

Article

Compound Optimum Designs for Clinical Trials in Personalized Medicine

Belmiro P. M. Duarte ^{1,2,3,*} , Anthony C. Atkinson ⁴ , David Pedrosa ^{5,6}  and Marlena van Munster ⁵ 

- ¹ Instituto Politécnico de Coimbra, Instituto Superior de Engenharia de Coimbra, Rua Pedro Nunes, 3030-199 Coimbra, Portugal
 - ² INESC Coimbra—Instituto de Engenharia de Sistemas e Computadores de Coimbra, Universidade de Coimbra, Rua Sílvio Lima—Pólo II, 3030-790 Coimbra, Portugal
 - ³ Research Center for Chemical Engineering and Renewable Resources for Sustainability, Universidade de Coimbra, Rua Sílvio Lima—Pólo II, 3030-790 Coimbra, Portugal
 - ⁴ Department of Statistics, London School of Economics, London WC2A 2AE, UK; a.c.atkinson@lse.ac.uk
 - ⁵ Department of Neurology, University Hospital Marburg, 35043 Marburg, Germany; david.pedrosa@staff.uni-marburg.de (D.P.); marlena.vanmunster@uni-marburg.de (M.v.M.)
 - ⁶ Centre of Brain, Mind and Behavior, Philipps-University Marburg, 35043 Marburg, Germany
- * Correspondence: bduarte@isec.pt

Abstract: We consider optimal designs for clinical trials when response variance depends on treatment and covariates are included in the response model. These designs are generalizations of Neyman allocation, and commonly employed in personalized medicine where external covariates linearly affect the response. Very often, these designs aim at maximizing the amount of information gathered but fail to assure ethical requirements. We analyze compound optimal designs that maximize a criterion weighting the amount of information and the reward of allocating the patients to the most effective/least risky treatment. We develop a general representation for static (a priori) allocation and propose a semidefinite programming (SDP) formulation to support their numerical computation. This setup is extended assuming the variance and the parameters of the response of all treatments are unknown and an adaptive sequential optimal design scheme is implemented and used for demonstration. Purely information theoretic designs for the same allocation have been addressed elsewhere, and we use them to support the techniques applied to compound designs.



Citation: Duarte, B.P.M.; Atkinson, A.C.; Pedrosa, D.; van Munster, M. Compound Optimum Designs for Clinical Trials in Personalized Medicine. *Mathematics* **2024**, *12*, 3007. <https://doi.org/10.3390/math12193007>

Academic Editor: Mudassir M. Rashid

Received: 25 August 2024
Revised: 19 September 2024
Accepted: 25 September 2024
Published: 26 September 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Keywords: compound optimal designs; Neyman allocation; covariates; information–regret designs; Parkinson’s disease

MSC: 62K05, 90C22

1. Introduction

This paper investigates methods for optimizing experimental designs in clinical trials, focusing on maximizing the information needed to predict responses to various treatments while considering ethical constraints relevant to personalized medicine. We extend existing approaches by addressing the challenge of patient allocation to treatments, aiming to balance statistical efficiency with the ethical imperative of minimizing the risk of administering suboptimal treatments to new patients. Our interest in personalized medicine leads to the modeling of patient–treatment interactions. Consequently, there is generally no single best treatment; the optimal choice depends on the patient’s covariates. Common patient covariates, also known as prognostic factors, include age, sex, prior illnesses, and baseline status on various indicators such as depression and susceptibility to certain conditions. These covariates may be represented by continuous metrics or discrete scales. Despite the differences in their nature, the challenge of optimally allocating patients remains a general problem, regardless of the type of covariates involved.

The main challenges in optimally allocating patients to effective treatments in clinical trials are (i) patient heterogeneity: variability among patients makes it difficult to predict individual responses to different treatments; (ii) ethical considerations: ensuring equitable access to potentially effective treatments for all participants while maintaining scientific rigor; (iii) bias and randomization: proper randomization is essential to avoid skewed results and the effects of uncontrolled factors, so as to ensure accurate assessments of treatment effectiveness; (iv) treatment effectiveness: identifying the most effective treatment for various patient subgroups can be challenging; (v) sample size and power: determining appropriate sample sizes is crucial for detecting meaningful differences between treatments; and (vi) dynamic changes in patient status: variations in patients' conditions over time can affect treatment effectiveness and influence trial conclusions [1]. In this study, we focus on points (ii) through (vi) and examine methods to address these challenges.

The problem of prescribing optimal designs in clinical trials is formalized as the maximization of a compound optimal design criterion that balances the information gained from the trial with the goal of maximizing patient outcomes by allocating them to the most effective treatments. While the first component of the objective function, related to statistical information, was discussed in Duarte and Atkinson [2], the second component addresses ethical concerns.

Typically, clinical trials are designed to either (i) maximize the information that enhances statistical inference about the response models; or (ii) satisfy ethical requirements by ensuring each patient is allocated to the most effective treatment as that becomes evident. These objectives often conflict with each other, leading to the formulation of an optimal allocation problem that may (i) maximize the information gained while imposing an upper limit on regret arising from allocation of patients to treatments which are not the best for them; or (ii) minimize regret while constraining the amount of information obtained, typically measured by a suitable convex function of the Fisher information matrix (FIM). When both objectives are convex (or concave), they can be combined using a linear combination with a Lagrange coefficient. This approach facilitates a balanced compromise between the incremental gain in information and the ethical considerations associated with treatment allocation [3,4].

In this paper, we consider the standard model commonly used in clinical trials, as described by Rosenberger and Lachin [5]. This model for the expected response to a given treatment includes a linear component to account for the effects of prognostic factors, or covariates, along with an additive constant representing the treatment effect [6]. The coefficients that link the covariates to the response for each treatment are either known a priori or can be sequentially estimated by fitting previous responses using least-squares methods. The response is assumed to follow a normal distribution, with each treatment associated with a distinct level of observational noise variance, which may be known beforehand or estimated from trial data. Such dependence of response variance on treatment is an extension of the standard model of Rosenberger and Lachin [5]. A second extension of the standard model is the existence of patient–treatment interactions.

Hu et al. [7] proposed a general family of covariate-adaptive, response-adaptive designs and derived the asymptotic properties of optimal allocation schemes. Additionally, Atkinson [8] presented optimal experimental designs for regression models—those with covariates—specifically when there are two treatments with differing variances. In this setup, a single normal response model is applied to all subjects. An up-to-date review of this topic can be found in Sverdlov et al. [9].

The use of compound criteria to identify sequential optimal designs is an area of active research, with the primary goal of reconciling the two key objectives of maximizing information and addressing ethical considerations. To incorporate ethical concerns into the optimal design criterion, a method based on randomly reinforced urn designs was applied, as detailed in May and Flournoy [10]. In a similar vein, Antognini and Giovagnoli [11] explored problems without covariates, integrating standardized information and ethical measures into the optimality criterion. Their approach dynamically adjusts the weight on

the two components based on the observed outcomes. Sverdlov et al. [12] examined responses related to time-to-event data, also without covariates. Furthermore, Metelkina and Pronzato [13] proposed a covariate-adaptive sequential allocation strategy that converges to the optimum, demonstrating the equivalence of this approach to penalized designs.

The numerical computation of compound designs that incorporate ethical considerations poses significant challenges, often hindering their practical application. In many cases, the computational complexity of these designs makes their implementation in real-world settings challenging. Consequently, decision makers (such as physicians and other agents) often opt for simpler methods that offer greater interpretability and can be more easily integrated into suitable information systems [14]. This paper addresses this issue by introducing a semidefinite programming formulation to derive asymptotic designs, along with a sequential allocation algorithm suited for real-world scenarios where prior knowledge is limited and the model parameters and response variance are extracted from individuals' responses.

Our approach automates the computation of standardized compound designs by using the D-optimality criterion to measure information and quantifying rewards based on the expected response. The framework assumes known treatment variances and determines the optimal (continuous) designs for patient allocation across treatments. The information component in our compound designs is grounded in the Neyman allocation method, incorporating covariates as previously explored in Duarte and Atkinson [2]. The ethics-based component is represented by the reward of the response, which is determined by the expectation calculated over a finitely large patient sample. When employing sequential designs, this expectation can also be iteratively approximated using a finite sample with known response values.

Further, we also address the more common scenario where treatment response variances and coefficients are initially unknown. In this case, we propose a sequential optimal design framework that updates the model after each patient response is obtained, subsequently allocating the next patient to the treatment that maximizes the compound objective.

This paper introduces three key innovations: (i) the development of a framework that enables the computation of compound designs incorporating standardized information measures and regret components; (ii) the formulation of semidefinite programming (SDP) methods to generalize the numerical computation of these designs; and (iii) the creation of an adaptive optimal allocation scheme that updates both the models for the responses and response variances in real time.

