

Treatment recommendation with distributional targets*

Anders Bredahl Kock

University of Oxford
 CREATES, Aarhus University
 10 Manor Rd, Oxford OX1 3UQ
anders.kock@economics.ox.ac.uk

David Preinerstorfer

ECARES, SBS-EM
 Université libre de Bruxelles
 50 Ave. F.D. Roosevelt CP 114/04, Brussels
david.preinerstorfer@ulb.be

Bezirgen Veliyev

CREATES, Department of Economics
 Aarhus University
 Fuglesangs Alle 10, 8210 Aarhus V.
bveliyev@econ.au.dk

May 25, 2020

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 regret-optimal policies. The first policy is static and thus applicable irrespectively of the subjects arriving sequentially or not in the course of the experimental 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.

JEL Classification: C18, C22, J68.

Keywords: Treatment allocation, best treatment identification, statistical decision theory, pure exploration, nonparametric multi-armed bandits.

*We thank Quentin Badolle for excellent research assistance in an early stage of this project.

1 Introduction

We consider a decision maker who wants to run an experiment to identify the “best” among a set of candidate treatments. To compare the treatments’ qualities, our framework lets the decision maker use a custom-built *distributional characteristic*, which summarizes the targeted properties of the outcome distributions, e.g., their poverty- or inequality-implications. 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. Note that the functional used to compare treatments may no longer be linear, or—more generally—may no longer be quasi-convex. That is, 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 of treatments is best.

A practically important aspect 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.

In the present article we establish theoretical results to assist a decision maker who needs to give a recommendation in this situation. 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 when the recommendation is based on an empirical-success rule.
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.

Our article draws on ideas from the multi-armed bandit literature and our framework resembles that of a “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. Thus, there is no exploration-exploitation trade-off. 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 (2019). 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.

Concerning functional targets, there are several contributions in the more traditional multi-armed bandit literature, which has the goal of balancing an exploration-exploitation trade-off. The recent papers Kock and Thyrgaard (2017), Cassel et al. (2018) and Kock et al. (2020a) are related to the present paper from a technical point-of-view. Note, however, that the results developed there consider a completely different problem, and require differently designed policies, as the regret function used involves *all* assignments made and does not target the quality of a final recommendation. Note also that the

just-mentioned articles do not allow the decision maker to target mixtures, or incorporate constraints.

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), Athey and Wager (2017), Kitagawa and Tetenov (2018). 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 (2019), who studied a class of rank dependent social welfare functions as targets. While 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.

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 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. All proofs are collected in Appendices A, B, C, and D.

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 $i \in \mathcal{I} := \{1, \dots, K\}$ 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.

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.*

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 $\mathbf{F} = (F^1, \dots, F^K)$. 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*

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

Given \mathbb{T} , the cdf $F \in D_{cdf}([a, b])$ is considered as being better than $G \in D_{cdf}([a, b])$ if $\mathbb{T}(F) > \mathbb{T}(G)$. Technically, the main assumption we impose on \mathbb{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 2.2. The functional $\mathsf{T} : D_{cdf}([a, b]) \rightarrow \mathbb{R}$ and the non-empty set $\mathcal{D} \subseteq D_{cdf}([a, b])$ satisfy

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

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 2.3 (Examples). Assumption 2.2 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 2.3-2.5 for some discussion, and see their Section 2.1 and Appendices C and E for a detailed discussion of examples.

If the decision maker knew F^1, \dots, F^K , the goal would be to target a vector of weights δ that maximizes $\mathsf{T}(\langle \delta, \mathbf{F} \rangle)$. Denoting by \mathcal{M}_K the set of weight vectors the decision maker is restricted to work with, the task would then be to find an element of¹

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

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

¹Note that the set in Equation (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 $D_{cdf}([a, b])$) and \mathcal{M}_K is a closed subset of the standard simplex of \mathbb{R}^K . This is guaranteed under Assumptions 2.1, 2.2, and 2.5.

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 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 2.4. 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

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

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 Equation (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 \mathbb{T} continuous on the convex set \mathcal{D}).

Assumption 2.5. *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 \mathbb{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 \mathbb{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 \mathbb{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 $\mathcal{M}_K^* \subsetneq \mathcal{M}_K$, 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 $\arg \max_{\delta \in \mathcal{M}_K^*} \mathbb{T}(\langle \delta, \mathbf{F} \rangle)$.

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

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

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).² 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 2.5.

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

2.2.2 Examples of functionals

Example 2.6. A functional for which Assumption 2.2 holds (with $C = 2(b - a)$ and $\mathcal{D} = D_{cdf}([a, b])$, cf. the discussion after Lemma C.9 in Kock et al. (2020a)) is the Gini-welfare measure

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

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 2.7. One important poverty measure, for which Assumption 2.2 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 C.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

$$F \mapsto -F(\mu(F)/2), \text{ where } \mu(F) = \int_{[a,b]} x dF(x). \quad (7)$$

²This follows from quasi-convexity implying that $\max_{\delta \in \mathcal{M}_K} \mathbb{T}(\langle \delta, \mathbf{F} \rangle) = \max_{\delta \in \text{conv}(\mathcal{M}_K)} \mathbb{T}(\langle \delta, \mathbf{F} \rangle)$, for $\text{conv}(\mathcal{M}_K)$ the convex hull of \mathcal{M}_K , together with $\max_{\delta \in \text{conv}(\mathcal{M}_K)} \mathbb{T}(\langle \delta, \mathbf{F} \rangle)$ 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 .

Here we multiply by -1 , as we aim at maximizing \mathbb{T} , and one typically wants to find the treatment combination leading to the smallest fraction of “poor” individuals. Other poverty measures that satisfy Assumption 2.2 but are not quasi-convex (in general) include the ones by Sen (1976) and Foster et al. (1984).

Example 2.8. 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 $\langle \delta, \mathbf{F} \rangle$ resembles this ideal cdf F^* as closely as possible. Then, one could work with the functional

$$F \mapsto -\|F - F^*\|_\infty. \quad (8)$$

By the inverse triangle inequality, Assumption 2.2 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 $\arg \max_{\delta \in \mathcal{M}_K} \mathbb{T}(\langle \delta, \mathbf{F} \rangle)$ 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: Besides learning \mathbf{F} 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 balanced 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 $\mathbb{T}(\langle \hat{\delta}, \mathbf{F} \rangle) \approx \max_{\delta \in \mathcal{M}_K} \mathbb{T}(\langle \delta, \mathbf{F} \rangle)$.

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.

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

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

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 $\mathbb{T}(\langle \pi_{n,n+1}(Z_n, G_{n+1}), \mathbf{F} \rangle)$ is close to $\max_{\delta \in \mathcal{M}_K} \mathbb{T}(\langle \delta, \mathbf{F} \rangle)$. Therefore, we will evaluate policies based on their *regret*

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

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 \mathbb{T} evaluated on \mathcal{D} (if \mathbb{T} is constant on \mathcal{D} the regret of any policy is constant equal to 0, an uninteresting situation).

Assumption 3.1. *The functional $\mathbb{T} : D_{cdf}([a, b]) \rightarrow \mathbb{R}$ satisfies Assumption 2.2, 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

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

As already noted in Kock et al. (2020b), where this assumption was introduced, we emphasize that Equation (11) in Assumption 3.1 is satisfied if, e.g., $\tau \mapsto \mathbb{T}(J_\tau)$ is continuously differentiable on $[0, 1]$ with an everywhere positive derivative.

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 2.5) as

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

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 2.5. The lower bound is as follows.

Theorem 3.2. *Suppose Assumptions 2.1, 2.2, 2.5 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*

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

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 $\{J_\tau : \tau \in [0, 1]\}$ 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 2.5) 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 $\mathcal{T} \subsetneq \mathcal{M}_K$, all $v \in [-1, 1]^K$, and all $w_i \in [-1, 1]$. It is easy to see that $\kappa(\mathcal{M}_K) \geq 0$ under Assumption 2.5.

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, 2.2, 2.5 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*

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

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

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

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 Equation (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 \mathbb{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 \mathbb{T} and operating over a corresponding line segment $\{J_\tau : \tau \in [0, 1]\}$ in \mathcal{D} along which \mathbb{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 $\mathbb{T}(F^1), \dots, \mathbb{T}(F^K)$ to determine $\arg \max_{\delta \in \mathcal{M}_K} \mathbb{T}(\langle \delta, \mathbf{F} \rangle)$, whereas in case, e.g., $\mathcal{M}_K = \mathcal{S}_K$, one needs to learn $\mathbb{T}(\langle \delta, \mathbf{F} \rangle)$ for *all* $\delta \in \mathcal{S}_K$, an uncountable set of potential maximizers. There is some hope, however, because learning the K cdfs F^1, \dots, F^K suffices in order to determine $\mathbb{T}(\langle \delta, \mathbf{F} \rangle)$ for all $\delta \in \mathcal{S}_K$. 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 Equation (15) shows that as long as

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

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 Equation (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 Equa-

tion (15) is 0, and furthermore also $\kappa(\mathcal{M}_K) = 0$ holds.³ 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 in 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 \mathbb{T} is continuous on \mathcal{D} , convex, and that \mathcal{M}_K satisfies Assumption 2.5.

An SA policy proceeds as follows: First, it allocates subject $t = 1, \dots, n$ 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

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

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 $\mathbb{T}(\langle \delta, \hat{\mathbf{F}}_{n,\Pi_n} \rangle)$ 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 \mathbb{T} , there is no guarantee that such a maximizer exists, because $\delta \mapsto \mathbb{T}(\langle \delta, \hat{\mathbf{F}}_{n,\Pi_n} \rangle)$ need not be continuous (even under Assumption 2.2). 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

$$\min \arg \max_{\delta \in \mathcal{M}_K^n} \mathbb{T}(\langle \delta, \hat{\mathbf{F}}_{n,\Pi_n} \rangle), \quad (17)$$

³That $\kappa(\mathcal{M}_K) = 0$ 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 $\max_{\delta \in \mathcal{M}_K} (v' \delta + w_j \delta_j) = \sup_{\delta \in \mathcal{M}_K \setminus \mathcal{T}} (v' \delta + w_j \delta_j)$.

