theta)\mathbf{M}^{-1}(\xi^*, \theta)] \leq B$ we obtain that $\mathbf{f}_{i,\theta}^\top(x)\mathbf{M}^{-1}(\xi^*, \theta)\mathbf{f}_{i,\theta}(x) \leq B$, $i = 1, \dots, m$. We have

$$\Lambda_{\min}[\mathbf{M}(\xi^*, \theta)] = \Lambda_{\max}^{-1}[\mathbf{M}^{-1}(\xi^*, \theta)] = [\max_{\|\mathbf{u}\|=1} \mathbf{u}^\top \mathbf{M}^{-1}(\xi^*, \theta) \mathbf{u}]^{-1}.$$

Consider the optimization problem defined by: maximize $\mathbf{u}^\top \mathbf{A}^\top \mathbf{A} \mathbf{u}$ with respect to \mathbf{A} and \mathbf{u} respectively in $\mathbb{R}^{n \times p}$ and \mathbb{R}^p , $n \leq p$, subject to the constraints $\|\mathbf{u}\| = 1$ and $\mathbf{f}_{i,\theta}^\top(x)\mathbf{A}^\top \mathbf{A} \mathbf{f}_{i,\theta}(x) \leq B$, $\forall x \in \mathcal{X}$ and $\forall i = 1, \dots, m$. The optimal solution is obtained for $\mathbf{A} = \mathbf{v}^\top \in \mathbb{R}^p$ such that $|\mathbf{f}_{i,\theta}^\top(x)\mathbf{v}| \leq \sqrt{B}$, $\forall x \in \mathcal{X}$, $\forall i = 1, \dots, m$, and $b = \mathbf{v}^\top \mathbf{v}$ is maximal. For x varying in \mathcal{X} the $\mathbf{f}_{i,\theta}(x)$'s span \mathbb{R}^p (since a nonsingular information matrix exists). Therefore, there exists a positive constant δ such that the optimal value for b is bounded by B/δ , and $\Lambda_{\min}[\mathbf{M}(\xi^*, \theta)] > \delta/B$. \blacksquare

Proof of Proposition 2.

Since $\text{trace}[\mu(x, \theta)\mathbf{M}^{-1}(\xi^*, \theta)] \leq p + \lambda[\phi(x, \theta) - \Phi(\xi^*, \theta)]$ for all $x \in \mathcal{X}$ when ξ^* is optimal, and $\int_{\mathcal{X}} \{\text{trace}[\mu(x, \theta)\mathbf{M}^{-1}(\xi^*, \theta)] - \lambda\phi(x, \theta)\} \xi^*(dx) = p - \lambda\Phi(\xi^*, \theta)$, we have $\text{trace}[\mu(\hat{x}, \theta)\mathbf{M}^{-1}(\xi^*, \theta)] = p + \lambda[\phi(\hat{x}, \theta) - \Phi(\xi^*, \theta)]$ at any \hat{x} support point of ξ^* . Suppose that λ is large enough so that there exists a design $\xi_\lambda \in \Xi$ satisfying $\Delta_\theta(\xi_\lambda) \geq p/\lambda$. We proceed as for Proposition 1-(ii) and construct a design $\tilde{\xi}_\lambda = (1 - \alpha)\xi_\lambda + \alpha\delta_{x^*}$ with $\alpha = 1 - p/[\lambda\Delta_\theta(\xi_\lambda)]$ so that $\Phi(\tilde{\xi}_\lambda, \theta) - \phi_\theta^* = p/\lambda$. With the same notation as in the proof of Proposition 1-(ii), we can write $\text{trace}[\mu(\hat{x}, \theta)\mathbf{M}^{-1}(\tilde{\xi}_\lambda, \theta)] = \sum_{i=1}^m \mathbf{f}_{i,\theta}^\top(\hat{x})\mathbf{M}^{-1}(\tilde{\xi}_\lambda, \theta)\mathbf{f}_{i,\theta}(\hat{x})$. We then follow the same approach as in Harman and Pronzato (2007) and define

$$\mathbf{H}(\xi^*, \xi, \theta) = \mathbf{M}^{1/2}(\xi^*, \theta)\mathbf{M}^{-1}(\xi, \theta)\mathbf{M}^{1/2}(\xi^*, \theta).$$

