

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconomTreatment recommendation with distributional targets[☆]Anders Bredahl Kock^a, David Preinerstorfer^{b,*}, Bezirgen Veliyev^c^a University of Oxford and CREATES, Aarhus University, 10 Manor Rd, Oxford OX1 3UQ, United Kingdom^b SEPS-SEW, University of St. Gallen, Varnbuelstrasse 14, 9000, St. Gallen, Switzerland^c CREATES, Department of Economics and Business Economics, Aarhus University, Fuglesangs Alle 4, 8210, Aarhus V., Denmark

ARTICLE INFO

Article history:

Received 23 August 2021

Received in revised form 15 August 2022

Accepted 20 August 2022

Available online 16 September 2022

JEL classification:

C18

C21

C44

Keywords:

Treatment allocation

Pure exploration

Best treatment identification

Statistical decision theory

Nonparametric multi-armed bandit

ABSTRACT

We study the problem of a decision maker who must provide the best possible treatment recommendation based on an experiment. The desirability of the outcome distribution resulting from the policy recommendation is measured through a functional capturing the distributional characteristic that the decision maker is interested in optimizing. This could be, e.g., its inherent inequality, welfare, level of poverty or its distance to a desired outcome distribution. If the functional of interest is not quasi-convex or if there are constraints, the optimal recommendation may be a *mixture* of treatments. This vastly expands the set of recommendations that must be considered. We characterize the difficulty of the problem by obtaining maximal expected regret lower bounds. Furthermore, we propose two (near) regret-optimal policies. The first policy is static and thus applicable irrespectively of the subjects arriving sequentially or not in the course of the experimentation phase. The second policy can utilize that subjects arrive sequentially by successively eliminating inferior treatments and thus spends the sampling effort where it is most needed.

© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

We consider a decision maker who wants to run an experiment to identify the “best” among a set of candidate treatments. While most previous work has focused on the case of targeting the distribution with the highest mean outcome, our framework lets the decision maker use a custom-built *distributional characteristic*. This characteristic summarizes the targeted properties of the outcome distributions, e.g., their poverty-, welfare-, or inequality-implications. Such generality is particularly crucial for socio-economic decision making, as targeting solely the highest mean outcome may result in excessive inequality or poverty, which can be avoided by using an appropriately designed functional. For a discussion concerning the importance of studying distributional effects of a policy beyond the mean in the context of welfare reform research we refer to [Bitler et al. \(2006\)](#).

Compared to targeting the treatment with the highest expectation, allowing for general distributional characteristics substantially increases the complexity of the decision maker’s problem. This is the case in particular if the functional used to compare treatments is not quasi-convex. Then, there is no guarantee that there always exists a single best treatment that is at least as good as any *mixture* of treatments. Thus, the decision maker has to determine which mixture

[☆] We thank the Co-Editor, an Associate Editor, and the Referees for helpful comments and suggestions. We also thank Quentin Badolle for excellent research assistance in an early stage of this project.

* Corresponding author.

E-mail addresses: anders.kock@economics.ox.ac.uk (A.B. Kock), david.preinerstorfer@unisg.ch (D. Preinerstorfer), bveliyev@econ.au.dk (B. Veliyev).

of treatments is best. Here, a vector of mixture weights corresponds to the proportion of the population to which each treatment is rolled out after the experimentation phase. Note that the functional is generally not quasi-convex as soon as one of its “components” is not quasi-convex. If, for example, the policy maker seeks to target a distribution that combines a high expectation with low poverty, the resulting functional is not quasi-convex if the poverty measure used is not quasi-convex.

A practically important aspect that we incorporate is that the decision maker may be obliged to respect *constraints* concerning the type of mixture that can be implemented. For example, there could be treatments that cannot be given to the whole population, or treatments that have to be given to at least a certain proportion of the population. Another example of restrictions arises when there are groups of treatments that are incompatible in the sense that they cannot be implemented jointly, because, e.g., they depend on different types of infrastructures that are too costly to maintain simultaneously. Even if the functional is quasi-convex, such as the mean, the decision maker can be confronted with constraints of the just-mentioned type. Then also in this case mixtures of treatments have to be considered.

Concerning functional targets, related papers are [Kock and Thyrgaard \(2017\)](#), [Cassel et al. \(2018\)](#) and [Kock et al. \(2020a\)](#). Note, however, that the results developed there consider a completely different problem of exploration–exploitation type. The regret in these papers is cumulative in the sense that (i) every subject not assigned to the best treatment contributes to the regret, and (ii) a loss incurred for one subject cannot be compensated by future assignments. This is in contrast to the present paper, where the goal is to roll out a mixture of treatments after an experimental phase, in which the decision maker can learn which mixture is best. Mistakes in assignments during the experimental phase do not contribute to the regret, which only measures the quality of the recommended mixture. Thus, it is solely the distribution across subjects in the roll-out phase that matters. Our new objective requires new policies. Indeed, it is known already for the case of targeting the mean functional that an *upper* bound on the exploration–exploitation regret of a policy implies a *lower* bound on its regret in the pure exploration problem (studied in the present work) targeting the quality of a final recommendation, cf. [Bubeck et al. \(2009\)](#). Thus, policies that work well in the exploration–exploitation settings of [Kock and Thyrgaard \(2017\)](#), [Cassel et al. \(2018\)](#) and [Kock et al. \(2020a\)](#) need not do so in our setting and new policies must be developed for our recommendation problem. This need for new policies is amplified by the fact that the just cited works do not allow the decision maker to target mixtures of treatments, nor to incorporate constraints (such as capacity or incompatibility constraints). Although incorporating constraints could also be interesting in the papers just cited, it is debatable whether targeting a mixture per se is even sensible in exploration–exploitation problems, where the goal is to assign every individual to the best treatment.

In the present article we establish theoretical results to assist a decision maker who needs to give a recommendation based on an experiment. We study optimality properties of policies in terms of their maximal expected regret. Here, regret is defined as the difference between the distributional characteristic of the optimal mixture of treatments (over the feasible set of weights considered) and the distributional characteristic of the weights recommended by the policy.

Before we discuss how our results are related to the literature, we summarize our contributions:

1. We develop general lower bounds on the maximal expected regret of static and sequential policies. Obtaining lower bounds allows us to understand how the size of the sample used for experimentation, the number of treatments, and structural aspects of the feasible set of mixture weights affect the difficulty of the problem. The size of these lower bounds depends intrinsically on the structure of the set of feasible mixture weights. Thus, establishing useful lower bounds is a delicate matter for severely restricted subsets.
2. We study static assignment policies and investigate under which conditions such policies are optimal.
3. We investigate optimality properties of a sequential assignment policy that attempts to eliminate “inferior” (groups of) treatments during the experimentation phase. This elimination strategy can help to target the sampling effort to where it is most needed, which we also illustrate in our numerical results.

1.1. Related literature

Our article draws on ideas from the multi-armed bandit literature. Important early contributions include [Thompson \(1933\)](#), [Robbins \(1952\)](#), [Gittins \(1979\)](#), and [Lai and Robbins \(1985\)](#); cf. [Cesa-Bianchi and Lugosi \(2006\)](#), [Bubeck and Cesa-Bianchi \(2012\)](#) and [Lattimore and Szepesvári \(2020\)](#) for introductions to the subject and for further references. A large part of this literature is focused on balancing an exploration–exploitation trade-off. In contrast, our framework resembles that of a (fixed budget) “pure exploration” problem, in that the goal is to give the best recommendation by the end of the experimentation phase. Suboptimal assignments made during this phase do not enter the regret function. The pure-exploration literature dates back at least to [Bubeck et al. \(2009\)](#), cf. also [Audibert et al. \(2010\)](#), [Karnin et al. \(2013\)](#), [Jamieson et al. \(2014\)](#), [Carpentier and Locatelli \(2016\)](#), [Kaufmann et al. \(2016\)](#), and [Kasy and Sautmann \(2021\)](#). In all articles just mentioned, the goal is to find the treatment with the highest expectation in unconstrained situations. An article that suggests a sequential batch-elimination policy in the pure exploration setting for distributional targets is [Tran-Thanh and Yu \(2014\)](#). The article does not study maximal expected regret optimality properties of the policy introduced nor related performance lower bounds. Furthermore, the policy introduced only allows targeting the best individual treatment, and does not allow the decision maker to target the best mixture of treatments, which is crucial for functionals that are not quasi-convex or for constrained problems.

The vector of mixture weights the decision maker can recommend takes its values in the subset of the unit simplex described by the imposed constraints. It is tempting to interpret this problem as a functional target version of the standard continuum (or \mathcal{X} -armed) bandit problem which is primarily concerned with targeting the mean (e.g., Agrawal (1995), Kleinberg et al. (2008) and Bubeck et al. (2011)). This, however, would completely overlook the delicate role played by the mixtures in the problem we are dealing with: The decision maker can always assign a subject to only one out of finitely many treatments, not to a mixture of those treatments. That is, in continuum-armed bandit parlance, the sampling in the experimentation phase can exclusively be done from the vertices of the simplex (even though these vertices may not satisfy the constraints imposed on the final recommendation). For general continuum-armed bandit problems, such a sampling scheme is obviously not enough to find the optimal recommendation, as there is no hope to learn a function in the interior of the simplex from its values on the vertices. However, the problem we consider is such that learning the underlying finitely many outcome distributions of the treatments permits us to solve the recommendation problem, because the target function we are optimizing is a (nonlinear) function of the mixture of the unknown treatment outcome distributions. Also note that, equipped with the special additional structure just explained, a standard continuum-armed bandit problem targeting the mean functional would reduce to a pure-exploration problem with finitely many arms. Thus, it is precisely the nonlinearity of the target functional and the presence of constraints that creates problems of the type considered in the present article.

Our paper is also related to the work on statistical treatment rules initiated by Manski (2004), and developed further by Dehejia (2005), Stoye (2009), Hirano and Porter (2009), Stoye (2012), Bhattacharya and Dupas (2012), Manski and Tetenov (2016), Kitagawa and Tetenov (2018), Athey and Wager (2021). The central question attacked in these papers is to learn, in a non-sequential setting, whom to assign to the treatment based on observed covariates. The non-asymptotic results in this literature are focused on designing treatment rules maximizing the conditional mean. One exception is Kitagawa and Tetenov (2021), who studied a class of rank dependent social welfare functions as targets. Although these social welfare functions cover interesting functionals such as the (extended) Gini-welfare measure, they are quasi-convex, which greatly simplifies the space of treatment rules as no mixtures need to be considered. In addition, the authors focused exclusively on static policies. However, Kitagawa and Tetenov (2021) allow the treatment choice to be based on a vector of covariates. In this sense our results are complimentary and none is a subcase of the other.¹

For many classes of distributional characteristics classical inference problems, such as constructing a non-sequential asymptotically efficient estimator or test, or partial identification issues concerning distributional effects have been studied; e.g., Biewen (2002), Davidson and Flachaire (2007), Barrett and Donald (2009), Rothe (2010, 2012), Stoye (2010), Dufour et al. (2019). Our goal is conceptually different, in that we analyze the specific decision problem of identifying the best treatment based on an experiment of a given size. As discussed in Manski and Tetenov (2016) this task is only weakly related to significance testing, which is typically focused on controlling Type I and Type II errors. From a technical perspective, we do not use asymptotic approximations, but we mainly use (finite-sample) concentration inequalities of plug-in estimators for the functional under consideration, and information-theoretic arguments to derive our maximal expected regret lower bounds.

The structure of the remaining article is as follows: We introduce the framework in Section 2, then we present lower bounds for maximal expected regret in Section 3. Static assignment policies are discussed in Section 4, whereas sequential policies are treated in Section 5. In Section 6 we summarize the results of three numerical experiments. We provide an informal discussion of how covariates can be dealt with in Section 7. All proofs are collected in Appendices A, B, C, D, and E.

2. Framework

In this section we formally describe the observational structure, the distributional target and the objective of the decision maker, as well as the policies we consider.

2.1. Observational structure

The potential outcome of assigning subject t to treatment $i \in \mathcal{I} := \{1, \dots, K\}$, $K \geq 2$, shall be denoted by $Y_{i,t}$. This potential outcome will be interpreted as a draw from an *unknown* cumulative distribution function (cdf) F^i , i.e., the cdf obtained by rolling out treatment i to an infinitely large population. We shall assume that $F^i \in \mathcal{D} \subseteq D_{cdf}([a, b])$ holds for $i = 1, \dots, K$, where $a < b$ are real numbers, $D_{cdf}([a, b])$ denotes the set of all cdfs F on \mathbb{R} such that $F(a-) = 0$ and $F(b) = 1$, and \mathcal{D} is the common “parameter space” of the unknown outcome distributions. As such, the set \mathcal{D} describes the assumptions one is willing to put on the unknown cdfs F^1, \dots, F^K . For example, \mathcal{D} could be a set of cdfs satisfying certain smoothness conditions. The potential outcome vector of subject t will be denoted by $Y_t := (Y_{1,t}, \dots, Y_{K,t})$. Furthermore, for every t , we let G_t be a random variable which can be used for randomization in assigning the t th subject. We think of the *randomization measure*, i.e., the distribution of G_t , as being fixed, e.g., the uniform distribution on $[0, 1]$.

Throughout we impose the following assumption.

¹ We provide a heuristic discussion of how to incorporate covariates into our framework in Section 7.

Assumption 2.1. The random vectors Y_t for $t \in \mathbb{N}$ are independent and identically distributed (i.i.d.); the sequence of random variables G_t for $t \in \mathbb{N}$ is i.i.d., and is independent of the sequence Y_t . Furthermore, \mathcal{D} is convex (and non-empty).

The assumption that one can sample independently from a target population is standard in the treatment assignment literature, and is reasonable if the target population is sufficiently large. Note that no assumptions are imposed concerning the dependence between the components of each random vector Y_t .

2.2. Target of the decision maker

Rolling out multiple treatments by randomly assigning treatment i with proportion δ_i leads to the population (mixture) cdf $\sum_{i=1}^K \delta_i F^i$, which we write as $\langle \delta, \mathbf{F} \rangle$ for $\mathbf{F} = (F^1, \dots, F^K)$. Here, $\delta_i \in [0, 1]$ and $\sum_{i=1}^K \delta_i = 1$. In order to judge which vector of proportions δ is “best”, we need to compare the quality of cdfs. To this end, we assume that the decision maker is interested in maximizing a certain characteristic of the cdf, measured by a pre-specified *functional*

$$T : D_{cdf}([a, b]) \rightarrow \mathbb{R}. \tag{1}$$

Given T , the cdf $F \in D_{cdf}([a, b])$ is considered as being better than $G \in D_{cdf}([a, b])$ if $T(F) > T(G)$.

Remark 2.2. We think of the treatment recommendation as being rolled out to a large population, in which a proportion δ_i of individuals is assigned to treatment i . Thus, every single subject is only assigned to one treatment, but not every subject need to be assigned to the same treatment.² The policy maker then wishes to choose a vector of proportions that maximizes $\delta \mapsto T(\langle \delta, \mathbf{F} \rangle)$, cf. (3).