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. 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**

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

end

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

4.1 Choosing a discretization

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

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

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 2.2 and for \mathcal{D} convex

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

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 assumed to be closed throughout, 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

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

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

$$\mathcal{M}_K^n := \mathcal{M}_K \cap \mathcal{S}_{K, \lfloor \varrho_n^{-1} \rfloor} = \bigcup_{i=1}^m \mathcal{S}_{A_i, K} \cap \mathcal{S}_{K, \lfloor \varrho_n^{-1} \rfloor} \quad (21)$$

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, 2.2, 2.5, and 4.2 hold. Then the SA policy $\hat{\pi}$ with partition Π_n , discretization \mathcal{M}_K^n and $\kappa_n := \min_{j=1, \dots, n} |\Pi_{n, j}|$ satisfies

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

Theorem 4.3 shows that working with a discretization having maximal optimization error $\varepsilon(n) \leq \sqrt{K/n}$ (cf. Equation (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 Equation (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. Remark 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 2.2 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) in 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 Theorem 5.1 and 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.* While 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 $\arg \max_{\delta \in \mathcal{M}_K} \mathbb{T}(\langle \delta, \mathbf{F} \rangle)$ have zero i -th coordinate, i.e., if

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

In this case treatment i does not contribute to any maximizer. Note that the maximum in the previous display is well defined if T is continuous on \mathcal{D} , convex, and \mathcal{M}_K is as in Assumption 2.5, which we shall assume throughout this section. Since \mathbf{F} is unknown and T is continuous, estimation uncertainty implies that an inferior treatment i is only empirically detectable if it is actually “strongly inferior” in the sense that

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

(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 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 2.4, which covers many situations of practical relevance. Since T is continuous on \mathcal{D} , convex, the condition

$$\max_{\delta \in \mathcal{M}_K} \mathsf{T}(\langle \delta, \mathbf{F} \rangle) = \max_{\delta \in \bigcup_{l \neq j} \mathcal{M}_{A_l, K}} \mathsf{T}(\langle \delta, \mathbf{F} \rangle) > \max_{\delta \in \mathcal{M}_{A_j, K}} \mathsf{T}(\langle \delta, \mathbf{F} \rangle) \quad (25)$$

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 $\bigcup_{l \neq j} \mathcal{M}_{A_l, K}$. In the special case $\mathcal{M}_K = \mathcal{E}_K$, amounting to $m = K$, the relation in (25) is actually equivalent to

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

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

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

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

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

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 2.4.

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 Equation (26). This is done as follows: at the end of round r we eliminate all treatments $i \in \mathcal{I}_{r-1}$ for which

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

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

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

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. 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_t , which is therefore suppressed in the notation.

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

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

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

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

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**

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

end

end

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

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 The general case

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 2.4. In this 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

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

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 Equation (31) holds by checking its discretized version

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

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 $\max_{\delta \in \bigcup_{l \neq j} \mathcal{M}_{A_l, K}^n} \mathbb{T}(\langle \delta, \mathbf{F} \rangle) > \mathbb{T}(\langle \gamma, \mathbf{F} \rangle)$; 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

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

where $u_\eta(r, n)$ is as defined in Equation (30), and $\hat{\mathbf{F}}_{t, n}$ was defined in Equation (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, 2.2, and 4.2 hold, and that \mathcal{M}_K is as in Example 2.4 and satisfies Assumption 2.5. 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 $D_{n, K} := \sqrt{K \log(17K)} \wedge \sqrt{\log(3.01 \times K(n+2))}$.