The paper is organized as follows. Section 2 provides the background and introduces the notation used for the problem formulation. Section 3 outlines the problem of optimal allocation using the compound-standardized information–regret (I-R) criterion. Section 4 presents the semidefinite programming (SDP) formulations designed to systematically address the problem; the construction of normalization values for computing normalized regret is detailed in Appendix A. Section 5 describes the algorithm for determining optimal sequential allocation designs, which is illustrated with an example. In Section 6, we apply the proposed methodologies to a real-world case, specifically focusing on the optimal allocation of Parkinson's disease patients to two alternative care treatments. Finally, in Section 7, we conclude with a brief discussion of our results.

2. Notation and Background

This section establishes the nomenclature and fundamental background used in the subsequent sections.

In our notation, boldface lowercase letters represent vectors, boldface capital letters represent continuous domains, blackboard bold capital letters denote discrete domains, and capital letters are used for matrices. Finite sets with ι elements are compactly represented by $\llbracket \iota \rrbracket \equiv \{1, \dots, \iota\}$. The transpose operation of a matrix or vector is denoted by "^T". The determinant of a matrix is represented by $\det(\bullet)$, and the trace is denoted by $\text{tr}(\bullet)$. The nomenclature used throughout the paper is provided at the end.

Let $k \in \llbracket K \rrbracket$ denote the covariates in the linear response model, $i \in \llbracket I \rrbracket$ represent the treatments, and $j \in \llbracket J \rrbracket$ indicate the levels each factor can have in the experiment. Here,

$K, I,$ and J denote the number of covariates, treatments, and factor levels in the response model, respectively.

When $i \in \{1, 2\}$ (indicating two treatments) and there are k covariates, the response model is given by

$$\eta_1(\mathbf{x}, \boldsymbol{\theta}) = \alpha_1 + \beta_{1,1}x_1 + \beta_{1,2}x_2 + \dots + \beta_{1,K}x_K, \quad \text{var}(\eta_1) = \sigma_1^2, \quad (1a)$$

$$\eta_2(\mathbf{x}, \boldsymbol{\theta}) = \alpha_2 + \beta_{2,1}x_1 + \beta_{2,2}x_2 + \dots + \beta_{2,K}x_K, \quad \text{var}(\eta_2) = \sigma_2^2, \quad (1b)$$

where $\mathbf{x} = (x_1, x_2, \dots, x_K) \in \mathbf{X} \subset \mathbb{R}^K$ is the vector of covariates, \mathbf{X} denotes the design space which contains all possible values of \mathbf{x} , $\boldsymbol{\theta}$ represents the set of parameters to be estimated (including all α 's and β 's), and $\Theta \subset \mathbb{R}^{n_\theta}$ is the domain of these parameters, with n_θ being the total number of parameters across all treatment models. The models $\eta_i(\mathbf{x}, \boldsymbol{\theta})$, for $i \in \llbracket 2 \rrbracket$, represent the expected responses, each characterized by an error term with variance $\text{var}(\eta_i)$. The p th realization of \mathbf{x} characterizing the p th patient is represented as \mathbf{x}_p .

The allocation is optimized to maximize the information extracted from patients' responses while considering the predictions from the model. Our goal is to determine the allocation of individuals to treatment 1 for estimating the parameters $(\alpha_1, \beta_{1,1}, \dots, \beta_{1,K})^\top$, and to treatment 2 for estimating the parameters of the second treatment model. In approximate optimal designs, w represents the fraction of individuals allocated to treatment 1, while $1 - w$ denotes the fraction allocated to treatment 2. This yields the design matrix

$$\xi = \begin{pmatrix} w & 1 - w \\ 1 & 2 \end{pmatrix}.$$

Here, the upper row of ξ indicates the fraction allocated to each treatment, while the lower row denotes the treatment number itself. In cases where the number of individuals is small and an exact design is required, rounding procedures can be used to determine the allocation, as described in Pukelsheim and Rieder [15].

Let $\tau = \sigma_2^2 / \sigma_1^2$, and without loss of generality, let $\sigma_1^2 = 1$. Consequently, the Fisher information matrix (FIM) for the experimental design ξ is a $2(K + 1) \times 2(K + 1)$ diagonal matrix given by Atkinson [8]:

$$M(\xi) = \begin{pmatrix} w & 0 & \dots & 0 & 0 \\ 0 & \frac{1-w}{\tau} & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & w & 0 \\ 0 & 0 & \dots & 0 & \frac{1-w}{\tau} \end{pmatrix}. \quad (2)$$

3. Standardized Compound Designs

In this section, we formulate standardized compound designs that balance information measurement criteria with regret. The response models for all treatments are continuous and exhibit linear dependence on K covariates. Moreover, the variances of the treatment responses are known a priori. Our goal is to establish a general conceptual framework for determining continuous optimal information–regret (I-R) designs, characterized by the fraction of patients allocated to each treatment.

We focus on the computation of I-R designs, which aim to reconcile the objectives of allocating individuals to the most promising treatments while optimizing parameter estimation efficiency across different response groups; see Hu et al. [7], Antognini and Giovagnoli [11], Metelkina and Pronzato [13]. In most applications, responses are either binary or follow a standard linear model. For both types, it is rational to evaluate allocation procedures based on the proportion of patients receiving the most effective treatment. However, in the presence of patient–treatment interactions, there is generally no single best treatment; the optimal choice depends on the patient's covariates. If the responses are significantly different, the choice of treatment is crucial. Conversely, if the responses are similar, the advantage of one treatment over another is minimal. Thus, we use the expected

response value based on previous allocations as our indicator for correct allocation. This approach represents an innovation compared to earlier methods.

This problem is particularly relevant as it mirrors the sequential allocation of individuals in clinical trials. It exploits prior information—characterized by both continuous and discrete covariates known a priori—to refine allocation strategies in later stages. Practically, optimization strategies offer significant advantages over pure randomization techniques, as discussed in Bertsimas et al. [16], Kallus [17].

In this context, we integrate standardized D-optimality, which quantifies the information content, with standardized regret, which measures the expected loss associated with not selecting the optimal treatment—the one corresponding to the highest response.

We first consider the information criterion. Assume the covariates x_k for $k \in \llbracket K \rrbracket$ are discrete, taking values in $\{-1, +1\}$. Consequently, the design space is defined as $\mathbf{X} \equiv \{-1, +1\}^K$. Let θ denote the set of parameters to be estimated, including the vectors α and β in Equation (1). In this simplified configuration, we consider categorical covariates with two levels: -1 (lower) and $+1$ (upper). According to the bound on the determinant of a Hadamard matrix [18,19], extreme observations provide the most information for the D-optimality criterion in first-order models.

We emphasize that using these limited support points for the designs helps reduce the computational burden when demonstrating the properties of our designs in a non-sequential context and provides valuable asymptotic allocations. In the application to constructing sequential clinical trials discussed in Section 5, covariate values are sampled from continuous distributions.

The criterion used for information-theoretical analysis in this work is D-optimality. This approach aims to maximize the determinant of the matrix $M(\xi)$, or equivalently, the $1/[2(K + 1)]$ power of the determinant, which represents the geometric mean of the diagonal elements of matrix B , as described in Duarte and Atkinson [2] (Equation 3.d).

To address our objective of integrating criteria that measure different aspects—specifically, the information content in the allocation and the regret associated with individual treatment assignments—we use standardized criteria, as detailed by Dette [20]. Hence, the information criterion we seek to maximize is given by

$$\Psi(\xi, \theta) = \left(\frac{\det[M(\xi, \theta)]}{\det[M(\xi^*, \theta)]} \right)^{1/n_\theta} \tag{3a}$$

$$M(\xi, \theta) = \sum_{i=1}^I w_i \mathcal{M}(\eta_i, \theta) \tag{3b}$$

$$\mathcal{M}(\eta_i, \theta) = \left(\frac{\partial \eta_i(\mathbf{x}, \theta)}{\partial \theta} \right)^\top \left(\frac{\partial \eta_i(\mathbf{x}, \theta)}{\partial \theta} \right), \tag{3c}$$

where $\mathcal{M}(\eta_i, \theta)$ represents the Fisher information matrix (FIM) for treatment i , and ξ^* denotes the optimal allocation when D-optimality is the sole criterion. For simplicity, we assume the models $\eta_i(\bullet)$ are linear in parameters, which makes the matrices $\mathcal{M}(\eta_i, \theta)$ independent of θ . The explicit inclusion of θ in the definition of $\Psi(\xi, \theta)$ serves two main purposes: (i) it clarifies the dependency on the number of parameters; and (ii) it aligns the functional with the regret representation, which does depend on θ .

We now introduce the regret criterion. The objective is to assign an individual, denoted as individual $m + 1$, who is characterized by covariate values \mathbf{x}_{m+1} , to the treatment that maximizes the expected response. For simplicity, we assume that higher response values are preferred, although the opposite preference could also be applied.