We obtain

$$\begin{aligned} \text{trace}[\mu(\hat{x}, \theta)\mathbf{M}^{-1}(\tilde{\xi}_\lambda, \theta)] &= \sum_{i=1}^m \mathbf{f}_{i,\theta}^\top(\hat{x})\mathbf{M}^{-1/2}(\xi^*, \theta) \mathbf{H}(\xi^*, \tilde{\xi}_\lambda, \theta)\mathbf{M}^{-1/2}(\xi^*, \theta)\mathbf{f}_{i,\theta}(\hat{x}) \\ &\geq \Lambda_{\min}[\mathbf{H}(\xi^*, \tilde{\xi}_\lambda, \theta)] \text{trace}[\mu(\hat{x}, \theta)\mathbf{M}^{-1}(\xi^*, \theta)] \\ &= \Lambda_{\min}[\mathbf{H}(\xi^*, \tilde{\xi}_\lambda, \theta)] \{p + \lambda[\phi(\hat{x}, \theta) - \Phi(\xi^*, \theta)]\} \\ &\geq \Lambda_{\min}[\mathbf{H}(\xi^*, \tilde{\xi}_\lambda, \theta)] \left\{p + \lambda[\phi(\hat{x}, \theta) - \Phi(\tilde{\xi}_\lambda, \theta)]\right\} \\ &= \Lambda_{\min}[\mathbf{H}(\xi^*, \tilde{\xi}_\lambda, \theta)] \lambda[\phi(\hat{x}, \theta) - \phi_\theta^*] \end{aligned}$$

where we used the property $\Delta_\theta(\xi^*) \leq p/\lambda = \Delta_\theta(\tilde{\xi}_\lambda)$, see (i). Therefore,

$$\text{trace}[\mu(\hat{x}, \theta)\mathbf{M}^{-1}(\xi_\lambda, \theta)] \geq (1 - \alpha)\text{trace}[\mu(\hat{x}, \theta)\mathbf{M}^{-1}(\tilde{\xi}_\lambda, \theta)] \quad (17)$$

$$\geq \frac{p \Lambda_{\min}[\mathbf{H}(\xi^*, \tilde{\xi}_\lambda, \theta)] [\phi(\hat{x}, \theta) - \phi_\theta^*]}{\Delta_\theta(\xi_\lambda)}. \quad (18)$$

The last step consists in deriving a lower bound on $\Lambda_{\min}[\mathbf{H}(\xi^*, \tilde{\xi}_\lambda, \theta)]$ that does not depend on ξ^* . Since $\text{trace}[\mathbf{H}^{-1}(\xi^*, \tilde{\xi}_\lambda, \theta)] = \int_{\mathcal{X}} \text{trace}[\mu(x, \theta)\mathbf{M}^{-1}(\xi^*, \theta)] \tilde{\xi}_\lambda(dx)$, the optimality of ξ^* implies

$$\begin{aligned} \text{trace}[\mathbf{H}^{-1}(\xi^*, \tilde{\xi}_\lambda, \theta)] &\leq p + \lambda \int_{\mathcal{X}} [\phi(x, \theta) - \Phi(\xi^*, \theta)] \tilde{\xi}_\lambda(dx) \\ &= p + \lambda [\Phi(\tilde{\xi}_\lambda, \theta) - \Phi(\xi^*, \theta)] \\ &\leq p + \lambda [\Phi(\tilde{\xi}_\lambda, \theta) - \phi_\theta^*] = 2p. \end{aligned} \quad (19)$$

Therefore, $\Lambda_{\min}[\mathbf{H}(\xi^*, \tilde{\xi}_\lambda, \theta)] \geq 1/(2p)$, which, together with (18), concludes the proof. ■

One might notice that, following Harman and Pronzato (2007), a tighter bound could be obtained in (8) by using

$$\text{trace}[\mathbf{H}(\xi^*, \tilde{\xi}_\lambda, \theta)] \leq \max_{x \in \mathcal{X}} [\mu(x, \theta)\mathbf{M}^{-1}(\tilde{\xi}_\lambda, \theta)] \leq (1 - \alpha)^{-1} \max_{x \in \mathcal{X}} [\mu(x, \theta)\mathbf{M}^{-1}(\xi_\lambda, \theta)]$$

in addition to (19). However, since $(1 - \alpha)^{-1} = \lambda\Delta_\theta(\xi_\lambda)/p$, the improvement obtained for the lower bound $1/(2p)$ on $\Lambda_{\min}[\mathbf{H}(\xi^*, \tilde{\xi}_\lambda, \theta)]$ is negligible for large λ .