Technically, the main assumption we impose on T is as follows, where for two cdfs F and G we denote $\|F - G\|_\infty = \sup_{x \in \mathbb{R}} |F(x) - G(x)|$ (the same symbol is used for the supremum norm on \mathbb{R}^K).

Assumption 2.3. The functional $T : D_{cdf}([a, b]) \rightarrow \mathbb{R}$ and the non-empty set $\mathcal{D} \subseteq D_{cdf}([a, b])$ satisfy

$$|T(F) - T(G)| \leq C \|F - G\|_\infty \quad \text{for every } F \in \mathcal{D} \text{ and every } G \in D_{cdf}([a, b]) \tag{2}$$

for some $C > 0$.

For expositional purposes it is best to keep T abstract. Several examples are briefly discussed in the following remark, and in more detail in Section 2.2.2 further below.

Remark 2.4 (Examples). Assumption 2.3 is satisfied for many inequality measures (e.g., the Schutz-coefficient, Gini-index, linear inequality measures, generalized entropy family, Atkinson-indices, or Kolm-indices), welfare measures based on these inequality measures (e.g., the Gini-welfare measure), poverty measures (e.g., the headcount ratio, or Sen- and Foster-families of poverty measures), quantiles, U-functionals, L-functionals, or generalized trimmed mean functionals. The assumption was introduced in Kock et al. (2020a); cf. their Remarks 2.2–2.4 for some discussion, and see their Section 2.1 and Appendices E and G for a detailed discussion of examples.³

Remark 2.5. In practice a policy maker often seeks to find a balance between several policy objectives. For example, one may wish to trade off a high expectation with low poverty. Such considerations are encompassed in the present framework: For functionals T_1, \dots, T_l on $D_{cdf}([a, b])$ and a function g defined on the range of these one can define $F \mapsto T(F)$ as

$$T(F) = g(T_1(F), \dots, T_l(F)).$$

If T_1, \dots, T_l satisfy Assumption 2.3 and g is Lipschitz continuous, then also T satisfies Assumption 2.3. Finally, note that T will generally not be quasi-convex as soon as one of the building blocks T_1, \dots, T_l is not quasi-convex (unless g has further special structure). For example, numerous poverty measures are not quasi-convex, cf. Example 2.9.

If the decision maker knew F^1, \dots, F^K , the goal would be to target a vector of weights δ that maximizes $T(\langle \delta, \mathbf{F} \rangle)$. Denoting by \mathcal{M}_K the (non-empty) set of weight vectors the decision maker is restricted to work with, the task would then be to find an element of⁴

$$\arg \max_{\delta \in \mathcal{M}_K} T(\langle \delta, \mathbf{F} \rangle). \tag{3}$$

² In particular, we do not think of the vector of proportions δ as meaning that every individual is assigned a “dose” of δ_i of treatment i – the treatments cannot be divided.

³ Note that if one samples independently from $F \in \mathcal{D}$, Assumption 2.3 implies that T applied to the empirical cdf converges to $T(F)$ at the parametric rate. Thus, functionals that cannot be estimated at that rate are not covered by our theory.

⁴ Note that the set in Eq. (3) is non-empty for every $\mathbf{F} \in \mathcal{D} \times \dots \times \mathcal{D}$ if, e.g., \mathcal{D} is convex, T is continuous on \mathcal{D} w.r.t. the supremum metric induced from $D_{cdf}([a, b])$ (that is for all $F \in \mathcal{D}$ and all $\varepsilon > 0$ there exists a $\delta > 0$ such that for all $G \in \mathcal{D}$ satisfying $\|F - G\|_\infty < \delta$ it holds that $|T(F) - T(G)| < \varepsilon$.) and \mathcal{M}_K is a (non-empty) closed subset of the standard simplex of \mathbb{R}^K . This is guaranteed under Assumptions 2.1, 2.3, and 2.7.

Leaving aside for a moment that \mathbf{F} is unknown to the decision maker and that solving the optimization problem in (3) is only one part of the problem, we now introduce some notation used throughout this article and discuss in more detail frequently encountered constraints imposed upon the decision maker.

In the “unrestricted” case the set \mathcal{M}_K of weights coincides with the standard simplex in \mathbb{R}^K , which we denote as

$$\mathcal{S}_K := \{\delta = (\delta_1, \dots, \delta_K)' \in [0, 1]^K : \delta_1 + \dots + \delta_K = 1\} \subseteq \mathbb{R}^K.$$

An important example of a constrained set is $\mathcal{M}_K = \mathcal{E}_K := \{e_1(K), \dots, e_K(K)\}$, where $e_i(K)$ denotes the i th element of the standard basis in \mathbb{R}^K . This set describes the situation in which the decision maker can only recommend a single treatment out of the K treatments. Other frequently encountered constraints are *capacity constraints*, imposing upper- or lower-bound restrictions on single weights δ_i or, more generally, imposing upper- or lower-bound restrictions on the total weight put on a combination of a subset of treatments; and *similarity constraints*, leading to upper bounds on the absolute difference between weights. It is clear that capacity constraints or similarity constraints can typically be incorporated by imposing linear inequality restrictions on δ .

In the following example we shall discuss *incompatibility constraints*, which are present whenever certain groups of treatments are “incompatible”, in the sense that treatments belonging to different groups cannot be rolled out jointly. This is relevant if, e.g., groups of treatments depend on different types of infrastructures, which cannot be administered together due to budget limitations.

Example 2.6. Let m be an integer such that $1 < m \leq K$, let $\{A_1, \dots, A_m\}$ be a partition of \mathcal{I} (i.e., the incompatible groups of treatments), and set

$$\mathcal{M}_K = \bigcup_{j=1}^m \mathcal{M}_{A_j, K}, \quad \text{where } \emptyset \neq \mathcal{M}_{A_j, K} \subseteq \mathcal{S}_K := \{\delta \in \mathcal{S}_K : \delta_k = 0 \text{ if } k \notin A_j\}. \tag{4}$$

As a special case, note that the partition $A_j = \{j\}$ for $j = 1, \dots, K$, leads to $\mathcal{M}_K = \mathcal{E}_K$, i.e., when each treatment constitutes its own group. Observe furthermore that in addition to the incompatibility constraints described by the partition, additional constraints might be imposed through the choices of $\mathcal{M}_{A_j, K}$.

The structure of \mathcal{M}_K described in Eq. (4) is already very flexible. For most results, however, we shall only need to impose the following assumption ensuring that the problem is non-trivial, and that the set in (3) is non-empty (for T continuous on the convex set \mathcal{D}).

Assumption 2.7. The set $\mathcal{M}_K \subseteq \mathcal{S}_K$ contains at least two elements and is closed.

Before we discuss policies that a decision maker may employ to learn an element of (3), we comment on how the restricted set \mathcal{M}_K imposed upon the decision maker can sometimes be reduced without loss under certain conditions on T and \mathcal{D} . Furthermore, we discuss three examples of functionals.

2.2.1. Special cases where \mathcal{M}_K can be reduced

A set of potential weights \mathcal{M}_K may sometimes without loss be reduced by exploiting properties of the specific functional T and parameter space \mathcal{D} , both of which are known to the decision maker (while the cdfs $\mathbf{F} = (F^1, \dots, F^K)$ are unknown). This is the case if T is such that for any $\mathbf{F} \in \mathcal{D} \times \dots \times \mathcal{D}$ the set of maximizers in (3) contains an element of $\mathcal{M}_K^* \subseteq \mathcal{M}_K$, where \mathcal{M}_K^* does not depend on \mathbf{F} . In such a case, nothing is lost by restricting \mathcal{M}_K to \mathcal{M}_K^* , i.e., by targeting $\arg \max_{\delta \in \mathcal{M}_K^*} T(\delta, \mathbf{F})$.

The leading example of such a reduction is the case where T restricted to the convex set \mathcal{D} is continuous and is quasi-convex, that is

$$\max(T(F), T(G)) \geq \max_{\alpha \in [0, 1]} T(\alpha F + (1 - \alpha)G), \quad \text{for every } F, G \in \mathcal{D}. \tag{5}$$

In this case one may restrict \mathcal{M}_K to the set of its extreme points (i.e., the subset of elements of \mathcal{M}_K that cannot be written as a strict convex combination of two other elements of \mathcal{M}_K).⁵ For example, if $\mathcal{E}_K \subseteq \mathcal{M}_K$, one can then replace \mathcal{M}_K by \mathcal{E}_K . That is, mixtures do not need to be taken into account and it suffices to consider the individual treatments. Note further that if $\mathcal{E}_K \not\subseteq \mathcal{M}_K$, e.g., because one or more treatments have capacity constraints, the set of extreme points of \mathcal{M}_K has a more complicated structure, and may even be infinite or not closed, in which case one may work with its closure to enforce Assumption 2.7.

Whether such reductions are possible depends on T and \mathcal{D} . We note that even if Eq. (5) holds, not reducing \mathcal{M}_K obviously does not lead to a loss concerning the maximal attainable value of T .

⁵ This follows from quasi-convexity implying that $\max_{\delta \in \mathcal{M}_K} T(\delta, \mathbf{F}) = \max_{\delta \in \text{conv}(\mathcal{M}_K)} T(\delta, \mathbf{F})$, for $\text{conv}(\mathcal{M}_K)$ the convex hull of \mathcal{M}_K , together with $\max_{\delta \in \text{conv}(\mathcal{M}_K)} T(\delta, \mathbf{F})$ being attained at an extreme point of $\text{conv}(\mathcal{M}_K)$, e.g., Theorem 4.6.3 in Cambini and Martein (2009), and thus at an extreme point of \mathcal{M}_K .

2.2.2. Examples of functionals

Example 2.8. A functional for which Assumption 2.3 holds (with $C = 2(b - a)$ and $\mathcal{D} = D_{cdf}([a, b])$, cf. the discussion after Lemma E.9 in Kock et al. (2020a)) is the Gini-welfare measure

$$F \mapsto \mu(F) - \frac{1}{2} \iint |x_1 - x_2| dF(x_1) dF(x_2), \text{ where } \mu(F) = \int x dF(x). \tag{6}$$

The functional can equivalently be written as $a + \int_a^b (1 - F(x))^2 dx$, from which it easily follows that it is quasi-convex with $\mathcal{D} = D_{cdf}([a, b])$.

Example 2.9. One important poverty measure, for which Assumption 2.3 holds (e.g., with $a = 0, b = 1, \mathcal{D}$ equal to the subset of cdfs in $D_{cdf}([0, 1])$ with a density bounded from above by $s > 0$, and $C = 1 + s/2$, cf. Lemma E.10 in Kock et al. (2020a)), but which is not quasi-convex in general, is the (negative) headcount ratio with poverty line equaling half the mean

$$F \mapsto -F(\mu(F)/2). \tag{7}$$

Here we multiply by -1 , as we aim at maximizing T , and one typically wants to find the treatment combination leading to the smallest fraction of “poor” individuals. Other poverty measures that satisfy Assumption 2.3 but are not quasi-convex (in general) include the ones by Sen (1976) and Foster et al. (1984).

Example 2.10. Consider finally a situation in which the decision maker has an “ideal” cdf $F^* \in \mathcal{D} \subseteq D_{cdf}([a, b])$ in mind, and intends to assign the treatments in such a way that (δ, \mathbf{F}) resembles this ideal cdf F^* as closely as possible. Then, one could work with the functional

$$F \mapsto -\|F - F^*\|_\infty. \tag{8}$$

By the inverse triangle inequality, Assumption 2.3 is seen to be satisfied with $C = 1$ (for any $\mathcal{D} \subseteq D_{cdf}([a, b])$, and any pair of real numbers $a < b$). Similar as in the previous example, it can be shown that (5) is not generally satisfied.

2.3. Policies

The decision maker’s targets $\arg \max_{\delta \in \mathcal{M}_K} T(\delta, \mathbf{F})$ are not readily accessible, because the cdfs \mathbf{F} are unknown. Therefore, before rolling out a certain combination of treatments, the decision maker needs to learn \mathbf{F} in order to reach a recommendation concerning which δ is best. The objective of the present paper is to devise optimal policies for this problem, i.e., optimal ways of obtaining such recommendations based on the results of assigning (statically or sequentially) a small number of subjects to treatments in the course of an experimentation phase.

Essentially, a policy is a prescription concerning two steps: the first step concerns exploring the efficacies of the available treatments based on an experiment involving n subjects. Here, a policy needs to describe how to assign these subjects to one out of the K treatments. These assignments may be *static*, e.g., based on an exogenous random draw from the set of all partitions of n subjects, or *sequential*, where the assignment of the t th subject depends on the outcomes of the previously observed subjects. Note that sequential assignments can be designed so as to adaptively focus their sampling effort to where it is needed most. In the second phase, after the outcomes of all n subjects have been observed, the policy prescribes which element $\hat{\delta} \in \mathcal{M}_K$ to recommend, aiming at a recommendation $\hat{\delta}$ satisfying $T(\hat{\delta}, \mathbf{F}) \approx \max_{\delta \in \mathcal{M}_K} T(\delta, \mathbf{F})$.

Formally, a policy π with recommendations in \mathcal{M}_K is a triangular array $\{\pi_{n,t} : n \in \mathbb{N}, n \geq K, t = 1, \dots, n + 1\}$ of measurable functions, such that for every natural number $n \geq K$

$$\begin{aligned} \pi_{n,t} &: ([a, b] \times \mathbb{R})^{t-1} \times \mathbb{R} \rightarrow \mathcal{I} = \{1, \dots, K\} \quad \text{for } t = 1, \dots, n, \\ \pi_{n,n+1} &: ([a, b] \times \mathbb{R})^n \times \mathbb{R} \rightarrow \mathcal{M}_K. \end{aligned}$$

For $t = 1, \dots, n$ the outcome of $\pi_{n,t}$ is to be interpreted as the assignment of subject t , whereas the outcome of $\pi_{n,n+1}$ is to be interpreted as the final recommendation. We emphasize that a subject can be assigned to a single treatment only. During the exploration phase, this assignment is governed by $\pi_{n,t}$ for $t = 1, \dots, n$. In the “roll-out phase” $\pi_{n,n+1}$ prescribes the proportion of subjects assigned to each treatment.

The input of $\pi_{n,t}$ is denoted by (Z_{t-1}, G_t) and is recursively defined, where $(Z_0, G_1) := (G_1)$ and one defines the $2(t - 1)$ dimensional random vector

$$Z_{t-1} = (Y_{\pi_{n,t-1}(Z_{t-2}, G_{t-1}), t-1}, G_{t-1}, \dots, Y_{\pi_{n,1}(G_1), 1}, G_1), \tag{9}$$

i.e., Z_{t-1} is the complete history of outcomes and randomizations observed before subject t arrives. Note that each assignment, and the final recommendation, can incorporate a draw from an exogenous random variable.

