Optimal designs for the development of personalized treatment rules

David Azriel^a, Yosef Rinott^b and Martin Posch^{c*}

October 11, 2021

 a Faculty of Industrial Engineering and Management, The Technion

^b Department of Statistics and Federmann Center for the Study of Rationality, The Hebrew University

 c Center for Medical Statistics, Informatics and Intelligent Systems, Medical University of Vienna

Abstract

We consider the development of personalized treatment rules that select treatments based on patients' individual covariates. In particular, we study the design of multi-armed parallel group clinical trials to estimate treatment rules that identify the best treatment for a given patient with given covariates. Assuming that the outcomes in each treatment arm are given by a homoscedastic linear model, with possibly different variances between treatment arms, and that the trial subjects are a random sample of an unselected overall population, we optimize the (possibly randomized) treatment allocation allowing the allocation rates to depend on the covariates. The trial designs minimize the regret for future patients assuming that they will arrive from the same population as the trial subjects. We find that, for the case of two treatments, the approximately optimal allocation rule does not depend on the value of the covariates but only on the variances of the responses. In contrast, for the case of three treatments or more, the optimal allocation design does depend on the covariates as we show for specific scenarios. In addition, the optimal allocation rule depends on the true regression coefficients in the case of three but not in the case of two treatment arms. Furthermore, in the case of three treatments, for small sample sizes randomized allocations are optimal while for very large studies almost deterministic allocation rules minimize the regret. The methods are illustrated with a recently published dietary clinical trial.

Key words: personalized medicine; minimal regret; optimal allocation; experimental design.

^{*}Corresponding author: Martin Posch; email: martin.posch@meduniwien.ac.at.

1 Introduction

We consider multi-armed clinical trials for selecting the best among several treatments, where the primary outcome may depend on patients' covariates. Modeling the outcome as a function of the covariates does not only reduce the error variance, but allows one also to estimate "personalized" treatment rules, that is, rules that assign the treatment with the best estimated outcome for a patient having given values of the covariates. This work concerns the study of optimal design of clinical trials with the goal of estimating personalized treatment rules. We assume that a sample of subjects from the relevant population is given, and we look for an optimal allocation of subjects participating in the trial to treatments as a function of their covariates, in order to minimize the regret for future patients. The regret is defined as the expected difference of outcomes under the optimal and the estimated treatment rule.

The development of statistical methods for personalized or precision medicine is a very active area in statistics. See, e.g., Kosorok and Moodie (2015) and Kosorok et al. (2021), both containing many recent references. Much of the literature deals with finding, on the basis of a given sample, the best treatment as a function of personal baseline covariates, which can be, for example, genetic or other biomarkers. Another common theme is adaptive or dynamic schemes, where treatment depends on the individual patient's course of disease. Kosorok and Laber (2019) describe two different settings: single- and multi-decision problems. Our focus here is on the former, where the treatment is only chosen at a single time point, and the choice may depend on personal baseline covariates. In this context, Minsker et al. (2016) and Zhao and LeBlanc (2020) study clinical trials where subjects can be selected in an optimal way in order to obtain an efficient design. In contrast, we assume that subjects arrive according to the population distribution and cannot be selected. Lee and Wason (2019) study optimal design for the purpose of comparing treatments, where treatment effects are estimated, and the design aims at minimizing the variances of the estimators. The covariates are assumed to be binary, and hence the design has a discrete nature. Zhang et al. (2021) study a similar regression model to the one considered here with two treatments and aim at finding the minimax design with respect to the variance of the decision rule. Both Zhang et al. (2021) and Lee and Wason (2019) differ from the present paper in the target function, formulated in the present paper as regret, which seems to be most natural, as well as the general approach: theirs is more computational, while we aim at finding explicit optimal designs, when possible.

In the related context of bandit problems with covariates, Goldenshluger and Zeevi (2013) studied a sequential allocation scheme where at each stage one out of two treatments needs to be assigned under a minimax framework. A high-dimensional version of this problem was studied in Bastani and Bayati (2015). The goal in the bandit formulation is to determine the best treatment while minimizing some loss or regret function for subjects in the trial sample. This is a different setting than the one considered

here, since rather than be concerned with optimal treatment for subjects in the trial, we look for designs that will allocate patients to treatments efficiently for the goal of finding the best treatment for future patients who may require one of the treatments under study.

Assuming that the outcomes in each treatment arm are given by a homoscedastic linear model (with possibly different variances between treatment arms) we study the regret via an approximation that assumes "ideal" conditions, such as normality instead of asymptotic normality of linear regression coefficient estimators. We show that this yields a useful approximation for quite small sample sizes. We also study an asymptotic regret. However, it turns out that when there are three or more treatments, the optimal designs for moderate sample sizes may substantially differ from the asymptotically optimal design.

For the case of two treatments, we show that the optimal trial design allocates patients in proportion to the standard deviations of the response under each treatment, and the optimal allocation probabilities do not depend on the covariates. This is shown for the case where the outcome in each treatment group depends on the covariates according to a *p*-dimensional linear regression model, as well as for the case of a single covariate and a polynomial regression model. We also formulate a minimax criterion and show that it is satisfied by designs with allocation probabilities that do not depend on the covariates values. For parameter spaces with the same bound on the variances of the two treatments, a balanced design (equal probability) is optimal under the minimax criterion.

For the case of three or more treatments and a linear regression model in a single covariate, we combine theoretical arguments and numerical calculations to study some examples of optimal designs. We show that the optimal allocation in general depends on the covariate. For specific cases, we show that the asymptotically optimal allocation is deterministic, i.e., the range of the covariate is partitioned into intervals and in each interval the allocation probability to a particular treatment is 1. Such allocation rules arise typically in very large studies.

We illustrate the optimal design in the setting of a dietary clinical trial where three diets were compared in a parallel group design (Ebbeling et al., 2018). The primary outcome in the original trial was averaged total energy expenditure, which is correlated with weight loss. An important patient covariate is insulin secretion and we consider the objective of determining a treatment rule that assigns for each patient the best diet based on his or her insulin secretion level measured at baseline. Here the utility of the diets is defined by the averaged total energy expenditure adjusted by diet costs. We derive the optimal experimental design, that is, the optimal allocation of subjects to treatments as a function of the measured insulin level, minimizing a suitable regret function. Based on the data in Ebbeling et al. (2018) we quantify the improvement in regret of the optimal design compared to designs with optimal allocation rates that do not depend on the covariate and the design with equal allocation rates that was used in the original trial (see Section 6 for details).

The paper is organized as follows. In Section 2 the estimation problem is formalized and the regret defined. In Section 3 an explicit formula that approximates the regret (dubbed the *ideal regret*) is derived. The optimal design for the case of two treatments is discussed in Section 4 after studying the rate of the approximating ideal regret to the true one. We also compute the asymptotic regret, give a numerical example, and discuss the case of polynomial regression (Section 4.3). In Section 5 the setting of K treatments and a single covariate is considered and the corresponding asymptotic regret is derived. In Section 5.1 we derive formula (14) for the ideal regret, and we compare the regret and ideal regret asymptotically. Section 5.2 discusses an algorithm based on (14) to obtain an approximate optimal design, and demonstrates the reduction in regret by such optimization. Furthermore, in Section 5.3 the computation of a lower bound for the optimal regret is discussed and an example is given where the optimal allocation rule minimizing the asymptotic regret can be derived. In Section 6 the procedure is illustrated with an example, and in Section 7 limitations and possible extensions of the approach are discussed.

2 Problem Statement

Consider K possible treatments T_1, \ldots, T_K . Let Y be a one-dimensional continuous response variable and let $\mathbf{X} \in \mathbb{R}^p$ be a vector of a subject's covariates. We assume that the joint distribution of the covariates is continuous with density denoted by $f(\mathbf{x})$. If some of the covariates are discrete, the density $f(\mathbf{x})$ should be taken with respect to a suitable measure, and our results apply as long as one of the covariates has a continuous distribution (conditioned on all other covariates). For simplicity of notation, we assume \mathbf{X} has a density with respect to Lebesgue measure. The expected response is $E(Y|\mathbf{X}, T_k) = g_k(\mathbf{X})$, where g_k is an unknown function. The optimal treatment for a subject with a covariate vector \mathbf{x} is $\delta^*(\mathbf{x}) = \arg \max_{k \in \{1,\ldots,K\}} g_k(\mathbf{x})$ (assuming that a higher response is better). If the above arg max contains more than one k, one arbitrary treatment is selected.

Suppose that a clinical trial with n subjects is performed in order to estimate δ^* . Let $\mathbf{X}_1, \ldots, \mathbf{X}_n$ denote a sample according to f of the covariates of the n subjects. The design we study here consists of allocating subjects to treatments, taking account of their covariates. Each subject i is allocated independently to treatment k with probability $\pi_k(\mathbf{X}_i)$, where $\pi_1(\mathbf{x}), \ldots, \pi_K(\mathbf{x})$ are non-negative functions satisfying for each $\mathbf{x}, \sum_{k=1}^K \pi_k(\mathbf{x}) = 1$. The allocation functions define densities for the covariates in the K treatment groups. Specifically, in treatment k, \mathbf{X} is sampled from the density $f_k(\mathbf{x}) := f(\mathbf{x})\pi_k(\mathbf{x})/\nu_k$ where $\nu_k := \int f(\mathbf{x})\pi_k(\mathbf{x})d\mathbf{x}$. Note that using the sampled $\mathbf{X}_1, \ldots, \mathbf{X}_n$ we can estimate the density f using a parametric or nonparametric estimator. In fact, we may have more observations of the covariates \mathbf{X}

than those *n* values which will enter the trial. Our purpose is to find the optimal allocation functions $\pi_1(\mathbf{x}), \ldots, \pi_K(\mathbf{x})$ that minimize the regret, which is defined next.

Let $\hat{\delta}$ denote the final estimate of the optimal rule δ^* based on the covariates, responses, and allocations of the subjects in the trial. Then the *regret* for a randomly chosen future patient is defined as

$$R(\pi_1, \dots, \pi_K) := E\left[g_{\delta^*(\widetilde{\mathbf{X}})}(\widetilde{\mathbf{X}}) - g_{\widehat{\delta}(\widetilde{\mathbf{X}})}(\widetilde{\mathbf{X}})\right],\tag{1}$$

where $\widetilde{\mathbf{X}}$ denotes the covariate vector of the chosen patient and the expectation is over $\widetilde{\mathbf{X}}$, $\widehat{\delta}$, and the randomizations in the allocations. In words, the regret is the difference between the expected response to the optimal treatment δ^* and its estimate $\widehat{\delta}$ for any independent future subject arising from the same population as those in the trial. A similar criterion appears, for examples, in Qian and Murphy (2011); see also Kosorok and Moodie (2015) and references therein. Our goal is to minimize the regret over all feasible allocations, that is, nonnegative functions $\pi_k(\cdot)$ that satisfy $\sum_{k=1}^{K} \pi_k(\mathbf{x}) = 1$ for every \mathbf{x} .

In this study we assume that we have some preliminary knowledge about the parameters of the problem, such as the functions g_k , and the variances of Y given **X**, T_k and the density f of the covariates. This approach concurs with existing literature on locally optimal designs; see, e.g., Chernoff (1953) and Silvey (2013, Chapter 6), where optimality is achieved for a given set of parameters. Sverdlov and Rosenberger (2013) review several methodological advances in (local) optimal allocation for clinical trials. Our goal is to find optimal allocation functions relative to the assumed parameter values, and obtain a treatment rule $\hat{\delta}$ based on an experiment according to our design. Knowledge about the parameters can arise from previous experience or theory, from a pilot study, or in some situations from earlier phases of the study.