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

Input: $n \in \mathbb{N}$, $K \geq n$, 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
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
     $\tilde{\pi}_{t+i,n}(Z_{t+i-1}) = j_i$ 
  end
end
 $\tilde{\pi}_{n,n+1}(Z_n) = \min \arg \max \{\mathbb{T}(\langle \delta, \hat{\mathbf{F}}_{n,n} \rangle) : \delta \in \bigcup_{j \in \mathcal{J}_r} \mathcal{M}_{A_j, K, r}^n\}$ 

```

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 $\delta \mapsto \mathbb{T}(\langle \delta, \mathbf{F} \rangle)$ over \mathcal{M}_K^n have been eliminated after n assignments, and that the last

term upper bounds the probability that all maximizers of $\delta \mapsto \mathbb{T}(\langle \delta, \mathbf{F} \rangle)$ 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 $\eta \geq 1$.

While 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 $\hat{\pi}$. We shall illustrate in the numerical results in the following section that if there are treatments that are sufficiently suboptimal compared to the best treatment, these are eliminated by the SE policy, and then lower expected regret can be obtained than for the SA policy.

6 Numerical results

We now illustrate the theoretical results established in this article by means of numerical examples. Simulations will be conducted for \mathbb{T} the Gini-welfare measure introduced in Example 2.6 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 2.7, 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 2.4, 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 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.

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 2.4. There, it is shown that this functional satisfies Assumption 2.2 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.

6.1.1 Results


The numerical results for the Gini-welfare measure are summarized in Figure 1. The top two panels in Figure 1 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 $\mathcal{M}_K = \mathcal{E}_K$.

Setting 1



Setting 2



Setting 3




Figure 1: 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.

6.2 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.⁴ 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 can not 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.

⁴In none of the situations we consider here, there are incompatibility constraints, or potentially strongly inferior treatments. If incompatibility constraints would be added, the SE policy could again be applied, and could potentially lead to lower average regret, as illustrated in the previous section. Since this has already been illustrated in detail, we do not study this further here.

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

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

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

6.2.1 Results

Figure 2 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.




Figure 2: Average regret of the SA policy for the mean functional under capacity and similarity constraints in Settings 1-3 described in the text.

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 Equation (7) of Example 2.7. 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:

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

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 0.3, 0.31, 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

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

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 2.4. It can be shown that $\delta = \left(\frac{736}{3171}, \frac{2435}{3171}, 0, 0, 0\right)'$ minimizes the headcount ratio over \mathcal{M}_K with a corresponding minimal value of 0.2739. Importantly, we note that the minimizer $\delta \notin \mathcal{E}_K$ even though $\mathcal{E}_K \subseteq \mathcal{M}_K$. The discretization used is

$$\mathcal{M}_K^n = \{\delta \in \mathcal{S}_5 : \delta_3 = \delta_4 = \delta_5 = 0, 100\delta_1 \in \mathbb{Z}\} \cup \{e_3(5)\} \cup \{e_4(5)\} \cup \{e_5(5)\},$$

and we set $C = 2.5$ in accordance with Example 2.7.

6.3.1 Results

The numerical results for the headcount ratio are contained in Figure 3. 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 (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.




Figure 3: Absolute and relative average regrets for Settings 1-3 for the headcount ratio. 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.

Online Appendices

A Auxiliary results

In this section, let $\{X_t\}_{t \in \mathbb{N}}$ be an i.i.d. sequence of random variables with cdf F ; furthermore, for every $t \in \mathbb{N}$, we denote by \hat{F}_t the empirical cdf based on X_1, \dots, X_t , i.e., $\hat{F}_t(x) := \frac{1}{t} \sum_{s=1}^t \mathbb{1}_{\{X_s \leq x\}}$ for every $x \in \mathbb{R}$.

We shall call (the distribution of) a random variable X *sub-Gaussian*, if there exist positive real numbers D and σ^2 such that $\mathbb{P}(|X| \geq x) \leq D e^{-\frac{x^2}{2\sigma^2}}$ for every $x > 0$. Note that we do not require X to have expectation 0; we only require a two-sided exponential tail bound.

Lemma A.1. *Let X be sub-Gaussian. Then, for every $s > 0$ and $k > 0$,*

$$\mathbb{E}[e^{sX}] \leq 1 + D\sqrt{2\pi}\sigma e^{\frac{s^2\sigma^2}{2}} \leq \left(1 + D\sqrt{k/e}\sqrt{2\pi}\right) e^{(\frac{1}{2} + \frac{1}{2k})s^2\sigma^2}. \quad (34)$$

Proof. Note that $\mathbb{E}(e^{sX}) \leq \mathbb{E}(e^{s|X|}) = \int_{(0,\infty)} \mathbb{P}(e^{s|X|} > x) dx$, which is bounded from above by

$$\int_{(0,\infty)} \mathbb{P}(|X| > \log(x)/s) dx \leq 1 + D \int_{(1,\infty)} e^{-(\log(x)/s)^2/(2\sigma^2)} dx,$$

which (after the change of variables $x \mapsto e^{sx}$) is seen to coincide with

$$1 + sD \int_{(1,\infty)} e^{-x^2/(2\sigma^2)+sx} dx = 1 + sD e^{\frac{s^2\sigma^2}{2}} \int_{(1,\infty)} e^{-\frac{1}{2\sigma^2}(x-s\sigma^2)^2} dx \leq 1 + D\sqrt{2\pi}\sigma e^{\frac{s^2\sigma^2}{2}}.$$

To prove the second inequality, use that for every $x \geq 0$ it holds that $\sqrt{x} \leq \sqrt{k/e} \exp(x/2k)$ [since $ex/k \leq \exp(x/k)$] with $x = s^2\sigma^2$. \square

We now summarize some consequences of Lemma A.1 (which hold for every $t \in \mathbb{N}$): Together with the Dvoretzky-Kiefer-Wolfowitz-Massart (DKWM)-inequality, the first inequality in (34) shows that (cf. also Shorack and Wellner (2009), p. 357)

$$\mathbb{E}\left(\exp(st\|\hat{F}_t - F\|_\infty)\right) \leq 1 + \sqrt{2\pi}\sigma\sqrt{t}e^{s^2t/8} \quad \text{for every } s > 0; \quad (35)$$

and the second inequality in (34) shows that for every $k > 0$

$$\mathbb{E}\left(\exp(st\|\hat{F}_t - F\|_\infty)\right) \leq \left(1 + 2\sqrt{\frac{2\pi k}{e}}\right) e^{(1+k^{-1})s^2t/8} \quad \text{for every } s > 0, \quad (36)$$

and thus in particular

$$\mathbb{E}\left(\exp(st\|\hat{F}_t - F\|_\infty)\right) \leq \left(1 + 2\sqrt{2\pi}\right) e^{(1+e^{-1})s^2t/8} \quad \text{for every } s > 0. \quad (37)$$

Doob's submartingale inequality, together with the inequality in Equation (36), delivers the upper bound in the following lemma.

Lemma A.2. For every $n \in \mathbb{N}$, every $x > 0$, and every $k > 0$ it holds that

$$\mathbb{P} \left(\max_{1 \leq t \leq n} t \|\hat{F}_t - F\|_\infty \geq x \right) \leq \left(1 + 2\sqrt{k/e}\sqrt{2\pi} \right) e^{-\frac{2x^2}{(1+1/k)n}}.$$

Proof. We fix $n \in \mathbb{N}$ and $s, x > 0$. It holds that

$$\mathbb{P} \left(\max_{1 \leq t \leq n} t \|\hat{F}_t - F\|_\infty \geq x \right) = \mathbb{P} \left(\max_{1 \leq t \leq n} e^{st} \|\hat{F}_t - F\|_\infty \geq e^{sx} \right).$$

We now verify that $\{e^{st} \|\hat{F}_t - F\|_\infty\}_{t \in \mathbb{N}}$ is a submartingale w.r.t. $\{\mathcal{F}_t\}_{t \in \mathbb{N}}$, the natural filtration generated by $\{X_t\}_{t \in \mathbb{N}}$. By Jensen's inequality for conditional expectations, it follows that for any pair of natural numbers $t_1 < t_2$,

$$\mathbb{E} \left(e^{st_2} \|\hat{F}_{t_2} - F\|_\infty \mid \mathcal{F}_{t_1} \right) \geq e^{s\mathbb{E}(t_2 \|\hat{F}_{t_2} - F\|_\infty \mid \mathcal{F}_{t_1})}. \quad (38)$$

Next, since

$$t_2 \|\hat{F}_{t_2} - F\|_\infty = \sup_{z \in \mathbb{R}} \left| \sum_{r=1}^{t_2} (\mathbf{1}\{X_r \leq z\} - F(z)) \right| \geq \left| \sum_{r=1}^{t_2} (\mathbf{1}\{X_r \leq z_0\} - F(z_0)) \right|,$$