The objective is to use the outcomes of the n subjects observed to give a recommendation $\pi_{n,n+1}$ in \mathcal{M}_K , such that $T(\pi_{n,n+1}(Z_n, G_{n+1}), \mathbf{F})$ is close to $\max_{\delta \in \mathcal{M}_K} T(\delta, \mathbf{F})$. Therefore, we will evaluate policies based on their *regret*

$$r_n(\pi, \mathcal{M}_K) = r_n(\pi, \mathcal{M}_K; \mathbf{F}; Z_n, G_{n+1}) := \max_{\delta \in \mathcal{M}_K} T(\delta, \mathbf{F}) - T(\pi_{n,n+1}(Z_n, G_{n+1}), \mathbf{F}). \tag{10}$$

Note that $r_n(\pi, \mathcal{M}_K)$ measures the “out-of-sample” performance of the recommendation $\pi_{n,n+1}(Z_n, G_{n+1})$. The regret we use in the present paper does not measure how “well” the n subjects used for obtaining the recommendation are assigned, it only incorporates the quality of the final recommendation. This makes sense in our framework, as n is assumed to be small compared to the whole population. In problems where n is relatively large, however, one may want to work with a different regret criterion also incorporating losses made during the experimentation phase. This problem is fundamentally different, and we refer the reader to [Kock et al. \(2020a\)](#), where an “in-sample” theory for a corresponding “individual-specific regret criterion” was developed.

3. Performance lower bounds

In this section we present lower bounds on maximal expected regret for policies as discussed in Section 2.3. The lower bounds are given under weak conditions on the structure of \mathcal{M}_K , allowing for rich forms of constraints. Besides delivering insights into the difficulty of the problem, the lower bounds will be instrumental in asserting that the specific policies introduced in later sections are optimal.

To rule out trivial situations, we need an assumption which guarantees a minimal amount of variation of the functional T evaluated on \mathcal{D} (if T is constant on \mathcal{D} the regret of any policy is constant equal to 0, an uninteresting situation).

Assumption 3.1. The functional $T : D_{cdf}([a, b]) \rightarrow \mathbb{R}$ satisfies [Assumption 2.3](#), and \mathcal{D} contains two elements H_1 and H_2 , such that

$$J_\tau := \tau H_1 + (1 - \tau)H_2 \in \mathcal{D} \quad \text{for every } \tau \in [0, 1],$$

and such that for some $c_- > 0$ we have

$$T(J_{\tau_2}) - T(J_{\tau_1}) \geq c_-(\tau_2 - \tau_1) \quad \text{for every } \tau_1 \leq \tau_2 \text{ in } [0, 1]. \tag{11}$$

As already noted in [Kock et al. \(2020b\)](#), where this assumption was introduced, we emphasize that [Eq. \(11\)](#) in [Assumption 3.1](#) is satisfied if, e.g., $\tau \mapsto T(J_\tau)$ is continuously differentiable on $[0, 1]$ with an everywhere positive derivative. Note that [Assumption 3.1](#) is weak, as one is free to choose H_1 and H_2 . In particular, it is typically satisfied as long as \mathcal{D} is reasonably large. For instance, as can be easily seen, this is the case for the examples given in [Section 2.2.2](#).

We start with a lower bound into which \mathcal{M}_K enters only via its squared diameter (with respect to the Euclidean norm $\|\cdot\|$), which we shall abbreviate (under [Assumption 2.7](#)) as

$$\text{diam}(\mathcal{M}_K) := \max\{\|v - \gamma\| : (v, \gamma) \in \mathcal{M}_K \times \mathcal{M}_K\}. \tag{12}$$

Obviously, $\text{diam}(\mathcal{M}_K) = 0$ if and only if \mathcal{M}_K is a singleton, an uninteresting case which is ruled out in [Assumption 2.7](#). The lower bound is as follows.

Theorem 3.2. Suppose [Assumptions 2.1, 2.3, 2.7 and 3.1](#) hold. Then there exists a constant $c > 0$, independent of K , n and \mathcal{M}_K , such that for every policy π with recommendations in \mathcal{M}_K , and any randomization measure \mathbb{P}_G , it holds that

$$\sup_{F^1, \dots, F^K \in \mathcal{U}_\tau : \tau \in [0, 1]} \mathbb{E}[r_n(\pi, \mathcal{M}_K)] \geq c \times \text{diam}^2(\mathcal{M}_K) / \sqrt{n}, \quad \text{for every } n \geq K, \tag{13}$$

where the supremum is taken over all potential outcome vectors with independent marginals and cdfs in $\{J_\tau : \tau \in [0, 1]\}$.

The lower bound in [Theorem 3.2](#) actually holds over the parametric subset $\{J_\tau : \tau \in [0, 1]\}$ (obtained from [Assumption 3.1](#)) of the potentially much larger and nonparametric set \mathcal{D} . This also holds for all other bounds in this section.

The main strength of [Theorem 3.2](#) is that it delivers a useful lower bound under minimal assumptions on \mathcal{M}_K : as soon as \mathcal{M}_K contains 2 elements that are bounded away from each other (uniformly over K), the theorem shows that there exists no policy with maximal expected regret decreasing at a rate faster than $1/\sqrt{n}$. The minimal assumption comes with a cost: namely that the theorem is silent about how the number of treatments affects the worst-case behavior of a policy. To make this more concrete, note for example that $\mathcal{M}_2 = \{e_1(2), e_2(2)\}$ and $\mathcal{M}_3 = \{e_1(3), e_2(3), e_3(3)\}$ both have diameter $\sqrt{2}$, and thus both lead to the same lower bound in the previous theorem, even though \mathcal{M}_3 obviously is more complex than \mathcal{M}_2 .

The remaining part of this section thus establishes lower bounds that also incorporate the dependence on K (in an optimal way, as the next section will show). It is clear that more refined structural properties of \mathcal{M}_K now need to enter the picture: the lower bounds need to reflect that the decision problem is more difficult if \mathcal{M}_K is very much “spread out” over \mathcal{S}_K instead of, e.g., being concentrated in a single corner of the simplex.

To formulate the second result in this section, we now define the quantity $\kappa(\mathcal{M}_K)$, which summarizes the structural properties of \mathcal{M}_K that enter into our general lower bound. The characteristic $\kappa(\mathcal{M}_K)$ captures, in a way tailored towards our method of proof, how “spread out” the set \mathcal{M}_K is over the simplex, and (under [Assumption 2.7](#)) is defined as

$$\kappa(\mathcal{M}_K) := \sup \left\{ \left[\max_{\delta \in \mathcal{M}_K} v' \delta - \sup_{\delta \in \mathcal{T}} v' \delta \right] \wedge \min_{i=1}^K \left[\max_{\delta \in \mathcal{M}_K} (v' \delta + w_i \delta_i) - \sup_{\delta \in \mathcal{M}_K \setminus \mathcal{T}} (v' \delta + w_i \delta_i) \right] \right\},$$

where the outer supremum is taken over all nonempty Borel sets $\mathcal{T} \subseteq \mathcal{M}_K$, all $v \in [-1, 1]^K$, and all $w_i \in [-1, 1]$ ($i = 1, \dots, K$). It is easy to see that $\kappa(\mathcal{M}_K) \geq 0$ under [Assumption 2.7](#).

Note that $\kappa(\mathcal{M}_K)$ is large if there exists a vector v , such that the maximum of $\delta \mapsto v'\delta$ decreases substantially upon imposing the restriction $\delta \in \mathcal{T}$, but where imposing the *complementary* restriction $\delta \notin \mathcal{T}$ decreases the maximum substantially upon suitably modifying *any* single coordinate of v . Since $\kappa(\mathcal{M}_K)$ is a rather abstract quantity, the following result also provides two lower bounds for $\kappa(\mathcal{M}_K)$ that are easy to interpret.

Theorem 3.3. *Suppose [Assumptions 2.1, 2.3, 2.7](#) and [3.1](#) hold. Then there exists a constant $c > 0$, independent of K, n and \mathcal{M}_K , such that for every policy π with recommendations in \mathcal{M}_K , and any randomization measure \mathbb{P}_G , it holds that*

$$\sup_{F^1, \dots, F^K \in \{J_\tau : \tau \in [0, 1]\}} \mathbb{E}[r_n(\pi, \mathcal{M}_K)] \geq c\kappa(\mathcal{M}_K)\sqrt{K/n}, \quad \text{for every } n \geq K, \tag{14}$$

where the supremum is taken over all potential outcome vectors with independent marginals and cdfs in $\{J_\tau : \tau \in [0, 1]\}$. Furthermore,

$$\kappa(\mathcal{M}_K) \geq \frac{1}{272} \times \max_{\delta \in \mathcal{M}_K} \|\delta\|^2 \times \min_{j=1}^K (\max_{\delta \in \mathcal{M}_K} \delta_j - \min_{\delta \in \mathcal{M}_K} \delta_j)^3; \tag{15}$$

and $\kappa(\mathcal{M}_K)$ is also lower bounded by the squared positive part of $\min_{j=1}^K \max_{\delta \in \mathcal{M}_K} \delta_j - 1/2$.

An important situation in [Theorem 3.3](#) is $\mathcal{M}_K \supseteq \mathcal{E}_K$. In this case, the second lower bound on $\kappa(\mathcal{M}_K)$ provided in the last sentence of the theorem implies $\kappa(\mathcal{M}_K) \geq 1/4$. Together with [Eq. \(14\)](#) this proves that no policy can have maximal expected regret decreasing at a faster rate than $\sqrt{K/n}$. The same is true for any \mathcal{M}_K as long as $\min_{j=1}^K \max_{\delta \in \mathcal{M}_K} \delta_j - 1/2$ is bounded away from 0 (uniformly in K).

For the special case in which the functional T is the mean functional, $a = 0, b = 1$, and $\mathcal{M}_K = \mathcal{E}_K$, a $\sqrt{K/n}$ lower bound is given in the distribution-free lower bound mentioned in [Bubeck et al. \(2009\)](#). The lower bound in [Theorem 3.3](#) is a non-trivial generalization of the lower bound in [Bubeck et al. \(2009\)](#): besides working with a general functional T and operating over a corresponding line segment $\{J_\tau : \tau \in [0, 1]\}$ in \mathcal{D} along which T varies sufficiently, we do not impose the condition $\mathcal{M}_K = \mathcal{E}_K$. This makes the situation substantially more delicate, as the recommendation then takes its values in a potentially complicated set \mathcal{M}_K .

At this point, one may wonder whether the lower bound obtained in [Theorem 3.3](#) can actually be attained in case $\mathcal{M}_K \neq \mathcal{E}_K$ for the following reason: Note that if $\mathcal{M}_K = \mathcal{E}_K$ the decision maker “only” needs to learn the K values $T(F^1), \dots, T(F^K)$ to determine $\arg \max_{\delta \in \mathcal{M}_K} T(\delta, \mathbf{F})$, whereas in case, e.g., $\mathcal{M}_K = \mathcal{I}_K$, one needs to learn $T(\delta, \mathbf{F})$ for *all* $\delta \in \mathcal{I}_K$. In the next section we shall see that lower bounds of rate $\sqrt{K/n}$ are attainable (up to logarithmic terms in K).

In situations where \mathcal{M}_K restricts a weight to $\delta_j \leq 1/2$, and the discussion right after [Theorem 3.3](#) thus does not apply, the lower bound for $\kappa(\mathcal{M}_K)$ given in [Eq. \(15\)](#) (and a simple argument) shows that

$$\max_{\delta \in \mathcal{M}_K} \|\delta\|^2 \times \min_{j=1}^K (\max_{\delta \in \mathcal{M}_K} \delta_j - \min_{\delta \in \mathcal{M}_K} \delta_j)^3 \geq \min_{j=1}^K (\max_{\delta \in \mathcal{M}_K} \delta_j - \min_{\delta \in \mathcal{M}_K} \delta_j)^5;$$

as long as this lower bound is bounded away from 0 (uniformly in K), no policy exists that has maximal expected regret decreasing at a faster rate than $\sqrt{K/n}$. This is a substantial generalization beyond the case already discussed after [Theorem 3.3](#), and in particular covers most cases of practical relevance where constraints are put on each treatment. Note in particular that if \mathcal{M}_K is such that $\max_{\delta \in \mathcal{M}_K} \delta_j - \min_{\delta \in \mathcal{M}_K} \delta_j > 0$ for every $j = 1, \dots, K$, i.e., if every δ_j can take at least two different values, then the lower bound for $\kappa(\mathcal{M}_K)$ given in [Eq. \(15\)](#) is strictly greater than 0.

Note also that in case $\max_{\delta \in \mathcal{M}_K} \delta_j - \min_{\delta \in \mathcal{M}_K} \delta_j = 0$ for a treatment j , this treatment always has to be assigned with the same weight. In this case the lower bound in [Eq. \(15\)](#) is 0, and furthermore also $\kappa(\mathcal{M}_K) = 0$ holds.⁶ Hence, [Theorem 3.3](#) does *not* deliver a $\sqrt{K/n}$ lower bound; it is nevertheless worth mentioning that [Theorem 3.2](#) still delivers a lower bound of order $1/\sqrt{n}$ as long as \mathcal{M}_K is not a singleton. In this situation, one could expect a lower bound to hold true that incorporates, instead of K , the number of treatments for which $\max_{\delta \in \mathcal{M}_K} \delta_j - \min_{\delta \in \mathcal{M}_K} \delta_j \neq 0$. Such a result can be established using [Theorem B.2](#) in [Appendix B.1](#), which fills this “gap” between [Theorems 3.2](#) and [3.3](#), but comes at the expense of more complicated notation. Since this case is very special, we shall not discuss it further.

4. Static assignment policies

We now consider a class of *static assignment* (SA) policies incorporating an “empirical-success” recommendation rule. Given balanced assignments in the experimentation phase, these policies will be shown to be minimax expected regret optimal, in the sense that they attain the $\sqrt{K/n}$ lower bound established in [Theorem 3.3](#) (up to $\sqrt{\log(K)}$). In the following discussion we assume that $n \geq K$, that T is continuous on the convex set \mathcal{D} , and that \mathcal{M}_K satisfies [Assumption 2.7](#).

⁶ That $\kappa(\mathcal{M}_K) = 0$ if $\max_{\delta \in \mathcal{M}_K} \delta_j - \min_{\delta \in \mathcal{M}_K} \delta_j = 0$ for treatment j follows upon noting that for any choice of $\mathcal{T}, v \in [-1, 1]^K$, and $w_j \in [-1, 1]$, either $\max_{\delta \in \mathcal{M}_K} v'\delta = \sup_{\delta \in \mathcal{T}} v'\delta$ or $\max_{\delta \in \mathcal{M}_K} (v'\delta + w_j\delta_j) = \sup_{\delta \in \mathcal{M}_K \setminus \mathcal{T}} (v'\delta + w_j\delta_j)$.