3 The regret under "ideal" conditions

Let $(Y_1, \mathbf{X}_1), \ldots, (Y_n, \mathbf{X}_n)$ be a sample obtained by certain allocation functions π_1, \ldots, π_n , given along with the treatment allocated for each \mathbf{X}_i . Assume $K \geq 2$ and $\mathbf{X} \in \mathbb{R}^p$, and under treatment T_k we have $Y = g_k(\mathbf{X}) + \varepsilon$, where $g_k(\mathbf{X}) = E(Y|\mathbf{X}, T_k) = \alpha_k + \beta_k^t \mathbf{X}$, and $\sigma_k^2 := Var(\varepsilon|\mathbf{X}, T_k)$. Let $\widehat{\alpha}_k, \widehat{\beta}_k$ denote the OLS estimators based on data from treatment $T_k, \widehat{g}_k(\mathbf{X})$ the corresponding estimated regression functions, and $\widehat{\delta}(\mathbf{x}) = \arg \max_k \widehat{g}_k(\mathbf{X})$ the decision rule. Then,

$$R(\pi_1, \dots, \pi_K) = \sum_{k=1}^K \int_{\mathbb{R}^p} P\left(\widehat{\delta}(\mathbf{x}) = k\right) \left[g_{\delta^*(\mathbf{x})}(\mathbf{x}) - g_k(\mathbf{x})\right] f(\mathbf{x}) d\mathbf{x}.$$
 (2)

For $k = 1, \ldots, K$, let

$$\boldsymbol{Q}_{k} := \int_{\mathbb{R}^{p}} \begin{pmatrix} 1 \\ \mathbf{x} \end{pmatrix} (1, \mathbf{x}^{t}) f_{k}(\mathbf{x}) d\mathbf{x}, \quad \boldsymbol{\Sigma}_{k} := \frac{1}{\nu_{k}} \sigma_{k}^{2} \boldsymbol{Q}_{k}^{-1}, \quad \boldsymbol{\xi}_{k}^{2}(\mathbf{x}) := (1, \mathbf{x}^{t}) \boldsymbol{\Sigma}_{k} \begin{pmatrix} 1 \\ \mathbf{x} \end{pmatrix}.$$
(3)

It is well known that under standard conditions (see, e.g., Hansen (2019) Hansen Chapter 7) we have $\sqrt{n}[(\widehat{\alpha}_k, \widehat{\beta}_k) - (\alpha_k, \beta_k)] \rightarrow N(0, \Sigma_k)$, and in particular Σ_k is the asymptotic variance of the OLS estimators.

In order to approximate the probability $P\left(\hat{\delta}(\mathbf{x}) = k\right)$, we first impose the "*ideal*" conditions, namely that $\hat{g}_k(\mathbf{x})$ possess their joint asymptotic distribution. We shall later drop this assumption . In particular, by considering the regression model $Y = \sum_{k=1}^{K} \{\alpha_k + \beta_k^t \mathbf{X}\} \mathbf{1}_k + \varepsilon$, where $\mathbf{1}_k$ is the indicator of treatment k, we obtain that under the ideal conditions, the OLS estimators are jointly normal with $Var(\hat{\alpha}_k, \hat{\beta}_k) = \mathbf{\Sigma}_k$ and hence $Var(\hat{g}_k(\mathbf{x})) = \xi_k^2(\mathbf{x})/n$. Also, under the ideal conditions the estimators $\hat{g}_k(\mathbf{x})$ are independent conditionally on $\mathbf{X}_1, \ldots, \mathbf{X}_n, T_1, \ldots, T_K$, with constant expectation $\alpha_k + \beta_k^t \mathbf{x}$. By the law of total covariance, conditional independence and constant conditional expectations imply that $\hat{g}_k(\mathbf{x})$ are uncorrelated, and by joint normality, they are independent.

Therefore, with P_I denoting the probability under the ideal conditions, we have

$$P_{I}\left(\widehat{\delta}(\mathbf{x})=k\right) = P_{I}\left(\max_{\ell=1,\dots,K,\ell\neq k}\widehat{g}_{\ell}(\mathbf{x}) < \widehat{g}_{k}(\mathbf{x})\right)$$

$$= P_{I}\left(\bigcap_{\ell\neq k}\left\{\frac{\sqrt{n}[\widehat{g}_{\ell}(\mathbf{x}) - g_{\ell}(\mathbf{x})]}{\xi_{\ell}(\mathbf{x})} < \frac{\sqrt{n}[\widehat{g}_{k}(\mathbf{x}) - g_{\ell}(\mathbf{x})]}{\xi_{\ell}(\mathbf{x})}\right\}\right)$$

$$= P\left(\bigcap_{\ell\neq k}\left\{Z_{\ell} < \frac{Z_{k}\xi_{k}(\mathbf{x}) + \sqrt{n}[g_{k}(\mathbf{x}) - g_{\ell}(\mathbf{x})]}{\xi_{\ell}(\mathbf{x})}\right\}\right)$$

$$= \int P\left(\bigcap_{\ell\neq k}\left\{Z_{\ell} < \frac{z\xi_{k}(\mathbf{x}) + \sqrt{n}[g_{k}(\mathbf{x}) - g_{\ell}(\mathbf{x})]}{\xi_{\ell}(\mathbf{x})}\right\} | Z_{k} = z\right)\varphi(z)dz$$

$$= \int\prod_{\ell=1,\dots,K,\ell\neq k}\Phi\left(\frac{z\xi_{k}(\mathbf{x}) + \sqrt{n}[g_{k}(\mathbf{x}) - g_{\ell}(\mathbf{x})]}{\xi_{\ell}(\mathbf{x})}\right)\varphi(z)dz, \qquad (4)$$

where P is with respect to $\{Z_m\}_{m=1}^K$, which are independent N(0,1), Z_m represents $\frac{\sqrt{n}[\hat{g}_m(\mathbf{x}) - g_m(\mathbf{x})]}{\xi_m(\mathbf{x})}$; φ and Φ denote the standard normal density and cumulative distribution function. The product in the last line was justified above by independence under the ideal conditions, which is also used to compute the conditional probability. We define the *ideal regret*, denoted by R_I , by

$$R_I(\pi_1,\ldots,\pi_K) := \sum_{k=1}^K \int_{\mathbb{R}^p} P_I\left(\widehat{\delta}(\mathbf{x}) = k\right) \left[g_{\delta^*(\mathbf{x})}(\mathbf{x}) - g_k(\mathbf{x})\right] f(\mathbf{x}) d\mathbf{x},\tag{5}$$

where $P_I\left(\widehat{\delta}(\mathbf{x}) = k\right)$ is given in (4).

4 The case of two treatments

4.1 Regret approximations and optimal design

We assume now that K = 2 and $Y = g_k(\mathbf{X}) + \varepsilon$ under treatment T_k for k = 1, 2, where $g_k(\mathbf{X})$ denotes the conditional mean of Y given \mathbf{X}, T_k , where $\mathbf{X} \in \mathbb{R}^p$ and

$$g_1(\mathbf{X}) = \alpha_1 + \boldsymbol{\beta}_1^t \mathbf{X}, \ g_2(\mathbf{X}) = \alpha_2 + \boldsymbol{\beta}_2^t \mathbf{X}, \ Var(\varepsilon | \mathbf{X}, T_1) = \sigma_1^2 \text{ and } Var(\varepsilon | \mathbf{X}, T_2) = \sigma_2^2.$$
(6)

We do not assume normality of the errors. However, we assume the existence of the moment generating function of ε when conditioned on \mathbf{X}, T_k . This assumption is needed for some large deviation exponential bounds used below, and could be relaxed to assuming finiteness of some high order moments instead. We also assume that \mathbf{X} is continuous with a bounded density $f(\mathbf{x})$ supported on $[0, 1]^p$.

Under the ideal conditions of Section 3 with K = 2, Equation (4) becomes

$$P_I\left(\widehat{\delta}(\mathbf{x}) = k\right) = \Phi\left(\frac{\sqrt{n}[g_k(\mathbf{x}) - g_l(\mathbf{x})]}{\sqrt{V(\mathbf{x})}}\right), \text{ where } V(\mathbf{x}) := \xi_1^2(\mathbf{x}) + \xi_2^2(\mathbf{x}) = (1, \mathbf{x}^t)\left(\mathbf{\Sigma}_1 + \mathbf{\Sigma}_2\right) \begin{pmatrix} 1\\ \mathbf{x} \end{pmatrix}$$
(7)

for $1 \le k \ne l \le 2$. The ideal regret (5) is

$$R_{I}(\pi_{1},\pi_{2}) = \int_{\{\mathbf{x}:g_{1}(\mathbf{x})>g_{2}(\mathbf{x})\}} \Phi\left(\frac{\sqrt{n}[g_{2}(\mathbf{x})-g_{1}(\mathbf{x})]}{\sqrt{V(\mathbf{x})}}\right) [g_{1}(\mathbf{x})-g_{2}(\mathbf{x})]f(\mathbf{x})d\mathbf{x}$$
$$+ \int_{\{\mathbf{x}:g_{1}(\mathbf{x})
$$= \int_{[0,1]^{p}} \Phi\left(-\frac{\sqrt{n}[g_{1}(\mathbf{x})-g_{2}(\mathbf{x})]}{\sqrt{V(\mathbf{x})}}\right) |g_{2}(\mathbf{x})-g_{1}(\mathbf{x})|f(\mathbf{x})d\mathbf{x}.$$
(8)$$

Note that R_I depends also on the parameters $\alpha_1, \alpha_2, \beta_1, \beta_2, \sigma_1^2, \sigma_2^2$, which are suppressed. By (8), minimization of $R_I(\pi_1, \pi_2)$ amounts to finding the allocation functions π_1, π_2 that minimize $V(\mathbf{x})$. The following theorem states that this can be done by minimizing $V(\mathbf{x})$ uniformly over \mathbf{x} .

Theorem 4.1. The allocation functions $\pi_1^0(\mathbf{x}) := \frac{\sigma_1}{\sigma_1 + \sigma_2}$ and $\pi_2^0(\mathbf{x}) := \frac{\sigma_2}{\sigma_1 + \sigma_2}$ minimize both $V(\mathbf{x})$ uniformly over \mathbf{x} , and $R_I(\pi_1, \pi_2)$.

Note that in order to implement the allocation functions $\pi_1^0(\mathbf{x})$ and $\pi_2^0(\mathbf{x})$ one only needs to know the ratio of the variances of the errors of the two treatments, and no other parameters of the model. The allocation functions also does not require knowledge of the density $f(\mathbf{x})$.

Recalling that $V(\mathbf{x})$ can be written as a quadratic form, see (7), Theorem 4.1 follows immediately from Lemma 4.2 below whose proof, as all other proofs, is given in the Appendix. Given matrices A and B, we write $A \succeq B$ if A - B is positive semi-definite. Let $\mathbf{Q} := \int_{[0,1]^p} {1 \choose \mathbf{x}} (1, \mathbf{x}^t) f(\mathbf{x}) d\mathbf{x}$, and recall the notation of (3). We have

Lemma 4.2. With the above definitions,

$$\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2 \succeq (\sigma_1 + \sigma_2)^2 \boldsymbol{Q}^{-1}.$$
(9)

Moreover, the lower bound is attained when $\pi_1(\mathbf{x}) = \pi_1^0(\mathbf{x})$ and $\pi_2(\mathbf{x}) = \pi_2^0(\mathbf{x})$.

To see the latter statement note that when $\pi_k(\mathbf{x})$ do not depend on \mathbf{x} , both matrices Q_k are proportional to Q and therefore both matrices Σ_k are proportional to Q^{-1} .

Instead of assuming that the regression parameters are known, consider now a minimax criterion for the ideal regret. Let γ denote the regression parameters, i.e., $\gamma := (\alpha_1, \alpha_2, \beta_1, \beta_2, \sigma_1^2, \sigma_2^2)$ and consider the parameter space Γ , where for some constant S, $\sigma_1^2, \sigma_2^2 \leq S$ and the rest of the parameters are unrestricted. The bound on the variances is required to make the problem nondegenerate. The allocation functions π_1^*, π_2^* are said to be minimax if

$$\sup_{\gamma \in \Gamma} R_I(\pi_1^*, \pi_2^*) \le \sup_{\gamma \in \Gamma} R_I(\pi_1, \pi_2) \text{ for all allocation functions } \pi_1, \pi_2.$$
(10)

By (8) and the definition of $V(\mathbf{x})$ in (7) it is easy to see that the regret is maximized for $\gamma \in \Gamma$ when $\sigma_1^2 = \sigma_2^2 = S$, in which case Theorem 4.1 implies that the optimal design is $\pi_1^*(\mathbf{x}) = \pi_2^*(\mathbf{x}) = 1/2$. This statement is true regardless of the value of the other parameters in γ . It follows that the balanced design satisfies (10) and is therefore minimax. The same argument shows that if the bound $\sigma_1^2, \sigma_2^2 \leq S$ is replaced by $\sigma_1^2 \leq S_1, \sigma_2^2 \leq S_2$, then the minimax design is $\pi_1^*(\mathbf{x}) = \frac{\sqrt{S_1}}{\sqrt{S_1} + \sqrt{S_2}}, \pi_2^*(\mathbf{x}) = \frac{\sqrt{S_2}}{\sqrt{S_1} + \sqrt{S_2}}$.