To align this criterion with the information criterion, it is necessary to normalize the responses to fall within the range $[0, 1]$. To achieve this normalization, we need to evaluate all I treatments and, for each treatment, consider all possible covariate values in $\Omega_m \equiv \{\mathbf{x}\}_1^m$. Here, Ω_m includes the m previous vectors of covariates that have already been allocated. Assuming perfect knowledge of the parameters θ for all models, let

$$y_{\Omega_m}^{\max} = \max_{i \in [I]} \max_{\mathbf{x} \in \Omega_m} \eta_i(\mathbf{x}, \boldsymbol{\theta})$$

denote the maximum response value achieved by the treatment that yields the highest response. Specifically, $y_{\Omega_m}^{\max}$ represents the highest value of the response over all models $i \in [I]$ and covariate values $\mathbf{x} \in \Omega_m$, obtained using model (1).

Conversely, the minimum response is denoted by $y_{\Omega_m}^{\min}$, which satisfies

$$y_{\Omega_m}^{\min} = \min_{i \in [I]} \min_{\mathbf{x} \in \Omega_m} \eta_i(\mathbf{x}, \boldsymbol{\theta}).$$

Let

$$\Delta = y_{\Omega_m}^{\max} - y_{\Omega_m}^{\min} \tag{4}$$

denote the difference in response expectations between the “optimal” and “least favorable” treatments for $\mathbf{x} \in \Omega_m$. Each of these values can be computed using the mixed-integer linear programming (MILP) procedure detailed in Appendix A. The parameter Δ ensures compatibility between the standardized regret and the standardized information criteria. To complete the scaling, we also need to determine the minimum response achieved (i.e., the baseline), denoted as

$$\delta_{\min} = y_{\Omega_m}^{\min}.$$

These problems determine the maximum (or minimum) response of a set of discrete points, considering the full set of response models. This approach is used to systematically find the extrema of possibly nonlinear or even discontinuous response models. We rely on a branch-and-cut large-scale solver [21], which ensures solvability even for larger numbers of discrete points and response models.

We now turn to the design of experiments. In the clinical trials motivating our work, the choice of covariate vector \mathbf{x} is not under our control. Each patient p arrives with a predetermined covariate vector. The central design question is determining the most suitable treatment for each patient.

The optimal treatment allocation for the p th individual is given by the treatment i_p^{\max} , which is defined as

$$i_p^{\max} = \arg \max_{i \in [I]} \eta_i(\mathbf{x}_p, \boldsymbol{\theta}),$$

where $\eta_i(\mathbf{x}_p, \boldsymbol{\theta})$ represents the response of treatment i given the covariate vector \mathbf{x}_p and the parameter vector $\boldsymbol{\theta}$.

In this context, regret measures the cost associated with not selecting the most effective treatment, as described by Rosenberger and Sverdlov [22]. When comparing different allocation strategies for I treatments across m individuals with known covariate information, regret is quantified as the average difference between the expected outcomes of the optimal allocation policy and those of an alternative policy.

Specifically, the non-standardized average regret, denoted as Γ^{ns} , is defined as

$$\Gamma_m^{\text{ns}}(\boldsymbol{\zeta}, \boldsymbol{\theta}) = \frac{1}{m} \sum_{j=1}^m \left[y_{\Omega_m}^{\max} - \eta_{\ell(j)}(\mathbf{x}_j, \boldsymbol{\theta}) \right], \tag{5}$$

where $y_{\Omega_m}^{\max}$ is the maximum expected response achievable across all treatments given the covariate information available, and $\eta_{\ell(j)}(\mathbf{x}_j, \boldsymbol{\theta})$ is the response of the treatment actually assigned to the j th individual. Here, $\boldsymbol{\ell}$ denotes a vector, where $\ell(j)$ specifies the treatment assigned to the j th individual.

The scale factor used to normalize the average regret to the unit interval, Δ , remains constant for optimal designs when the model coefficients of treatments are known. However, it must be re-estimated in sequential allocation schemes as new data become available. Given that

$$y_{\Omega_m}^{\max} = \Delta + \delta_{\min},$$

the standardized average regret, referred to as “regret” (without superscript) in the following sections, is given by

$$\Gamma_m(\xi, \theta) = \frac{\Gamma_m^{ns}(\xi, \theta) + \delta_{\min}}{\Delta}.$$

Alternatively, $\Gamma_m(\xi, \theta)$ can be expressed as

$$\Gamma_m(\xi, \theta) = 1 - \frac{1}{m} \sum_{j=1}^m \frac{\eta_{\ell(j)}(\mathbf{x}_j, \theta) - \delta_{\min}}{\Delta}, \tag{6}$$

where $\eta_{\ell(j)}(\mathbf{x}_j, \theta)$ represents the response of the treatment assigned to the j th individual.

Minimizing regret in the allocation process corresponds to selecting the treatment that maximizes the following objective function (up to constant terms):

$$\max_{\mathbf{w}} [-\Gamma_m(\xi, \theta)] \equiv \max_{i \in [I]} \frac{1}{m} \sum_{j=1}^m \frac{\eta_{\ell(j)=i}(\mathbf{x}_j, \theta) - \delta_{\min}}{\Delta},$$

where $\mathbf{w} \in \mathbb{R}^I$ denotes the vector of proportions of patients allocated to each treatment, subject to the constraint $\sum_{i=1}^I w_i = 1$.

Let m_i (with $m_i < m$) denote the number of subjects previously assigned to treatment i . Consequently, we have

$$\frac{1}{m} \sum_{i=1}^I \sum_{j=1}^{m_i} \frac{\eta_{\ell(j)=i}(\mathbf{x}_j, \theta) - \delta_{\min}}{\Delta} = \sum_{i=1}^I w_i \mathbb{E}_m \left(\frac{\eta_i(\mathbf{x}, \theta) - \delta_{\min}}{\Delta} \right), \tag{7}$$

where $\mathbb{E}_m \left(\frac{\eta_i(\mathbf{x}, \theta) - \delta_{\min}}{\Delta} \right)$ represents the expectation of the standardized loss for the m_i individuals previously allocated to treatment i . For a detailed derivation of Equation (7), refer to Appendix B.

The I-R-optimality criterion is formulated by integrating the information-theoretic criterion from Equation (3) with the regret criterion from Equation (7). The I-R-optimal design seeks to address the following optimization problem:

$$\max_{\mathbf{w} \in [0,1]^I} (1 - \gamma) \Psi(\xi, \theta) - \gamma \Gamma_m(\xi, \theta)$$

which is equivalently expressed as

$$\max_{\mathbf{w} \in [0,1]^I} (1 - \gamma) \left(\frac{\det[M(\xi, \theta)]}{\det[M(\xi^*, \theta)]} \right)^{1/n_\theta} + \gamma \sum_{i=1}^I w_i \mathbb{E}_m \left(\frac{\eta_i(\mathbf{x}, \theta) - \delta_{\min}}{\Delta} \right), \tag{8a}$$

$$\text{s.t. } \sum_{i=1}^I w_i = 1 \tag{8b}$$

where $\gamma \in [0, 1]$ represents the relative weight assigned to the regret criterion, which is determined in advance. In practice, the decision maker selects this parameter to reflect the relative importance of ethical considerations versus inferential accuracy when developing response models for both groups. For instance, if γ is set to 0.95, it suggests that ethical concerns are given 19 times more significance than the information criterion. This high value of γ is typically used in scenarios where patients are well-informed, and the priority is to minimize regret by allocating them to the treatment with the least potential regret.

Problem (8) is formulated similarly to (Metelkina and Pronzato [13], Equation 2.9) and can also be represented as a semidefinite program. However, it is important to note that the multiplicative compound criterion proposed by Hu et al. [7] cannot be expressed in semidefinite form due to the lack of semidefinite representability of the functions involved. Additionally, when $\gamma = 0$, the I-R-optimality criterion simplifies to D-optimality, as dis-

cussed in (Duarte and Atkinson [2], §4). On the other hand, when $\gamma = 1$, the criterion aligns with concepts from multi-armed bandit theory, such as those explored by Villar et al. [23].

4. Optimal I-R Designs via SDP

We now present a formulation to address problem (8) using semidefinite programming. Solving this problem via SDP involves several key steps: (i) determining the value of $\det[M(\xi^*, \theta)]$ in Equation (3a), which can be precomputed using the procedure outlined in (Duarte and Atkinson [2], A1.a–A1.e). This corresponds to calculating the determinant for the optimal design when the objective is the D-optimality criterion. (ii) Estimating the average response of each treatment, a process that relies on the parameter vector θ derived from previous allocations and is inherently subject to noise. (iii) Computing the quantities $y_{\Omega_m}^{\min}$, $y_{\Omega_m}^{\max}$, Δ , and δ_{\min} , as obtained from Equations (4), (A1) and (A2). (iv) Incorporating the weight γ , as specified by the user.