An SA policy proceeds as follows: First, it allocates subjects $t = 1, \dots, n$ to treatment j according to whether $t \in \Pi_{n,j}$, where $\Pi_n := (\Pi_{n,1}, \dots, \Pi_{n,K})$ is a partition of $\{1, \dots, n\}$. For simplicity, we treat the partition as fixed; random partitions that are obtained by an exogenous randomization mechanism can be easily accommodated via conditioning. Once all n outcomes are observed, one estimates every F^j by the corresponding empirical cdf

$$\hat{F}_{n,\Pi_n}^j(\cdot) := |\Pi_{n,j}|^{-1} \sum_{t \in \Pi_{n,j}} \mathbb{1}\{Y_{j,t} \leq \cdot\}, \tag{16}$$

where $|\Pi_{n,j}|$ denotes the cardinality of $\Pi_{n,j}$. The obvious way to proceed would now be to search for a maximizer of $T(\langle \delta, \hat{\mathbf{F}}_{n,\Pi_n} \rangle)$ over \mathcal{M}_K , where $\hat{\mathbf{F}}_{n,\Pi_n} := (\hat{F}_{n,\Pi_n}^1, \dots, \hat{F}_{n,\Pi_n}^K)$. However, without further assumptions on T , there is no guarantee that such a maximizer exists, because $\delta \mapsto T(\langle \delta, \hat{\mathbf{F}}_{n,\Pi_n} \rangle)$ need not be continuous (even under Assumption 2.3; noting that the coordinates of $\hat{\mathbf{F}}_{n,\Pi_n}$ are not necessarily elements of \mathcal{D}). To avoid adding additional assumptions, we shall work with an approximate maximum, based on a discretization of \mathcal{M}_K , i.e., a non-empty and finite set $\mathcal{M}_K^n \subseteq \mathcal{M}_K$. Based on the discretization \mathcal{M}_K^n the SA policy then recommends

$$\min_{\delta \in \mathcal{M}_K^n} \arg \max T(\langle \delta, \hat{\mathbf{F}}_{n,\Pi_n} \rangle), \tag{17}$$

where the minimum in (17) is taken as a concrete way of breaking ties and where the ambient set \mathcal{S}_K is equipped with the lexicographic order.⁷ SA policies will be denoted by the generic symbol $\hat{\pi}$, and are summarized in Policy 1 for later reference. Note that the assignments of the subjects $t = 1, \dots, n$ neither depend on Z_{t-1} (i.e., the policy is static) nor incorporate external randomization G_t , which are therefore dropped as arguments from the policy in the summary in Policy 1. Note also that the recommendation given does not incorporate an external randomization G_{n+1} , which we therefore drop from the notation in Policy 1 as well.

Policy 1: Static Assignment Policy $\hat{\pi}$

Input: $n \in \mathbb{N}$, $K \leq n$, partition Π_n of $\{1, \dots, n\}$, discretization \mathcal{M}_K^n of \mathcal{M}_K

for $t = 1, \dots, n$ **do**

$$\quad | \hat{\pi}_{n,t} = \sum_{i=1}^K i \mathbb{1}_{\Pi_{n,i}}(t)$$

end

$$\hat{\pi}_{n,n+1}(Z_n) = \min \arg \max_{\delta \in \mathcal{M}_K^n} T(\langle \delta, \hat{\mathbf{F}}_{n,\Pi_n} \rangle)$$

4.1. Choosing a discretization

Associate to any discretization $\mathcal{M}_K^n \subseteq \mathcal{M}_K$ its (worst case) “optimization error”

$$\varepsilon(n) = \varepsilon(n, \mathcal{D}, \mathcal{M}_K, \mathcal{M}_K^n) := \sup_{\mathbf{F} \in \mathcal{D} \times \dots \times \mathcal{D}} \left| \max_{\delta \in \mathcal{M}_K} T(\langle \delta, \mathbf{F} \rangle) - \max_{\delta \in \mathcal{M}_K^n} T(\langle \delta, \mathbf{F} \rangle) \right|, \tag{18}$$

i.e., the maximal loss possible by optimizing over \mathcal{M}_K^n rather than \mathcal{M}_K . In case \mathcal{M}_K has finitely many elements, one may choose $\mathcal{M}_K^n = \mathcal{M}_K$, implying $\varepsilon(n, \mathcal{D}, \mathcal{M}_K, \mathcal{M}_K^n) = 0$. If \mathcal{M}_K is not finite (and cannot be reduced to a finite set, cf. Section 2.2.1), or if \mathcal{M}_K is finite but large, controlling the optimization error is often more conveniently done by choosing a fine enough discretization. Note that under Assumption 2.3 and for \mathcal{D} convex

$$\varepsilon(n, \mathcal{D}, \mathcal{M}_K, \mathcal{M}_K^n) \leq CK \sup_{\delta \in \mathcal{M}_K} \inf_{\gamma \in \mathcal{M}_K^n} \|\delta - \gamma\|_\infty =: CK \varrho(\mathcal{M}_K, \mathcal{M}_K^n), \tag{19}$$

and we refer to $\varrho(\mathcal{M}_K, \mathcal{M}_K^n)$ as the (worst-case) “resolution error” of the discretization \mathcal{M}_K^n of \mathcal{M}_K . Since \mathcal{M}_K is throughout assumed to be closed, a finite $\mathcal{M}_K^n \subseteq \mathcal{M}_K$ with $\varrho(\mathcal{M}_K, \mathcal{M}_K^n)$ (and hence $\varepsilon(n)$) as small as one wishes always exists. To give a concrete example, a common way of constructing discretizations is discussed next for the unrestricted case $\mathcal{M}_K = \mathcal{S}_K$ and for sets \mathcal{M}_K incorporating incompatibility constraints (but are otherwise unrestricted).

Example 4.1. Let \mathcal{M}_K be as in (4) (potentially with $m = 1$) with $\mathcal{M}_{A_j,K} = \mathcal{S}_{A_j,K}$ for $j = 1, \dots, m$. For $P \in \mathbb{N}$ define

$$\mathcal{S}_{K,P} := \{\delta \in \mathcal{S}_K : P\delta \in \mathbb{Z}^K\}, \tag{20}$$

i.e., the intersection of the simplex \mathcal{S}_K and $P^{-1}\mathbb{Z}^K$. Theorem 7 of Bomze et al. (2014) implies that for every $x \in \mathcal{S}_K$ there exists a $y \in \mathcal{S}_{K,P}$, such that $\|x - y\|_\infty \leq P^{-1}(1 - 1/K)$. The same theorem, but applied to the simplex $\mathcal{S}_{A_i,K}$, shows that for

⁷ Here, for two vectors $a \neq b$ of real numbers, a is lexicographically smaller than b if $a_{i^*} < b_{i^*}$ for the smallest index i^* such that $a_{i^*} \neq b_{i^*}$.

every $x \in \mathcal{S}_{A_i, K}$ there exists a $y \in \mathcal{S}_{A_i, K} \cap \mathcal{S}_{K, P}$ such that $\|x - y\|_\infty \leq P^{-1}(1 - 1/|A_i|)$. It thus follows that for every $\varrho_n > 0$ the discretization

$$\mathcal{M}_K^n := \mathcal{M}_K \cap \mathcal{S}_{K, \lceil \varrho_n^{-1} \rceil} = \bigcup_{i=1}^m \mathcal{S}_{A_i, K} \cap \mathcal{S}_{K, \lceil \varrho_n^{-1} \rceil} \tag{21}$$

satisfies $\varrho(\mathcal{M}_K, \mathcal{M}_K^n) \leq \varrho_n$ and $\mathcal{M}_K^n \subseteq \mathcal{M}_K$.

4.2. Maximal expected regret upper bound for $\hat{\pi}$

We now prove an upper bound on the maximal expected regret of $\hat{\pi}$. One question we have not discussed so far is the measurability of $\hat{\pi}$, which holds under the following weak condition as shown in the proof of [Theorem 4.3](#) in Appendix C.

Assumption 4.2. For every $m \in \mathbb{N}$, and every $\delta \in \mathcal{S}_m$ the function on $[a, b]^m$ defined via $x \mapsto \mathbb{T}(\sum_{j=1}^m \delta_j \mathbb{1}\{x_j \leq \cdot\})$ is Borel measurable.

Theorem 4.3. Suppose [Assumptions 2.1, 2.3, 2.7, and 4.2](#) hold. Then the SA policy $\hat{\pi}$ with partition Π_n , discretization \mathcal{M}_K^n and $\beta_n := \min_{j=1, \dots, K} |\Pi_{n,j}|$ satisfies

$$\sup_{\substack{F^i \in \mathcal{D} \\ i=1, \dots, K}} \mathbb{E} [r_n(\hat{\pi}, \mathcal{M}_K)] \leq \varepsilon(n) + 3.01 \times C \sqrt{\log(K)/\beta_n}, \text{ for every } n \geq K. \tag{22}$$

[Theorem 4.3](#) shows that working with a discretization having maximal optimization error $\varepsilon(n) \leq \sqrt{K/n}$ (cf. Eq. (18)) together with “balanced partitions”, i.e., partitions satisfying $|\Pi_{n,j}| \geq \lfloor n/K \rfloor$ for every $n \in \mathbb{N}$ and every $j = 1, \dots, K$, results in an SA policy that attains the lower bound established in [Theorem 3.3](#), up to a $\sqrt{\log(K)}$ factor (and multiplicative constants). Hence, for such a choice of partition and discretization, the SA policy is (near) minimax expected regret optimal. Note that the optimality statement also holds in the high-dimensional regime where the number of treatments grows with n . By Eq. (19), a discretization satisfying $\varepsilon(n) \leq \sqrt{K/n}$ can be obtained by choosing \mathcal{M}_K^n such that its resolution error $\varrho(\mathcal{M}_K, \mathcal{M}_K^n) \leq \sqrt{K/n}/(CK)$, cf. [Example 4.1](#) for specific constructions.

In the special case where \mathbb{T} is the mean functional (which is quasi-convex) and $\mathcal{M}_K = \mathcal{E}_K$, the SA policy $\hat{\pi}$ with $\mathcal{M}_K^n = \mathcal{E}_K$ and balanced assignment reduces to the “uniform allocation policy” in [Bubeck et al. \(2009\)](#) and to the empirical-success rule studied in [Manski \(2004\)](#) and [Manski and Tetenov \(2016\)](#). In this special case [Assumption 2.3](#) is satisfied with $C = b - a$, $\varepsilon(n) = 0$, and [Theorem 4.3](#) delivers an upper bound with the same dependence on K and n as the upper bound in the just mentioned articles, but with an additional factor 3.01. This additional factor is due to our different method of proof that needs to deal with the possibility that in general $\mathcal{E}_K \neq \mathcal{M}_K$ (and that \mathbb{T} is not the mean functional).

5. Sequential elimination policies

The static assignment policy $\hat{\pi}$ assigns each subject $t = 1, \dots, n$ to a treatment in \mathcal{I} according to the partition Π_n . This partition is fixed by the decision maker before the first assignment is made. Thus, the policy $\hat{\pi}$ is static. In the present section we consider policies that are sequential, in the sense that the assignment of subject $t \in \{1, \dots, n\}$ can depend on the outcomes of all previously assigned subjects $1, \dots, t - 1$. Intuitively, this opens up the following opportunity: by sequentially monitoring the performance of the treatments, one can target the sampling effort to where it is most useful, rather than deciding up front to assign each treatment, e.g., equally often.

We already know from [Theorem 4.3](#) that the maximal expected regret lower bound from [Theorem 3.3](#) is attainable (up to a $\sqrt{\log(K)}$ term) within the class of static policies. Therefore, not much can be gained in terms of worst-case expected regret from sequential policies. Nevertheless, it is plausible that a policy which exploits that subjects arrive sequentially can improve on the static policy $\hat{\pi}$ for many potential outcome distributions $\mathbf{F} = (F^1, \dots, F^K) \in \mathcal{D} \times \dots \times \mathcal{D}$, without having a higher worst-case expected regret. For the policies suggested in the present article this is confirmed by [Theorems 5.1 and 5.4](#) as well as the numerical results in Section 6.

Essentially, we propose sequential policies which are based on the following rationale: *stop assigning “inferior” treatments as soon as possible, and do not include treatments in the recommendation once eliminated.* Although it is clear in the special case of $\mathcal{M}_K = \mathcal{E}_K$ that all treatments i not contained in $\arg \max_{j \in \mathcal{I}} \mathbb{T}(F^j)$ are inferior, it is less clear in the general case what an inferior treatment is. We shall consider treatment i to be inferior, if *all* elements of $\arg \max_{\delta \in \mathcal{M}_K} \mathbb{T}((\delta, \mathbf{F}))$ have zero i th coordinate, i.e., if

$$\arg \max_{\delta \in \mathcal{M}_K} \mathbb{T}((\delta, \mathbf{F})) \cap \{\delta \in \mathcal{M}_K : \delta_i > 0\} = \emptyset. \tag{23}$$

In this case treatment i does not contribute to any maximizer. Note that the set of maximizers in the previous display is non-empty if \mathbb{T} is continuous on \mathcal{D} , convex, and \mathcal{M}_K is as in [Assumption 2.7](#), which we shall assume throughout this section. Since \mathbf{F} is unknown and \mathbb{T} is continuous, estimation uncertainty implies that an inferior treatment i is only empirically

detectable (while controlling the error probability of prematurely concluding inferiority of treatments) if it is actually “strongly inferior” in the sense that

$$\arg \max_{\delta \in \mathcal{M}_K} T((\delta, \mathbf{F})) \cap \overline{\{\delta \in \mathcal{M}_K : \delta_i > 0\}} = \emptyset \tag{24}$$

(that is $\{\delta \in \mathcal{M}_K : \delta_i > 0\}$ in (23) is replaced by its closure). It is easy to see that structural assumptions need to be put on \mathcal{M}_K in order to guarantee that strongly inferior treatments exist for some \mathbf{F} . As an example, no strongly inferior treatments exist for $\mathcal{M}_K = \mathcal{S}_K$, since then $\overline{\{\delta \in \mathcal{S}_K : \delta_i > 0\}} = \mathcal{S}_K$ for every i . In such cases, attempting to eliminate treatments based on assessing (24) can never result in efficiency gains over $\hat{\pi}$, because (24) is then never satisfied.

In the construction of our sequential policies it thus only makes sense to consider \mathcal{M}_K that do not a priori rule out the existence of strongly inferior treatments. We shall therefore mainly focus on $\mathcal{M}_K = \bigcup_{j=1}^m \mathcal{M}_{A_j, K}$ for a partition A_1, \dots, A_m of \mathcal{I} as in Example 2.6, which covers many situations of practical relevance (a sequential policy that also works without this structural assumption is considered in Section 5.3). Since T is continuous on \mathcal{D} , convex, the condition

$$\max_{\delta \in \mathcal{M}_K} T((\delta, \mathbf{F})) = \max_{\delta \in \bigcup_{i \neq j} \mathcal{M}_{A_i, K}} T((\delta, \mathbf{F})) > \max_{\delta \in \mathcal{M}_{A_j, K}} T((\delta, \mathbf{F})) \tag{25}$$

is then sufficient for all treatments in A_j to be strongly inferior, because the closure of $\mathcal{M}_{A_j, K}$ is contained in $\mathcal{S}_{A_j, K}$, which is disjoint with $\bigcup_{i \neq j} \mathcal{M}_{A_i, K}$. In the special case $\mathcal{M}_K = \mathcal{E}_K$, amounting to $m = K$, the relation in (25) is actually equivalent to

$$\max_{i \in \mathcal{I}} T(F^i) = \max_{i \neq j} T(F^i) > T(F^j). \tag{26}$$

The policies introduced in the present section eliminate strongly inferior treatments based on verifying an empirical equivalent of (25). Once the data firmly suggest (25), all treatments in A_j are eliminated. Our policy carefully needs to avoid concluding (25) prematurely, in case this inequality is false. Whether or not the validity of (25) can actually be confirmed depends on the difference between the two maxima in (25) in relation to sample size. In case all treatments in A_1 and all treatments in A_2 are strongly inferior, it can for example happen that (25) can be concluded for $j = 1$ at a much earlier stage than for $j = 2$. Therefore, in a sequential policy, the set of indices over which (25) is assessed evolves during the sampling process, and depends on the previously eliminated treatments.