We now discuss the regret (2). In Theorem 4.3 we approximate the regret by the ideal regret, and then we use the above results for the ideal regret to find an asymptotically optimal design for the regret itself. We later demonstrate numerically that the latter design provides a good approximation for finite sample sizes. The regret is given by

$$R(\pi_{1},\pi_{2}) = \int_{\{\mathbf{x}:g_{1}(\mathbf{x})>g_{2}(\mathbf{x})\}} P(\widehat{g}_{2}(\mathbf{x})>\widehat{g}_{1}(\mathbf{x}))(g_{1}(\mathbf{x})-g_{2}(\mathbf{x}))f(\mathbf{x})d\mathbf{x} + \int_{\{\mathbf{x}:g_{2}(\mathbf{x})>g_{1}(\mathbf{x})\}} P(\widehat{g}_{1}(\mathbf{x})>\widehat{g}_{2}(\mathbf{x}))(g_{2}(\mathbf{x})-g_{1}(\mathbf{x}))f(\mathbf{x})d\mathbf{x}.$$
 (11)

Theorem 4.3. Under model (6) with $\nu_1, \nu_2 > 0$, we have for any $\varepsilon > 0$

$$n^{3/2-\varepsilon} |R(\pi_1,\pi_2) - R_I(\pi_1,\pi_2)| \to 0 \text{ as } n \to \infty.$$

To facilitate the representation, we state the next result in the one-dimensional case first. In this case g_1 and g_2 are functions of a single continuous variable, and we assume that their intersection point $\theta := \frac{\alpha_1 - \alpha_2}{\beta_2 - \beta_1}$ is in [0,1]. Assume without loss of generality, that $\beta_2 > \beta_1$. It follows that for $x < \theta$, $g_1(x) > g_2(x)$, that is, T_1 is the better treatment, and when $x > \theta$, T_2 is better. In the case p = 1 the limit of the regret is simple, given by the following theorem.

Theorem 4.4. Under model (6) with p = 1, if $\nu_1, \nu_2 > 0$, then $\lim_{n \to \infty} nR(\pi_1, \pi_2) = \frac{V(\theta)f(\theta)}{2(\beta_2 - \beta_1)}I\{\theta \in [0, 1]\}.$

It follows that in order to minimize $\lim_{n\to\infty} nR(\pi_1,\pi_2)$ we have to minimize $V(\theta)$, which by Theorem 4.1 is achieved by taking $\pi_1(\mathbf{x}) = \pi_1^0(\mathbf{x})$ and $\pi_2(\mathbf{x}) = \pi_2^0(\mathbf{x})$. We remark that by a delta method calculation, the asymptotic variance of $\hat{\theta} = \frac{\hat{\alpha}_1 - \hat{\alpha}_2}{\hat{\beta}_2 - \hat{\beta}_1}$ is $\frac{V(\theta)}{(\beta_2 - \beta_1)^2}$, which is proportional to the limit of $nR(\pi_1, \pi_2)$. Thus, minimizing the latter limit is equivalent to minimizing the asymptotic variance of the estimator of the intersection point θ .

We return to general dimension p. Set $\beta_k^t := (\beta_{k,1}, \ldots, \beta_{k,p}), k = 1, 2, \beta_{k,-1}^t := (\beta_{k,2}, \ldots, \beta_{k,p})$ and without loss of generality assume that X_1 is a continuous variable given the rest of the covariates, and that $\beta_{2,1} > \beta_{1,1}$. For given $\mathbf{x}_{-1} = (x_2, \ldots, x_p)$ define $\theta_1 := \theta_1(\mathbf{x}_{-1}) := \frac{\alpha_1 - \alpha_2 + (\beta_{1,-1} - \beta_{2,-1})^t \mathbf{x}_{-1}}{\beta_{2,1} - \beta_{1,1}}$. If $\theta_1 \in [0,1]$, treatment T_1 is better for a covariate vector \mathbf{x} satisfying $x_1 < \theta_1 = \theta_1(\mathbf{x}_{-1})$, and otherwise T_2 is better. The limit of the regret is given in the following theorem.

Theorem 4.5. Assume model (6) and $\nu_1, \nu_2 > 0$, then

$$\lim_{n \to \infty} nR(\pi_1, \pi_2) = \frac{1}{2(\beta_{2,1} - \beta_{1,1})} \int_{[0,1]^{p-1}} V(\theta_1, \mathbf{x}_{-1}) f(\theta_1, \mathbf{x}_{-1}) I\{\theta_1 \in [0,1]\} d\mathbf{x}_{-1}.$$
 (12)

Thus, asymptotically, $nR(\pi_1, \pi_2)$ is proportional to the integral of the variance of the estimate of the intersection curve $\theta_1(\mathbf{x}_{-1})$, with weights proportional to the density at the intersection points, and by (12) $\lim_{n\to\infty} nR(\pi_1, \pi_2)$ is minimized by the allocation $\pi_1^0(\mathbf{x}), \pi_2^0(\mathbf{x})$ since this allocation minimizes V uniformly. By Theorems 4.3 and 4.4 we thus have

Theorem 4.6. The allocation $\pi_1^0(\mathbf{x}), \pi_2^0(\mathbf{x})$ is optimal in the sense of minimizing $\lim_{n\to\infty} nR(\pi_1, \pi_2)$ for any number of covariates p.

Moreover, for any $\varepsilon > 0$ there exists C > 0 such that for any allocation functions π_1, π_2 with $\nu_1, \nu_2 > 0$

$$R(\pi_1, \pi_2) \ge R(\pi_1^0, \pi_2^0) - \frac{C}{n^{3/2-\varepsilon}}.$$

In particular, the optimal design that minimizes $R(\pi_1, \pi_2)$ can improve the allocation π_1^0, π_2^0 by at most $C/n^{3/2-\varepsilon}$.

The error term is meaningful as it is of smaller order than the regret, which is 1/n according to Theorem 4.5.

In model (6) separate regressions models are considered for each of the treatments and we used the optimal (best linear unbiased estimators) estimators for each model, which are the least squares estimators, to compute $\hat{\delta}$. Alternatively, one can consider a joint regression model $Y = \sum_{k=1}^{2} \{\alpha_k + \beta_k^t \mathbf{X}\} \mathbf{1}_k + \varepsilon$, where $\mathbf{1}_k$ is the indicator of treatment k. Since we did not assume a common variance among the treatments, the variance of ε in this model depends on $\mathbf{1}_k$. If the variances σ_1^2, σ_2^2 are known or can be estimated accurately enough, the weighted-least squares estimator with weights proportional to $1/\sigma_1^2, 1/\sigma_2^2$ is optimal. This amounts to making the variances equal, and then the above results imply that a balanced design (i.e., $\pi_1(\mathbf{x}) = 1/2$) is asymptotically optimal.

4.2 A numerical example

To illustrate the asymptotic approximations of the regret and the optimal allocations derived in Section 4.1 consider an example of model (6) with p = 1, $\alpha_1 = 0.2$, $\alpha_2 = 0$, $\beta_1 = 0.5$, $\beta_2 = 1$, $\sigma_1^2 = 0.1$, $\sigma_2^2 = 0.2$, and $X \sim U[0, 1]$. The optimal allocation is $\pi_1(\mathbf{x}) = \nu_1 = 0.414$. The σ 's are chosen such that R^2 in both regressions is about 0.2. Figure 1(a) shows the regression lines. For x smaller than $\theta = 0.4$, treatment 1 is better and otherwise treatment 2 is preferred. The regret (11) is evaluated using simulations (with 10^5 replications) and is compared to the ideal regret (8). Two scenarios are considered for the residuals in the regression model: normal and centered exponential. Figures (b) and (c) show the regret and ideal regret for n = 100 under the allocation function $\pi_1(x) = \nu_1$ where ν_1 vary from 0.2 to 0.6. The allocation optimizing the ideal regret, it seems that the optimal allocations, with respect to both, are very close. Thus, the approximation of Theorem 4.3 works well also for small n in this example. Figure 1 (d) shows that the regret slowly converges to its limit, and it is close to the limit only when n is quite large.





(d) Ideal regret and asymptotic regret

Figure 1: Figure (a): the regression lines. Figures (b) and (c) show the regret (computed by a simulation with 95% confidence intervals), and the ideal regret for n = 100 and different allocation ratios, where in (b) the residual is normal and in (c) it is exponential (centered). The regret is calculated for the allocation function $\pi_1(x) = \nu_1$ where ν_1 varies from 0.2 to 0.6. The optimal allocation is marked by the vertical line. Figure (d) shows n times the ideal regret (blue line) and its limit (gray line).

4.3 Polynomial Regression

We now consider the model of (6) where the functions $g_k(x)$ are assumed to be polynomials of a single continuous covariate X with density f, i.e., $g_k(X) = \alpha_k + \sum_{j=1}^J \beta_{jk} X^j$, where J denotes the degree of the polynomial. We will show that $\pi_1^0(x), \pi_2^0(x)$ are asymptotically optimal also in this setting.

Let $\theta_1 < \cdots < \theta_L$, where $L \leq J$, be the crossing points of $g_1(x)$ and $g_2(x)$, and define $\theta_0 = 0$ and $\theta_{L+1} = 1$. Assume without loss of generality that $g_1(x) > g_2(x)$ for $x \in (\theta_\ell, \theta_{\ell+1})$ if ℓ is odd, and the reverse inequality holds when ℓ is even. Therefore, the regret is

$$\begin{split} R(\pi_1, \pi_2) &= \sum_{\ell \text{ is odd}} \int_{\theta_\ell}^{\theta_{\ell+1}} P(\widehat{g}_2(x) > \widehat{g}_1(x)) (g_1(x) - g_2(x)) f(x) dx \\ &+ \sum_{\ell \text{ is even}} \int_{\theta_\ell}^{\theta_{\ell+1}} P(\widehat{g}_1(x) > \widehat{g}_2(x)) (g_2(x) - g_1(x)) f(x) dx. \end{split}$$

Notice that Theorems 4.3 and 4.5, which provide results for multivariate \mathbf{X} , do not apply in this case because the random vector $\mathbf{X} = (X, X^2, \dots, X^J)$ does not have a joint density. However, using the arguments of Theorem 4.3 (under similar notation and regularity conditions) it can be shown that $n^{3/2-\varepsilon}|R(\pi_1, \pi_2) - R_I(\pi_1, \pi_2)| \to 0$, for any $\varepsilon > 0$, where

$$R_I(\pi_1, \pi_2) = \int_0^1 \Phi\left(\frac{-\sqrt{n}|g_2(x) - g_1(x)|}{\sqrt{V(x)}}\right) |g_2(x) - g_1(x)|f(x)dx,\tag{13}$$

and now $\mathbf{x}^t = (1, x, \dots, x^J)$, $\boldsymbol{Q}_k = \int_{\mathbb{R}^p} \mathbf{x} \mathbf{x}^t f_k(x) dx$ and $\boldsymbol{\Sigma}_k = \frac{1}{\nu_k} \sigma_k^2 \boldsymbol{Q}_k^{-1}$, k = 1, 2, and $V(x) = \mathbf{x}^t (\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2) \mathbf{x}$, which is the asymptotic variance of $\hat{g}_1(x) - \hat{g}_2(x)$.

Theorem 4.6 implies that the design where $\pi_1(x) = \pi_1^0(x)$ minimizes V(x) uniformly over all x's. It follows that this design is asymptotically optimal also for this problem.