In practical terms, this design can be interpreted as an approximate, infinite-size local design intended to estimate a specific parameter vector θ . The primary goal of this formulation is to establish an asymptotic framework for resolving local designs within sequential procedures, as outlined by Pocock [24]. The formulation is presented below.

$$\max_{\xi, B, t} (1 - \gamma) \frac{t}{\{\det[M(\xi^*, \theta)]\}^{1/[2(K+1)]}} + \gamma \sum_{i=1}^I w_i \mathbb{E}_m \left(\frac{\eta_i(\mathbf{x}, \theta) - \delta_{\min}}{\Delta} \right) \tag{9a}$$

$$\text{s.t.} \begin{pmatrix} M(\xi, \theta) & B^T \\ B & \text{diag}(B) \end{pmatrix} \succeq 0_{4(K+1)} \tag{9b}$$

$$t \leq \prod_{j=1}^{2(K+1)} B_{j,j}^{1/[2(K+1)]} \tag{9c}$$

$$\sum_{i=1}^I w_i = 1 \tag{9d}$$

$$0 \leq w_i \leq 1, \quad i \in \llbracket I \rrbracket. \tag{9e}$$

Here, B is a lower triangular matrix of size $2(K + 1) \times 2(K + 1)$, $B_{j,j}$, $j \in \llbracket 2(K + 1) \rrbracket$ the diagonal elements of B , and t the hypograph of the determinant function.

To illustrate the application of Problem (9), consider a scenario similar to the one presented in Section 4 of Duarte and Atkinson [2]. In this setup, there are two treatments and two covariates (denoted as x_1 and x_2), each with two levels. The design space is defined as $\mathbf{X} \equiv \{-1, +1\}^2$. The local parameter values θ for treatment 1 are $(\alpha_1, \beta_{1,1}, \beta_{1,2}) = (0.5, 0.2, 0.4)$, and for treatment 2 they are $(\alpha_2, \beta_{2,1}, \beta_{2,2}) = (1.0, 0.8, 1.0)$. By solving programs (A1) and (A2), we obtain the following values: $y_{\Omega_m}^{\min} = \delta_{\min} = -0.8$, $y_{\Omega_m}^{\max} = 2.8$, and $\Delta = 3.6$. The expected values for each model are $\mathbb{E}_m[\eta_1(\mathbf{x}, \theta)] = 0.5$ and $\mathbb{E}_m[\eta_2(\mathbf{x}, \theta)] = 1.0$, respectively. This setup demonstrates that the second treatment has a significantly higher expected value, making it the preferable choice to minimize regret. However, it is crucial to recognize that decision making is also influenced by the uncertainty in the responses, as quantified by τ , which in turn impacts the information-based criterion.

We systematically vary the parameter τ to conduct a comparative analysis between I-R-optimal designs and purely information-theoretical designs, as presented in Duarte and Atkinson [2]. For each value of τ , we use the optimal $\det[M(\xi^*, \theta)]$ values from Table 1 in Duarte and Atkinson [2], which were derived for the case $K = 2$. Additionally, we explore different values of γ in the set $\{0.90, 0.95\}$ to assess the impact of the weighting parameter, where a higher γ emphasizes the importance of regret in the allocation process.

The performance of these designs depends on the values of \mathbf{x} . In our calculations, we assume an equal occurrence of each \mathbf{x} value, resulting in balanced 2^K experiments for each treatment. The corresponding results are presented in Table 1 (the third and fifth columns) and align well with our expected trends. To simplify the interpretation of the results, we denote w_1 as the proportion of individuals allocated to the first treatment, with $w_2 = 1 - w_1$

representing the proportion allocated to the second treatment. “Opt” refers to the objective function value at convergence. As anticipated, for treatment 1, which has a lower expected response, a smaller fraction of individuals is allocated. The weights consistently remain below a half, reflecting the allocations observed in purely information-theoretical designs, specifically those based on the D-optimality criterion, as shown in Table 1 of Duarte and Atkinson [2]. As expected, the fraction of patients allocated to treatment 1 decreases with increasing γ , illustrating the growing influence of the ethical component in decision making. The optimum decreases as τ increases, with this effect being more pronounced for smaller values of γ . This reduction in the optimum with increasing τ results from heightened uncertainty, which consequently diminishes the amount of information gathered from the experimental plan (see Table 1 in Duarte and Atkinson [2] for an analysis based solely on information-based allocation). To evaluate the influence of γ on the optimum, we must first consider the allocation driven purely by regret minimization. In this case, the allocation is minimally affected by τ , as individuals are primarily assigned to the treatment with the higher expected response (treatment 2), regardless of its variance. As γ decreases, the standardized information component in $\Psi(\xi, \theta)$ (see Equation (8a)) becomes more critical, and its reduction less pronounced than the decrease in the standardized regret.

Table 1. I-R-optimal designs for model (1) (including two treatments, two covariates and $x_k \in \{-1, +1\}$, $k \in \llbracket K \rrbracket$, and $w_2 = 1 - w_1$).

σ_1^2	σ_2^2	I-R-Opt. Design [†]		I-R-Opt. Design [‡]	
		w_1	Opt	w_1	Opt
1.0000	0.2000	0.1548	0.5046	0.2941	0.5392
1.0000	0.4000	0.1067	0.4929	0.2417	0.5085
1.0000	0.6000	0.0833	0.4882	0.2093	0.4950
1.0000	0.8000	0.0690	0.4856	0.1864	0.4871
1.0000	1.0000	0.0593	0.4839	0.1689	0.4819
1.0000	1.2500	0.0507	0.4824	0.1519	0.4773
1.0000	1.6667	0.0412	0.4809	0.1310	0.4722
1.0000	2.5000	0.0304	0.4793	0.1043	0.4665
1.0000	5.0000	0.0178	0.4774	0.0672	0.4597

[†] $\gamma = 0.95$; [‡] $\gamma = 0.90$.

To apply the results presented in Table 1 effectively in practical scenarios, researchers should follow these steps:

- (i) Use the given values of σ_1^2 and σ_2^2 to compute τ ;
- (ii) Set the value of γ ;
- (iii) Find the corresponding value of w for the compound optimality criterion based on τ from step (i);
- (iv) Allocate $\lceil Nw_i \rceil$ individuals to each treatment i .

5. Sequential Optimal Design Based on the I-R Rule

We now investigate the optimal sequential design based on the I-R rule, as represented by (9). Our approach is consistent with sequential allocation algorithms similar to those used by Atkinson [6] and Atkinson [25]. Specifically, (i) the values of the variances and coefficients within the response models for all treatments, denoted as θ , are unknown; (ii) the value of $\det[M(\xi^*, \theta)]$ is determined in advance by solving the D-optimal design problem; (iii) the weighting parameter γ remains fixed throughout the analysis; and (iv) we have information about the underlying distributions of the covariates.

Without loss of generality, we assume that the covariates are uniformly distributed within the range $[-1, +1]$ (i.e., $x_k \sim \mathcal{U}[-1, +1]$ for $k \in \llbracket K \rrbracket$). Additionally, the response noise is modeled as a normal distribution with mean zero and standard deviation σ_i for treatment $i \in \llbracket I \rrbracket$. While these σ_i values are initially unknown, estimates of them are updated iteratively as each patient is allocated and their response is recorded. The updated

estimates, denoted as s_i for $i \in \llbracket I \rrbracket$, represent the residual mean square errors from the fitted models for the I treatments and reflect the effectiveness of the iterative fitting procedure in capturing the noise in the responses. For $I = 2$, the estimate $\hat{\tau}$ is given by $\hat{\tau} = s_2^2/s_1^2$. As in previous sections, I represents the number of treatments, and $p \in \llbracket P \rrbracket$ indexes the patients enrolled in the trial. The covariate x_k denotes an individual covariate, and $\mathbf{x}_p \equiv \{x_{k,p} : k \in \llbracket K \rrbracket\}$ represents the vector of covariates for patient p .

In the context of sequential treatment allocation, we have access to the covariates and allocations for the initial p patients, forming a matrix G_p that includes both allocations and explanatory variables. For patient $p + 1$, who has a vector \mathbf{x}_{p+1} of covariates, when treatment i is assigned, the combined vector of allocation and explanatory variables for this patient is represented as $g_{i,p+1}$ for $i \in \llbracket I \rrbracket$.

Previous work on optimal experimental designs, particularly those focusing on information criteria (e.g., Atkinson [6] and (Smith [26], §10)), has shown that minimizing the variance of the estimate after $p + 1$ observations can be achieved by selecting the treatment for which the sensitivity function reaches its maximum value. This approach corresponds to a specific case of optimal design theory aimed at reducing the variance of a single parameter estimate within a model with multiple nuisance parameters, known as D_s -optimality. For more details, refer to (Atkinson et al. [27], §10.3) for the case with $s = 1$. However, in our approach, which uses a compound (convex) criterion, we take a different approach. Specifically, we compute the value of Equation (8) for each available treatment for patient $p + 1$. We then allocate the patient to the treatment that maximizes the overall performance according to this criterion.