for all $z_0 \in \mathbb{R}$, another application of Jensen's inequality for conditional expectations yields

$$\mathbb{E} \left(t_2 \|\hat{F}_{t_2} - F\|_\infty \mid \mathcal{F}_{t_1} \right) \geq \left| \sum_{r=1}^{t_2} \mathbb{E} \left(\mathbf{1}\{X_r \leq z_0\} - F(z_0) \mid \mathcal{F}_{t_1} \right) \right| = \left| \sum_{r=1}^{t_1} (\mathbf{1}\{X_r \leq z_0\} - F(z_0)) \right|,$$

the equality following from the independence of the random variables X_t . Thus, we conclude that

$$\mathbb{E} \left(t_2 \|\hat{F}_{t_2} - F\|_\infty \mid \mathcal{F}_{t_1} \right) \geq \sup_{z \in \mathbb{R}} \left| \sum_{r=1}^{t_1} (\mathbf{1}\{X_r \leq z\} - F(z)) \right| = t_1 \|\hat{F}_{t_1} - F\|_\infty,$$

which together with (38) establishes that $\{e^{st} \|\hat{F}_t - F\|_\infty\}_{t \in \mathbb{N}}$ is a (positive) submartingale. Hence, Doob's submartingale inequality together with Equation (36) shows for all $s > 0$ that

$$\mathbb{P} \left(\max_{1 \leq t \leq n} t \|\hat{F}_t - F\|_\infty \geq x \right) \leq \frac{\mathbb{E} e^{sn} \|\hat{F}_n - F\|_\infty}{e^{sx}} \leq \left(1 + 2\sqrt{k/e}\sqrt{2\pi} \right) e^{(\frac{1}{2} + \frac{1}{2k}) \frac{n}{4} s^2 - sx}.$$

Minimizing the right hand side of the above display over $s > 0$ yields that

$$\mathbb{P} \left(\max_{1 \leq t \leq n} t \|\hat{F}_t - F\|_\infty \geq x \right) \leq \left(1 + 2\sqrt{k/e}\sqrt{2\pi} \right) e^{-\frac{2x^2}{(1+1/k)n}}.$$

□

Lemma A.3. *Let X and Y be independent sub-Gaussian random variables with the same parameters D and σ^2 . Then, for every $z > 0$ and every $k > 0$,*

$$\mathbb{P}(X + Y \geq z) \leq \left(1 + D\sqrt{k/e}\sqrt{2\pi}\right)^2 e^{-\frac{z^2}{4(1+1/k)\sigma^2}}.$$

Proof. Note that by the independence of X and Y , for all $s > 0$, $k > 0$, and $z > 0$,

$$\mathbb{P}(X + Y \geq z) = \mathbb{P}(e^{s(X+Y)} \geq e^{sz}) \leq \frac{\mathbb{E}e^{sX}\mathbb{E}e^{sY}}{e^{sz}} \leq \left(1 + D\sqrt{k/e}\sqrt{2\pi}\right)^2 e^{(1+\frac{1}{k})s^2\sigma^2 - sz},$$

where the second estimate follows from Lemma A.1. Minimizing the right-hand side of the above display in $s > 0$ yields that

$$\mathbb{P}(X + Y \geq z) \leq \left(1 + D\sqrt{k/e}\sqrt{2\pi}\right)^2 e^{-\frac{z^2}{4(1+1/k)\sigma^2}}.$$

□

Corollary A.4. *Let $K \in \mathbb{N}$, and let $\{X_{1,t}\}_{t \in \mathbb{N}}, \dots, \{X_{K,t}\}_{t \in \mathbb{N}}$ be independent sequences of random variables. Suppose further that for every $i = 1, \dots, K$ the random variables $\{X_{i,t}\}_{t \in \mathbb{N}}$ are i.i.d. with cdf F_i . Denote by $\hat{F}_{i,n}$ the empirical cdf based on $X_{i,1}, \dots, X_{i,n}$. Then, for any $A \subseteq \{1, \dots, K\}$, any $x > 0$, $k > 0$ and $n \in \mathbb{N}$*

$$\mathbb{P}\left(\max_{i \in A} \|\hat{F}_{i,n} - F_i\|_\infty + \max_{i \notin A} \|\hat{F}_{i,n} - F_i\|_\infty \geq x\right) \leq \left(1 + K\sqrt{k/e}\sqrt{8\pi}\right)^2 e^{-\frac{x^2 n}{(1+1/k)}}.$$

Proof. By a union bound and the DKWM inequality $\max_{i \in B} \|\hat{F}_{i,n} - F_i\|_\infty$ for $B \in \{A, A^c\}$ are sub-Gaussian with $D = 2K$ and $\sigma^2 = 1/(4n)$. As the two random variables are also independent, Lemma A.3 yields the desired estimate. □

B Proofs and additional results for Section 3

B.1 A general lower bound

In this section, we start with a lower bound result that will be instrumental for proving Theorem 3.2 and the first part of Theorem 3.3. We need to introduce some further notation. Assume that \mathcal{M}_K satisfies Assumption 2.5, i.e., \mathcal{M}_K is closed and contains at least two distinct elements. Given $\mathcal{H} = \{\mathcal{H}_1, \dots, \mathcal{H}_l\}$, a partition of $\{1, \dots, K\}$, we define for every $i = 1, \dots, l$ the $(K \times |\mathcal{H}_i|)$ -dimensional matrix $B_i = (e_j(K))_{j \in \mathcal{H}_i}$, and set $\kappa(\mathcal{H}, \mathcal{M}_K)$ equal to

$$\sup \left\{ \left[\max_{\delta \in \mathcal{M}_K} v'\delta - \sup_{\delta \in \mathcal{T}} v'\delta \right] \wedge \min_{i=1}^l \left[\max_{\delta \in \mathcal{M}_K} (v + B_i w_i)'\delta - \sup_{\delta \in \mathcal{M}_K \setminus \mathcal{T}} (v + B_i w_i)'\delta \right] \right\}, \quad (39)$$

where the supremum is taken over all nonempty Borel sets $\mathcal{T} \subsetneq \mathcal{M}_K$, all $v \in [-1, 1]^K$, and all $w_i \in [-1, 1]^{|\mathcal{H}_i|}$. Note that for the partition $\mathcal{H} = \{\{1\}, \dots, \{K\}\}$ the expression $\kappa(\mathcal{H}, \mathcal{M}_K)$ coincides with $\kappa(\mathcal{M}_K)$. Note further that $\kappa(\mathcal{H}, \mathcal{M}_K) \geq 0$ (under Assumption 2.5). The following lemma is crucial for establishing Theorem 3.2; the main result in this section is given subsequently.

Lemma B.1. *Let \mathcal{M}_K satisfy Assumption 2.5. Then*

$$2 \times \kappa(\{\{1, \dots, K\}\}, \mathcal{M}_K) \geq \text{diam}^2(\mathcal{M}_K) := \max \{ \|\nu - \gamma\|^2 : (\nu, \gamma) \in \mathcal{M}_K \times \mathcal{M}_K \}.$$

Proof. Let (ν, γ) solve $\max_{(\nu, \gamma) \in \mathcal{M}_K \times \mathcal{M}_K} \|\nu - \gamma\|^2$, and set $v = (\nu - \gamma)/2$ and $w = -2v$. The Cauchy-Schwarz inequality then shows that ν solves $\max_{\delta \in \mathcal{M}_K} v'(\delta - \gamma) = \|\nu - \gamma\|^2/2$, and γ solves $\max_{\delta \in \mathcal{M}_K} (-v)'(\delta - \nu) = \|\nu - \gamma\|^2/2$. For the non-empty Borel set $\mathcal{T} := \{\delta \in \mathcal{M}_K : v'(\delta - \gamma) \leq 0\} \subsetneq \mathcal{M}_K$, we obtain

$$\begin{aligned} \max_{\delta \in \mathcal{M}_K} v'\delta - \sup_{\delta \in \mathcal{T}} v'\delta &= \max_{\delta \in \mathcal{M}_K} v'(\delta - \gamma) - \sup_{\delta \in \mathcal{T}} v'(\delta - \gamma) \geq \|\nu - \gamma\|^2/2, \\ \max_{\delta \in \mathcal{M}_K} (v + w)'\delta - \sup_{\delta \in \mathcal{M}_K \setminus \mathcal{T}} (v + w)'\delta &= \max_{\delta \in \mathcal{M}_K} (-v)'(\delta - \nu) - \sup_{\delta \in \mathcal{M}_K \setminus \mathcal{T}} (-v)'(\delta - \nu) \geq \|\nu - \gamma\|^2/2, \end{aligned}$$

which shows that for this choice of \mathcal{T} , v , and $w = B_1 w_1 = w_1$ (and for the partition $H = \{\{1, \dots, K\}\}$) the term in braces in (39) is not smaller than $\|\nu - \gamma\|^2/2$. \square

Theorem B.2. *Suppose Assumptions 2.1, 2.2, 2.5 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*

$$\sup_{F^1, \dots, F^K \in \{J_\tau : \tau \in [0, 1]\}} \mathbb{E}[r_n(\pi, \mathcal{M}_K)] \geq c \max_{\mathcal{H} \in \{\text{Partitions of } \{1, \dots, K\}\}} \kappa(\mathcal{H}, \mathcal{M}_K) \sqrt{|\mathcal{H}|/n}, \quad \forall n \geq K,$$

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

Proof. We show that there exists a constant c as in the statement of the theorem, such that for every policy π with recommendations in \mathcal{M}_K , any randomization measure \mathbb{P}_G , and any partition $\mathcal{H} = \{\mathcal{H}_1, \dots, \mathcal{H}_l\}$ of $\{1, \dots, K\}$ it holds that

$$\sup_{F^1, \dots, F^K \in \{J_\tau : \tau \in [0, 1]\}} \mathbb{E}[r_n(\pi, \mathcal{M}_K)] \geq c \kappa(\mathcal{H}, \mathcal{M}_K) \sqrt{l/n}, \quad \text{for every } n \geq K. \quad (40)$$

Fix a partition \mathcal{H} , and abbreviate $\kappa := \kappa(\mathcal{H}, \mathcal{M}_K)$ in what follows. If $\kappa = 0$ there is nothing to show (recall that $\kappa \geq 0$). We therefore assume throughout that $\kappa > 0$. Let $\varepsilon := 2/\sqrt{17} < 1/2$. Arguing as in Step 0 of the proof of Theorem 3.9 in Kock et al.

(2020b), we note that Assumption 3.1 implies the existence of a $\zeta > 0$, such that $H_v := J_{1/2+v}$ satisfies

$$\text{KL}^{1/2}(\mu_{H_v}, \mu_{H_w}) \leq \zeta|v - w| \quad \text{for every } v, w \text{ in } [-\varepsilon, \varepsilon], \quad (41)$$

where μ_{H_v} denotes the (unique) probability measure corresponding to the cdf H_v , and

$$c_-(w - v) \leq \mathsf{T}(H_w) - \mathsf{T}(H_v) \leq C(w - v) \quad \text{for every } v \leq w \text{ in } [-\varepsilon, \varepsilon], \quad (42)$$

implying in particular that $w \mapsto \mathsf{T}(H_w)$ is continuous and strictly increasing on $[-\varepsilon, \varepsilon]$.