By Proposition 8.1 (in the appendix) we have

$$\lim_{n \to \infty} nR(\pi_1, \pi_2) = \sum_{\ell=1}^{L} \frac{f(\theta_\ell) V(\theta_\ell)}{2|(\beta_2 - \beta_1)^t \boldsymbol{\zeta}_\ell|},$$

where $\boldsymbol{\zeta}_{\ell} := (1, 2\theta_{\ell}, ..., J\theta_{\ell}^{J-1})^t$.

A careful inspection of the proofs of the above results shows that the optimality of the design with $\pi_1(x) = \pi_1^0(x)$ generalizes to regression functions of the form $g_k(X) = \sum_j \beta_{jk} h_j(X)$ for any functions h_j having bounded derivatives.

5 K Treatments and one Covariate

5.1 Regret and ideal regret

We consider the case p = 1 and $g_k(X) = \alpha_k + \beta_k X$, $Var(Y|X, T_k) = \sigma_k^2$, k = 1, ..., K. As in Section 4, we assume the existence of the moment generating function of ε when conditioned on (X, T_k) . We also assume that X is continuous with density f(x) and is supported on [0, 1]. Let θ_k , k = 0, ..., K be in increasing order, and such that treatment k is best in the interval (θ_{k-1}, θ_k) ; we assume that each treatment is best in some open interval, or equivalently that the intervals are nonempty.

Equations (2) and (5) for the regret and ideal regret imply immediately

$$R(\pi_1, \dots, \pi_K) = \sum_{k=1}^K \sum_{m=1}^K \int_{\theta_{m-1}}^{\theta_m} P\left(\widehat{g}_k(x) > \max_{\ell \neq k} \widehat{g}_\ell(x)\right) [g_m(x) - g_k(x)] f(x) dx$$

and the ideal regret is

$$R_{I}(\pi_{1},...,\pi_{K}) = \sum_{k=1}^{K} \sum_{m=1}^{K} \int_{\theta_{m-1}}^{\theta_{m}} \left(\int \prod_{\ell=1,...,K,\ell \neq k} \Phi\left(\frac{z\xi_{k}(x) + \sqrt{n}[g_{k}(x) - g_{\ell}(x)]}{\xi_{\ell}(x)}\right) \varphi(z) dz \right) [g_{m}(x) - g_{k}(x)] f(x) dx.$$
(14)

Similar to the results in Section 4, we give asymptotic results on the rate of convergence of the regret to the ideal regret R_I as well as the limit of the regret. Given allocation functions $\pi_k(x)$, we have by (3) that the distribution of the estimated regression coefficients $(\hat{\alpha}_k, \hat{\beta}_k)$ is approximately bivariate normal with means (α_k, β_k) and covariance

$$\Sigma_{k} = \frac{\sigma_{k}^{2}}{\nu_{k}} \begin{pmatrix} \frac{\tau_{k}^{2} + \mu_{k}^{2}}{\tau_{k}^{2}} & \frac{-\mu_{k}}{\tau_{k}^{2}} \\ \frac{-\mu_{k}}{\tau_{k}^{2}} & \frac{1}{\tau_{k}^{2}} \end{pmatrix} , \qquad (15)$$

where we assume positivity of $\nu_k := \int f(x)\pi_k(x)dx$, and denote the mean and variance of the covariate in group k by $\mu_k := \int_0^1 x f_k(x) dx$ and $\tau_k^2 := \int_0^1 (x - \mu_k)^2 f_k(x) dx$, respectively, $k = 1, \ldots, K$. Parallel to Theorems 4.3 and 4.4 we have

Theorem 5.1. Under the assumptions in the beginning of Section 5.1 we have for any $\varepsilon > 0$

$$\lim_{n \to \infty} n^{3/2 - \varepsilon} |R(\pi_1, \dots, \pi_K) - R_I(\pi_1, \dots, \pi_K)| = 0.$$

Theorem 5.2. Under the assumptions in the beginning of Section 5.1, we have

$$n \lim_{n \to \infty} R(\pi_1, \dots, \pi_K) = \sum_{m=1}^{K-1} \frac{V_m(\theta_m) f(\theta_m)}{2|\beta_{m+1} - \beta_m|},$$
(16)

where $V_m(x) = (1, x) \left(\boldsymbol{\Sigma}_m + \boldsymbol{\Sigma}_{m+1} \right) {\binom{1}{x}}.$

Using (15) it is possible to write $V_m(\theta_m)$ explicitly as follows

$$V_m(\theta_m) = \frac{\sigma_m^2}{\nu_m} \left[1 + \frac{(\theta_m - \mu_m)^2}{\tau_m^2} \right] + \frac{\sigma_{m+1}^2}{\nu_{m+1}} \left[1 + \frac{(\theta_m - \mu_{m+1})^2}{\tau_{m+1}^2} \right].$$

Theorem 5.2 implies that asymptotically the optimal allocation problem reduces to minimize a weighted average of the variances of $\hat{\theta}_1, \ldots, \hat{\theta}_{K-1}$ (see the discussion after Theorem 4.4), which are the estimates of the intersection points. The weights depend on the β 's and on the density f at the intersection points. Note also that $V_m(\theta_m)$ is the sum of the asymptotic variances of $\hat{\alpha}_m + \hat{\beta}_m \theta_m$ and $\hat{\alpha}_{m+1} + \hat{\beta}_{m+1} \theta_m$. Theorem 4.6 continues to hold for K treatments and thus, optimizing the ideal regret approximately optimizes the regret itself.

5.2 Approximate numeric optimization of allocation rules

In this section we demonstrate the utility of (14) in computing optimal designs when K > 2, and the reduction in the regret that can be achieved. The focus here is on the ideal regret. By Theorem 5.1 and the numerical results of Section 4.2, a reduction in the ideal regret implies a similar reduction in the regret itself.

To numerically find allocation rules minimizing the regret (14), we can use a finite dimensional parametrization of the set of allocation probability functions $\pi_k(x)$, $k = 1, \ldots, K$. One option is to use suitably rescaled polynomials of degree $M \ge 0$: Let $\mathbf{A} = (a_{k,m})$ denote a $(K-1) \times (M+1)$ matrix and set $h_k(x) = \exp(\sum_{m=0}^M a_{k,m}x^m)$ for $k = 1, \ldots, K-1$. Then, the allocation probabilities are defined by $\pi_k(x) = \frac{h_k(x)}{1+\sum_{j=1}^{K-1}h_j(x)}$, $k = 1, \ldots, K-1$ and $\pi_K(x) = 1 - \sum_{j=1}^{K-1}\pi_j(x)$. Note that $\mathbf{A} = \mathbf{0}$ corresponds to equal allocation $\pi_k(x) = 1/K$, $k = 1, \ldots, K$. Now, we can optimize the ideal regret by plugging $\pi_k(x)$ into (14) and numerically minimizing with respect to \mathbf{A} (e.g., by using the algorithms provided with the R function optim).

For M = 0 we obtain (approximately) optimal fixed allocation probabilities that do not depend on x. The larger M the more flexible allocation functions in x are fitted. However, because the target function (14) is not necessarily convex, we cannot guarantee that the numerical optimization will converge to the global minimum regret. Because the spaces of allocation probability functions are nested for increasing M, a heuristic optimization strategy is to iteratively optimize (14) starting with M = 0 and then subsequently increase M using the solution of the last iteration step as starting values for the next iteration (adding a column of zeroes to **A**). Note that this approach can also be used to optimize allocation rules minimizing the asymptotic regret (16) instead of the ideal regret (14).

To illustrate this approach we consider a specific setting where n = 200, K = 3, $(\alpha_1, \alpha_2, \alpha_3) = (0.0, -0.1, -1.2)$ and $(\beta_1, \beta_2, \beta_3) = (0.2, 0.5, 2.0)$; the left panel of Figure 2 plots the regression lines. We looked at two options for the variances: heteroscedastic - where $(\sigma_1, \sigma_2, \sigma_3) = \frac{1}{\sqrt{12}}(\beta_1, \beta_2, \beta_3)$ (so that $\sigma_k^2 = Var(\beta_k X)$ when X is uniform for each k, i.e., $R^2 \approx 1/2$ for each regression) and homoscedastic - where $\sigma_1 = \sigma_2 = \sigma_3 = \frac{\beta_2}{\sqrt{12}}$. We consider four settings for the distribution of X: Uniform(0,1) = Beta(1,1), Beta(0.5, 0.5) (U-shaped), Beta(2,2) (symmetric, unimodal), Beta(2,5) (asymmetric).

Starting with the balanced design, Table 1 shows the percentage of regret reduced by using the optimized fixed allocation design (M = 0) and by the optimized design with M = 4. The allocation probabilities when M = 0 are also given. It is clear from the table that the regret depends heavily on the distribution of X and on the σ 's. The table demonstrates that the reduction of the regret is quite significant. Furthermore, the optimal design where the allocation depends on x improves significantly over the fixed design (M = 0). An exception is the homoscedastic case and $X \sim Beta(2,5)$, in which the density function is small when $g_3(x)$ is maximal (see the left panel of Figure 2) and additionally σ_3 is small compared to β_3 . This makes this case close to K = 2 where fixed allocation designs are optimal. Another finding, which is not given in the table, is that the difference, in terms of the regret, between M = 1 and M = 4 is small; it is at most 2.3% in the eight scenarios we considered.

		heter		homoscedastic						
	% reduction		fixed alloc. $(M = 0)$		% reduction		fixed alloc. $(M = 0)$			
distribution	M = 0	M=4	π_1	π_2	π_3	M = 0	M=4	π_1	π_2	π_3
U(0,1)	24.0	40.9	0.12	0.33	0.55	11.7	25.4	0.40	0.43	0.17
Beta(0.5, 0.5)	25.1	41.9	0.13	0.34	0.53	16.5	29.6	0.41	0.44	0.15
Beta(2,2)	20.5	34.2	0.11	0.32	0.57	7.0	18.8	0.38	0.43	0.19
Beta(2,5)	11.5	22.8	0.18	0.45	0.37	21.3	25.8	0.45	0.46	0.09

Table 1: Results for the setting mentioned in the text under different distributions of X and for the heteroscedastic and homoscedastic models: the percentage of reduction achieved by the optimized designs with M = 0 and M = 4 relative to a balanced design, and the fixed allocation probabilities for M = 0.

5.3 A lower bound for the asymptotic regret (16)

By (14), (15), and (16) (see also (3)) both the ideal and the asymptotic regret depend on the functions $\pi_k(x)$ only via the quantities ν_k , μ_k , τ_k^2 , $k = 1, \ldots, K$. Therefore, we can minimize (14) and (16) in these parameters to obtain a lower bound for the regret. Let $\mu := \int_0^1 x f(x) \, dx$, $\tau^2 := \int_0^1 (x - \mu)^2 f(x) \, dx$ and

recall that $f(x) = \sum_{k=1}^{K} \nu_k f_k(x)$. Then, by the relations for central moments of mixture distributions we have

$$\sum_{k=1}^{K} \nu_k = 1, \quad \sum_{k=1}^{K} \nu_k \mu_k = \mu, \quad \sum_{k=1}^{K} \nu_k (\tau_k^2 + \mu_k^2) = \tau^2 + \mu^2.$$
(17)

Minimizing (14) or (16) in ν_k , μ_k , τ_k , k = 1, ..., K subject to the constraints (17) and ν_k , $\tau_k \ge 0, k = 1, ..., K$ yields a lower bound for the achievable regret. To obtain the minimum regret (instead of a lower bound) one needs to add additional constraints on μ_k , τ_k , k = 1, ..., K to restrict optimization to values for which there exist mixture components f_k with weighted sum f that assume these moments.

We now consider an example with a uniform f and K = 3, and minimize the asymptotic regret (16). Since now $f(x) \leq 1$ the constraint $f(x) = \sum_{k=1}^{3} \nu_k f_k(x)$ implies that $\nu_k f_k(x) \leq 1$, k = 1, 2, 3, $x \in [0, 1]$. By Lemma 1 below this implies that the variance of $f_k(x)$, k = 1, 2, 3 is bounded from below by $\nu_k^2/12$.