In the sequential development of designs, especially when employing linear response models and information criteria, calculating the sensitivity function is a straightforward task as it directly depends on the information matrix and can be readily computed for each proposed treatment allocation. However, for I-R-optimal designs, this process requires an alternative approach due to the complexity of establishing the sensitivity function for the regret component. To circumvent this challenge, we adopt a different strategy. We compute the metric that represents the compound criterion for I distinct scenarios, each one constructed by considering the allocation of patient $p + 1$ to a different treatment $i \in \llbracket I \rrbracket$. In other words, we evaluate $v[\xi_i^{(p+1)}]$ given by

$$v[\xi_i^{(p+1)}] = (1 - \gamma) \left(\frac{\det[M(\xi, \theta)]}{\det[M(\xi^*, \theta)]} \right)^{1/n_\theta} + \gamma \sum_{i=1}^I w_i \mathbb{E}_m \left(\frac{\eta_i(\mathbf{x}, \theta) - \delta_{\min}}{\Delta} \right), \quad i \in \llbracket I \rrbracket \quad (10)$$

which represents the I-R performance of the design ξ constructed after treatments have been allocated to p individuals (with the superscript denoting the number of individuals with available information), and treatment i is assigned to patient $p + 1$. Here, the subscript aids in tracking the provisional treatment allocation for subject $p + 1$.

To calculate the design $\xi_i^{(p+1)}$, several steps are required. First, we need to update the weights, \mathbf{w} , and their corresponding Fisher information matrices (FIMs) following the tentative allocation of treatment i to subject $p + 1$. Additionally, we must compute the metrics $v[\xi_i^{(p+1)}]$ for each treatment i . The updates involve the following:

1. Parameter estimates: Recalculate $\hat{\theta}$, the estimates of s_1^2 and s_2^2 , and so $\hat{\tau}$, based on the available data;
2. Updated metrics: Reconstruct $\hat{\delta}_{\min}$ and $\hat{\Delta}$ using the responses from the p individuals, whose treatment allocations and responses are known;
3. Weights vector: Recalculate the vector of weights, where $w_i = p_i/p$ represents the fraction of individuals allocated to each treatment.

Once these updates are complete, we compute the normalized ratios as follows:

$$\pi_i = \frac{v[\xi_i^{(p+1)}]}{\sum_{j=1}^I v[\xi_j^{(p+1)}]}, \quad i \in \llbracket I \rrbracket. \quad (11)$$

To determine the allocation, we use the following rule:

$$\begin{cases} \text{Allocate patient } p + 1 \text{ to treatment 1} & \text{if } \pi_1 \geq 0.5, \\ \text{Allocate patient } p + 1 \text{ to treatment 2} & \text{otherwise.} \end{cases} \tag{12}$$

This rule dictates that patient $p + 1$ will be allocated to treatment 1 if the ratio π_1 is greater than or equal to 0.5, and to treatment 2 otherwise.

In the context of tentatively allocating treatment i to subject $p + 1$, we predict the response of this subject using the i th treatment model from Equation (1), denoted as $\eta_i(\mathbf{x}_{p+1}, \boldsymbol{\theta})$. Here, $\boldsymbol{\theta}$ is estimated based on data from the first p_i individuals who were allocated to this specific treatment. The predicted response is then used to update two key elements:

1. Parameter estimates: Recalculate $\hat{\boldsymbol{\theta}}$ for the model associated with the treatment being considered for allocation;
2. Metrics update: Update $\hat{\delta}_{\min}$ and $\hat{\Delta}$ based on the new response information.

Additionally, the experimental design itself is updated. This involves recalculating the weights using the formula

$$w_i^{(p+1)} = \frac{p_i + 1}{p + 1}, \quad i \in \llbracket I \rrbracket,$$

where p_i represents the number of patients previously allocated to treatment i . The updated design is denoted as $\tilde{\zeta}^{(p+1)}$. The treatment for subject $p + 1$ is then selected based on the one that maximizes the objective function.

This procedure is repeated iteratively for a total of P individuals, where P can be predetermined or adjusted based on ongoing trial results using a specified criterion. To assess the properties of our procedure for a given value of P , we use simulations. To minimize the impact of randomness, we replicate the entire procedure n_s times. Algorithm 1 outlines the treatment allocation process.

Initially, $k + 1$ individuals are allocated to each treatment to ensure a non-singular Fisher information matrix (FIM). Parameter estimates for the model are derived from the responses of these $2(K + 1)$ individuals. The procedure is then systematically repeated until the allocation of the P th individual. After each new patient is admitted, their covariates are considered in the decision-making process to optimize treatment allocation. Once the response is available, it is integrated into the dataset, prompting updates to the model parameters, the FIM, and the expectation estimates. In the algorithm $\mathbb{1}_{i,j}$ is an indicator variable with value 1 if the j th individual is allocated to treatment i and 0 otherwise.

Algorithm 2 systematically outlines the procedure for sequential allocation based on the proposed I-R criterion. To illustrate the mechanics of the procedure, we consider the allocation of the p th individual. The implementation requires the following quantities:

- (i) The number of competing treatments, denoted as I ;
- (ii) The number of covariates, denoted as K ;
- (iii) The vector of covariates for the p th individual, denoted as \mathbf{x}_p ;
- (iv) The Fisher information matrices (FIMs) for each treatment, and the global FIM after allocating the $(p - 1)$ th individual, denoted as $\mathcal{M}^{(p-1)}(\tilde{\zeta})$;
- (v) The estimates of $\hat{\boldsymbol{\theta}}^{(p-1)}$;
- (vi) The values of $\Delta^{(p-1)}$, $\delta_{\min}^{(p-1)}$, γ , and $\det[M(\tilde{\zeta}^*, \boldsymbol{\theta})]$;
- (vii) The estimated error variances for all treatments, denoted as $\hat{\sigma}_i^2$, obtained after allocating the $(p - 1)$ th individual.

This iterative process is applied to all individuals, $p \in \llbracket P \rrbracket$, as they enter the trial.

Algorithm 1 Algorithm to randomize the allocation ($k = 2$).

```

procedure RANDOMIZEALLOCATION( $I, n_s, k, \gamma, \det[M(\zeta^*, \theta)]$ )
  Allocate  $2(K + 1)$  individuals to both treatments ( $\mathbf{w} = 1/2$ )
   $j \leftarrow 2(K + 1)$ 
  Construct  $\hat{\theta}^{(j)}, \hat{\delta}_{\min}, \hat{\Delta}$ 
  for  $m = 1$  to  $n_s$  do ▷ sampling loop
    for  $j = 2(K + 1) + 1$  to  $P$  do ▷ individuals loop
      for  $i = 1$  to  $k$  do ▷ treatments loop
        Allocate  $j$  individual to treatment  $i$ 
        Update  $\mathbf{w}^{(j+1)}, \zeta_i^{(j+1)}$  from  $\mathbf{w}^{(j)}$  and  $\zeta_i^{(j)}$ 
        Update  $M(\zeta_i^{(j+1)})$ 
        Update  $\mathbb{E}_m \left[ (\eta_i(\mathbf{x}, \theta^{(j)}) - \hat{\delta}_{\min}) / \hat{\Delta} \right]$ 
        Compute  $v[\zeta_i^{(j+1)}]$  with  $\mathbf{x}_{j+1}$  using (10)
      end for
       $\pi_1 \leftarrow v[\zeta_1^{(j+1)}] / \sum_{i=1}^I v[\zeta_i^{(j+1)}]$ 
      Allocate  $j$  applying rule (12)
      Update the set of responses obtained from treatment  $i$ 
      Update  $\hat{\theta}^{(j+1)}, s_i^{(j+1)}, i \in [2]$  and  $\hat{\tau}$  via LS
      Update  $\hat{\delta}_{\min}$  and  $\hat{\Delta}$ 
      Update  $\mathbf{w}^{(j+1)}$  ( $w_i^{(j+1)} \leftarrow \sum_{m=1}^j \mathbb{1}_{i,j} / j$ )
       $W^{(m,j+1)} \leftarrow \mathbf{w}^{(j+1)}, j \in [I]$ 
    end for
  end for
   $\mathbb{E}_{n_s}[\mathbf{w}^{(j)}] \leftarrow \sum_{m=1}^{n_s} W^{(m,j)} / n_s, j \in [P]$ 
end procedure