Before we proceed to the policies we suggest, we need some notation. Given a policy π , $n \in \mathbb{N}$ and $t = 1, \dots, n$, we denote the number of times treatment i has been assigned up to time t by

$$S_{i,n}(t) := \sum_{s=1}^t \mathbb{1}\{\pi_{n,s}(Z_{s-1}, G_s) = i\}. \tag{27}$$

On the event $\{S_{i,n}(t) > 0\}$ we define the empirical cdf based on the outcomes of all subjects in $\{1, \dots, t\}$ that have been assigned to treatment i

$$\hat{F}_{i,t,n}(z) := S_{i,n}^{-1}(t) \sum_{\substack{1 \leq s \leq t \\ \pi_{n,s}(Z_{s-1}, G_s) = i}} \mathbb{1}\{Y_{i,s} \leq z\}, \quad \text{for every } z \in \mathbb{R}; \tag{28}$$

we leave $\hat{F}_{i,t,n}$ undefined on $\{S_{i,n}(t) = 0\}$. Note that the random sampling times s such that $\pi_{n,s}(Z_{s-1}, G_s) = i$ depend on previously observed outcomes. Finally, we define $\hat{\mathbf{F}}_{t,n} = (\hat{F}_{1,t,n}, \dots, \hat{F}_{K,t,n})$ on the event $\{S_{i,n}(t) > 0 \text{ for every } i = 1, \dots, K\}$.

Equipped with this notation, we can now introduce and discuss sequential policies. To ease the exposition, we start with the simplest case $\mathcal{M}_K = \mathcal{E}_K$, which is then extended to all \mathcal{M}_K as in Example 2.6. Finally, we shall go beyond incompatibility constraints and consider a sequential policy, which can be applied without \mathcal{M}_K possessing any particular structure (whether anything can be gained in practice compared to a non-sequential policy then depends on whether the set \mathcal{M}_k allows for the existence of strongly inferior treatments).

5.1. The case $\mathcal{M}_K = \mathcal{E}_K$

We shall refer to the sequential policy introduced as the *sequential elimination* (SE) policy $\tilde{\pi}$. In the SE policy, the treatments are assigned in rounds. In every round r , all treatments $\mathcal{I}_{r-1} \subseteq \mathcal{I}$, say, that have not been eliminated in one of the previous rounds are assigned exactly once. Elimination is based on checking whether the data observed so far firmly suggest Eq. (26). This is done as follows: at the end of round r we eliminate all treatments $i \in \mathcal{I}_{r-1}$ for which

$$\max_{j \in \mathcal{I}_{r-1}} T(\hat{F}_{j,t,n}) > T(\hat{F}_{i,t,n}) + u_\eta(r, n); \tag{29}$$

i.e., one defines \mathcal{I}_r as the subset of elements of \mathcal{I}_{r-1} that do not satisfy (29). Here $u_\eta(r, n) : \mathbb{N}^2 \rightarrow [0, \infty)$ is a threshold function giving a “critical value” determining whether a treatment can be eliminated. We shall base our policies upon the threshold

$$u_\eta(r, n) = C \sqrt{\frac{1 + \eta}{r} [0.5 \log(n) + \log(rK)]}, \quad \text{for a tuning parameter } \eta > 0, \tag{30}$$

and where C is the constant from [Assumption 2.3](#). Note that the more rounds r have been completed, the better the unknown cdfs can be estimated, and thus the more aggressively $\tilde{\pi}$ eliminates treatments.

The decision maker may not want to check [\(29\)](#) and update \mathcal{I}_r after every single round. For example, subjects may arrive in batches. We allow for this by assuming that the decision maker a priori decides on a subset of “elimination rounds” $\mathcal{R} \subseteq \{1, \dots, \lfloor \frac{n}{2} \rfloor\}$, after which [\(29\)](#) is assessed, and treatments are potentially eliminated. Note that there are at most $\lfloor n/2 \rfloor$ rounds at the beginning of which there are at least two treatments left. We shall denote by $\underline{r} := \min \mathcal{R}$ the first round after which elimination takes place, and assume throughout that $\underline{r} \leq \lfloor n/K \rfloor$, since otherwise an elimination round can never be reached. To provide some examples and motivation consider the following relevant special cases:

1. $\mathcal{R} = \{1, \dots, \lfloor \frac{n}{2} \rfloor\}$: This amounts to checking after every round whether any of the remaining treatments can be eliminated.
2. $\mathcal{R} = \{\underline{r}, \dots, \lfloor \frac{n}{2} \rfloor\}$ for some $1 < \underline{r} \leq \lfloor n/K \rfloor$: This amounts to allowing for a “burn-in” phase before checking (after every subsequent round) whether treatments can be eliminated. This is relevant in case one does not want to eliminate any treatments based on very few assignments.
3. $\mathcal{R} = k\mathbb{N} \cap \{1, \dots, \lfloor \frac{n}{2} \rfloor\}$ for some $k \in \{1, \dots, \lfloor \frac{n}{K} \rfloor\}$: This choice amounts to checking after every k th round whether treatments can be eliminated.

A detailed description of the policy $\tilde{\pi}$, concretizing the explanation above, is given in [Policy 2](#). Here, we also need to take care of the possibility that after round r has been completed, there might be less subjects left than treatments in \mathcal{I}_r , not allowing for a further complete round of assignments. This is why the policy is separated in an outer “while”-loop and an outer “if”-expression. As long as enough subjects are left to assign all remaining treatments once, the “while”-loop proceeds in rounds as discussed above, where after each round in \mathcal{R} treatments may be eliminated. Once there are more treatments than subjects left, all remaining subjects (if there are any left) are assigned to a subset of the remaining treatments. This happens within the outer “if”-expression. Finally, the recommendation is based on an empirical-success rule, where the minimum is taken as a concrete way of breaking ties. The policy does not use external randomization G_r , which is therefore suppressed in the notation.

Policy 2: Sequential Elimination Policy $\tilde{\pi}$

Input: $n \in \mathbb{N}, K \leq n$, set of elimination rounds $\mathcal{R} \subseteq \{1, \dots, \lfloor \frac{n}{2} \rfloor\}, \eta > 0$

Set: $t \leftarrow 0, r \leftarrow 0, \mathcal{I}_0 \leftarrow \mathcal{I}$

while $n - t \geq |\mathcal{I}_r|$ **do**

for $i \in \mathcal{I}_r$ **do**

$t \leftarrow t + 1$

$\tilde{\pi}_{n,t}(Z_{t-1}) = i$

end

$r \leftarrow r + 1$

if $r \in \mathcal{R}$ **then**

$\mathcal{I}_r \leftarrow \{j \in \mathcal{I}_{r-1} : \max_{i \in \mathcal{I}_{r-1}} \mathbb{T}(\hat{F}_{i,t,n}) - \mathbb{T}(\hat{F}_{j,t,n}) \leq u_\eta(r, n)\}$

end

else

$\mathcal{I}_r \leftarrow \mathcal{I}_{r-1}$

end

end

if $t < n$ **then**

Let $(j_1, \dots, j_{|\mathcal{I}_r|})$ be the elements of \mathcal{I}_r , ordered from smallest to largest.

for $i = 1, \dots, (n - t)$ **do**

$\tilde{\pi}_{n,t+i}(Z_{t+i-1}) = j_i$

end

end

$\tilde{\pi}_{n,n+1}(Z_n) = \min \arg \max \{ \mathbb{T}(\hat{F}_{i,n,n}) : i \in \mathcal{I}_r \}$

Most importantly, our theoretical results in the general case in the next section imply that [Policy 2](#) is minimax regret optimal up to logarithmic factors. This follows from the upper bound on expected maximal regret established in [Theorem 5.1](#) together with the lower bound established in [Theorem 3.3](#). We abstain from formulating a theorem for the special case $\mathcal{M}_K = \mathcal{E}_K$.

In pure exploration problems targeting the mean and $\mathcal{M}_K = \mathcal{E}_K$, sequential policies have been studied in [Audibert et al. \(2010\)](#) and [Karnin et al. \(2013\)](#). Furthermore, [Tran-Thanh and Yu \(2014\)](#) studied these policies for a family of quasi-convex functionals in the case of $\mathcal{M}_K = \mathcal{E}_K$. The just mentioned policies, however, fix a set of elimination times upfront at which a pre-specified (nonzero) number of treatments *must* be eliminated. [Policy 2](#), on the other hand, decides in a data-driven way if and when a treatment may be eliminated.

5.2. General incompatibility constraints

We shall now extend the policy from $\mathcal{M}_K = \mathcal{E}_K$ to the more general case where $\mathcal{M}_K = \bigcup_{j=1}^m \mathcal{M}_{A_j, K}$ for a partition A_1, \dots, A_m of \mathcal{I} , cf. [Example 2.6](#). In this more general case, the policy proceeds in rounds in the same way as detailed in the previous section. Again, the decision maker needs to a priori specify a subset of “elimination rounds” \mathcal{R} , in which treatments can potentially be eliminated. Checking whether treatments can be eliminated is a bit more involved and explained in the following.

Recall from the discussion around [\(25\)](#) that the treatments in A_j are strongly inferior if

$$\max_{\delta \in \bigcup_{l \neq j} \mathcal{M}_{A_l, K}} T(\delta, \mathbf{F}) > \max_{\delta \in \mathcal{M}_{A_j, K}} T(\delta, \mathbf{F}). \tag{31}$$

In order to empirically check whether this inequality holds (and whether the set of treatments A_j can be eliminated) we again rely on a discretization of \mathcal{M}_K . To this end fix a discretization $\mathcal{M}_K^n = \bigcup_{j=1}^m \mathcal{M}_{A_j, K}^n$ such that $\emptyset \neq \mathcal{M}_{A_j, K}^n \subseteq \mathcal{M}_{A_j, K}$ for every $j = 1, \dots, m$. Given \mathcal{M}_K^n of this form, our policy aims to verify whether [Eq. \(31\)](#) holds by checking its discretized version

$$\max_{\delta \in \bigcup_{l \neq j} \mathcal{M}_{A_l, K}^n} T(\delta, \mathbf{F}) > \max_{\delta \in \mathcal{M}_{A_j, K}^n} T(\delta, \mathbf{F}). \tag{32}$$

Intuitively, we empirically check this condition in two steps: (i) for every j and in every elimination round we “remove” those elements $\gamma \in \mathcal{M}_{A_j, K}^n$, for which the data firmly suggest $\max_{\delta \in \bigcup_{l \neq j} \mathcal{M}_{A_l, K}^n} T(\delta, \mathbf{F}) > T(\gamma, \mathbf{F})$; and (ii) we eliminate all treatments in A_j , once all $\gamma \in \mathcal{M}_{A_j, K}^n$ have been removed. As in the previous subsection, if a treatment is eliminated, it is no longer assigned, and does not contribute to the final recommendation.

To make this two-step check more precise, we denote by $\mathcal{M}_{A_j, K, r}^n$ the subset of elements of $\mathcal{M}_{A_j, K}^n$, which have not been removed after r rounds. After each elimination round $r \in \mathcal{R}$, we remove all $\gamma \in \mathcal{M}_{A_j, K, r}^n$ for which

$$\max_{\delta \in \bigcup_{l \neq j} \mathcal{M}_{A_l, K, r}^n} T(\delta, \hat{\mathbf{F}}_{t, n}) > T(\gamma, \hat{\mathbf{F}}_{t, n}) + u_\eta(r, n), \tag{33}$$

where $u_\eta(r, n)$ is as defined in [Eq. \(30\)](#), and $\hat{\mathbf{F}}_{t, n}$ was defined in [Eq. \(28\)](#) (in case the set over which the maximum in the previous display is taken is empty, no removal takes place). If all elements in \mathcal{M}_{K, A_j}^n are removed, we eliminate all treatments in A_j , and no longer assign them.

A detailed description of $\tilde{\pi}$ (for general partitions) is given in [Policy 3](#). The structure is the same as in [Policy 2](#), but elimination is based on the two-step procedure just explained. The recommendation is again based on an empirical-success rule. Furthermore, the policy does not depend on external randomization, which is therefore dropped from the notation. It is easy to verify that for $m = K$, i.e., $\mathcal{M}_K = \mathcal{E}_K$, [Policy 3](#) reduces to [Policy 2](#).

Now that we have defined the policy in the general case, we are ready to state a theorem regarding its maximal expected regret. The sequential nature of the SE policy makes the analysis more involved than that of the SA policy, which leads to a slightly more complicated upper bound than the one from [Theorem 4.3](#).

Theorem 5.1. *Suppose [Assumptions 2.1, 2.3, and 4.2](#) hold, and that \mathcal{M}_K is as in [Example 2.6](#) and satisfies [Assumption 2.7](#). Then, the SE policy $\tilde{\pi}$ as in [Policy 3](#) satisfies*

$$\sup_{\substack{F^i \in \mathcal{D} \\ i=1, \dots, K}} \mathbb{E} [r_n(\tilde{\pi}, \mathcal{M}_K)] \leq \varepsilon(n) + 4C \sqrt{\frac{K}{n}} D_{n, K} + 4C \frac{K^{1 - \frac{3\eta}{4+2\eta}}}{\sqrt{n}} \times (1 + 18.5\eta^{-1})^2, \quad \forall n \geq 2K,$$

where $D_{n, K} := \sqrt{K \log(17K)} \wedge \sqrt{\log(3.01 \times K(n+2))}$.

The upper bound on maximal expected regret is a sum of three terms. The first term is the optimization error resulting from targeting the best recommendation in \mathcal{M}_K^n , rather than the best recommendation in \mathcal{M}_K . Inspection of the proof shows that the second term is an upper bound on the expected regret on the event where not all maximizers of $\delta \mapsto T(\delta, \mathbf{F})$ over \mathcal{M}_K^n have been eliminated after n assignments, and that the last term upper bounds the probability that all maximizers of $\delta \mapsto T(\delta, \mathbf{F})$ over \mathcal{M}_K^n are eliminated by $\tilde{\pi}$ after n assignments. We note that in [Theorem D.1](#) in [Appendix D](#) we prove a slightly more general result, from which [Theorem 5.1](#) follows. The generality is bought at the expense of more complicated notation, but it provides a stronger bound valid for all $n > K$, which shows in more detail how the choice of \underline{r} affects the upper bound.

For K fixed and any choice of $\eta > 0$, the bound in [Theorem 5.1](#) achieves the minimax optimal dependence in n if one chooses a discretization with optimization error $\varepsilon(n)$ of no larger order than $n^{-1/2}$ (cf. the discussion after [Theorem 3.3](#)). For regimes where K grows with n , the dependence of the upper bound in [Theorem 5.1](#) on n and K is optimal (up to a factor of order $\sqrt{\log(nK)}$) as long as $\eta \geq 1$.