As an example consider the setting of Section 5.2 where $(\alpha_1, \alpha_2, \alpha_3) = (0.0, -0.1, -1.2), (\beta_1, \beta_2, \beta_3) = (0.2, 0.5, 2.0)$, and the residual variances σ_k^2 , k = 1, 2, 3 are equal across treatment groups. Figure 2 shows the resulting scenario and optimal allocation probabilities (whose computation is explained below). Then, numerically (see details below) minimizing (16) in ν_k , $\mu_k, \tau_k, k = 1, \ldots, K$ subject to the constraints (17) as well as $\tau_k^2 \ge \nu_k^2/12$, k = 1, 2, 3 gives $\nu_k = 0.346, 0.444, 0.210, \mu_k = 0.342, 0.512, 0.735, \tau_k^2 = 0.00997, 0.132, 0.00369$, respectively, and the minimized value of (16) is 12.128. At this point we know it is a lower bound to (16) since only part of the constraints were applied. However, we now demonstrate a design which achieves this lower bound.

We note that $\tau_k^2 = \nu_k^2/12$, for k = 1, 3. Applying Lemma 1 for $c = 1/\nu_k$ for k = 1, 3, the only density $f_k(x)$ with mean μ_k and variance τ_k^2 and satisfying the constraint $f_k(x) \leq 1/\nu_k$, $x \in [0, 1]$ is the uniform distributions on $[\mu_k - \nu_k/2, \mu_k + \nu_k/2]$ with densities $f_k(x) = 1_{[\mu_k - \nu_k/2, \mu_k + \nu_k/2]}(x)/\nu_k$. Furthermore, we set

$$f_2(x) = \frac{1 - \nu_1 f_1(x) - \nu_3 f_3(x)}{\nu_2}$$

Because the supports of f_1 , f_3 are disjoint and $f_k(x) \leq 1/\nu_k$, k = 1, 3, $f_2(x)$ is a valid density. By construction, the mean and variance of $f_2(x)$ are μ_2, τ_2^2 and the three densities satisfy $f(x) = \sum_{k=1}^3 \nu_k f_k(x)$. The resulting allocation probabilities $\pi_k(x) = f_k(x)/\nu_k$ achieve the lower bound computed above, and hence, assuming the above minimization algorithm yielded the correct minimum (which we verified in several ways, including the fact that it is unique, see below), it is an optimal allocation (see Figure 2 for a plot). It follows that in this case the above constraints suffice. Furthermore, as f_1, f_3 are unique, they are a unique optimal allocation (under the assumption that the minimization in $\nu_k, \mu_k, \tau_k, k = 1, \ldots, K$ gives a unique solution). It also follows that the optimal solution leads to a deterministic allocation.

Some comments: (1) To implement the numerical optimization of (16) subject to the constraints, we



Figure 2: The functions g_k for the three treatments (left) and the asymptotically optimal allocation probabilities minimizing (16) (as a stacked area chart, right) for the example in Section 5.3. For instance, for x's under the black area, treatment 1 is assigned with probability $\pi_1(x) = 1$.

use the R command optim with the following reparametrization, which allows one to apply optimization algorithms for unrestricted optimization. Let $A_i = (a_{i,1}, a_{i,2}) \in \mathbb{R}^2$, i = 1, 2, 3. First we transform the vectors A_i into proportions setting $a'_{ik} = \exp(a_{i,k}) / \sum_{m=1}^3 \exp(a_{i,m})$, i, k = 1, 2, 3 where $a_{i,3} \equiv 0$, and set $\nu_k = a'_{1k}, \ \mu_k = a'_{2k} \mu / \nu_k$, and

$$\tau_k^2 = \frac{\nu_k^2}{12} + \frac{a'_{3,k} \left(\mu^2 + \tau^2 - \sum_{m=1}^3 (\nu_m \mu_m^2 + \nu_m^3 / 12)\right)}{\nu_k}, \ k = 1, 2, 3.$$

As long as the second summand remains positive, this parametrization ensures that $\tau_k^2 \ge \nu_k^2/12$ and that the constraints of (17) hold. To confirm the robustness of the solution, the optimization was performed repeatedly with randomly chosen initial conditions. (2) If the intersection points $\theta_{k,k+1}$ of the lines g_k are moved closer together by a different specification of the parameters, so will the intervals of x's where all patients are allocated to Treatment 1 or Treatment 3 (the blue and black regions in Figure 2). For scenarios where they overlap, we can no longer use the above approach to obtain an optimal solution and the solutions seem to become more complex. This is the case, for example, for $\alpha_k = 0.0, -0.1, -.8, \beta_k =$ 0.2, 0.5, 2.0, k = 1, 2, 3.

Lemma 1 Let c > 0. Among all continuous distributions f with (i) mean μ such that (ii) $f(x) \leq c$ for all $x \in \mathbb{R}$, the uniform distribution on $[\mu - 1/(2c), \mu + 1/(2c)]$ has the smallest variance. This minimum variance is $1/(c^2 12)$.

6 A clinical trial to determine personalized diets

To demonstrate our approach, we consider a study by Ebbeling et al. (2018) in which three diets were compared in a parallel group design. We fitted a model to the data of that trial and derived an alternative optimal design based on the estimated model parameters.

Subjects were randomly assigned to three diet groups. The diets differ in their carbohydrate content, high (60%), moderate (40%) or low (20%) and were followed for 20 weeks. The total number of subjects allocated to these three treatment groups was 54, 53 and 57, respectively, so the design is (almost) balanced. The main covariate is insulin secretion (insulin concentration 30 minutes after oral glucose). The primary outcome in the original trial was averaged total energy expenditure over two measurements, in the middle and the end of the trial. Total energy expenditure is measured by doubly labeled water; for details on this measure see Hills et al. (2014).

Of the 1685 subjects initially screened, 234 participated in a run-in phase which preceded the trial itself. Of these, 164 achieved a target of 12% ($\pm 2\%$) weight loss and were qualified for the trial; they were randomly assigned to one of the three diets.

To account for the cost of the diets, either in terms of money or in terms of inconvenience to the dieter, we consider the utility of the diets as an outcome variable which is defined as the energy expenditure minus the diet cost. We assume that the lower the carbon content of the diet, the higher the cost (see Hagberg et al. (2019) for a similar definition) but otherwise our choice of the cost for this example is arbitrary. On the basis of Figure 4 in Ebbeling et al. (2018), we assume that under treatment T_k , k = 1, 2, 3,

utility =
$$\alpha_k + \beta_k X + \varepsilon_k - \cos t_k$$
, (18)

where $\cos t_k = 0$, 150, 300 for k = 1, 2, 3, respectively, X is insulin secretion and its distribution is Gamma(3.12, 0.02), and the standard deviations of $\varepsilon_1, \varepsilon_2, \varepsilon_3$ are 190, 150, 130, respectively. The design we propose is based on the assumed values of α_k , which are 40, -80, -240, and of β_k , which are -0.8, 0.1, and 0.8 for k = 1, 2, 3, respectively. These values and the parameters of the gamma distribution of X and the standard deviation of ε 's are all based on Figure 4 of Ebbeling et al. (2018), which plots the energy expenditure as a function of the insulin secretion for the three diets in the data, and roughly shows the distribution of insulin secretion. In particular, the parameters of the Gamma distribution are chosen to match empirical quantiles of the X's.

In order to numerically find the optimal allocation probabilities $\pi_k(x)$ that minimize the ideal regret (14), we use the approach outlined in subsection 5.2 setting M = 4 and using the R function optim with the Nelder-Mead method.

Figure 3(a) plots the regression lines (18) for the three diets and the assumed distribution of the covariate x. The high, moderate, and low carbohydrate diets are optimal when x belongs to the intervals $(0, 133), (133, 229), (229, \infty)$, respectively. The difference between the treatments is more pronounced for extreme values of x. Hence, people with very small or very large x would benefit more from a personalized treatment choice compared to those with medium values of x.

Figures 3(b)-(d) present the optimal allocation probabilities, as a stacked area chart, when $n = 164, 1000, \text{ and } \infty$. When the sample size is larger, the optimal allocation rule is closer to being deterministic. For n = 164, which was the actual trial size, the optimal trial allocation rule tends to allocate subjects to T_3 (respectively, T_2) for large (respectively, small) values of x. Compared to equal allocation, which was the actual design used, there is a reduction of about 11% of the regret. A saving of 11% may be quite significant if the treatments are applied to a large number of future patients. Furthermore, the regret of the balanced design with n = 164 can be achieved under the optimal design with n = 143, which represents a reduction of 13% of the sample size. Taking into account that in this experiment only about 10% of recruited patients passed the screening process, a saving of about 20 subjects amounts to a reduction of about 200 subjects that would have to be recruited and pretested. Also, the duration of 20 weeks, and the cost involved in the experiment, imply that an optimal design is desirable.

To investigate if the reduction in the regret is due to unbalanced allocation or the covariate dependent allocation, we also computed the regret for the optimal fixed allocation that is restricted to allocation probabilities that do not depend on the covariate. This can be achieved as above by setting M = 0. We found numerically the optimal fixed design for n = 164 leads to reduction of the regret by about 8%, whereas optimizing with allocation functions that depend on the covariate, leads to a reduction by an additional 3%. For larger sample sizes, it appears that the gain by optimal covariate dependent allocation probabilities increases (see Table 2). As the algorithm only gives an approximate optimal solution, we also used the approach in Section 5.3 to compute a lower bound for the limit nR_I for $n \to \infty$. The obtained lower bound of 735.1 is close to 738.2, the corresponding regret obtained by the derived allocation function depicted in Figure 3d.

It should be noticed that all these designs are based on the true values of the parameters and therefore require a preliminary study, or an adaptive design.

		% rec	luction	fixed allocation $(M = 0)$				
n	nR_I	M=4	M = 0	high	moderate	low		
164	842.8	10.9	7.9	0.45	0.37	0.18		
1000	829.6	13.1	4.4	0.37	0.40	0.23		
∞	738.2	17.1	3.2	0.35	0.40	0.25		

Table 2: Diets example: nR_I , the percentage of reduction achieved by the optimized design with M = 0 and M = 4 relative to a balanced design, and the fixed allocation probabilities for M = 0.



Figure 3: Figure (a) plots the regression models (18) and the assumed density of the covariate. Figures (b)-(d) present the optimal allocation probabilities (as a stacked area chart) when $n = 164, 1000, \infty$.

7 Discussion

In this paper we derived optimal trial designs to determine treatment rules that are based on patients' individual covariates, assuming that the outcomes in each treatment arm are given by a linear model. For two treatments the optimal allocation rule depends only on the variances of the error terms in the linear models and does not depend on other model parameters, nor on the covariates. In contrast, for three or more treatments the optimal rule depends on all the linear model parameters, and the covariates. In the latter case one faces locally optimal designs whose computation requires prior knowledge on the model parameters. If such information is not available, extensions of the framework can be considered. For example, sequential designs that start with equal allocation and then compute optimal allocation rules for further stages based on interim estimates of the model parameters can be used. Another option are Bayesian designs, where a prior distribution for the model parameters is specified and the regret, averaged over this prior is minimized. Furthermore, the sequential and Bayesian approaches can be combined by computing posterior distributions of the model parameters in interim analyses and optimizing the allocation ratios based on the regret averaged over these posterior distributions. Alternatively, a minimax approach can be used and the maximal regret across a range of plausible scenarios can be minimized. The (approximate) optimal minimax design was computed in Section 4.1 for two treatments. The case of more than two treatment is deferred for future work.

A trial of the kind we study ends with recommendations on treatments, amounting to a claim of causality. Here we allow the allocation to treatments to be a function of the covariate vector \mathbf{X} only, and in particular there are no other confounding variables which affects both the allocation and the response, and causality can be deduced; see, e.g., Hernán and Robins (2020) and Stock and Watson (2012). This remains true even if the allocation functions are "deterministic", taking only the values 0 and 1. We assume that a linear model holds true (see Section 3) rather than taking the approach that "all models are wrong." This allows us to make consistent estimation of the model parameters for any allocation, barring a degenerate design which concentrates on a single \mathbf{X} . Assuming a model is necessary if one wants to avoid a separate trial for each level of \mathbf{X} , which seems inefficient. A linear model assumption may be reasonable if the range of values of \mathbf{X} is restricted, which is often very natural. As usual, linear models include a polynomial regression of some degree rather than a simple linear model (see Section 4.3), and interaction terms may be included.