```

Algorithm 2 Allocation of the p th individual.

```

procedure ALLOCATEINDIVIDUAL(Input:  $I, K, \mathbf{x}_p, \mathcal{M}^{(p-1)}(\zeta), \hat{\sigma}_i^2, \Delta^{(p-1)}, \gamma, \delta_{\min}^{(p-1)}, \det[M(\zeta^*, \theta)], \mathbb{E}_{p-1} \left( \frac{\eta_i(\mathbf{x}, \theta) - \delta_{\min}}{\Delta} \right)$ )
  Construct the vector  $\mathbf{z}_p \equiv (1, \mathbf{x}_p)$ 
  for  $i$  in  $[I]$  do
    Compute the function  $\max_i v[\zeta_i^{(p)}]$  for  $\mathbf{z}_p$ 
  end for
  Allocate the individual to the treatment prescribed by rules (11–12)
  Update  $\mathbf{w}^{(p)}, \mathcal{M}^{(p)}(\zeta), \hat{\sigma}_i^2, \Delta^{(p)}, \delta_{\min}^{(p)}, \hat{\theta}^{(p)}, \mathbb{E}_p \left( \frac{\eta_i(\mathbf{x}, \theta) - \delta_{\min}}{\Delta} \right)$ 
end procedure

```

To evaluate the algorithm, we consider a scenario with two treatments ($I = 2$), two covariates ($K = 2$), $n_s = 2000$ simulations, and $P = 1200$ individuals. The parameter values used to generate the response data for the treatments are $\theta_1 = (\alpha_1, \beta_{1,1}, \beta_{1,2})^T = (0.5, 0.2, 0.4)$ for treatment 1 and $\theta_2 = (\alpha_2, \beta_{2,1}, \beta_{2,2})^T = (1.0, 0.8, 1.0)$ for treatment 2, as detailed in Section 4. We set $\tau = 1.6667$, with corresponding variances of $\sigma_1^2 = 1.0$ and $\sigma_2^2 = 1.6667$, and examine two different values of γ , specifically $\gamma \in \{0.90, 0.95\}$. Additionally, we use the value of $(\det[M(\zeta^*, \theta)])^{1/[2(K+1)]} = 0.0968$, which was found by (Duarte and Atkinson [2], Table 1, line 7) for $K = 2$ and $\tau = 1.6667$. The table gives properties of D- (and other) optimal designs for the information component of our compound design.

In the upper panel of Figure 1, we observe the fraction w_1 , which represents the proportion of individuals allocated to treatment 1. Two distinct patterns emerge:

1. Convergence below 0.5: w_1 converges to a value below 0.5 for both values of γ . This indicates that fewer individuals are allocated to the treatment with the lower average expectation, aligning with our intended allocation strategy;
2. Overestimation of optimal allocation: The values of w_1 tend to overestimate the optimal allocation values obtained through the SDP approach, especially when considering an infinite sample size (as shown in Table 1, line 7 for I-R-optimal designs

with $\gamma \in \{0.90, 0.95\}$). This discrepancy arises primarily due to differences in expectation construction. In the SDP formulation, the expectation is established a priori using four points, whereas in the randomization procedure, it is adaptively constructed by considering the continuous design space of regressors for sampling. For example, the SDP-based optimal designs yield $w_1 = 0.0412$ for $\gamma = 0.95$ and $w_1 = 0.1310$ for $\gamma = 0.90$. In contrast, the adaptive design achieves $w_1 = 0.1491$ and $w_1 = 0.2163$ at convergence.

It is also observed that w_1 decreases as γ increases, reflecting the increased weight given to the regret term. As a result, more individuals are allocated to the treatment with a higher expectation (treatment 2), despite its larger variance. The difference in allocation between the two treatments at convergence is 0.0728. In the lower panel of Figure 1, we examine the variance estimates s_i^2 for both treatments ($i \in \llbracket I \rrbracket$). Both estimates converge to the actual variance of the responses. Notably, s_1^2 shows more pronounced fluctuations, mainly because it is estimated from a smaller fraction of patients, whereas s_2^2 is estimated from the larger remaining fraction.




Figure 1. Results for $\gamma \in \{0.90, 0.95\}$: (i) upper panel— w_1 ; (ii) lower panel— s^2 .

Figure 2 illustrates the progression of parameter estimates for both treatments over the course of the numerical experiment. It is noticeable that all of them converge towards the true values of θ . However, a pattern resembling that observed for the variance estimates s_i^2 is also evident. Specifically, the estimation of the parameters for $\eta_1(\mathbf{x}, \theta)$ relies on a relatively small number of individuals, causing any new data point to have a more pronounced impact on the existing estimates. In contrast, the parameter estimates for $\eta_2(\mathbf{x}, \theta)$ exhibit smoother trajectories due to the larger number of responses considered in the fitting process.

Figure 3 illustrates the evolution of w_1 across a range of γ values from 0.50 to 0.99. The observed trends align with our expectations: as the weight assigned to the regret component increases, the values of w_1 decrease, and the corresponding curves become flatter. This downward trend in w_1 continues as γ increases. Notably, when γ reaches 1.0, corresponding to an allocation based solely on regret, w_1 approaches zero. Consequently, the Fisher information matrix (FIM) becomes nearly singular, and $\det[M(\xi, \theta)]$ approaches zero. This suggests that the information component of the objective function in Equation (8) becomes prone to numerical instability. This phenomenon occurs because when the FIM approaches singularity it becomes ill-conditioned, meaning its condition number increases significantly. This leads to numerical instability in the computation of its determinant, as

small errors in the input data or round-off errors during calculations can result in large inaccuracies in the output.

To address these challenges while still evaluating the algorithm’s performance at extreme γ values, we conducted tests with $\gamma = 0.99$. The results remained consistent with the broader trends observed for other γ values.




Figure 2. Evolution of the model parameters for both treatments. The smoother upper trajectories are for treatment 2, which is allocated more frequently.




Figure 3. Evolution of w_1 with γ . As $\gamma \rightarrow 1$, there is a decreasing emphasis on parameter estimation and increasing on the allocation to the most promising treatment (the arrow represents the increase of γ).

6. A Real Case Application

In this section, we apply the methodologies developed in Sections 4 and 5 to a real-world case study. Specifically, we investigate the clinical trial described by Atkinson et al. [28], which focuses on the randomization of Parkinson’s disease patients. The primary aim of this study was to devise a randomization protocol for assigning patients to one of two treatment groups:

1. Treatment 1: A novel procedure exploiting digital technologies.
2. Treatment 2: A conventional treatment and monitoring approach.

To evaluate treatment efficacy, we use the quality of life (QoL) as the response variable. This is measured by the Parkinson’s Disease Questionnaire with 8 items (PDQ-8) [29], with QoL, denoted as $pdq8$, quantified as a percentage (ranging from 0% to 100%), where higher values indicate a lower quality of life. Several prognostic factors are known to influence QoL, with two demonstrating particularly strong correlations: (i) disease duration and stage, represented as $h&q$, and measured on the Hoehn and Yahr scale [30,31]; and (ii) psychological well-being and neuropsychiatric symptoms (e.g., depression and anxiety), denoted as bdi , and assessed using Beck’s Depression Inventory (BDI) [32], based on responses to a structured questionnaire.

The models used for preliminary results for both treatments follow the structure outlined in Equation (1). The first prognostic factor, x_1 , corresponds to $h&q$, while the second, x_2 , represents bdi . In the dataset used for modeling the two responses to these two variables $h&q$ and bdi are constrained to integer values. Specifically, $h&q$ ranges from 1 to 5 (i.e., $h&q \in [1, 5]$), and bdi ranges from 1 to 35 (i.e., $bdi \in [1, 35]$). Consequently, the design space considered is denoted as $\mathbf{X} \equiv [1, 5] \times [1, 35]$, with covariate values randomly, but not uniformly, drawn from this space.

To mimic the empirical distribution of covariates we simulated from continuous distributions, which were then discretized. The variable $h&q$ was generated by sampling from a normal distribution with a mean of 2.3837 and a standard deviation of 0.2857. Additionally, bdi was sampled from a Gamma distribution with a shape parameter of 1.7678 and a scale parameter of 6.8145. These parameters were estimated from historical data. The model parameters for each treatment are specified as follows:

- For treatment 1: $(\alpha_1, \beta_{1,1}, \beta_{1,2}) = (2.0941, 3.9057, 1.1089)$, with $\sigma_1^2 = 207.9267$.
- For treatment 2: $(\alpha_2, \beta_{2,1}, \beta_{2,2}) = (-0.1015, 4.1312, 1.5117)$, with $\sigma_2^2 = 107.8993$.

The expected values for each model are (i) $\mathbb{E}_m[\eta_1(\mathbf{x}, \boldsymbol{\theta})] = 33.7714$; and (ii) $\mathbb{E}_m[\eta_2(\mathbf{x}, \boldsymbol{\theta})] = 39.5027$. Additionally, the values for δ_{\min} , $y_{\Omega_m}^{\max}$, and Δ are (i) $y_{\Omega_m}^{\min} = \delta_{\min} = 5.5414$; (ii) $y_{\Omega_m}^{\max} = 73.4640$; and (iii) $\Delta = 67.9226$.