Although the policy $\tilde{\pi}$ cannot improve on $\hat{\pi}$ in terms of worst-case expected regret (and there is not much room for improvement given our lower bound results), there are many empirically relevant potential outcome distributions for which the policy $\tilde{\pi}$ is more efficient than $\hat{\pi}$. A prominent example of such distributions is the following:

Policy 3: Sequential Elimination $\tilde{\pi}$

Input: $n \in \mathbb{N}$, $K \leq n$, partition $\{A_1, \dots, A_m\}$ of \mathcal{I} such that $m \geq 2$, discretization $\mathcal{M}_K^n = \bigcup_{j=1}^m \mathcal{M}_{A_j, K}^n$, such that $\emptyset \neq \mathcal{M}_{A_j, K}^n \subseteq \mathcal{M}_{A_j, K}$ for every $j = 1, \dots, m$, set of elimination rounds \mathcal{R} , $\eta > 0$

Set: $t \leftarrow 0$, $r \leftarrow 0$, $\mathcal{I}_0 \leftarrow \mathcal{I}$, $\mathcal{J}_0 \leftarrow \{1, \dots, m\}$, $\mathcal{M}_{A_j, K, 0}^n \leftarrow \mathcal{M}_{A_j, K}^n$ for $j = 1, \dots, m$.

```

while  $n - t \geq |\mathcal{I}_r|$  do
  for  $i \in \mathcal{I}_r$  do
     $t \leftarrow t + 1$ 
     $\tilde{\pi}_{n,t}(Z_{t-1}) = i$ 
  end
   $r \leftarrow r + 1$ 
   $\mathcal{I}_r \leftarrow \mathcal{I}_{r-1}$ 
   $\mathcal{J}_r \leftarrow \mathcal{J}_{r-1}$ 
  for  $j \in \mathcal{J}_r$  do
    if  $r \in \mathcal{R}$  and  $|\mathcal{J}_r| \geq 2$  then
       $h \leftarrow \max\{\mathbb{T}(\langle \delta, \hat{\mathbf{F}}_{t,n} \rangle) : \delta \in \mathcal{M}_{A_l, K, r-1}^n, l \in \mathcal{J}_r \setminus \{j\}\}$ 
       $\mathcal{M}_{A_j, K, r}^n \leftarrow \{\gamma \in \mathcal{M}_{A_j, K, r-1}^n : h - \mathbb{T}(\langle \gamma, \hat{\mathbf{F}}_{t,n} \rangle) \leq u_\eta(r, n)\}$ 
      if  $\mathcal{M}_{A_j, K, r}^n = \emptyset$  then
         $\mathcal{I}_r \leftarrow \mathcal{I}_r \setminus A_j$ 
         $\mathcal{J}_r \leftarrow \mathcal{J}_r \setminus \{j\}$ 
      end
    end
    else
       $\mathcal{M}_{A_j, K, r}^n \leftarrow \mathcal{M}_{A_j, K, r-1}^n$ 
    end
  end
end
end
if  $t < n$  then
  Let  $(j_1, \dots, j_{|\mathcal{I}_r|})$  be the elements of  $\mathcal{I}_r$ , ordered from smallest to largest.
  for  $i = 1, \dots, (n - t)$  do
     $\tilde{\pi}_{n,t+i}(Z_{t+i-1}) = j_i$ 
  end
end
 $\tilde{\pi}_{n,n+1}(Z_n) = \min \arg \max \{\mathbb{T}(\langle \delta, \hat{\mathbf{F}}_{n,n} \rangle) : \delta \in \bigcup_{j \in \mathcal{J}_r} \mathcal{M}_{A_j, K, r}^n\}$ 

```

Example 5.2. For simplicity, assume that the functional of interest \mathbb{T} is quasi-convex and that $\mathcal{M}_K = \mathcal{E}_K$ such that no mixtures need to be studied. Consider $K > 2$ potential outcome distributions F^1, \dots, F^K for which it holds that

$$\mathbb{T}(F^1) \approx \mathbb{T}(F^2) \quad \text{and} \quad \mathbb{T}(F^3) = \dots = \mathbb{T}(F^K) \ll \mathbb{T}(F^1),$$

that is there are two top treatments (Treatments 1 and 2) that are approximately equally good and $K - 2$ very inferior treatments. Instead of assigning all K treatments n/K times as in the basic version of the SA policy (assuming here that n/K is an integer), it is often possible to identify that Treatments 3 to K are inferior based on much fewer than n/K assignments to each of these. This allows one to assign Treatments 1 and 2 more often, which is beneficial since these are the two real candidates for the best treatment. In particular if K is large this can lead to better recommendations as one can sample much more often from the two distributions that are difficult to distinguish than in the static policy. Indeed, this is the idea underlying the third simulation setting in Section 6.1 where the outcome distribution with maximal Gini-welfare is targeted. There, it is illustrated numerically that in such a setting the sequential elimination policy can result in much lower regret than the static assignment policy.

5.3. Beyond incompatibility constraints

So far, we have considered the case where \mathcal{M}_K satisfies an incompatibility constraint, which we exploited in the construction of the sequential policy. As our final investigation, we now ask whether a sequential policy with optimal expected regret guarantees can also be designed without exploiting that type of constraint. Recall that we have already argued that strongly inferior treatments should exist in order that a sequential policy can eliminate treatments. Because \mathcal{M}_K is known to the decision maker, this can be checked before deciding on which policy to use.




Fig. 1. Incompatibility constraint (A) vs. an entry cost constraint (B).

Example 5.3. One example where strongly inferior treatments can exist, without an incompatibility constraint being satisfied, are capacity constrained sets \mathcal{M}_K , where, for some index i , it holds for every $\delta \in \mathcal{M}_K$ that either $\delta_i = 0$ or $\delta_i \geq \epsilon$ for some $\epsilon \in (0, 1)$. This corresponds to “entry cost” restrictions, where treatment i can only be rolled out if a certain minimal proportion ϵ is assigned to that treatment. Then, if all optimal mixture vectors are such that their i th coordinate is 0, treatment i is clearly strongly inferior. An illustration of such a situation in comparison to an incompatibility constraint is given in Fig. 1, where there are three treatments. The set \mathcal{M}_K is colored red, whereas its complement in the simplex (i.e., those weights that are ruled out) is shown in gray. In Part (A) of that figure, we show a situation corresponding to an incompatibility constraint, where Treatments 1 and 3 are elements of the same group and Treatment 2 constitutes the second group. In Part (B) of that figure we show an “entry cost” constraint on Treatment 2. Note that in (A) Treatments 1 and 3, or Treatment 2 can be strongly inferior treatments, whereas in (B), only Treatment 2 can be strongly inferior.

The policy we propose proceeds in rounds (analogous to the sequential policies already discussed). In every round, treatments that have not been eliminated in previous rounds are assigned once. The decision maker needs to a priori specify a subset of “elimination rounds” \mathcal{R} , in which treatments can potentially be eliminated. Fix a discretization \mathcal{M}_K^n of \mathcal{M}_K , which is updated after each round, starting with $\mathcal{M}_{K,0}^n = \mathcal{M}_K^n$. In each round r our policy determines $\mathcal{M}_{K,r}^n$ by dropping from that set all $\gamma \in \mathcal{M}_{K,r-1}^n$ for which

$$\max_{\delta \in \mathcal{M}_{K,r-1}^n} T(\langle \delta, \hat{\mathbf{F}}_{t,n} \rangle) > T(\langle \gamma, \hat{\mathbf{F}}_{t,n} \rangle) + 2u_\eta(r, n), \tag{34}$$

cf. Eq. (30) where $u_\eta(r, n)$ was introduced. If $\mathcal{M}_{K,r}^n$ does not contain a single weights vector with positive i th coordinate, treatment i is eliminated, i.e., is no longer assigned in subsequent rounds. Note that compared to the previously introduced sequential policies (that were tailored to incompatibility constraints), the elimination threshold $2u_\eta(r, n)$ is now *more conservative*, because we can no longer exploit any particular structure of \mathcal{M}_K . This is the price we have to pay for obtaining a policy that can be applied under weaker structural assumptions on \mathcal{M}_K .

A detailed description of $\check{\pi}$ is given in Policy 4. The recommendation is based on an empirical-success rule and does not depend on external randomization, which is therefore dropped from the notation. Note also that one can only eliminate treatments, if the user chooses the discretization $\mathcal{M}_K^n \subseteq \mathcal{M}_K$ in such a way that it contains “sparse” elements (i.e., vectors with coordinates equaling 0), because otherwise the policy will never eliminate treatments.

Theorem 5.4. Suppose Assumptions 2.1, 2.3, and 4.2 hold, and that \mathcal{M}_K satisfies Assumption 2.7. Then, the SE policy $\check{\pi}$ as in Policy 4 satisfies the same regret upper bound as the policy $\tilde{\pi}$ in Theorem 5.1.

One advantage of the policy $\check{\pi}$ over $\tilde{\pi}$ is that it can be run and has optimality guarantees regardless of whether \mathcal{M}_K is incompatibility constrained or not, because $\check{\pi}$ does not exploit any structural properties of \mathcal{M}_K . This is achieved through a more conservative constant $2u_\eta(r, n)$.

6. Numerical results

We now illustrate the theoretical results established in this article by means of numerical examples. Simulations will be conducted for T the Gini-welfare measure introduced in Example 2.8 with $\mathcal{M}_K = \mathcal{E}_K$, for the mean functional with \mathcal{M}_K incorporating different types of capacity and similarity constraints, and for the headcount ratio discussed in Example 2.9, a widely used poverty measure. In all cases we study the expected regret of the static assignment policy $\hat{\pi}$ and, for those types of \mathcal{M}_K as in Example 2.6, we also study the performance of the sequential elimination policy $\tilde{\pi}$. For simplicity in the simulations, the SA policy is implemented assigning the K treatments available cyclically in the course of the experimentation phase (i.e., with the partition $\Pi_{n,i} = \cup_{s=0}^\infty \{i + Ks\} \cap \{1, \dots, n\}$), which leads to balanced groups. Note that considering cyclical assignments is without loss of generality here, as the performance of the SA policy only depends

Policy 4: Sequential Elimination $\tilde{\pi}$

Input: $n \in \mathbb{N}$, $K \leq n$, discretization \mathcal{M}_K^n , set of elimination rounds \mathcal{R} , $\eta > 0$

Set: $t \leftarrow 0$, $r \leftarrow 0$, $\mathcal{I}_0 \leftarrow \mathcal{I}$, $\mathcal{M}_{K,0}^n \leftarrow \mathcal{M}_K^n$.

```

while  $n - t \geq |\mathcal{I}_r|$  do
  for  $i \in \mathcal{I}_r$  do
     $t \leftarrow t + 1$ 
     $\tilde{\pi}_{n,t}(Z_{t-1}) = i$ 
  end
   $r \leftarrow r + 1$ 
   $\mathcal{I}_r \leftarrow \mathcal{I}_{r-1}$ 
  if  $r \in \mathcal{R}$  and  $|\mathcal{I}_r| \geq 2$  then
     $h \leftarrow \max\{\mathbb{T}(\langle \delta, \hat{\mathbf{F}}_{t,n} \rangle) : \delta \in \mathcal{M}_{K,r-1}^n\}$ 
     $\mathcal{M}_{K,r}^n \leftarrow \{\gamma \in \mathcal{M}_{K,r-1}^n : h - \mathbb{T}(\langle \gamma, \hat{\mathbf{F}}_{t,n} \rangle) \leq 2u_\eta(r, n)\}$ 
     $\mathcal{I}_r \leftarrow \{i \in \mathcal{I}_r : \text{there exists a } \gamma \in \mathcal{M}_{K,r}^n \text{ with } \gamma_i > 0\}$ 
  end
end
if  $t < n$  then
  Let  $(j_1, \dots, j_{|\mathcal{I}_r|})$  be the elements of  $\mathcal{I}_r$ , ordered from smallest to largest.
  for  $i = 1, \dots, (n - t)$  do
     $\tilde{\pi}_{n,t+i}(Z_{t+i-1}) = j_i$ 
  end
end
 $\tilde{\pi}_{n,n+1}(Z_n) = \min \arg \max \{\mathbb{T}(\langle \delta, \hat{\mathbf{F}}_{n,n} \rangle) : \delta \in \mathcal{M}_{K,r}^n\}$ 

```

on the partition via the number of subjects assigned to every treatment. The SE policy is implemented with $\eta = 0.0001$, and with elimination rounds $\mathcal{R} = \{1, \dots, \lfloor n/2 \rfloor\}$, i.e., allowing for elimination after each round.⁸

We study the expected regret of the policies for all n that are integer multiples of 100 between 100 and 5,000 and approximate the expected regret by the average regret (i.e., the arithmetic mean) over 5,000 replications for each sample size.

A brief summary of the findings is as follows. The SA and SE policy both generally incur low average regret, confirming our theoretical optimality results for these policies. For sufficiently small sample sizes, the recommendations of the two policies are identical, as no treatments are eliminated by the SE policy in this case. However, for larger samples sizes, the SE policy eliminates clearly suboptimal treatments and thus incurs a lower average regret. In this sense, our numerical results indicate that in terms of regret there is nothing lost from using the SE policy for such \mathcal{M}_K , as it always performs at least as well as the SA policy and sometimes strictly better. The SE policy is, however, computationally more burdensome, as it must compare the values of the functional after every elimination round. Note, however, that the number of elimination rounds can be chosen by the user to lighten the numerical cost of using the SE policy.

6.1. Gini-welfare

Recall the Gini-welfare measure from [Example 2.6](#). There, it is shown that this functional satisfies [Assumption 2.3](#) with $\mathcal{D} = D_{cdf}([0, 1])$ and $C = 2$, and is quasi-convex. We know from the discussion in [Section 2.2.1](#) that any \mathcal{M}_K can thus without loss of generality be replaced by the set of its extreme points. In situations where \mathcal{M}_K contains \mathcal{E}_K , it therefore suffices to work with \mathcal{E}_K , the case we focus on here (and which allows an application of the SE policy). We consider $K = 5$ treatments, and where $Y_{i,t}$ for $i = 1, \dots, K$ are (independently) distributed with F^i the cdf of a Beta-distribution with parameters α_i and 1. Three settings of parameters are considered.

1. *Approximately equal Gini-welfare:* $\alpha_1 = 5$ and $\alpha_i = 4.5$ for $i = 2, \dots, 5$. These distributions have a Gini welfare of 0.76 and 0.74, respectively.
2. *Approximately equal Gini-welfare within three classes:* (i) $\alpha_1 = 5$ and $\alpha_2 = 4.5$, (ii) $\alpha_3 = \alpha_4 = 2$ and (iii) $\alpha_5 = 1$. These distributions have a Gini-welfare of 0.76, 0.74, 0.53 and 0.33, respectively.
3. *Two strong treatments and three inferior ones:* $\alpha_1 = 5$, $\alpha_2 = 4.5$ and $\alpha_i = 1$ for $i = 3, \dots, 5$. These distributions have a Gini-welfare of 0.76, 0.74 and 0.33, respectively.