References

- Bastani, H. and M. Bayati (2015). Online decision-making with high-dimensional covariates. Available at SSRN 2661896.
- Chen, L. H., L. Goldstein, and Q.-M. Shao (2010). Normal Approximation by Stein's Method. Springer Science & Business Media.
- Chernoff, H. (1953). Locally optimal designs for estimating parameters. Annals of Mathematical Statistics 24 (4), 586–602.

- Ebbeling, C. B., H. A. Feldman, G. L. Klein, J. M. W. Wong, L. Bielak, S. K. Steltz, P. K. Luoto, R. R. Wolfe, W. W. Wong, and D. S. Ludwig (2018). Effects of a low carbohydrate diet on energy expenditure during weight loss maintenance: randomized trial. *BMJ 363*.
- Goldenshluger, A. and A. Zeevi (2013). A linear response bandit problem. *Stochastic Systems* 3(1), 230–261.
- Hagberg, L., A. Winkvist, H. K. Brekke, F. Bertz, E. H. Johansson, and E. Huseinovic (2019). Costeffectiveness and quality of life of a diet intervention postpartum: 2-year results from a randomized controlled trial. *BMC Public Health* 19(1), 38.
- Hansen, B. (2019). Econometrics. University of Wisconsin: Online.
- Hernán, M. A. and J. M. Robins (2020). Causal Inference: What If. CRC Boca Raton, FL.
- Hills, A. P., N. Mokhtar, and N. M. Byrne (2014). Assessment of physical activity and energy expenditure: an overview of objective measures. *Frontiers in Nutrition* 1, 5.
- Kosorok, M. R. and E. B. Laber (2019). Precision medicine. Annual review of statistics and its application 6, 263–286.
- Kosorok, M. R., E. B. Laber, D. S. Small, and D. Zeng (2021). Introduction to the theory and methods special issue on precision medicine and individualized policy discovery. *Journal of the American Statistical Association* 116(533), 159–161.
- Kosorok, M. R. and E. E. Moodie (2015). Adaptive Treatment Strategies in Practice: Planning Trials and Analyzing Data for Personalized Medicine. SIAM.
- Lee, K. M. and J. Wason (2019). Design of experiments for a confirmatory trial of precision medicine. Journal of Statistical Planning and Inference 199, 179–187.
- Minsker, S., Y.-Q. Zhao, and G. Cheng (2016). Active clinical trials for personalized medicine. Journal of the American Statistical Association 111 (514), 875–887.
- Moore, M. (1973). A convex matrix function. The American Mathematical Monthly 80(4), 408–409.
- Qian, M. and S. A. Murphy (2011). Performance guarantees for individualized treatment rules. The Annals of Statistics 39(2), 1180 – 1210.
- Silvey, S. (2013). Optimal Design: an Introduction to the Theory for Parameter Estimation, Volume 1. Springer Science & Business Media.

Stock, J. H. and M. W. Watson (2012). Introduction to Econometrics, Volume 3. Pearson New York.

- Sverdlov, O. and W. F. Rosenberger (2013). On recent advances in optimal allocation designs in clinical trials. Journal of Statistical Theory and Practice 7(4), 753–773.
- Tropp, J. A. (2012). User-friendly tail bounds for sums of random matrices. Foundations of Computational Mathematics 12(4), 389–434.
- Zhang, Q., A. Khademi, and Y. Song (2021+). Min-max optimal design of two-armed trials with side information. INFORMS Journal on Computing, In press.
- Zhao, Y.-Q. and M. L. LeBlanc (2020). Designing precision medicine trials to yield a greater population impact. *Biometrics* 76(2), 643–653.

8 Appendix: Proofs

Proof of Lemma 4.2

We use the convexity relation $\lambda A^{-1} + (1 - \lambda)B^{-1} \succeq [\lambda A + (1 - \lambda)B]^{-1}$ (Moore, 1973) where A and B are positive definite matrices. We have

$$\begin{split} \boldsymbol{\Sigma}_{1} + \boldsymbol{\Sigma}_{2} &= \sigma_{1} \left(\frac{\nu_{1} \boldsymbol{Q}_{1}}{\sigma_{1}} \right)^{-1} + \sigma_{2} \left(\frac{\nu_{2} \boldsymbol{Q}_{2}}{\sigma_{2}} \right)^{-1} = (\sigma_{1} + \sigma_{2}) \left[\frac{\sigma_{1}}{\sigma_{1} + \sigma_{2}} \left(\frac{\nu_{1} \boldsymbol{Q}_{1}}{\sigma_{1}} \right)^{-1} + \frac{\sigma_{2}}{\sigma_{1} + \sigma_{2}} \left(\frac{\nu_{2} \boldsymbol{Q}_{2}}{\sigma_{2}} \right)^{-1} \right] \\ & \succeq (\sigma_{1} + \sigma_{2}) \left[\frac{\sigma_{1}}{\sigma_{1} + \sigma_{2}} \left(\frac{\nu_{1} \boldsymbol{Q}_{1}}{\sigma_{1}} \right) + \frac{\sigma_{2}}{\sigma_{1} + \sigma_{2}} \left(\frac{\nu_{2} \boldsymbol{Q}_{2}}{\sigma_{2}} \right) \right]^{-1} = (\sigma_{1} + \sigma_{2})^{2} \boldsymbol{Q}^{-1}, \end{split}$$

where in the last equality we used the fact that $\nu_1 Q_1 + \nu_2 Q_2 = Q$. Equality holds when the two matrices to the right of the \succeq sign are equal, which happens when $\pi_1(\mathbf{x}) = \frac{\sigma_1}{\sigma_1 + \sigma_2}$ and $\pi_2(\mathbf{x}) = \frac{\sigma_2}{\sigma_1 + \sigma_2}$.

Proof of Theorem 4.3

We start with the case of p = 1. It is enough to show (as the other part of the integral, from θ to 1, is symmetric) that

$$n^{3/2-\varepsilon} \int_0^\theta \left| P(\hat{g}_2(x) > \hat{g}_1(x)) - \Phi\left(\frac{-\sqrt{n}[g_1(x) - g_2(x)]}{\sqrt{V(x)}}\right) \right| [g_1(x) - g_2(x)]f(x)dx \to 0.$$
(19)

The idea of the proof is quite simple. For $x \in \theta \pm 1/\sqrt{n}$ the quantity in absolute value converges to zero at rate $1/\sqrt{n}$, and both the range of the integral and $g_1(x) - g_2(x)$ are of order $1/\sqrt{n}$ each. For other x's both the probability and Φ in the integrand are exponentially small in n.

We will prove (19) by showing that

$$n^{3/2-\varepsilon} \int_0^\theta \left| P(\hat{g}_2(x) > \hat{g}_1(x)) - E\Phi\left(\frac{-\sqrt{n}(\beta_2 - \beta_1)(\theta - x)}{S(x)}\right) \right| [g_1(x) - g_2(x)]f(x)dx \to 0 \text{ and}$$
(20)

$$n^{3/2-\varepsilon} \int_0^\theta \left| E\Phi\left(\frac{-\sqrt{n}(\beta_2 - \beta_1)(\theta - x)}{S(x)}\right) - \Phi\left(\frac{-\sqrt{n}[g_1(x) - g_2(x)]}{\sqrt{V(x)}}\right) \right| [g_1(x) - g_2(x)]f(x)dx \to 0, \quad (21)$$

where

$$S^{2}(x) = n(1,x) \left\{ \sigma_{1}^{2} \left(\sum_{i \in T_{1}} {1 \choose X_{i}} (1,X_{i}) \right)^{-1} + \sigma_{2}^{2} \left(\sum_{i \in T_{2}} {1 \choose X_{i}} (1,X_{i}) \right)^{-1} \right\} {1 \choose x},$$
(22)

and the expectation in (20) and (21) is with respect to X_1, \ldots, X_n which appear in $S^2(x)$ and the random sets T_k . We start with (21). Since θ is the intersection point, $\alpha_1 + \beta_1 \theta = \alpha_2 + \beta_2 \theta$; therefore,

 $g_1(x) - g_2(x) = \alpha_1 + \beta_1 x - (\alpha_2 + \beta_2 x) = (\beta_2 - \beta_1)(\theta - x). \text{ Hence, (21) can be written as}$ $n^{3/2-\varepsilon} \int_0^\theta \left| E\Phi\left(\frac{-\sqrt{n}(\beta_2 - \beta_1)(\theta - x)}{S(x)}\right) - \Phi\left(\frac{-\sqrt{n}(\beta_2 - \beta_1)(\theta - x)}{\sqrt{V(x)}}\right) \right| (\beta_2 - \beta_1)(\theta - x)f(x)dx \to 0.$

Fix $c_n = n^{\varepsilon/5}$. We will show that

$$n^{3/2-\varepsilon} \int_0^{\theta-c_n/\sqrt{n}} \left| E\Phi\left(\frac{-\sqrt{n}(\beta_2-\beta_1)(\theta-x)}{S(x)}\right) - \Phi\left(\frac{-\sqrt{n}(\beta_2-\beta_1)(\theta-x)}{\sqrt{V(x)}}\right) \right| (\beta_2-\beta_1)(\theta-x)f(x)dx \to 0,$$
(23)

and that

$$n^{3/2-\varepsilon} \int_{\theta-c_n/\sqrt{n}}^{\theta} \left| E\Phi\left(\frac{-\sqrt{n}(\beta_2-\beta_1)(\theta-x)}{S(x)}\right) - \Phi\left(\frac{-\sqrt{n}(\beta_2-\beta_1)(\theta-x)}{\sqrt{V(x)}}\right) \right| (\beta_2-\beta_1)(\theta-x)f(x)dx \to 0.$$

$$(24)$$

We show (23) by arguing that the probabilities in (23) are exponentially small. With notation defined in Section 3 we have that

$$V(x) \le (1+x^2) \left(\frac{\sigma_1^2}{\nu_1 \lambda_{min}(\boldsymbol{Q}_1)} + \frac{\sigma_2^2}{\nu_2 \lambda_{min}(\boldsymbol{Q}_2)} \right) \le C,$$

for a constant C > 0 (since x is bounded). With the bound $\Phi(-t) \le \exp(-t^2/2)$ for t > 1 and $(\theta - x)\sqrt{n} \le c_n$ we have

$$\Phi\left(\frac{-(\beta_2-\beta_1)(\theta-x)\sqrt{n}}{\sqrt{V(x)}}\right) \le \exp\left(-\frac{(\beta_2-\beta_1)^2 c_n^2}{2C}\right),$$

which shows that the second normal probability $\Phi(\cdot)$ in (23) is exponentially small. We now argue that the expectation in (23) is also exponentially small. Consider the event

$$A_{n} := \left\{ \lambda_{min} \left(\frac{1}{n} \sum_{i \in T_{1}} {\binom{1}{X_{i}}} (1, X_{i}) \right) \le \nu_{1} \frac{\lambda_{min}(\boldsymbol{Q}_{1})}{2} \right\}$$
$$\bigcup \left\{ \lambda_{min} \left(\frac{1}{n} \sum_{i \in T_{2}} {\binom{1}{X_{i}}} (1, X_{i}) \right) \le \nu_{2} \frac{\lambda_{min}(\boldsymbol{Q}_{2})}{2} \right\}. \quad (25)$$

Writing $\frac{1}{n} \sum_{i \in T_1} {1 \choose X_i} (1, X_i) = \frac{1}{n} \sum_{i=1}^n {1 \choose X_i} (1, X_i) I(i \in T_1)$, a sum of iid terms, we can apply Theorem 5.1 of Tropp (2012) and the union bound to conclude that $P(A_n) \leq C \exp(-Cn)$ for some constant C > 0.

Going back to (23), write

$$E\Phi\left(\frac{-\sqrt{n}(\beta_2-\beta_1)(\theta-x)}{S(x)}\right)$$
$$= E\Phi\left(\frac{-\sqrt{n}(\beta_2-\beta_1)(\theta-x)}{S(x)}\right)I(A_n) + E\Phi\left(\frac{-\sqrt{n}(\beta_2-\beta_1)(\theta-x)}{S(x)}\right)I(A_n^c)$$
$$\leq P(A_n) + E\Phi\left(\frac{-\sqrt{n}(\beta_2-\beta_1)(\theta-x)}{S(x)}\right)I(A_n^c).$$

When A_n^c occurs, the denominator $S^2(x)$ is bounded and the same argument as above shows that this probability is exponentially small. We conclude that both probabilities in (23) are exponentially small, which implies that (23) is correct.