The D-optimal design is uniformly distributed, with $w_1 = w_2 = 0.5000$. The normalization factor for the D-optimal design used, $\det[M(\zeta^*, \boldsymbol{\theta})]^{1/[2(K+1)]}$, is 0.4203. It is important to note that given the characteristics of the response variable our objective is to minimize the second term in Equation (10), which is equivalent to maximizing its negative. Consequently, treatment 1 should be preferred due to its lower expectation, despite having a higher level of uncertainty. Visual inspection of the graphical representations of both models shows that in a specific region of the design space, defined as $\Omega \equiv \{(h&q, bdi) \in [1, 5] \times [1, 35] \mid 0.2251h&q + 0.4028bdi \leq 2.1956\}$, where both bdi and $h&q$ are low, treatment 2 becomes the preferred choice.

The compound designs derived from the SDP formulation (9) for $\gamma = \{0.90, 0.95\}$ are summarized in Table 2. These results align with the key trends identified in Section 4. Specifically, the treatment with the higher expectation is assigned to more than half of the individuals, despite also having a larger variance. This approach is consistent with our goal of maximizing the normalized regret. Notably, treatment 1 emerges as the most favorable option across almost the entire design domain. Moreover, as the weight of the regret component increases, the proportion of individuals allocated to treatment 1 also rises.

Table 2. I-R-optimal designs for model (1) considering the response is $pdq8$ model ($w_2 = 1 - w_1$).

σ_1^2	σ_2^2	I-R-Opt. Design [†]		I-R-Opt. Design [‡]	
		w_1	Opt	w_1	Opt
207.9267	107.8993	0.8120	−0.3711	0.6769	−0.3052

[†] $\gamma = 0.95$; [‡] $\gamma = 0.90$.

We now apply the treatment allocation algorithm discussed in Section 5 using the I-R rule. For this simulation, we set $n_s = 2000$ and $P = 1200$ (i.e., 1200 individuals). The values for $h&q$ and bdi are generated randomly from a uniform distribution over the ranges $[1, 5]$ and $[1, 35]$, respectively.

The upper panel of Figure 4 presents the weights w_1 for the two values of γ . These weights closely match the results obtained using SDP, as shown in Table 2. The difference in w_1 at convergence between $\gamma = 0.90$ and $\gamma = 0.95$ is approximately 0.0699. As expected, the weight assigned to treatment 1, which has the higher expectation, exceeds 0.5 in both cases, reaching around 0.80. This indicates that the allocation favors the treatment with the higher expectation.

The lower panel of Figure 4 displays the estimated variances of the response models for both treatment groups. The variances converge to a stable value, consistent with the test case. Notably, the group with the larger weight shows less variability in its variance estimate.




Figure 4. Results for $\gamma \in \{0.90, 0.95\}$: (i) upper panel— w_1 ; (ii) lower panel— s^2 . For both values of γ treatment 1 is allocated more frequently, and consequently, has a more precise estimate of variance.

7. Conclusions

This paper demonstrates the application of compound design criteria for patient allocation in personalized medicine. Building upon the framework proposed in Duarte and Atkinson [2], which focused on patient allocation without considering ethical aspects, this study integrates ethical considerations into the allocation process. Specifically, the aim is to balance the acquisition of information for parametric response models with ethical considerations to ensure that each patient is assigned to the most effective available treatment.

The proposed approach is designed for response models that are linearly dependent on prognostic factors and characterized by continuous responses with distinct noise variances. A concave compound design criterion is developed that combines D-optimality with regret minimization, as detailed in Section 3. This dual criterion reflects the necessity of both maximizing the precision of parameter estimates and minimizing patient harm or regret, emphasizing the importance of fairness in clinical decision making, considering a weighted sum objective function.

To compute continuous optimal designs for the I-R criterion within a discrete design space, semidefinite programming is considered, as outlined in Section 4. Previous work in Duarte and Atkinson [2] addressed the numerical computation of designs maximizing information about parameters in a similar setup, and this paper builds upon those formulations. This work assumes that the variances of the treatments and model parameters are known, facilitating the construction of the expectation of each treatment response.

In scenarios where information about variances and model parameters is unavailable, a sequential adaptive optimal design algorithm is presented in Section 5. This algorithm iteratively updates estimates of parameters and variances as responses are collected. A systematic method for computing sequential adaptive optimal allocation schemes is proposed, and its application is illustrated with a concrete example. Although the approach aligns with reinforcement learning algorithms described in studies such as Zhao et al. [33] and Matsuura et al. [34], it differs by utilizing a known parametric structure a priori rather than a model-free representation.

The proposed computational methods are evaluated through two distinct scenarios: first, using a simplified illustrative example; and second, applying them to a real-world problem involving Parkinson's disease patients. The treatments considered are (i) treatment 1, which employs a cutting-edge digital technology-centered procedure; and (ii) treatment 2, encompassing a traditional treatment and monitoring protocol.

Beyond computational efficiency, the proposed methods thoughtfully integrate ethical considerations into traditionally information-based designs in personalized medicine. This helps mitigate the ethical dilemma of balancing data collection in experimental settings with ensuring patients receive the best possible care. By incorporating regret minimization, the approach ensures that even patients not initially assigned to the optimal treatment experience reduced harm during the early stages of the experiment.

Moreover, this approach could be extended to non-medical domains such as marketing, economics, or policy-making, where optimal allocation strategies must balance information acquisition with fairness or ethical considerations. The framework's flexibility also opens up the potential for application to treatments with binary or categorical outcomes, further broadening its relevance as well as models of linear and nonlinear classes.

The computational methods proposed are robust and versatile, offering practical insights into the allocation process. Future research could explore alternative randomization techniques, as discussed in Atkinson [25], which address trade-offs between efficiency losses due to randomization and bias reduction. Additionally, developing theoretical results to support the convergence of these methods would be beneficial. The approach introduced could also be adapted to models beyond the linear regression framework considered here, with careful attention to selecting appropriate regret measures and scaling for these alternative models.

Finally, while this study centers on known model parameters and variances, exploring scenarios with more complex model structures or unknown variance assumptions remains a promising direction. Introducing Bayesian methods or robust optimization techniques could offer a means of addressing the uncertainties inherent in patient responses or treatment variabilities. In particular, incorporating patient heterogeneity more explicitly into the model could further enhance the framework's applicability to real-world clinical trials where patient responses may not follow identical parametric structures. The inclusion of such considerations could lead to even more refined and ethically sound allocation strategies, reinforcing the value of this approach in personalized medicine.

Author Contributions: B.D. conceived the methodology, implemented the methods computationally and wrote the initial version of the manuscript. A.C.A. edited and revised the manuscript and validated the methods. D.P. and M.v.M. provided the data set and validated the results. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The code for one of the examples considered herein is available on request.

Conflicts of Interest: The authors declare no conflict of interest.

Nomenclature

i	Counter for treatments
I	Number of treatments
j	Counter for factor levels
J	Number of factor levels
k	Counter for covariates
K	Number of covariates
i^{\max}	Treatment with maximum response
$v(\bullet)$	Compound criterion
\mathbf{x}	Vector of covariates
\mathbf{X}	Design space
\mathbf{w}	Vector of treatment weights
\mathbf{z}	Vector of covariates including 1
\mathbb{E}_m	Expectation based on m -individuals
$M(\bullet)$	Global FIM
$\mathcal{M}(\bullet)$	Treatment FIM
$\boldsymbol{\theta}$	Parameter vector
Θ	Parameter domain
Δ	Response range
δ_{\min}	Minimum (reference) response
Γ^{ns}	Non-standardized regret function
Γ	Standardized regret function
γ	Weight relating regret and information
Ψ	Standardized D-optimality criterion

Appendix A. Formulations to Determine the Bounds of Response Expectation

Here, we set the formulations used to estimate the bounds of the expectation for response to treatment. First, the upper bound problem:

$$\delta_{\max} = \max_{\mathbf{x}, \delta, \zeta, x} \sum_{i=1}^n \zeta_i \tag{A1a}$$

$$\text{s.t } \mu_i(\mathbf{x}) = \delta_i, \quad i \in \llbracket n \rrbracket \tag{A1b}$$

$$\delta_i \leq \zeta_i \tag{A1c}$$

$$\zeta_i \leq w_i M^{\text{big}} \tag{A1d}$$

$$\sum_{i=1}^n \omega_i = 1, \quad \mathbf{X} \in \{0, 1\}, \quad \mathbf{x} \in \Omega, \tag{A1e}$$

And now, the lower bound problem:

$$\delta_{\min} = \min_{\mathbf{x}, \delta, \zeta, x} \sum_{i=1}^n \zeta_i \tag{A2a}$$

$$\text{s.t } \mu_i(\mathbf{x}) = \delta_i, \quad i \in \llbracket n \rrbracket \tag{A2b}$$

$$\delta_i \geq \zeta_i \tag{A2c}$$

$$\zeta_i \leq w_i M^{\text{big}} \tag{A2d}$$

$$\sum_{i=1}^n \omega_i = 1, \quad \mathbf{X} \in \{0, 1\}, \quad \mathbf{x} \in \Omega. \tag{A2e}$$