Now, let π be a policy with recommendations in \mathcal{M}_K . Fix $n \geq K$, and denote $\pi_{n,t} = \pi_t$ for $t = 1, \dots, n + 1$. For $t = 1, \dots, n$ and every $u \in [-\varepsilon, \varepsilon]^K$ we denote by $\mathbb{P}_{\pi,u}^t$ the distribution induced by the random vector Z_t (see Equation (9)) on the Borel sets of \mathbb{R}^{2t} when G_t are i.i.d. with distribution \mathbb{P}_G , and $Y_t = (Y_{1,t}, \dots, Y_{K,t})$ with independent coordinates and marginal cdfs

$$F^i = H_{u_i} \text{ for } i = 1, \dots, K; \quad (43)$$

correspondingly, we denote $\mathbf{F}_u := (H_{u_1}, \dots, H_{u_K})$. We also denote the product measure $\mathbb{P}_{\pi,u}^t \otimes \mathbb{P}_G$ by $\tilde{\mathbb{P}}_{\pi,u}^t$, and write $\mathbb{E}_{\pi,u}^t$ and $\tilde{\mathbb{E}}_{\pi,u}^t$ for the expectation operators corresponding to $\mathbb{P}_{\pi,u}^t$ and $\tilde{\mathbb{P}}_{\pi,u}^t$, respectively.

For every $\delta \in \mathcal{S}_K$ and every $u \in [-\varepsilon, \varepsilon]^K$, we observe that

$$\langle \delta, \mathbf{F}_u \rangle = \sum_{j=1}^K \delta_j J_{1/2+u_j} = J_{1/2+\delta'u} = H_{\delta'u}. \quad (44)$$

Recalling the definition of κ and the notation introduced before the statement of Theorem B.2, we choose $v \in [-1, 1]^K$ and $w_i \in [-1, 1]^{|H_i|}$ for $i = 1, \dots, l$ and a Borel set \mathcal{T} satisfying $\emptyset \neq \mathcal{T} \subsetneq \mathcal{M}_K$, such that

$$\min \left(\max_{\delta \in \mathcal{M}_K} v'\delta - \sup_{\delta \in \mathcal{T}} v'\delta, \min_{i=1}^l \left[\max_{\delta \in \mathcal{M}_K} (v + B_i w_i)'\delta - \sup_{\delta \in \mathcal{M}_K \setminus \mathcal{T}} (v + B_i w_i)'\delta \right] \right) > \kappa/2 > 0. \quad (45)$$

Abbreviate $\epsilon := \sqrt{l/n}\varepsilon$. Define $v^* = (\epsilon/2)v \in [-\varepsilon, \varepsilon]^K$. Because the functions π_1, \dots, π_n map into \mathcal{I} and since $\mathcal{H}_1, \dots, \mathcal{H}_l$ is a partition of \mathcal{I} , it must hold that $\sum_{j=1}^l \sum_{t=1}^n \mathbb{1}\{\pi_t \in \mathcal{H}_j\} = n$. Hence, there exists an index $j^* \in \{1, \dots, l\}$, such that

$$\tilde{\mathbb{E}}_{\pi, v^*}^{n-1} \left[\sum_{t=1}^n \mathbb{1}\{\pi_t \in \mathcal{H}_{j^*}\} \right] \leq \frac{n}{l}. \quad (46)$$

Fix such an index j^* , define $\bar{w} := v + B_{j^*} w_{j^*}$, and rescale this vector to obtain $w^* := (\epsilon/2)\bar{w} \in [-\epsilon, \epsilon]^K$ (recalling that v and w_{j^*} have coordinates in $[-1, 1]$). By construction $v_j^* = w_j^*$ for every $j \notin \mathcal{H}_{j^*}$.

Equation (44) implies

$$\max_{\delta \in \mathcal{M}_K} \mathbb{T}(\langle \delta, \mathbf{F}_{v^*} \rangle) - \sup_{\delta \in \mathcal{T}} \mathbb{T}(\langle \delta, \mathbf{F}_{v^*} \rangle) = \max_{\delta \in \mathcal{M}_K} \mathbb{T}(H_{\delta' v^*}) - \sup_{\delta \in \mathcal{T}} \mathbb{T}(H_{\delta' v^*}). \quad (47)$$

As observed after Equation (42), the function $a \mapsto \mathbb{T}(H_a)$ is continuous and strictly increasing on $[-\epsilon, \epsilon]$. Furthermore, $\max_{\delta \in \mathcal{M}_K} [v^*]' \delta - \sup_{\delta \in \mathcal{T}} [v^*]' \delta > \epsilon \kappa / 4$ as a consequence of (45). It thus follows from the lower bound in Equation (42) that

$$\max_{\delta \in \mathcal{M}_K} \mathbb{T}(\langle \delta, \mathbf{F}_{v^*} \rangle) - \sup_{\delta \in \mathcal{T}} \mathbb{T}(\langle \delta, \mathbf{F}_{v^*} \rangle) \geq c_- \epsilon \kappa / 4.$$

Likewise, it follows that

$$\max_{\delta \in \mathcal{M}_K} \mathbb{T}(\langle \delta, \mathbf{F}_{w^*} \rangle) - \sup_{\delta \notin \mathcal{T}} \mathbb{T}(\langle \delta, \mathbf{F}_{w^*} \rangle) = \max_{\delta \in \mathcal{M}_K} \mathbb{T}(H_{\delta' w^*}) - \sup_{\delta \notin \mathcal{T}} \mathbb{T}(H_{\delta' w^*}) \geq c_- \epsilon \kappa / 4.$$

Therefore, on the event $\{\pi_{n+1} \in \mathcal{T}\}$ we have $r_n(\pi, \mathcal{M}_K; \mathbf{F}_{v^*}; Z_n, G_{n+1}) \geq c_- \epsilon \kappa / 4$; whereas on the event $\{\pi_{n+1} \notin \mathcal{T}\}$ we have $r_n(\pi, \mathcal{M}_K; \mathbf{F}_{w^*}; Z_n, G_{n+1}) \geq c_- \epsilon \kappa / 4$. We conclude that

$$\begin{aligned} \sup_{F^1, \dots, F^K \in \{J_\tau; \tau \in [0, 1]\}} \mathbb{E}[r_n(\pi, \mathcal{M}_K)] &\geq \frac{1}{2} \left[\tilde{\mathbb{E}}_{\pi, v^*}^n(r_n(\pi, \mathcal{M}_K)) + \tilde{\mathbb{E}}_{\pi, w^*}^n(r_n(\pi, \mathcal{M}_K)) \right] \\ &\geq \frac{c_- \epsilon \kappa}{8} \left[\tilde{\mathbb{E}}_{\pi, v^*}^n(\mathbb{1}\{\pi_{n+1} \in \mathcal{T}\}) + \tilde{\mathbb{E}}_{\pi, w^*}^n(\mathbb{1}\{\pi_{n+1} \notin \mathcal{T}\}) \right] \\ &\geq \frac{c_- \epsilon \kappa}{8} \exp\left(-\text{KL}(\mathbb{P}_{\pi, v^*}^n, \mathbb{P}_{\pi, w^*}^n)\right), \end{aligned}$$

where for the third inequality we used that the term in brackets is the sum of the Type I and Type II errors of the test $\mathbb{1}\{\pi_{n+1} \in \mathcal{T}\}$ for testing $\tilde{\mathbb{P}}_{\pi, v^*}^n$ against $\tilde{\mathbb{P}}_{\pi, w^*}^n$, together with, e.g., Theorem 2.2(iii) in Tsybakov (2009) and $\text{KL}(\tilde{\mathbb{P}}_{\pi, v^*}^n, \tilde{\mathbb{P}}_{\pi, w^*}^n) = \text{KL}(\mathbb{P}_{\pi, v^*}^n, \mathbb{P}_{\pi, w^*}^n)$, the latter following, e.g., from the chain rule for the Kullback-Leibler divergence (see, e.g., Lemma A.1 in Kock et al. (2020a) for a suitable statement).

Next note that (cf. also Equation (9)) for every $u \in [-\epsilon, \epsilon]^K$ we can write

$$\mathbb{P}_{\pi, u}^n = \left(\sum_{j=1}^K \mu_{H_{u_j}} \mathbb{1}\{\pi_n = j\} \right) \otimes \tilde{\mathbb{P}}_{\pi, u}^{n-1}, \quad (48)$$

as the term in parentheses defines a stochastic kernel on $\mathcal{B}(\mathbb{R}) \times (\mathbb{R}^{2(n-1)} \times \mathbb{R})$. It thus follows from the chain rule for the Kullback-Leibler divergence, and since $v_j^* = w_j^*$ holds for all $j \notin \mathcal{H}_{j^*}$, that

$$\text{KL}(\mathbb{P}_{\pi, v^*}^n, \mathbb{P}_{\pi, w^*}^n) \leq \text{KL}(\mathbb{P}_{\pi, v^*}^{n-1}, \mathbb{P}_{\pi, w^*}^{n-1}) + \tilde{\mathbb{E}}_{\pi, v^*}^{n-1}(\mathbb{1}\{\pi_n \in \mathcal{H}_{j^*}\}) \max_{j \in \mathcal{H}_{j^*}} \text{KL}(\mu_{H_{v_j^*}}, \mu_{H_{w_j^*}}).$$

By induction and Equations (41) and (46), it hence follows that

$$\text{KL}(\mathbb{P}_{\pi, v^*}^n, \mathbb{P}_{\pi, u^*}^n) \leq \frac{n}{l} \zeta^2 \max_{j \in \mathcal{H}_j^*} (v_j^* - w_j^*)^2 \leq \frac{n}{l} \zeta^2 \epsilon^2.$$

We thus arrive at

$$\sup_{F^1, \dots, F^K \in \{J_\tau : \tau \in [0, 1]\}} \mathbb{E}[r_n(\pi)] \geq \frac{c - \epsilon \kappa}{8} \exp\left(-\frac{n}{l} \zeta^2 \epsilon^2\right) = \sqrt{l/n} \frac{c - \epsilon \kappa}{8} \exp(-\zeta^2 \epsilon^2),$$

where the equality follows from $\epsilon = \sqrt{l/n} \epsilon$. \square

B.2 Proof of Theorem 3.2

This follows immediately from Theorem B.2 and Lemma B.1.

B.3 Proof of Theorem 3.3

Equation (14) follows immediately from Theorem B.2. It remains to show that

$$\kappa(\mathcal{M}_K) \geq \frac{1}{272} \times \max_{\delta \in \mathcal{M}_K} \|\delta\|^2 \times \min_{j=1}^K (\max_{\delta \in \mathcal{M}_K} \delta_j - \min_{\delta \in \mathcal{M}_K} \delta_j)^3; \quad (49)$$

and that

$$\kappa(\mathcal{M}_K) \geq (\min_{j=1}^K \max_{\delta \in \mathcal{M}_K} \delta_j - 1/2)_+^2, \quad (50)$$

where for a real number x we denote its positive part $\max(x, 0)$ by $(x)_+$.

We collect some notation used in this proof. For $j = 1, \dots, K$ define

$$\underline{\delta}_j := \min_{\delta \in \mathcal{M}_K} \delta_j, \quad \bar{\delta}_j := \max_{\delta \in \mathcal{M}_K} \delta_j, \quad \hat{\delta}_j := (\underline{\delta}_j + \bar{\delta}_j)/2, \quad \text{and } d_j := \frac{\bar{\delta}_j - \underline{\delta}_j}{2} = \bar{\delta}_j - \hat{\delta}_j = \hat{\delta}_j - \underline{\delta}_j > 0.$$

By Assumption 2.5 the minima and maxima in the previous display are all well defined.

We organize the proofs of the statements in Equations (49) and (50) in the following two subsections, respectively.

B.3.1 Proof of Equation (49)