Before we proceed to the results, we note that since \mathcal{E}_K is finite, we used the discretization $\mathcal{M}_K^n = \mathcal{E}_K$ in both policies.

⁸ Experiments (not reported) yielded that lower values of $\eta > 0$, i.e., more aggressive elimination of treatments, resulted in lower regret for the SE policy which we therefore recommend for practice.

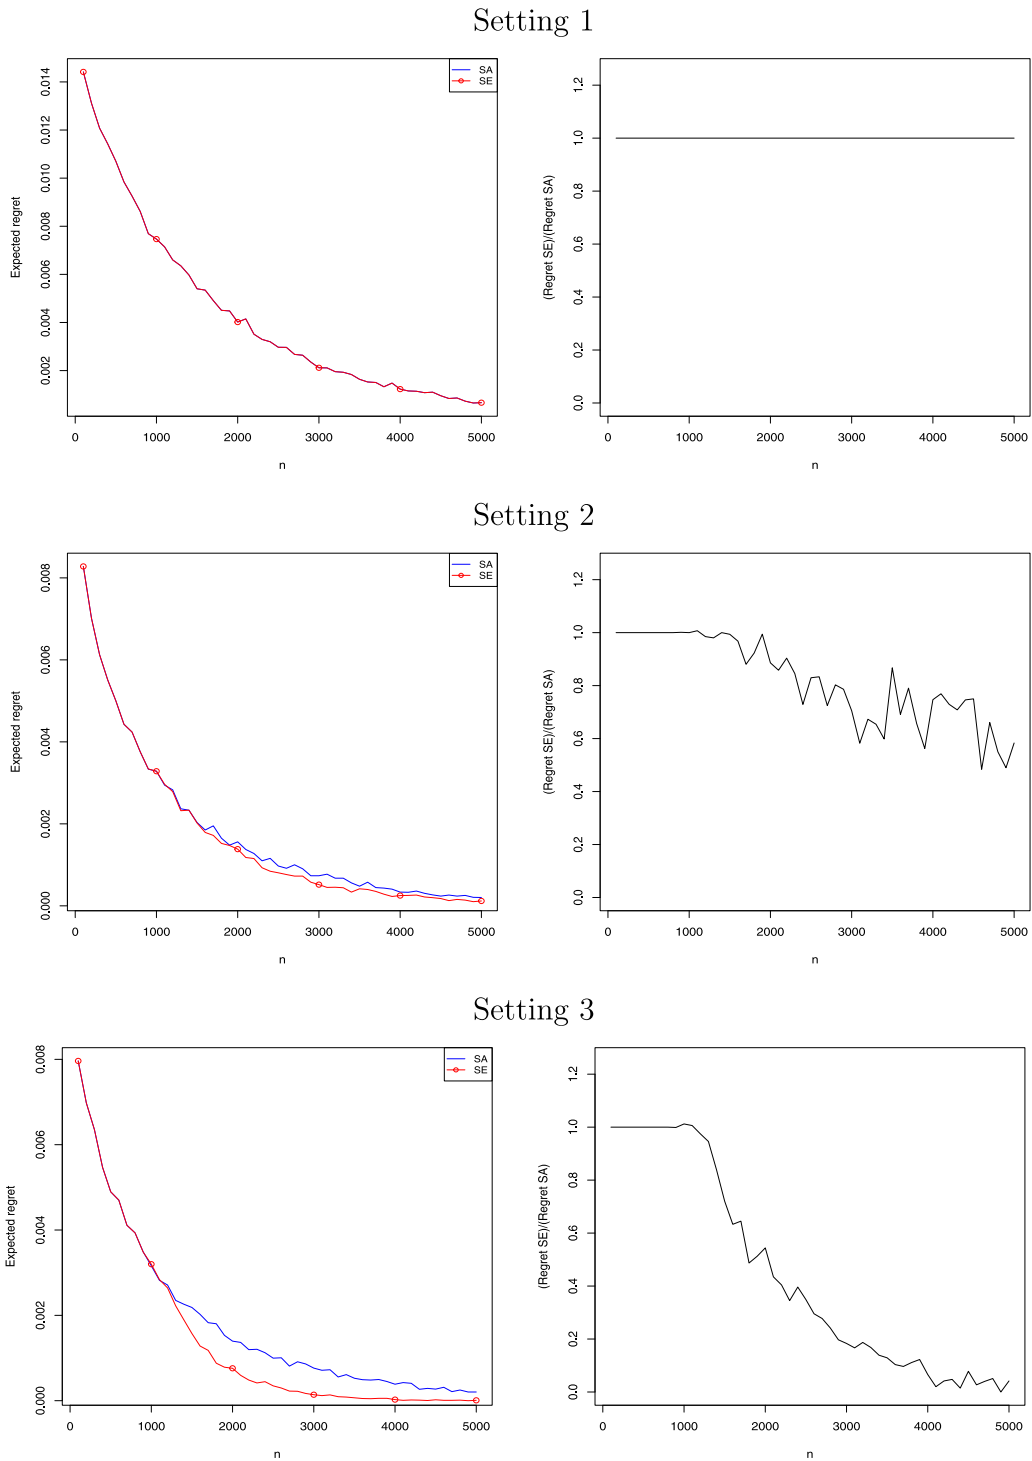


Fig. 2. Absolute and relative average regrets for Settings 1-3 for the Gini-welfare measure. Each row corresponds to a setting. The left column contains the absolute average regrets, while the right column contains the relative average regret of the SE to the SA policy. Thus, numbers less than one for the relative average regret mean that the SE has a lower average regret than the SA policy.

6.1.1. Results

The numerical results for the Gini-welfare measure are summarized in Fig. 2. The top two panels in Fig. 2 contain the results for Setting 1 where all treatments are approximately equally good. The Gini welfares of the five treatments are so

close to each other that no treatment is ever eliminated by the SE policy for any of the sample sizes considered. Thus, the SA and SE policies are identical and incur the same average regret as witnessed by the left and right panel alike. From the left panel it is seen that even for the smallest sample sizes considered, the SA and SE policies both incur rather small average regret.

The middle two panels contain the results for Setting 2, where there are three classes of treatments. The most important difference to Setting 1 is that the regrets of the SA and SE policies now differ. In fact, by eliminating the very inferior Treatment 5 after (on average) around 250 rounds (and the moderately inferior Treatments 3 and 4 after around 850 rounds), the SE policy is able to allocate more assignments to Treatments 1 and 2, which are the ones that are hard to distinguish. As a result of this, the SE policy recommends the superior Treatment 1 more often than the SA policy and hence generally incurs a lower average regret. This gain is almost uniform in n except for at $n = 1, 100$ where the relative regret is 1.007. This slightly higher average regret is not due to the SE policy eliminating the best treatment – it never did in our simulations. Rather it is a consequence of approximating expected regret by average regret. This also explains why the average regret of the SE policy is not monotonically decreasing in n in any of the three settings. The right panel shows that the relative gains of the SE over the SA policy can be more than 50% in Setting 2.

The bottom panel contains the results for Setting 3 where there are two strong treatments and three inferior ones. Both policies have a small average regret. However, the SE policy is able to eliminate the inferior treatments after (on average) about 250 rounds. This results in the SE policy incurring an average regret that is lower by several orders of magnitude than the one of the SE policy for the largest sample sizes considered.

Summarizing, the numerical results for Gini-welfare and $\mathcal{M}_K = \mathcal{E}_K$ suggest that the SE policy is never inferior to the SA policy in terms of average regret, and can be superior whenever there are treatments that are not too similar to the best (in relation to sample size). The computational burden of the SE policy is furthermore negligible in case $\mathcal{M}_K = \mathcal{E}_K$. Thus, if a sequential assignment scheme is practically feasible, the SE policy should be used for the Gini-welfare and $\mathcal{M}_K = \mathcal{E}_K$.

6.2. Mean functional with capacity and similarity constraints

When a decision maker targets the treatment with the highest expectation, and is free to roll out a single best treatment, the situation is very similar (apart from the functional used) to the situation in the previous section. However, if the decision maker needs to take capacity or similarity constraints into consideration, this is no longer true, as such constraints typically imply restrictions of the type $\mathcal{E}_K \not\subseteq \mathcal{M}_K$. In this section, we study the performance of the SA policy under such restrictions.⁹ Again, in all of the settings we consider, there are $K = 5$ treatments, and F^i is the cdf of a Beta-distribution with parameters α_i and 1. We consider the parameters $\alpha_1 = 1.2$, $\alpha_2 = 1.1$, $\alpha_3 = 1$, $\alpha_4 = 0.9$ and $\alpha_5 = 0.8$, the corresponding Beta-distributions having means 0.55, 0.52, 0.5, 0.47 and 0.44, respectively. The restrictions $\mathcal{M}_K^{(1)} \supseteq \mathcal{M}_K^{(2)} \supseteq \mathcal{M}_K^{(3)}$ we study are as follows.

1. The capacities of all treatments are equally restricted:

$$\mathcal{M}_K^{(1)} = \{\delta \in \mathcal{S}_K : 0.1 \leq \delta_i \leq 0.5, \text{ for } i = 1, \dots, K\}.$$

Here, the optimal weights vector is $\delta = (0.5, 0.2, 0.1, 0.1, 0.1)'$. As expected, this puts as much weight as possible on the treatments with the highest means.

2. The joint capacity of the first two treatments is further restricted:

$$\mathcal{M}_K^{(2)} = \{\delta \in \mathcal{M}_K^{(1)} : 2\delta_1 + \delta_2 \leq 0.5\}.$$

Now the total weight on the best two treatments cannot be too high. In particular, Treatment 1 is “expensive”. The optimal weights vector is $\delta = (0.1, 0.3, 0.4, 0.1, 0.1)'$. Observe that now Treatment 1 is assigned as little as possible, despite having the highest mean.

3. A similarity constraint concerning Treatments 3 and 5 is added:

$$\mathcal{M}_K^{(3)} = \{\delta \in \mathcal{M}_K^{(2)} : |\delta_3 - \delta_5| \leq 0.1\}.$$

The optimal weights vector here is $\delta = (0.1, 0.3, 0.2, 0.3, 0.1)'$. Now the weights of Treatment 3, which were highest in the previous setting, must be close to those of the individual treatment with the lowest mean.

Before we proceed to the numerical results, we note that since the constraints are all linear, the recommendation in the SA policy can be found by solving (for $i = 1, 2, 3$) the linear program

$$\max_{\delta \in \mathcal{M}_K^{(i)}} T(\delta, \hat{\mathbf{F}}_n, \Pi_n) = \max_{\delta \in \mathcal{M}_K^{(i)}} \sum_{j=1}^K \delta_j \hat{\mu}_j$$

via the simplex-algorithm; here $\hat{\mu}_j$ denotes the mean of the empirical cdf based on the observations assigned to treatment j .

⁹ In none of the situations we consider here, there are strongly inferior treatments. Hence, we do not implement any version of the SE policy.




Fig. 3. Average regret of the SA policy for the mean functional under capacity and similarity constraints in Settings 1–3 described in the text.

6.2.1. Results

Fig. 3 contains the results for Settings 1–3 just described. The regret decreases with n in all three settings in accordance with the theoretical guarantees described in Theorem 4.3. Furthermore, even for small n , the regret is “low” for all three sets $\mathcal{M}_K^{(i)}$ considered. The figure also illustrates that the smaller the feasible set is, the lower the regret is for small sample sizes. This is in accordance with the lower bounds in Theorem 3.3, which involve smaller constants for smaller feasible sets.

6.3. Headcount ratio

We finally study a situation in which the decision maker targets a common poverty measure: namely the (negative) headcount ratio as defined in Eq. (7) of Example 2.9. In contrast to the Gini welfare measure and the mean, this functional is not quasi-convex. Thus, even when $\mathcal{E}_K \subseteq \mathcal{M}_K$, the mixture that minimizes the headcount ratio need not be an element of \mathcal{E}_K . One example of this type is the following simple setting:

$$\begin{aligned}
 F^1(x) &= 1.25x\mathbb{1}_{\{x < 0.25\}} + 0.3125\mathbb{1}_{\{0.25 \leq x < 0.3125\}} + x\mathbb{1}_{\{x \geq 0.3125\}}, \\
 F^2(x) &= 3x\mathbb{1}_{\{x < 0.1\}} + 0.3\mathbb{1}_{\{0.1 \leq x < 0.15\}} + 2x\mathbb{1}_{\{0.15 \leq x < 1/2\}} + \mathbb{1}_{\{x \geq 1/2\}}, \\
 F^i(x) &= a_i\mathbb{1}_{\{x < 1/2\}} + (p_i x + 1 - p_i)\mathbb{1}_{\{x \geq 1/2\}}, \quad \text{for } i \in \{3, 4, 5\},
 \end{aligned}
 \tag{35}$$

with $p_i = 2(1 - a_i)$ and $a_3 = 0.5$, $a_4 = 0.7$, $a_5 = 0.9$, respectively. It can be verified that the headcount ratios of these five cdfs are 0.31, 0.3, 0.5, 0.7 and 0.9, respectively (mentioned in the order of F^1 to F^5). Intuitively, the larger a_i is, the larger is the mass at zero of F^i , i.e. the larger is the proportion of very poor people. We let

$$\mathcal{M}_K = \{\delta \in \mathcal{S} : \delta_3 = \delta_4 = \delta_5 = 0\} \cup \{e_3(5)\} \cup \{e_4(5)\} \cup \{e_5(5)\},$$

such that mixtures of F^1 and F^2 are allowed, while F^3 , F^4 , and F^5 can only be assigned on their own. This is an instance of an incompatibility constraint as described in Example 2.6. It can be shown that $\delta = (\frac{736}{3171}, \frac{2435}{3171}, 0, 0, 0)'$ minimizes the headcount ratio over \mathcal{M}_K with a corresponding minimal value of 0.2739. The values of the headcount ratios corresponding to mixtures of F^1 and F^2 are shown in Fig. 4 (which nicely illustrates that the lowest headcount ratio is obtained by combining two treatments). Importantly, we note that the minimizer $\delta \notin \mathcal{E}_K$ even though $\mathcal{E}_K \subseteq \mathcal{M}_K$. The discretization used is

$$\mathcal{M}_K^n = \{\delta \in \mathcal{S} : \delta_3 = \delta_4 = \delta_5 = 0, 100\delta_1 \in \mathbb{Z}\} \cup \{e_3(5)\} \cup \{e_4(5)\} \cup \{e_5(5)\},$$

and we set $C = 2.5$ in accordance with Example 2.9.

6.3.1. Results

The numerical results for the headcount ratio with outcome distributions as in (35) are contained in Fig. 5. The left panel, which contains the absolute average regret for the SA and the SE policies, shows that both of these generally incur a low average regret and that this is decreasing in the sample size. Observe also that the SE policy is, up to simulation error, uniformly better (in n) than the SA policy. The right panel contains the relative average regret of the SE policy to that of the SA policy. It can be seen that for $n \geq 3,000$ the average regret of the SE policy is always at least 25% lower




Fig. 4. Headcount ratios for mixtures of the form $\delta_1 F^1 + (1 - \delta_1) F^2$ in dependence on $\delta_1 \in [0, 1]$.