Acknowledgements. The author wishes to thank V. Fedorov for his comments and N. Flournoy for her careful reading of the paper, her detailed comments and advice. The careful reading of the two referees is gratefully acknowledged.

References

- Chaudhuri, P., Mykland, P., 1993. Nonlinear experiments: optimal design and inference based likelihood. *Journal of the American Statistical Association* 88 (422), 538–546.
- Cook, D., Fedorov, V., 1995. Constrained optimization of experimental design (invited discussion paper). *Statistics* 26, 129–178.
- Cook, D., Wong, W., 1994. On the equivalence between constrained and compound optimal designs. *Journal of the American Statistical Association* 89 (426), 687–692.
- Dragalin, V., Fedorov, V., 2006. Adaptive designs for dose-finding based on efficacy-toxicity response. *Journal of Statistical Planning and Inference* 136, 1800–1823.
- Dragalin, V., Fedorov, V., Wu, Y., 2008. Adaptive designs for selecting drug combinations based on efficacy-toxicity response. *Journal of Statistical Planning and Inference* 138, 352–373.
- Durham, S., Flournoy, N., 1994. Random walks for quantile estimation. In: Berger, J., Gupta, S. (Eds.), *Statistical Decision Theory and Related Topics V*. Springer, New York, pp. 467–476.
- Fedorov, V., Hackl, P., 1997. *Model-Oriented Design of Experiments*. Springer, Berlin.
- Gautier, R., Pronzato, L., 1998. Sequential design and active control. In: Flournoy, N., Rosenberger, W., Wong, W. (Eds.), *New Developments and Applications in Experimental Design, Lecture Notes — Monograph Series*, vol. 34. IMS, Hayward, pp. 138–151.
- Gautier, R., Pronzato, L., 2000. Adaptive control for sequential design. *Discussiones Mathematicae, Probability & Statistics* 20 (1), 97–114.
- Giovagnoli, A., Pintacuda, N., 1998. Properties of frequency distributions induced by general ‘up-and-down’ methods for estimating quantiles. *Journal of Statistical Planning and Inference* 74, 51–63.
- Haines, L., Perevozskaya, I., Rosenberger, W., 2003. Bayesian optimal designs in Phase I clinical trials. *Biometrics* 59, 591–600.
- Hardwick, J., Stout, Q., 2001. Optimizing a unimodal response function for binary variables. In: Atkinson, A., Bogacka, B., Zhigljavsky, A. (Eds.), *Optimum Design 2000*. Kluwer, Dordrecht, Ch. 18, pp. 195–210.
- Harman, R., Pronzato, L., 2007. Improvements on removing non-optimal support points in D-optimum design algorithms. *Statistics & Probability Letters* 77, 90–94.

- Ivanova, A., 2003. A new dose-finding design for bivariate outcomes. *Biometrics* 59, 1001–1007.
- Kiefer, J., Wolfowitz, J., 1960. The equivalence of two extremum problems. *Canadian Journal of Mathematics* 12, 363–366.
- Kpamegan, E., Flournoy, N., 2001. An optimizing up-and-down design. In: Atkinson, A., Bogacka, B., Zhigljavsky, A. (Eds.), *Optimum Design 2000*. Kluwer, Dordrecht, Ch. 19, pp. 211–224.
- Mikulecká, J., 1983. On a hybrid experimental design. *Kybernetika* 19 (1), 1–14.
- Pronzato, L., 2000. Adaptive optimisation and D -optimum experimental design. *Annals of Statistics* 28 (6), 1743–1761.
- Pronzato, L., 2009a. Asymptotic properties of nonlinear least squares estimates in stochastic regression models over a finite design space. Application to self-tuning optimisation. In: *Proc. 15th IFAC Symposium on System Identification, Saint-Malo, France*.
- Pronzato, L., 2009b. One-step ahead adaptive D -optimal design on a finite design space is asymptotically optimal. *Metrika* (to appear).
- Pukelsheim, F., 1993. *Optimal Experimental Design*. Wiley, New York.