Write the lower bound in Equation (49) as $\lambda := \|\mu\|^2 \min_{i=1}^K \frac{d_i^3}{34}$ for $\mu \in \arg \max_{\delta \in \mathcal{M}_K} \|\delta\|$. Since $\kappa(\mathcal{M}_K) \geq 0$ always holds, there is nothing to show if $\lambda = 0$, and we shall thus assume $\lambda > 0$. To show $\kappa(\mathcal{M}_K) \geq \lambda$, we now construct a non-empty Borel set $\mathcal{T} \subsetneq \mathcal{M}_K$, a vector $v \in [-1, 1]^K$, and weights $w_i \in [-1, 1]$ such that

$$\max_{\delta \in \mathcal{M}_K} v' \delta - \sup_{\delta \in \mathcal{T}} v' \delta \geq \lambda, \quad (51)$$

$$\max_{\delta \in \mathcal{M}_K} (v' \delta + w_i \delta_i) - \sup_{\delta \in \mathcal{M}_K \setminus \mathcal{T}} (v' \delta + w_i \delta_i) \geq \lambda \text{ for } i = 1, \dots, K. \quad (52)$$

We set $v := \tau\mu \in [0, 1]^K$ for $\tau := \frac{8}{17} \min_{i=1, \dots, K} d_i$, define the (Borel) set

$$\mathcal{T} := \{\delta \in \mathcal{M}_K : \|\mu\|^2 - \mu'\delta \geq \tau^{-1}\lambda\} = \{\delta \in \mathcal{M}_K : v'\mu - v'\delta \geq \lambda\},$$

and, for $i = 1, \dots, K$, set w_i to -1 if $\mu_i \geq \hat{\delta}_i$, and to 1 if $\mu_i < \hat{\delta}_i$. To proceed (and to show that \mathcal{T} is a non-empty strict subset of \mathcal{M}_K , in particular), we need an auxiliary result: We claim that for every $i = 1, \dots, K$

$$\|\mu\|^2 - \mu'\gamma \geq \|\mu\|^2 d_i^2 / 16, \quad \forall \gamma \in \{\delta \in \mathcal{M}_K : w_i \delta_i \geq w_i [2\hat{\delta}_i + (1 - w_i)\underline{\delta}_i + (1 + w_i)\bar{\delta}_i] / 4\}. \quad (53)$$

Note that the set in the previous display is non-empty, which follows upon choosing γ with i -th coordinate equal to $\underline{\delta}_i$ (equal to $\bar{\delta}_i$) if $w_i = -1$ (if $w_i = 1$). To prove the claim fix i , let $\gamma \in \mathcal{M}_K$ be such that $w_i \gamma_i \geq w_i (2\hat{\delta}_i + (1 - w_i)\underline{\delta}_i + (1 + w_i)\bar{\delta}_i) / 4$, and bound

$$\frac{d_i^2}{4} \leq (\mu_i - \gamma_i)^2 \leq \|\mu - \gamma\|^2 = \|\mu\|^2 + \|\gamma\|^2 - 2\|\mu\|\|\gamma\|\cos(\theta),$$

for θ such that $\cos(\theta) = \mu'\gamma / (\|\mu\|\|\gamma\|)$. Noting that $\|\mu\|\|\gamma\| \leq 1$, as μ and γ are elements of \mathcal{S}_K , we obtain for $x := \|\mu\|/\|\gamma\| \geq 1$ that

$$\cos(\theta) \leq \frac{1}{2} (x + x^{-1} - d_i^2 / 4).$$

If $1 \leq x \leq 1 + \frac{d_i^2}{8}$, then the upper bound just derived can be further upper bounded by $1 - \frac{d_i^2}{16}$, which allows us to conclude that

$$\mu'\gamma - \mu'\mu = \cos(\theta)\|\gamma\|\|\mu\| - \|\mu\|^2 \leq -\|\mu\|^2 d_i^2 / 16. \quad (54)$$

Multiplying by -1 we obtain the first inequality in (53). If $x > 1 + d_i^2/8$, then we obtain

$$\mu'\gamma - \mu'\mu \leq \|\mu\|\|\gamma\| - \|\mu\|^2 = \|\mu\|^2 (x^{-1} - 1) \leq -\frac{\|\mu\|^2}{8} \left(\frac{d_i^2}{1 + d_i^2/8} \right),$$

which is bounded from above by the upper bound in (54), and which proves the claim.

The equality $\tau^{-1}\lambda = \|\mu\|^2 \min_{i=1}^K d_i^2 / 16$ implies that \mathcal{T} contains the non-empty set in (53) (for every i). Hence \mathcal{T} is non-empty. Furthermore, $\mu \notin \mathcal{T}$, as $\tau^{-1}\lambda > 0$. This establishes that \mathcal{T} is strictly contained in \mathcal{M}_K . By definition of \mathcal{T} , this now delivers (51).

We now prove (52), and shall use that for every $\delta \in \mathcal{M}_K$ such that $\delta \neq \mu$, the Cauchy-Schwarz inequality implies

$$v'\delta = \tau\mu'\delta < \tau\|\mu\|\|\delta\| \leq \tau\|\mu\|^2 = v'\mu, \quad (55)$$

in particular $\max_{\delta \in \mathcal{M}_K} v'\delta = v'\mu$. Now, fix $i \in \{1, \dots, K\}$. Note that the definition of \mathcal{T} , $\tau^{-1}\lambda = \|\mu\|^2 \min_{i=1}^K d_i^2 / 16$ and the statement in (53) implies

$$w_i \delta_i < w_i \left(2\hat{\delta}_i + (1 - w_i)\underline{\delta}_i + (1 + w_i)\bar{\delta}_i \right) / 4 \text{ for every } \delta \in \mathcal{M}_K \setminus \mathcal{T}.$$

Suppose first that $w_i = -1$. Let $\eta \in \mathcal{M}_K$ be such that $\eta_i = \underline{\delta}_i$, and let $\delta \in \mathcal{M}_K \setminus \mathcal{T}$. Then the previous display implies $\delta_i > (\hat{\delta}_i + \underline{\delta}_i)/2$, which together with (55) gives

$$(\eta'v + w_i\eta_i) - (\delta'v + w_i\delta_i) = \eta'v + (\delta_i - \underline{\delta}_i) - \delta'v \geq (\hat{\delta}_i/2 - \underline{\delta}_i/2) - \mu'v = \frac{d_i}{2} - \tau\|\mu\|^2,$$

which by the definition of τ and $\|\mu\|^2 \leq 1$ is bounded from below by

$$\frac{d_i}{2}(1 - 16/17) = d_i/34 \geq \lambda, \quad (56)$$

which proves Equation (52). Assume next that $w_i = 1$, let $\eta \in \mathcal{M}_K$ be such that $\eta_i = \bar{\delta}_i$, and $\delta \in \mathcal{M}_K \setminus \mathcal{T}$. Then one has $\delta_i < (\hat{\delta}_i + \bar{\delta}_i)/2$, which gives

$$(\eta'v + w_i\eta_i) - (\delta'v + w_i\delta_i) \geq \eta'v + (\bar{\delta}_i/2 - \hat{\delta}_i/2) - \delta'v \geq \frac{d_i}{2} - \tau\|\mu\|^2,$$

which again proves Equation (52).

B.3.2 Proof of Equation (50)

If $\min_{j=1}^K \bar{\delta}_j - 1/2 \leq 0$, there is nothing to show. Hence, we assume from now on that $\min_{j=1}^K \bar{\delta}_j - 1/2 =: \varrho > 0$, and thus in particular

$$\bar{\delta}_j - 1/2 \geq \varrho, \text{ for every } j = 1, \dots, K. \quad (57)$$

It remains to verify that $\kappa(\mathcal{M}_K) \geq \varrho^2$. Assuming (57) there must exist an index $j^* \in \{1, \dots, K\}$, such that

$$\underline{\delta}_{j^*} \leq \frac{1}{2} - \varrho. \quad (58)$$

Now, set $v = (0, \dots, 0, \varrho, 0, \dots, 0)' \in [-1, 1]^K$ (ϱ being at the j^* -th coordinate), and define for $i = 1, \dots, K$ the weights $w_i = 1$ for $i \neq j^*$ and $w_{j^*} = -1$. Finally, define $\mathcal{T} := \{\delta \in \mathcal{M}_K : \delta_{j^*} \leq 1/2\}$ (a non-empty Borel set, that does not coincide with \mathcal{M}_K).

To verify $\kappa(\mathcal{M}_K) \geq \varrho^2$ we now show

$$\min \left(\max_{\delta \in \mathcal{M}_K} v'\delta - \sup_{\delta \in \mathcal{T}} v'\delta, \min_{i=1}^K \left[\max_{\delta \in \mathcal{M}_K} (v'\delta + w_i\delta_i) - \sup_{\delta \in \mathcal{M}_K \setminus \mathcal{T}} (v'\delta + w_i\delta_i) \right] \right) \geq \varrho^2. \quad (59)$$

We start with

$$\max_{\delta \in \mathcal{M}_K} v'\delta - \sup_{\delta \in \mathcal{T}} v'\delta = \max_{\delta \in \mathcal{M}_K} \varrho\delta_{j^*} - \sup_{\delta \in \mathcal{M}_K : \delta_{j^*} \leq 1/2} \varrho\delta_{j^*} \geq \varrho\bar{\delta}_{j^*} - \varrho/2 \geq \varrho^2,$$

where the second inequality follows from (57). Next, consider an index $i \neq j^*$ and write

$$\max_{\delta \in \mathcal{M}_K} (v'\delta + w_i\delta_i) - \sup_{\delta \in \mathcal{M}_K \setminus \mathcal{T}} (v'\delta + w_i\delta_i) = \max_{\delta \in \mathcal{M}_K} (\delta_{j^*}\varrho + \delta_i) - \sup_{\delta \in \mathcal{M}_K : \delta_{j^*} > 1/2} (\delta_{j^*}\varrho + \delta_i).$$

The maximum to the right can trivially be bounded from below by $\bar{\delta}_i \geq 1/2 + \varrho$; the supremum can be upper bounded by (observing that $\varrho \in (0, 1/2]$)

$$\sup_{\delta \in \mathcal{M}_K: \delta_{j^*} > 1/2} (\delta_{j^*} \varrho + \delta_i) \leq \sup_{\delta \in \mathcal{M}_K: \delta_{j^*} > 1/2} [\delta_{j^*} \varrho + (1 - \delta_{j^*})] \leq 1 + (\varrho - 1)/2 = 1/2 + \varrho/2.$$

Combining what we have just observed yields

$$\max_{\delta \in \mathcal{M}_K} (v' \delta + w_i \delta_i) - \sup_{\delta \in \mathcal{M}_K \setminus \mathcal{T}} (v' \delta + w_i \delta_i) \geq 1/2 + \varrho - [1/2 + \varrho/2] = \varrho/2 \geq \varrho^2.$$

Finally, consider

$$\max_{\delta \in \mathcal{M}_K} (v' \delta + w_{j^*} \delta_{j^*}) - \sup_{\delta \in \mathcal{M}_K \setminus \mathcal{T}} (v' \delta + w_{j^*} \delta_{j^*}) = \max_{\delta \in \mathcal{M}_K} \delta_{j^*} (\varrho - 1) - \sup_{\delta \in \mathcal{M}_K: \delta_{j^*} > 1/2} \delta_{j^*} (\varrho - 1).$$

Since $\varrho - 1$ is negative, the maximum to the right equals $\underline{\delta}_{j^*} (\varrho - 1)$. For the same reason, the supremum is bounded from above by $(\varrho - 1)/2$. Together this implies the lower bound

$$(\varrho - 1)(\underline{\delta}_{j^*} - 1/2) = (1 - \varrho)(1/2 - \underline{\delta}_{j^*}) \geq \varrho(1 - \varrho) \geq \varrho^2,$$