Fig. 5. Absolute and relative average regrets for the headcount ratio with outcome distributions as in (35). The left plot contains the absolute average regrets, while the right plot contains the relative average regret of the SE to the SA policy. Thus, numbers less than one for the relative average regret mean that the SE has a lower average regret than the SA policy.

(and up to 38% lower) than that of the SA policy. As was the case for the Gini welfare measure, this is due to the SE policy correctly eliminating inferior treatments and dedicating more sampling effort to, in this case, Treatments 1 and 2 on which the best mixture puts positive weights.

6.3.2. Mixing inferior and superior treatments

In the setting of (35) all mixtures of Treatments 1 and 2 are superior to the remaining Treatments 3, 4 and 5. This may make the elimination of Treatments 3, 4 and 5 too easy. We therefore consider now a setting where this is no longer the case. The potential outcome distributions are in the same family as those used in (35). To be precise, we let

$$\begin{aligned}
 G^1(x) &= 3x\mathbb{1}_{\{x < 0.1\}} + 0.3\mathbb{1}_{\{0.1 \leq x < 0.15\}} + 2x\mathbb{1}_{\{0.15 \leq x < 1/2\}} + \mathbb{1}_{\{x \geq 1/2\}}, \\
 G^i(x) &= b_i\mathbb{1}_{\{x < 1/2\}} + (q_i x + 1 - q_i)\mathbb{1}_{\{x \geq 1/2\}}, \quad \text{for } i \in \{2, 3, 4, 5\};
 \end{aligned}
 \tag{36}$$

with $q_i = 2(1 - b_i)$ and $b_2 = 0.7$, $b_3 = 0.32$, $b_4 = 0.7$, $b_5 = 0.9$, respectively. It can be verified that the headcount ratios of these five cdfs are 0.3, 0.7, 0.32, 0.7 and 0.9, respectively (mentioned in the order of G^1 to G^5). We allow for the superior Treatment 1 to be mixed with the dominated Treatment 2 while the remaining treatments can only be assigned on their own, that is

$$\mathcal{M}_K = \{\delta \in \mathcal{S}_5 : \delta_3 = \delta_4 = \delta_5 = 0\} \cup \{e_3(5)\} \cup \{e_4(5)\} \cup \{e_5(5)\}.$$




Fig. 6. Absolute and relative average regrets for the headcount ratio with outcome distributions as in (36). The left plot contains the absolute average regrets, while the right plot contains the relative average regret of the SE to the SA policy. Thus, numbers less than one for the relative average regret mean that the SE has a lower average regret than the SA policy.

It can now be shown that $\delta = (1, 0, 0, 0, 0)$ minimizes the headcount ratio over \mathcal{M}_K with a corresponding value of 0.3 (the headcount ratio corresponding to assigning Treatment 1 on its own). The discretization of \mathcal{M}_K that we use is

$$\mathcal{M}_K^n = \{ \delta \in \mathcal{S} : \delta_3 = \delta_4 = \delta_5 = 0, 100\delta_1 \in \mathbb{Z} \} \cup \{e_3(5)\} \cup \{e_4(5)\} \cup \{e_5(5)\},$$

and we again set $C = 2.5$ in accordance with Example 2.9.

6.3.3. Additional results

The results for the headcount ratio with outcome distributions as in (36) are contained in Fig. 6. The findings are qualitatively like those in Fig. 5 – as long as no treatments are eliminated the SA and SE policy make the same recommendation but once the SE policy begins eliminating inferior treatments it more often recommends the regret minimizing mixture $(1, 0, 0, 0, 0)$.

We also experimented with settings in which the headcount ratio of the second best treatment (Treatment 3) is much closer to (further from) that of the best treatment (Treatment 1). In the former case, the recommendation problem can become so difficult that neither the SA nor SE policy can distinguish between the two top treatments at the studied sample sizes and the relative regret is essentially one. If, on the other hand, the headcount ratio of the second best treatment is much higher than that of the best treatment then the recommendation problem can become so easy that neither the SA nor the SE policy make any mistakes and again they are equally good. Let us stress that even in such settings nothing is lost from using the SE policy and the above simulations show that there are settings in which one obtains substantially lower regret by using the SE rather than the SA policy.

7. Incorporating covariates

One relevant extension of the setup we considered is the situation in which for each subject one also observes a vector of covariates $X_t \in \mathcal{X}$ with (unknown) distribution P_X , say. While we do not spell out any theoretical results in this setup, we shall heuristically outline in the following how a decision maker can incorporate covariate information using a natural extension of our plug-in based approach. Denoting the conditional distribution of the potential outcome vector Y_t given X_t by $\mathbf{F}(y|x)$ (a distribution on \mathbb{R}^K), any mixture $\delta : \mathcal{X} \rightarrow \mathcal{M}_K$ (in principle \mathcal{M}_K could then also depend on the covariates but we do not highlight this possibility here) now induces a population cdf via

$$\int \langle \delta(x), \mathbf{F}(\cdot|x) \rangle dP_X(x).$$

Note that in a setting with covariates the weights vector δ now actually is a vector-valued function taking as an argument the covariate vector. The goal of the policy maker here is to determine an optimal function δ , i.e., a maximizer of

$$\delta \mapsto \mathbb{T} \left(\int \langle \delta(x), \mathbf{F}(\cdot|x) \rangle dP_X(x) \right). \tag{37}$$

Following our plug-in approach, given an estimate of the conditional distribution $\hat{\mathbf{F}}(\cdot|x)$ and an estimate \hat{P}_X of P_X , the mixture recommended at the end of the experimental phase would be a maximizer of

$$\delta \mapsto \mathbb{T} \left(\int \langle \delta(x), \hat{\mathbf{F}}(\cdot|x) \rangle d\hat{P}_X(x) \right).$$

To illustrate this in some detail, consider a specific approach based on partitioning the covariate space, an approach that has been used in related settings, e.g., [Rigollet and Zeevi \(2010\)](#), [Perchet and Rigollet \(2013\)](#) or [Kock et al. \(2020b\)](#). Essentially, one would partition the covariate space into B clusters, say, obtain the estimated conditional cdfs \hat{F}_j (a cdf on \mathbb{R}^K) for $j = 1, \dots, B$ by aggregating outcomes within each cluster, and use these estimates as plug-ins to maximize

$$(\delta_1, \dots, \delta_B) \mapsto T \left(\sum_{j=1}^B \hat{p}_j(\delta^{(j)}, \hat{F}_j) \right). \quad (38)$$

Here \hat{p}_j denotes the relative frequency of observations falling into cluster j , and each $\delta^{(j)}$ is an element of \mathcal{M}_K . Optimizing (38) in $(\delta^{(1)}, \dots, \delta^{(B)})$ one obtains the cluster-specific allocation vectors $(\hat{\delta}^{(1)}, \dots, \hat{\delta}^{(B)})$. The optimal size of the clusters (ensuring that (38) well approximates (37)) depends on the sample size n and on smoothness assumptions imposed on $F(y|\mathcal{X})$. In the important case where \mathcal{X} is finite and not too large no clustering is necessary at all, i.e., B would coincide with the number of elements of \mathcal{X} . Static and sequential assignment schemes are possible and would be implemented on the cluster level (but would need to be adjusted so as to take into account the global objective in (38) in the recommendation and in eliminating treatments on the cluster level).

8. Conclusions

This paper has studied the problem of a policy maker who conducts a sequential experiment to recommend the treatment that is best according to a desirable functional characteristic of the outcome distribution of the treatments. Our theory allows the policy maker to target a wide class of functionals. The problem is particularly intricate when targeting a functional that is not quasi-convex or when the set of treatments that can be rolled out is restricted as one must then learn the best mixture of treatments. We have characterized the difficulty of this decision problem and shown how it depends on the type of constraints faced. When it is feasible (and sensible, i.e., the constraints do not rule out strongly inferior treatments), we recommend to assign subjects using one of our sequential policies due to their superior regret behavior.

Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jeconom.2022.08.003>.

References

- Agrawal, R., 1995. The continuum-armed bandit problem. *SIAM J. Control Optim.* 33 (6), 1926–1951.
- Athey, S., Wager, S., 2021. Policy learning with observational data. *Econometrica* 89 (1), 133–161.
- Audibert, J.-Y., Bubeck, S., Munos, R., 2010. Best arm identification in multi-armed bandits. In: COLT-23th Conference on Learning Theory-2010. pp. 13–p.
- Barrett, G.F., Donald, S.G., 2009. Statistical inference with generalized Gini indices of inequality, poverty, and welfare. *J. Bus. Econom. Statist.* 27 (1), 1–17.
- Bhattacharya, D., Dupas, P., 2012. Inferring welfare maximizing treatment assignment under budget constraints. *J. Econometrics* 167 (1), 168–196.
- Biewen, M., 2002. Bootstrap inference for inequality and poverty measurement. *J. Econometrics* 108 (2), 317–342.
- Bitler, M.P., Gelbach, J.B., Hoynes, H.W., 2006. What mean impacts miss: Distributional effects of welfare reform experiments. *Amer. Econ. Rev.* 96 (4), 988–1012.
- Bomze, I.M., Gollwitzer, S., Yildirim, E.A., 2014. Rounding on the standard simplex: Regular grids for global optimization. *J. Global Optim.* 59 (2–3), 243–258.
- Bubeck, S., Cesa-Bianchi, N., 2012. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Found. Trends[®] Mach. Learn.* 5 (1), 1–122.
- Bubeck, S., Munos, R., Stoltz, G., 2009. Pure exploration in multi-armed bandits problems. In: *International Conference on Algorithmic Learning Theory*. Springer, pp. 23–37.
- Bubeck, S., Munos, R., Stoltz, G., Szepesvári, C., 2011. X-Armed bandits. *J. Mach. Learn. Res.* 12 (5).
- Cambini, A., Martein, L., 2009. *Generalized Convexity and Optimization*. In: *Lecture Notes in Economics and Mathematical Systems*, vol. 616, Berlin, Springer.
- Carpentier, A., Locatelli, A., 2016. Tight (lower) bounds for the fixed budget best arm identification bandit problem. In: *Conference on Learning Theory*. pp. 590–604.
- Cassel, A., Mannor, S., Zeevi, A., 2018. A general approach to multi-armed bandits under risk criteria. arXiv preprint [arXiv:1806.01380](https://arxiv.org/abs/1806.01380).
- Cesa-Bianchi, N., Lugosi, G., 2006. *Prediction, Learning, and Games*. Cambridge University Press.
- Davidson, R., Flachaire, E., 2007. Asymptotic and bootstrap inference for inequality and poverty measures. *J. Econometrics* 141 (1), 141–166.
- Dehejia, R.H., 2005. Program evaluation as a decision problem. *J. Econometrics* 125 (1), 141–173.
- Dufour, J.-M., Flachaire, E., Khalaf, L., 2019. Permutation tests for comparing inequality measures. *J. Bus. Econom. Statist.* 37 (3), 457–470.
- Foster, J., Greer, J., Thorbecke, E., 1984. A class of decomposable poverty measures. *Econometrica* 52 (3), 761–766.
- Gittins, J.C., 1979. Bandit processes and dynamic allocation indices. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 41 (2), 148–164.
- Hirano, K., Porter, J.R., 2009. Asymptotics for statistical treatment rules. *Econometrica* 77 (5), 1683–1701.
- Jamieson, K., Malloy, M., Nowak, R., Bubeck, S., 2014. Lill'ucb: An optimal exploration algorithm for multi-armed bandits. In: *Conference on Learning Theory*. pp. 423–439.
- Karnin, Z., Koren, T., Somekh, O., 2013. Almost optimal exploration in multi-armed bandits. In: *International Conference on Machine Learning*. pp. 1238–1246.
- Kasy, M., Sautmann, A., 2021. Adaptive treatment assignment in experiments for policy choice. *Econometrica* 89 (1), 113–132.

- Kaufmann, E., Cappé, O., Garivier, A., 2016. On the complexity of best-arm identification in multi-armed bandit models. *J. Mach. Learn. Res.* 17 (1), 1–42.
- Kitagawa, T., Tetenov, A., 2018. Who should be treated? Empirical welfare maximization methods for treatment choice. *Econometrica* 86 (2), 591–616.
- Kitagawa, T., Tetenov, A., 2021. Equality-minded treatment choice. *J. Bus. Econom. Statist.* 39, 561–574.
- Kleinberg, R., Slivkins, A., Upfal, E., 2008. Multi-armed bandits in metric spaces. In: *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*. pp. 681–690.
- Kock, A.B., Preinerstorfer, D., Veliyev, B., 2020a. Functional sequential treatment allocation. *J. Am. Stat. Assoc.* <http://dx.doi.org/10.1080/01621459.2020.1851236>, Forthcoming.
- Kock, A.B., Preinerstorfer, D., Veliyev, B., 2020b. Functional sequential treatment allocation with covariates. arXiv preprint [arXiv:2001.10996](https://arxiv.org/abs/2001.10996).
- Kock, A.B., Thyrgaard, M., 2017. Optimal sequential treatment allocation. arXiv preprint [arXiv:1705.09952](https://arxiv.org/abs/1705.09952).
- Lai, T.L., Robbins, H., 1985. Asymptotically efficient adaptive allocation rules. *Adv. Appl. Math.* 6 (1), 4–22.
- Lattimore, T., Szepesvári, C., 2020. *Bandit Algorithms*. Cambridge University Press, Cambridge.
- Manski, C.F., 2004. Statistical treatment rules for heterogeneous populations. *Econometrica* 72 (4), 1221–1246.
- Manski, C.F., Tetenov, A., 2016. Sufficient trial size to inform clinical practice. *Proc. Natl. Acad. Sci.* 113 (38), 10518–10523.
- Perchet, V., Rigollet, P., 2013. The multi-armed bandit problem with covariates. *Ann. Statist.* 693–721.
- Rigollet, P., Zeevi, A., 2010. Nonparametric bandits with covariates. *Proceedings of COLT*.
- Robbins, H., 1952. Some aspects of the sequential design of experiments. *Bull. Am. Math. Soc.* 58 (5), 527–535.
- Rothe, C., 2010. Nonparametric estimation of distributional policy effects. *J. Econometrics* 155 (1), 56–70.
- Rothe, C., 2012. Partial distributional policy effects. *Econometrica* 80 (5), 2269–2301.
- Sen, A., 1976. Poverty: an ordinal approach to measurement. *Econometrica* 219–231.
- Stoye, J., 2009. Minimax regret treatment choice with finite samples. *J. Econometrics* 151 (1), 70–81.
- Stoye, J., 2010. Partial identification of spread parameters. *Quantit. Econ.* 1 (2), 323–357. <http://dx.doi.org/10.3982/QE24>.
- Stoye, J., 2012. Minimax regret treatment choice with covariates or with limited validity of experiments. *J. Econometrics* 166 (1), 138–156.
- Thompson, W.R., 1933. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* 25 (3/4), 285–294.
- Tran-Thanh, L., Yu, J.Y., 2014. Functional bandits. arXiv preprint [arXiv:1405.2432](https://arxiv.org/abs/1405.2432).