We now prove (24). By the mean value theorem, for $x \in (\theta - c_n/\sqrt{n}, \theta)$ and for some $V^*(x)$ between $S^2(x)$ and V(x), we have for some C > 0

$$\begin{aligned} \left| \Phi\left(\frac{-\sqrt{n}(\beta_2 - \beta_1)(\theta - x)}{\sqrt{S^2(x)}}\right) - \Phi\left(\frac{-\sqrt{n}(\beta_2 - \beta_1)(\theta - x)}{\sqrt{V(x)}}\right) \right| \\ &= |S^2(x) - V(x)|\varphi\left(\frac{-\sqrt{n}(\beta_2 - \beta_1)(\theta - x)}{\sqrt{V^*(x)}}\right) \frac{\sqrt{n}(\beta_2 - \beta_1)(\theta - x)}{2}(V^*(x))^{-3/2} \\ &\leq C|S^2(x) - V(x)|c_n(\min\{S^2(x), V(x)\})^{-3/2}, \text{ and} \end{aligned}$$

$$\{S^{2}(x) - V(x)\} = (1, x) \left\{ \frac{\sigma_{1}^{2}}{\nu_{1}} \left(\mathbf{M}_{1}^{-1} - \boldsymbol{Q}_{1}^{-1} \right) + \frac{\sigma_{2}^{2}}{\nu_{2}} \left(\mathbf{M}_{2}^{-1} - \boldsymbol{Q}_{2}^{-1} \right) \right\} \begin{pmatrix} 1\\ x \end{pmatrix},$$
(26)

where $\mathbf{M}_k = \frac{1}{n\nu_k} \sum_{i \in T_k} {\binom{1}{X_i}}(1, X_i), k = 1, 2$. We have that $\mathbf{M}_1^{-1} - \mathbf{Q}_1^{-1} = \mathbf{M}_1^{-1} (\mathbf{Q}_1 - \mathbf{M}_1) \mathbf{Q}_1^{-1}$, implying that the difference of the inverses is small when the difference of the matrices is small and the inverse is bounded. Define the event

$$\widetilde{A}_n := \{ \max\{ \|\boldsymbol{Q}_1 - \mathbf{M}_1\|_{\infty}, \|\boldsymbol{Q}_2 - \mathbf{M}_2\|_{\infty} \} \ge c_n/\sqrt{n} \},\$$

where $\|\cdot\|_{\infty}$ denotes the maximum (entry-wise) norm. By Hoeffding's inequality (since the random variables are bounded), $P(\tilde{A}_n) \leq C \exp(-Cc_n^2)$ (which is exponentially small). Therefore, as before, this event can be ignored. On the complement of \tilde{A}_n , \mathbf{M}_1^{-1} , \mathbf{M}_2^{-1} , $S^2(x)$ are all bounded, and by similar arguments $S^2(x)$ and V(x) are bounded below. Hence, when \tilde{A}_n^c occurs,

$$\left|\Phi\left(\frac{-\sqrt{n}(\beta_2-\beta_1)(\theta-x)}{S(x)}\right) - \Phi\left(\frac{-\sqrt{n}(\beta_2-\beta_1)(\theta-x)}{\sqrt{V(x)}}\right)\right| \le C\frac{c_n^2}{\sqrt{n}}$$

Therefore the integral in (24) is bounded by

$$Cn^{1-\varepsilon}c_n^2 \int_{\theta-c_n/\sqrt{n}}^{\theta} (\theta-x)f(x)dx \le Cn^{1/2-\varepsilon}c_n^3 \int_{\theta-c_n/\sqrt{n}}^{\theta} f(x)dx \le Cc_n^4/n^{\varepsilon} = n^{-\varepsilon/5} \to 0.$$

which completes the proof of (24) and hence of (21).

We now show (20) by dividing the integral into two ranges: (a) $[\theta - c_n/\sqrt{n}, \theta]$, and (b) $[0, \theta - c_n/\sqrt{n}]$. We have that

$$P\left(\widehat{g}_{2}(x) > \widehat{g}_{1}(x)\right) = P\left(\left[\left(\frac{\widehat{\alpha}_{1} - \alpha_{1}}{\widehat{\beta}_{1} - \beta_{1}}\right) - \left(\frac{\widehat{\alpha}_{2} - \alpha_{2}}{\widehat{\beta}_{2} - \beta_{2}}\right)\right]^{t} \begin{pmatrix}1\\x\end{pmatrix} > (\beta_{2} - \beta_{1})(x - \theta)\right)$$
$$= P\left(\left[\frac{\mathbf{M}_{1}^{-1}}{n\nu_{1}}\sum_{i\in T_{1}} \begin{pmatrix}\varepsilon_{i}\\X_{i}\varepsilon_{i}\end{pmatrix} - \frac{\mathbf{M}_{2}^{-1}}{n\nu_{2}}\sum_{i\in T_{2}} \begin{pmatrix}\varepsilon_{i}\\X_{i}\varepsilon_{i}\end{pmatrix}\right]^{t} \begin{pmatrix}1\\x\end{pmatrix} > (\beta_{2} - \beta_{1})(x - \theta)\right)$$
$$= P\left(\sum_{i=1}^{n} a_{i}\varepsilon_{i} > (\beta_{2} - \beta_{1})(x - \theta)\right),$$

where $a_i = \begin{cases} (1, X_i) \frac{\mathbf{M}_1^{-1}}{n\nu_1} \begin{pmatrix} 1 \\ x \end{pmatrix} & i \in T_1 \\ -(1, X_i) \frac{\mathbf{M}_2^{-1}}{n\nu_2} \begin{pmatrix} 1 \\ x \end{pmatrix} & i \in T_2 \end{cases}$, $\mathbf{M}_1, \mathbf{M}_2$ are defined in (26) and ε_i are the regression error terms. Notice that $Var(\varepsilon_i)$ is equal to σ_1^2 if $i \in T_1$ and to σ_2^2 if $i \in T_2$. We are going to condition on X_1, \ldots, X_n obtained in the experiment and their allocation, denoted together by D. Then $Var(\sum_{i=1}^n a_i \varepsilon_i | D) = S^2(x)/n$, where $S^2(x)$ is defined in (22).

We aim to apply the Berry Esseen Theorem, for which we need some bounds. The values na_i are bounded below since the matrices \mathbf{M}_k are bounded above. In order to obtain upper bounds, recall the event A_n defined in (25). On the set A_n^c we have that na_i are bounded above. Since $P(A_n)$ is exponentially small, we can ignore it, and assume that na_i are bounded with probability one. We can apply the Berry-Esseen Theorem for the non-identically distributed case (for a convenient reference see Chen et al. (2010) (3.27)) to $\sum_{i=1}^n \frac{\sqrt{na_i}}{S(x)} \varepsilon_i \equiv \sum_{i=1}^n V_i$ where $\sum_{i=1}^n Var(V_i|D) = 1$ and $E(|V_i|^3 | D) \leq$ $max\{na_i/S(x)E(|\varepsilon_i|^3|D): 1 \leq i \leq n\}n^{-3/2}$ to obtain by unconditioning on D

$$\left| P\left(\sum_{i=1}^{n} a_i \varepsilon_i > (\beta_2 - \beta_1)(x - \theta)\right) - E\Phi\left(\frac{-\sqrt{n}(\beta_2 - \beta_1)(\theta - x)}{S(x)}\right) \right| \le C/\sqrt{n}$$

Therefore,

$$n^{3/2-\varepsilon} \int_{\theta-c_n/\sqrt{n}}^{\theta} \left| P(\hat{g}_2(x) > \hat{g}_1(x)) - E\Phi\left(\frac{-\sqrt{n}(\beta_2 - \beta_1)(\theta - x)}{S(x)}\right) \right| (\beta_2 - \beta_1)(\theta - x)f(x)dx$$

$$\leq n^{3/2-\varepsilon} \frac{C}{\sqrt{n}} \int_{\theta-c_n/\sqrt{n}}^{\theta} (\beta_2 - \beta_1)(\theta - x)f(x)dx \leq n^{3/2-\varepsilon} \frac{C}{\sqrt{n}} \cdot \frac{c_n}{\sqrt{n}} \int_{\theta-c_n/\sqrt{n}}^{\theta} f(x)dx$$

$$\leq C \frac{c_n^2}{n^{\varepsilon}} = C \frac{n^{2\varepsilon/5}}{n^{\varepsilon}} \to 0. \quad (27)$$

By a standard large deviation bound, the integral of (20) from 0 to $\theta - c_n/\sqrt{n}$ multiplied by $n^{3/2-\varepsilon}$ is easily shown to be of order $n^{3/2-\varepsilon}O(exp(-c_n\sqrt{n})) \to 0$. This proves the theorem for covariates in \mathbb{R} (p=1).

We now consider the case of covariates in \mathbb{R}^p for $p \geq 2$. Suppose, without loss of generality, that $\beta_{2,1} > \beta_{1,1}$. For given $\mathbf{x}_{-1} = (x_2, \ldots, x_p)$ define $\theta_1(\mathbf{x}_{-1}) := \frac{\alpha_1 - \alpha_2 + (\beta_{1,-1} - \beta_{2,-1})^t \mathbf{x}_{-1}}{\beta_{2,1} - \beta_{1,1}}$. For \mathbf{x} such that $\theta_1(\mathbf{x}_{-1}) \notin [0,1]$ the probability of making a mistake (i.e., $\hat{g}_2(\mathbf{x}) - \hat{g}_1(\mathbf{x})$ has the wrong sign) is exponentially small. Therefore,

$$\begin{aligned} R(\pi_1, \pi_2) &= \int_{[0,1]^{p-1}} \int_0^{\theta_1} P(\widehat{g}_2(\mathbf{x}) > \widehat{g}_1(\mathbf{x})) (g_1(\mathbf{x}) - g_2(\mathbf{x})) f(\mathbf{x}) dx_1 I(\theta_1 \in [0,1]) d\mathbf{x}_{-1} \\ &+ \int_{[0,1]^{p-1}} \int_{\theta_1}^1 P(\widehat{g}_2(\mathbf{x}) > \widehat{g}_1(\mathbf{x})) (g_1(\mathbf{x}) - g_2(\mathbf{x})) f(\mathbf{x}) dx_1 I(\theta_1 \in [0,1]) d\mathbf{x}_{-1} + a_n, \end{aligned}$$

where a_n is exponentially small. By a slight variation on the above one-dimensional case, the inner integral equals

$$\int_{0}^{1} \Phi\left(\frac{-\sqrt{n}|g_{2}(\mathbf{x}) - g_{1}(\mathbf{x})|}{\sqrt{V(\mathbf{x})}}\right) |g_{2}(\mathbf{x}) - g_{1}(\mathbf{x})|f(\mathbf{x})dx_{1}I(\theta_{1} \in [0,1]) + o(1/n^{3/2-\varepsilon}),$$
(28)

where the error term o is uniform in \mathbf{x}_{-1} . Taking the outer integral, the theorem follows.