where we used Equation (58) and $\varrho \in (0, 1/2]$.

C Proof of Theorem 4.3

Before we prove the theorem, we verify that the functions $\pi_{n,t}$ defined in Policy 1 are measurable under Assumption 4.2: It is obvious that for any $n \in \mathbb{N}$ the functions $\hat{\pi}_{n,t}$ for $t = 1, \dots, n$ are measurable, simply because they are constant. Assumption 4.2 guarantees, for every $n \in \mathbb{N}$ and every $\delta \in \mathcal{S}_K$, the measurability of $y \mapsto \mathbb{T}(\langle \delta, \hat{\mathbf{F}}_{n, \Pi_n} \rangle)$ (interpreted as a function on \mathbb{R}^n). Thus, the measurability of the selection $\hat{\pi}_{n,n+1}$ immediately follows, noting that for every $\bar{\delta} \in \mathcal{M}_K^n$ we may write $\{\hat{\pi}_{n,n+1} = \bar{\delta}\}$ as the intersection

$$\bigcap_{\delta \in \mathcal{M}_K^n: \delta \geq \bar{\delta}} \{\mathbb{T}(\langle \bar{\delta}, \hat{\mathbf{F}}_{n, \Pi_n} \rangle) \geq \mathbb{T}(\langle \delta, \hat{\mathbf{F}}_{n, \Pi_n} \rangle)\} \cap \bigcap_{\delta \in \mathcal{M}_K^n: \delta < \bar{\delta}} \{\mathbb{T}(\langle \bar{\delta}, \hat{\mathbf{F}}_{n, \Pi_n} \rangle) > \mathbb{T}(\langle \delta, \hat{\mathbf{F}}_{n, \Pi_n} \rangle)\}.$$

Now, we move on to the proof of Theorem 4.3: Fix $n \geq K$, and denote $\hat{\pi}_{n,t} = \hat{\pi}_t$ for $t = 1, \dots, n+1$. We start with the observation that, using convexity of \mathcal{D} and of $D_{cdf}([a, b])$ together with Assumption 2.2 for the second inequality,

$$\begin{aligned} \left| \max_{\delta \in \mathcal{M}_K^n} \mathbb{T}(\langle \delta, \hat{\mathbf{F}}_{n, \Pi_n} \rangle) - \max_{\delta \in \mathcal{M}_K^n} \mathbb{T}(\langle \delta, \mathbf{F} \rangle) \right| &\leq \max_{\delta \in \mathcal{M}_K^n} |\mathbb{T}(\langle \delta, \hat{\mathbf{F}}_{n, \Pi_n} \rangle) - \mathbb{T}(\langle \delta, \mathbf{F} \rangle)| \\ &\leq C \max_{\delta \in \mathcal{M}_K^n} \|\langle \delta, \hat{\mathbf{F}}_{n, \Pi_n} \rangle - \langle \delta, \mathbf{F} \rangle\|_\infty \\ &\leq C \max_{j=1}^K \|\hat{F}_{n, \Pi_n}^j - F^j\|_\infty. \end{aligned}$$

Now, $\hat{\pi}_{n,n+1} = \min \arg \max_{\delta \in \mathcal{M}_K^n} \mathsf{T}(\langle \delta, \hat{\mathbf{F}}_{n, \Pi_n} \rangle)$ and (18) shows that the regret $r_n(\hat{\pi}, \mathcal{M}_K) = \max_{\delta \in \mathcal{M}_K} \mathsf{T}(\langle \delta, \mathbf{F} \rangle) - \mathsf{T}(\langle \hat{\pi}_{n+1}, \mathbf{F} \rangle)$ is upper bounded by

$$\max_{\delta \in \mathcal{M}_K^n} \mathsf{T}(\langle \delta, \mathbf{F} \rangle) - \max_{\delta \in \mathcal{M}_K^n} \mathsf{T}(\langle \delta, \hat{\mathbf{F}}_{n, \Pi_n} \rangle) + \mathsf{T}(\langle \hat{\pi}_{n+1}, \hat{\mathbf{F}}_{n, \Pi_n} \rangle) - \mathsf{T}(\langle \hat{\pi}_{n+1}, \mathbf{F} \rangle) + \varepsilon(n).$$

From $\hat{\pi}_{n+1} \in \mathcal{M}_K^n$ we obtain

$$\mathsf{T}(\langle \hat{\pi}_{n+1}, \hat{\mathbf{F}}_{n, \Pi_n} \rangle) - \mathsf{T}(\langle \hat{\pi}_{n+1}, \mathbf{F} \rangle) \leq \max_{\delta \in \mathcal{M}_K^n} |\mathsf{T}(\langle \delta, \hat{\mathbf{F}}_{n, \Pi_n} \rangle) - \mathsf{T}(\langle \delta, \mathbf{F} \rangle)|,$$

for which an upper bound has already been developed above. Summarizing, we obtain

$$r_n(\hat{\pi}, \mathcal{M}_K) \leq \varepsilon(n) + 2C \max_{j=1}^K \|\hat{F}_{n, \Pi_n}^j - F^j\|_\infty. \quad (60)$$

Denote $M_j := \|\hat{F}_{n, \Pi_n}^j - F^j\|_\infty$. Using Jensen's inequality, Equation (35) and $\kappa_n \leq |\Pi_{n,j}|$, for every j , we get for every $t > 0$ that

$$\exp \left[t \mathbb{E} \left(\max_{j=1}^K M_j \right) \right] \leq \sum_{j=1}^K \mathbb{E} \left(e^{\frac{t}{|\Pi_{n,j}|} |\Pi_{n,j}| M_j} \right) \leq K \left(1 + \sqrt{2\pi/\kappa_n} t e^{t^2/(8\kappa_n)} \right), \quad (61)$$

or equivalently

$$\mathbb{E} \left(\max_{j=1}^K M_j \right) \leq t^{-1} \log \left(K \left(1 + \sqrt{2\pi/\kappa_n} t e^{t^2/(8\kappa_n)} \right) \right).$$

Upon inserting $t = \sqrt{8\kappa_n \log(K) a}$ for some $a \geq 1$, the upper bound becomes

$$\left(\frac{1}{\sqrt{8a}} + \sqrt{a} \frac{\log(1 + 4\sqrt{\pi \log(K^a)} K^a)}{\sqrt{8} \log(K^a)} \right) \sqrt{\log(K)/\kappa_n},$$

which, by Lemma C.1 below (with $c = 4\sqrt{\pi}$, thus $e^{e/c^2} < 2^a$), is bounded from above by

$$\frac{1}{\sqrt{8}} \left(\frac{1}{\sqrt{a}} + \sqrt{a} \frac{\log(1 + 4\sqrt{\pi \log(2^a)} 2^a)}{\log(2^a)} \right) \sqrt{\log(K)/\kappa_n}.$$

This yields

$$\mathbb{E} \left(\max_{j=1}^K M_j \right) \leq \inf_{a \geq 1} \left(\frac{\log(2 + 8\sqrt{\pi \log(2^a)} 2^a)}{\log(2) \sqrt{8a}} \right) \sqrt{\frac{\log(K)}{\kappa_n}} \leq 1.505 \times \sqrt{\frac{\log(K)}{\kappa_n}}, \quad (62)$$

where the last inequality is obtained by setting $a = 3$. Together with Equation (60), this proves the theorem.

Lemma C.1. *Let $c > 0$. Then $x \mapsto \frac{\log(1+c\sqrt{\log(x)x})}{\log(x)}$ is strictly decreasing on $(e^{e/c^2}, \infty)$.*

Proof. It suffices to show that $x \mapsto \log(1+cx e^{x^2})/(x^2)$ is strictly decreasing on $(\sqrt{e}/c, \infty)$. The derivative of this function at $x \in (\sqrt{e}/c, \infty)$ equals

$$\frac{2cx^2 e^{x^2} + ce^{x^2}}{x^2(cx e^{x^2} + 1)} - \frac{2 \log(cx e^{x^2} + 1)}{x^3},$$

which is negative if and only if $2 \log(cx e^{x^2} + 1) + ce^{x^2} x(2 \log(cx e^{x^2} + 1) - 2x^2 - 1)$ is positive, which holds because of $2 \log(cx e^{x^2} + 1) > 2x^2 + 1$, a consequence of $cx \geq e^{1/2}$. \square

D Proof of Theorem 5.1

Instead of proving Theorem 5.1, we establish the following slightly stronger theorem, where we denote $f_{n,K}(u) := \frac{1}{2} + \frac{u}{n} - \frac{uK}{2n} + \frac{1}{n}$ (for $n > 0$).

Theorem D.1. *Suppose Assumptions 2.1, 2.2, and 4.2 hold, and that \mathcal{M}_K is as in Example 2.4 and satisfies Assumption 2.5. Then, the SE policy $\tilde{\pi}$ satisfies*

$$\sup_{\substack{F^i \in \mathcal{D} \\ i=1, \dots, K}} \mathbb{E} [r_n(\tilde{\pi}, \mathcal{M}_K)] \leq \varepsilon(n) + C \times (2 \min(A_{n,K}, B_{n,K}) + C_{n,K}), \text{ for every } n > K, \quad (63)$$

where

$$\begin{aligned} A_{n,K} &:= \frac{1+e}{\sqrt{2e(1-K/n)}} \sqrt{\log(K [1 + (1 + \sqrt{8\pi})\sqrt{2\pi}])} \sqrt{f_{n,K}(\underline{r})} \frac{K}{\sqrt{n}} \\ &\leq \frac{1}{(1-K/n)} \sqrt{\log(17 \times K)} \frac{K}{\sqrt{n}} \end{aligned}$$

$$\begin{aligned} B_{n,K} &:= \sqrt{\frac{1+e}{2e(1-K/n)} \log\left(K \left[\underline{r} + \left\lceil \frac{n - \underline{r}K}{2} \right\rceil - \lfloor n/K \rfloor + 1\right] (1 + \sqrt{8\pi})\right)} \frac{K}{n} \\ &\leq \sqrt{\frac{1}{(1-K/n)} \log\left(6.02 \times K \left[3 + \frac{n}{2} - \frac{n}{K}\right]\right)} \frac{K}{n} \end{aligned}$$

and

$$\begin{aligned} C_{n,K} &:= \min\left(\frac{(1 + K\sqrt{16\pi}/(\eta e))^2}{(\sqrt{n}K)^{1+\eta/(2+\eta)}} \sum_{r \in \mathcal{R}} \frac{1}{r^{1+\eta/(2+\eta)}}, \frac{2K^{(1-\eta)/2}}{n^{(1+\eta)/4}} \sum_{r \in \mathcal{R}} \frac{1}{r^{(1+\eta)/2}}\right) \\ &\leq \min\left(\frac{(1 + K\sqrt{16\pi}/(\eta e))^2}{(\sqrt{n}K)^{1+\eta/(2+\eta)}} \frac{2(1+\eta)}{\eta}, \frac{2K^{(1-\eta)/2}}{n^{(1+\eta)/4}} \sum_{r \in \mathcal{R}} \frac{1}{r^{(1+\eta)/2}}\right), \end{aligned}$$

where

$$\frac{(1 + K\sqrt{16\pi}/(\eta e))^2}{(\sqrt{n}K)^{1+\eta/(2+\eta)}} \frac{2(1+\eta)}{\eta} \leq \frac{1 + 18.5 \times K^2/\eta}{(\sqrt{n}K)^{1+\eta/(2+\eta)}} \frac{4(1+\eta)}{\eta} \leq \frac{K^{1-\frac{3\eta}{4+2\eta}}}{\sqrt{n}} \times 4(1 + 18.5\eta^{-1})^2.$$

To prove Theorem D.1, fix $n \in \mathbb{N}$, $n > K$, a partition $\{A_1, \dots, A_m\}$ of \mathcal{I} such that $m \geq 2$, a discretization $\mathcal{M}_K^n = \bigcup_{j=1}^m \mathcal{M}_{A_j, K}^n$ (finite), such that $\emptyset \neq \mathcal{M}_{A_j, K}^n \subseteq \mathcal{M}_{A_j, K}$ for every $j = 1, \dots, m$, a set of elimination rounds \mathcal{R} , and $\eta > 0$. Abbreviate $\tilde{\pi}_{n,t} = \tilde{\pi}_t$ for $t = 1, \dots, n+1$.