Appendix B. Construction of the Regret Based on Previous Responses

Here, we construct the metric used to measure the regret based on the expectation of the responses available.

$$\begin{aligned}
\frac{1}{m} \sum_{i=1}^I \sum_{j=1}^{m_i} \frac{\eta_i(\mathbf{x}_j, \boldsymbol{\theta}) - \delta_{\min}}{\Delta} &= \frac{1}{m} \sum_{i=1}^I \sum_{j=1}^{m_i} \frac{\eta_i(\mathbf{x}_j, \boldsymbol{\theta}) - \delta_{\min}}{\Delta} \\
&= \frac{1}{m} \sum_{i=1}^I m_i \sum_{j=1}^{m_i} \frac{\eta_i(\mathbf{x}_j, \boldsymbol{\theta}) - \delta_{\min}}{m_i \Delta} \\
&= \frac{1}{m} \sum_{i=1}^I m_i \mathbb{E}_m \left(\frac{\eta_i(\mathbf{x}, \boldsymbol{\theta}) - \delta_{\min}}{\Delta} \right) \\
&= \sum_{i=1}^I \frac{m_i}{m} \mathbb{E}_m \left(\frac{\eta_i(\mathbf{x}, \boldsymbol{\theta}) - \delta_{\min}}{\Delta} \right) \\
&= \sum_{i=1}^I w_i \mathbb{E}_m \left(\frac{\eta_i(\mathbf{x}, \boldsymbol{\theta}) - \delta_{\min}}{\Delta} \right) \tag{A3}
\end{aligned}$$

Note that the first equality results from assuming that each patient is allocated to only one treatment.

References

- Piantadosi, S. *Clinical Trials: A Methodologic Perspective*, 2nd ed.; Wiley-Blackwell: Hoboken, NJ, USA, 2005. [\[CrossRef\]](#)
- Duarte, B.P.M.; Atkinson, A.C. Optimum designs for clinical trials in personalized medicine when response variance depends on treatment. *J. Biopharm. Stat.* **2024**, *accepted*. [\[CrossRef\]](#) [\[PubMed\]](#)
- Dragalin, V.; Fedorov, V. Adaptive designs for dose-finding based on efficacy-toxicity response. *J. Stat. Plan. Inference* **2006**, *136*, 1800–1823. [\[CrossRef\]](#)
- Pronzato, L. On the sequential construction of optimum bounded designs. *J. Stat. Plan. Inference* **2006**, *136*, 2783–2804. [\[CrossRef\]](#)
- Rosenberger, W.F.; Lachin, J.L. *Randomization in Clinical Trials: Theory and Practice*, 2nd ed.; Wiley: New York, NY, USA, 2016.
- Atkinson, A.C. Optimum biased coin designs for sequential clinical trials with prognostic factors. *Biometrika* **1982**, *69*, 61–67. [\[CrossRef\]](#)
- Hu, J.; Zhu, H.; Hu, F. A unified family of covariate-adjusted response-adaptive designs based on efficiency and ethics. *J. Am. Stat. Assoc.* **2015**, *110*, 357–367. [\[CrossRef\]](#)
- Atkinson, A.C. Optimum designs for two treatments with unequal variances in the presence of covariates. *Biometrika* **2015**, *102*, 494–499. [\[CrossRef\]](#)
- Sverdlov, O.; Ryzhnik, Y.; Wong, W.K. On optimal designs for clinical trials: An updated review. *J. Stat. Theory Pract.* **2020**, *14*, 1–29. [\[CrossRef\]](#)
- May, C.; Flournoy, N. Asymptotics in response-adaptive designs generated by a two-color, randomly reinforced urn. *Ann. Stat.* **2009**, *37*, 1058–1078. [\[CrossRef\]](#)
- Antognini, A.B.; Giovagnoli, A. Compound optimal allocation for individual and collective ethics in binary clinical trials. *Biometrika* **2010**, *97*, 935–946. [\[CrossRef\]](#)
- Sverdlov, O.; Tymofeyev, Y.; Wong, W.K. Optimal response-adaptive randomized designs for multi-armed survival trials. *Stat. Med.* **2011**, *30*, 2890–2910. [\[CrossRef\]](#)
- Metelkina, A.; Pronzato, L. Information-regret compromise in covariate-adaptive treatment allocation. *Ann. Stat.* **2017**, *45*, 2046–2073. [\[CrossRef\]](#)
- Chao, Y.S.; Wu, C.J.; Wu, H.C.; McGolrick, D.; Chen, W.C. Interpretable trials: Is interpretability a reason why clinical trials fail? *Front. Med.* **2021**, *8*, 541405. [\[CrossRef\]](#) [\[PubMed\]](#)
- Pukelsheim, F.; Rieder, S. Efficient rounding of approximate designs. *Biometrika* **1992**, *79*, 763–770. [\[CrossRef\]](#)
- Bertsimas, D.; Johnson, M.; Kallus, N. The power of optimization over randomization in designing experiments involving small samples. *Oper. Res.* **2015**, *63*, 868–876. [\[CrossRef\]](#)
- Kallus, N. On the optimality of randomization in experimental design: How to randomize for minimax variance and design-based inference. *J. R. Stat. Soc. Ser. B (Stat. Methodol.)* **2021**, *83*, 404–409. [\[CrossRef\]](#)
- Hadamard, J. Résolution d’une question relative aux déterminants. *Bull. Sci. Math.* **1893**, *2*, 240–246.
- Brenner, J. The Hadamard maximum determinant problem. *Am. Math. Mon.* **1972**, *79*, 626–630. [\[CrossRef\]](#)
- Detle, H. Designing experiments with respect to “standardized” optimality criteria. *J. R. Stat. Soc. Ser. B (Stat. Methodol.)* **1997**, *59*, 97–110. [\[CrossRef\]](#)
- IBM ILOG. *IBM ILOG CPLEX Optimization Studio*; IBM: Armonk, NY, USA, 2024.
- Rosenberger, W.F.; Sverdlov, O. Handling covariates in the design of clinical trials. *Stat. Sci.* **2008**, *23*, 404–419. [\[CrossRef\]](#)
- Villar, S.S.; Bowden, J.; Wason, J. Multi-armed bandit models for the optimal design of clinical trials: Benefits and challenges. *Stat. Sci. Rev. J. Inst. Math. Stat.* **2015**, *30*, 199. [\[CrossRef\]](#)
- Pocock, S.J. Group sequential methods in the design and analysis of clinical trials. *Biometrika* **1977**, *64*, 191–199. [\[CrossRef\]](#)

25. Atkinson, A.C. The comparison of designs for sequential clinical trials with covariate information. *J. R. Stat. Soc. Ser. A* **2002**, *165*, 349–373. [[CrossRef](#)]
26. Smith, R.L. Sequential treatment allocation using biased coin designs. *J. R. Stat. Soc. Ser. B* **1984**, *46*, 519–543. [[CrossRef](#)]
27. Atkinson, A.C.; Donev, A.N.; Tobias, R.D. *Optimum Experimental Designs, with SAS*; Oxford University Press: Oxford, UK, 2007.
28. Atkinson, A.C.; Duarte, B.P.; Pedrosa, D.J.; van Munster, M. Randomizing a clinical trial in neuro-degenerative disease. *Contemp. Clin. Trials Commun.* **2023**, *33*, 101140. [[CrossRef](#)] [[PubMed](#)]
29. Peto, V.; Jenkinson, C.; Fitzpatrick, R. PDQ-39: A review of the development, validation and application of a Parkinson's disease quality of life questionnaire and its associated measures. *J. Neurol.* **1998**, *245*, S10–S14. [[CrossRef](#)]
30. Hoehn, M.M.; Yahr, M.D. Parkinsonism: Onset, progression, and mortality. *Neurology* **1967**, *17*, 427–442. [[CrossRef](#)]
31. Hoehn, M.M.; Yahr, M.D. Parkinsonism: Onset, progression, and mortality. 1967. *Neurology* **1998**, *50*, 318–333. [[CrossRef](#)]
32. Beck, A.T.; Steer, R.A.; Pompili, M. *BHS, Beck Hopelessness Scale: Manual*; Psychological Corporation: San Antonio, TX, USA, 1988.
33. Zhao, Y.; Kosorok, M.R.; Zeng, D. Reinforcement learning design for cancer clinical trials. *Stat. Med.* **2009**, *28*, 3294–3315. [[CrossRef](#)]
34. Matsuura, K.; Honda, J.; El Hanafi, I.; Sozu, T.; Sakamaki, K. Optimal adaptive allocation using deep reinforcement learning in a dose-response study. *Stat. Med.* **2022**, *41*, 1157–1171. [[CrossRef](#)]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.