Proof of Theorem 4.4

As computed two lines below (22), $g_1(x) - g_2(x) = (\beta_2 - \beta_1)(\theta - x)$. We will show that for continuous t

$$\lim_{t \to \infty} t \int_0^\theta \Phi\left(\frac{-\sqrt{t}(\beta_2 - \beta_1)(\theta - x)}{\sqrt{V(x)}}\right) (\beta_2 - \beta_1)(\theta - x)f(x)dx = \frac{V(\theta)f(\theta)}{4(\beta_2 - \beta_1)}.$$
(29)

The integral from θ to 1 is similar, yielding the desired limit. By L'Hôpital's rule the limit in (29) equals

$$\lim_{t \to \infty} \frac{t^2}{2} \int_0^\theta \varphi\left(\frac{\sqrt{t}(\beta_2 - \beta_1)(\theta - x)}{\sqrt{V(x)}}\right) \frac{(\beta_2 - \beta_1)(\theta - x)}{\sqrt{V(x)}} \frac{1}{\sqrt{t}} (\beta_2 - \beta_1)(\theta - x)f(x)dx$$

Substitution of x by $y = (\theta - x)\sqrt{t}$ in the integral yields

$$\lim_{t \to \infty} \frac{1}{2} \int_0^{\theta \sqrt{t}} \varphi \left(\frac{(\beta_2 - \beta_1)y}{\sqrt{V(\theta - y/\sqrt{t})}} \right) \frac{(\beta_2 - \beta_1)^2}{\sqrt{V(\theta - y/\sqrt{t})}} y^2 f(\theta - y/\sqrt{t}) dy.$$

The limit is

$$\frac{1}{2} \int_0^\infty \varphi\left(\frac{(\beta_2 - \beta_1)y}{\sqrt{V(\theta)}}\right) \frac{(\beta_2 - \beta_1)^2}{\sqrt{V(\theta)}} y^2 f(\theta) dy.$$

Substitution of y with $z = \frac{(\beta_2 - \beta_1)y}{\sqrt{V(\theta)}}$ in the integral yields the claimed limit $\frac{f(\theta)V(\theta)}{2(\beta_2 - \beta_1)} \int_0^\infty \varphi(z) z^2 dz = \frac{f(\theta)V(\theta)}{4(\beta_2 - \beta_1)}$.

Proof of Theorem 4.5

Consider the inner integral in (28). By Theorem 4.4, the limit of the inner integral times n is

$$\frac{V(\theta_1, \mathbf{x}_{-1}) f(\theta_1, \mathbf{x}_{-1})}{2(\beta_{2,1} - \beta_{1,1})}$$

By taking the outer integral the result follows.

The limit of the regret for polynomial regression

Suppose that $g_1(x) = \alpha_1 + (x, \dots, x^J)^t \beta_1$ and $g_2(x) = \alpha_2 + (x, \dots, x^J)^t \beta_2$. Let θ be an intersection point, i.e., $g_1(\theta) = g_2(\theta)$. We have

$$g_2(x) - g_1(x) = g_2(x) - g_2(\theta) - [g_1(x) - g_1(\theta)] = \eta^t (\beta_2 - \beta_1),$$

where $\boldsymbol{\eta} := (x - \theta, x^2 - \theta^2, \dots, x^J - \theta^J)^t$. The following proposition is parallel to Theorem 4.4 (in the one-dimensional case) and provides the limit of the regret; see (13). The proof is similar, although some modifications are required.

Proposition 8.1. For a continuous argument s

$$\lim_{s \to \infty} s \int_0^\theta \Phi\left(\frac{-\sqrt{t}(\boldsymbol{\beta}_2 - \boldsymbol{\beta}_1)^t \boldsymbol{\eta}}{\sqrt{V(x)}}\right) (\boldsymbol{\beta}_2 - \boldsymbol{\beta}_1)^t \boldsymbol{\eta} f(x) dx = \frac{f(\theta)V(\theta)}{4|(\boldsymbol{\beta}_2 - \boldsymbol{\beta}_1)^t \boldsymbol{\zeta}|},\tag{30}$$

where $\boldsymbol{\zeta} := (1, 2\theta, ..., J\theta^{J-1})^t$.

Proof. By L'Hôpital's rule the limit in (30) equals

$$\lim_{s \to \infty} \frac{s^2}{2} \int_0^\theta \varphi\left(\frac{\sqrt{t}(\boldsymbol{\beta}_2 - \boldsymbol{\beta}_1)^t \boldsymbol{\eta}}{\sqrt{V(x)}}\right) \frac{(\boldsymbol{\beta}_2 - \boldsymbol{\beta}_1)^t \boldsymbol{\eta}}{\sqrt{V(x)}} \frac{1}{\sqrt{t}} (\boldsymbol{\beta}_2 - \boldsymbol{\beta}_1)^t \boldsymbol{\eta} f(x) dx.$$

Substituting $y = (\theta - x)\sqrt{s}$ we get

$$\sqrt{s}\boldsymbol{\eta} = \sqrt{s} \begin{pmatrix} x-\theta\\ x^2-\theta^2\\ \vdots\\ x^J-\theta^J \end{pmatrix} = \sqrt{s} \begin{pmatrix} x-\theta\\ (x-\theta)(x+\theta)\\ \vdots\\ (x-\theta)(x^{J-1}+x^{J-2}\theta+\ldots+\theta^{J-1}) \end{pmatrix} = y\tilde{\boldsymbol{\eta}},$$

where

$$\begin{split} \widetilde{\boldsymbol{\eta}} &:= (1, x + \theta, \dots, x^{J-1} + x^{J-2}\theta + \dots + \theta^{J-1})^t \\ &= (1, 2\theta - y/\sqrt{s}, \dots, (\theta - y/\sqrt{s})^{J-1} + (\theta - y/\sqrt{s})^{J-2}\theta + \dots + \theta^{J-1})^t. \end{split}$$

The integral reads

$$\lim_{s \to \infty} \frac{1}{2} \int_0^{\theta \sqrt{s}} \varphi \left(y \frac{(\beta_2 - \beta_1)^t \widetilde{\boldsymbol{\eta}}}{\sqrt{V(\theta - y/\sqrt{s})}} \right) \frac{y^2 [(\beta_2 - \beta_1)^t \widetilde{\boldsymbol{\eta}}]^2}{\sqrt{V(\theta - y/\sqrt{s})}} f(\theta - y/\sqrt{t}) dy.$$

The limit is

$$\frac{1}{2} \int_0^\infty \varphi\left(\frac{y(\boldsymbol{\beta}_2 - \boldsymbol{\beta}_1)^t \boldsymbol{\zeta}}{\sqrt{V(\boldsymbol{\theta})}}\right) \frac{[y(\boldsymbol{\beta}_2 - \boldsymbol{\beta}_1)^t \boldsymbol{\zeta}]^2}{\sqrt{V(\boldsymbol{\theta})}} f(\boldsymbol{\theta}) dy.$$

Substitution of y with $z = \frac{y(\beta_2 - \beta_1)^t \zeta}{\sqrt{V(\theta)}}$ yields

$$\frac{f(\theta)V(\theta)}{2|(\boldsymbol{\beta}_2-\boldsymbol{\beta}_1)^t\boldsymbol{\zeta}|}\int_0^\infty \varphi\left(z\right)z^2dz=\frac{f(\theta)V(\theta)}{4|(\boldsymbol{\beta}_2-\boldsymbol{\beta}_1)^t\boldsymbol{\zeta}|},$$

which completes the proof.

For the proofs of Theorems 5.1 and 5.2 we need the lemma below and some notation. For every m and k, let $I_{m,k} := \int_{\theta_{m-1}}^{\theta_m} P\left(\widehat{g}_k(x) > \max_{\ell \neq k} \widehat{g}_\ell(x)\right) \left[g_m(x) - g_k(x)\right] f(x) dx$. We then have $R(\pi_1, \ldots, \pi_K) = \sum_{k=1}^K \sum_{m=1}^K I_{m,k}$.

Lemma 8.2. Under the assumptions in the beginning of Section 5.1, the integral $I_{m,k}$ is exponentially

small for $k \neq m-1, m+1$, and for every $\varepsilon > 0$

$$\lim_{n \to \infty} n^{3/2-\varepsilon} \left| I_{m,m-1} - \int_{\theta_{m-1}}^{\theta_{m-1}+\frac{c_n}{\sqrt{n}}} \Phi\left(\frac{-\sqrt{n}\{g_m(x) - g_{m-1}(x)\}}{\sqrt{V_{m-1}(x)}}\right) \left[g_m(x) - g_{m-1}(x)\right] f(x) dx \right| = 0, \quad (31)$$
$$\lim_{n \to \infty} n^{3/2-\varepsilon} \left| I_{m,m+1} - \int_{\theta_m}^{\theta_m} \Phi\left(\frac{-\sqrt{n}\{g_m(x) - g_{m+1}(x)\}}{\sqrt{V_m(x)}}\right) \left[g_m(x) - g_{m+1}(x)\right] f(x) dx \right| = 0,$$

where $c_n = n^{\varepsilon/5}$ and $V_m(x) = (1, x) \left(\Sigma_m + \Sigma_{m+1} \right) {1 \choose x}$. Moreover, for every m and k

$$n^{3/2-\varepsilon} \left| I_{m,k} - \int_{\theta_{m-1}}^{\theta_m} \left(\int \prod_{\ell=1,\dots,K, l \neq k} \Phi\left(\frac{z\xi_k(x) + \sqrt{n}[g_k(x) - g_\ell(x)]}{\xi_\ell(x)} \right) \varphi(z) dz \right) [g_m(x) - g_k(x)] f(x) dx \right| \to 0,$$

as $n \to \infty$.

Proof of Lemma 8.2

The proof of the first claim in Lemma 8.2 is similar to that of Theorem 4.3. Here is a sketch. For $x \in I_{m,k}$ and bounded away from θ_{m-1} and θ_m , and $k \neq m$ both the quantities

$$P\left(\widehat{g}_k(x) > \max_{\ell \neq k} \widehat{g}_\ell(x)\right) \text{ and } \int \prod_{\ell=1,\dots,K, l \neq k} \Phi\left(\frac{z\xi_k(x) + \sqrt{n}[g_k(x) - g_\ell(x)]}{\xi_\ell(x)}\right) \varphi(z) dz$$

are exponentially small (for k = m the regret is zero). The same holds true for every x and when $k \neq m-1, m+1$.

We now prove (31) concerning $I_{m,m-1}$ (where m > 1). The other relation for $I_{m,m+1}$ is similar. Fix $c_n = n^{\varepsilon/5}$ and consider x such that $x \in (\theta_{m-1}, \theta_{m-1} + c_n/\sqrt{n})$. For such x we have by a standard large deviations argument that

$$P\left(\hat{g}_{m-1}(x) > \max_{\ell \neq m-1} \hat{g}_{\ell}(x)\right) = P(\hat{g}_{m-1}(x) > \hat{g}_{m}(x)) + a_{n}, \text{ and}$$

$$\int \prod_{\ell=1,\dots,K, l \neq m-1} \Phi\left(\frac{z\xi_{m-1}(x) + \sqrt{n}[g_{m-1}(x) - g_{\ell}(x)]}{\xi_{\ell}(x)}\right) \varphi(z)dz = \Phi\left(\frac{-\sqrt{n}[g_{m}(x) - g_{m-1}(x)]}{\sqrt{V_{m-1}(x)}}\right) + b_{n},$$
(32)

where a_n and b_n are exponentially small (uniformly in x). By a Berry-Esseen type bound applied to the first part of the integrand, and the smallness of the other part for x is the given range of the integral, and the smallness of the range itself, as in (27), we obtain

$$\int_{\theta_{m-1}}^{\theta_{m-1}+\frac{c_n}{\sqrt{n}}} \left[P(\widehat{g}_{m-1}(x) > \widehat{g}_m(x)) - \Phi\left(\frac{-\sqrt{n}[g_m(x) - g_{m-1}(x)]}{\sqrt{V_m(x)}}\right) \right] [g_m(x) - g_{m-1}(x)] f(x) dx = o(1/n^{3/2-\varepsilon}),$$

and (31) follows. This, together with (32), implies the last part of the theorem when k = m - 1.

Proof of Theorem 5.1

This follows directly from the last part of Lemma 8.2 and the fact that $R(\pi_1, \ldots, \pi_K) = \sum_{k,m=1}^K I_{m,k}$. Note that we can replace that latter sum by $\sum_{m=1}^K I_{m,m-1}$ since $I_{m,k}$ is exponentially small for $k \neq m-1, m+1$.

Proof of Theorem 5.2

The first part of Lemma 8.2 implies that

$$\sum_{k=1}^{K} \sum_{m=1}^{K} I_{m,k} = \sum_{m=1}^{K-1} \int_{\theta_m - \frac{c_n}{\sqrt{n}}}^{\theta_m + \frac{c_n}{\sqrt{n}}} \Phi\left(\frac{-\sqrt{n}|g_m(x) - g_{m-1}(x)|}{\sqrt{V_{m-1}(x)}}\right) |g_m(x) - g_{m-1}(x)| f(x) dx + o(1/n).$$

The limit of the latter integrals can be calculated using Theorem 4.4 implying the result.