Before we establish the upper bound in Theorem D.1, we verify that the functions $\tilde{\pi}_{n,t}$ (implicitly) defined in Policy 3 are Borel measurable. For $t = 1, \dots, K$ the function $\tilde{\pi}_t$ is constant, and hence Borel measurable. Consider next the case where $K+1 \leq t \leq n$. Since the policy does not depend on external randomization, we need to verify that $\{z \in [a, b]^{t-1} : \tilde{\pi}_t(z) = s\}$ is a Borel set for every $s \in \{1, \dots, K\}$. Fix such an s . It is tedious but not difficult to see that $\{z \in [a, b]^{t-1} : \tilde{\pi}_t(z) = s\}$ can be written as a finite number of intersections and unions of sets (and their complements) of the form

$$A(\mathcal{S}, \Pi, c, \gamma) = \{z \in [a, b]^{t-1} : \max_{\delta \in \mathcal{S}} \mathbb{T}(\langle \delta, \mathbf{F}_\Pi(z) \rangle) - \mathbb{T}(\langle \gamma, \mathbf{F}_\Pi(z) \rangle) \leq c\}; \quad (64)$$

where \mathcal{S} is a finite subset of \mathcal{S}_K ; $\Pi = (\Pi_1, \dots, \Pi_K)$ where the Π_i are disjoint and non-empty subsets of $\{1, \dots, t-1\}$; where $\mathbf{F}_\Pi(z) = (F_{\Pi_1}(z), \dots, F_{\Pi_K}(z))$ with $F_{\Pi_j}(\cdot) = |\Pi_j|^{-1} \sum_{i \in \Pi_j} \mathbf{1}\{z_i \leq \cdot\}$; $c > 0$; and $\gamma \in \mathcal{S}_K$. That every such $A(\mathcal{S}, \Pi, c, \gamma)$ is Borel measurable follows immediately from Assumption 4.2, which implies that $\tilde{\pi}_t$ is measurable. The measurability of $\tilde{\pi}_{n,n+1}$ is shown analogously, observing that for every $\delta \in \mathcal{M}_K^n$ the set $\{z \in [a, b]^n : \tilde{\pi}_{n,n+1}(z) = \delta\}$ can be written as a finite union and intersection of sets (and their complements) as in the previous display (but with $t-1$ replaced by n).

To establish the upper bound in (63), denote by r_C the last round completed, and denote the last elimination round completed by $\bar{r} := \max\{r \in \mathcal{R} : r \leq r_C\}$. We have $\bar{r} \leq \min(r_C, \lfloor n/2 \rfloor)$, since $r \leq \lfloor n/2 \rfloor$ for all $r \in \mathcal{R}$. Note that r_C and \bar{r} are random variables (Borel measurability can be verified as in the beginning of the proof; we suppress the dependence of r_C and \bar{r} on $\omega \in \Omega$, for $(\Omega, \mathcal{A}, \mathbb{P})$ the underlying probability space), and note that $\mathcal{I}_{\bar{r}} = \mathcal{I}_{r_C}$, $\mathcal{J}_{\bar{r}} = \mathcal{J}_{r_C}$, and $\mathcal{M}_{A_j, K, \bar{r}}^n = \mathcal{M}_{K, A_j, r_C}^n$ for every $j \in \mathcal{J}_{\bar{r}}$, since no elimination takes place after round \bar{r} . Set

$$\mathcal{M}_{K, \bar{r}}^n := \bigcup_{j \in \mathcal{J}_{\bar{r}}} \mathcal{M}_{A_j, K, \bar{r}}^n.$$

Define the event (Borel measurability can again be verified as in the argument in the beginning of the proof) where not all $\delta \in \arg \max_{\delta \in \mathcal{M}_K^n} \mathbb{T}(\langle \delta, \mathbf{F} \rangle)$ have been eliminated after \bar{r} rounds

$$\mathcal{G} := \{\omega \in \Omega : \mathcal{M}_{K, \bar{r}}^n \cap \arg \max_{\delta \in \mathcal{M}_K^n} \mathbb{T}(\langle \delta, \mathbf{F} \rangle) \neq \emptyset\}.$$

Consider $\omega \in \mathcal{G}$ in this paragraph: if $|\mathcal{I}_{\bar{r}}| = 1$, i.e., when there is only a single treatment left after the last completed elimination round, then the weights vector which puts mass 1

on that treatment must be an element of $\arg \max_{\delta \in \mathcal{M}_K^n} \mathbb{T}(\langle \delta, \mathbf{F} \rangle)$, and the regret is at most $\varepsilon(n)$. Assume that ω is such that $|\mathcal{I}_{\bar{r}}| = |\mathcal{I}_{r_C}| \geq 2$. Then, $\omega \in \mathcal{G}$ implies

$$r_n(\tilde{\pi}, \mathcal{M}_K) = \max_{\delta \in \mathcal{M}_K} \mathbb{T}(\langle \delta, \mathbf{F} \rangle) - \mathbb{T}(\langle \tilde{\pi}_{n+1}, \mathbf{F} \rangle) \leq \max_{\delta \in \mathcal{M}_{K, \bar{r}}^n} \mathbb{T}(\langle \delta, \mathbf{F} \rangle) - \mathbb{T}(\langle \tilde{\pi}_{n+1}, \mathbf{F} \rangle) + \varepsilon(n).$$

The right-hand side of the above display can also be written as

$$\max_{\delta \in \mathcal{M}_{K, \bar{r}}^n} \mathbb{T}(\langle \delta, \mathbf{F} \rangle) - \mathbb{T}(\langle \tilde{\pi}_{n+1}, \hat{\mathbf{F}}_{n,n} \rangle) + \mathbb{T}(\langle \tilde{\pi}_{n+1}, \hat{\mathbf{F}}_{n,n} \rangle) - \mathbb{T}(\langle \tilde{\pi}_{n+1}, \mathbf{F} \rangle) + \varepsilon(n),$$

which, since by definition $\tilde{\pi}_{n+1} = \min \arg \max \{ \mathbb{T}(\langle \delta, \hat{\mathbf{F}}_{n,n} \rangle) : \delta \in \mathcal{M}_{K, \bar{r}}^n \}$, equals

$$\max_{\delta \in \mathcal{M}_{K, \bar{r}}^n} \mathbb{T}(\langle \delta, \mathbf{F} \rangle) - \max_{\delta \in \mathcal{M}_{K, \bar{r}}^n} \mathbb{T}(\langle \delta, \hat{\mathbf{F}}_{n,n} \rangle) + \mathbb{T}(\langle \tilde{\pi}_{n+1}, \hat{\mathbf{F}}_{n,n} \rangle) - \mathbb{T}(\langle \tilde{\pi}_{n+1}, \mathbf{F} \rangle) + \varepsilon(n).$$

Next observe that $\max_{\delta \in \mathcal{M}_{K, \bar{r}}^n} \mathbb{T}(\langle \delta, \mathbf{F} \rangle) - \max_{\delta \in \mathcal{M}_{K, \bar{r}}^n} \mathbb{T}(\langle \delta, \hat{\mathbf{F}}_{n,n} \rangle)$ and $\mathbb{T}(\langle \tilde{\pi}_{n+1}, \hat{\mathbf{F}}_{n,n} \rangle) - \mathbb{T}(\langle \tilde{\pi}_{n+1}, \mathbf{F} \rangle)$ are bounded from above by

$$\max_{\delta \in \mathcal{M}_{K, \bar{r}}^n} |\mathbb{T}(\langle \delta, \hat{\mathbf{F}}_{n,n} \rangle) - \mathbb{T}(\langle \delta, \mathbf{F} \rangle)| \leq C \max_{\delta \in \mathcal{M}_{K, \bar{r}}^n} \|\langle \delta, \hat{\mathbf{F}}_{n,n} \rangle - \langle \delta, \mathbf{F} \rangle\|_\infty \leq C \max_{i \in \mathcal{I}_{\bar{r}}} \|\hat{F}_{i,n,n} - F^i\|_\infty,$$

where the first inequality follows from convexity of \mathcal{D} and Assumption 2.2. Note that $\underline{r} := \min \mathcal{R}$ is the number of times each treatment in \mathcal{I} has been assigned by the end of the first elimination round. Together with $|\mathcal{I}_{\bar{r}}| = |\mathcal{I}_{r_C}| \geq 2$ this implies

$$\lfloor n/K \rfloor \leq S_i(n) \leq \underline{r} + \left\lceil \frac{n - \underline{r}K}{2} \right\rceil =: R, \text{ for every } i \in \mathcal{I}_{\bar{r}},$$

(note that R is non-random, as $\underline{r} = \min \mathcal{R} \leq n/K$ is fixed), from which it follows that

$$\begin{aligned} \max_{i \in \mathcal{I}_{\bar{r}}} \|\hat{F}_{i,n,n} - F^i\|_\infty &\leq \max_{i=1}^K \left[\|\hat{F}_{i,n,n} - F^i\|_\infty \mathbf{1}\{S_i(n) \in \{\lfloor n/K \rfloor, \dots, R\}\} \right] \\ &= \max_{i=1}^K \left[\sum_{s=\lfloor n/K \rfloor}^R \|\hat{F}_{i,n,n} - F^i\|_\infty \mathbf{1}\{S_i(n) = s\} \right]. \end{aligned}$$

Using that $r_n(\tilde{\pi}, \mathcal{M}_K) \leq C$ for $\omega \in \Omega \setminus \mathcal{G}$ by Assumption 2.2, we conclude that for every $\omega \in \Omega$ it holds that

$$r_n(\tilde{\pi}, \mathcal{M}_K) \leq 2C \max_{i=1}^K \left[\sum_{s=\lfloor n/K \rfloor}^R \|\hat{F}_{i,n,n} - F^i\|_\infty \mathbf{1}\{S_i(n) = s\} \right] + \varepsilon(n) + C \mathbf{1}\{\Omega \setminus \mathcal{G}\}.$$

The theorem now follows from Parts 3 and 4 of the auxiliary Lemma D.2 given next (the remaining upper bounds on $A_{n,K}$, $B_{n,K}$ and $C_{n,K}$ being obvious).

While we continue to use the notation and notational conventions introduced already, to formulate and prove Lemma D.2, we define the sequence of measurable functions f_t as:

$$f_t := \begin{cases} \tilde{\pi}_{n,t} & \text{for } 1 \leq t \leq n \\ (t \bmod K) + 1 & \text{for } t > n, \end{cases}$$

and define $\tilde{Z}_t = Z_t$ for $t = 1, \dots, n$, and recursively set $\tilde{Z}_{t+1} = (Y_{f_t(Z_{t-1})}, Z_t)$ for $t \geq n + 1$. For every $i \in \mathcal{I}$ and $t \in \mathbb{N}$ denote the random variable $\tau_{i,t}$ by

$$\tau_{i,t} := \min \left\{ r \in \mathbb{N} : \sum_{j=1}^r \mathbf{1} \{ f_j(\tilde{Z}_{j-1}) = i \} = t \right\}.$$

Lemma D.2. *It holds that*

1. *The sequences of random variables $(Y_{i,\tau_{i,t}})_{t \in \mathbb{N}}$ and $(Y_{i,t})_{t \in \mathbb{N}}$ have the same distribution for every $i \in \mathcal{I}$.*
2. *The sequences of random variables $(Y_{1,\tau_{1,t}})_{t \in \mathbb{N}}, \dots, (Y_{K,\tau_{K,t}})_{t \in \mathbb{N}}$ are mutually independent.*
3. *It holds that*

$$\mathbb{E} \max_{i=1}^K \left[\sum_{s=\lfloor n/K \rfloor}^R \|\hat{F}_{i,n,n} - F^i\|_\infty \mathbf{1} \{ S_i(n) = s \} \right] \leq A_{n,K} \wedge B_{n,K}.$$

4. *$1 - \mathbb{P}(\mathcal{G})$ is bounded from above by*

$$\frac{(1 + K\sqrt{16\pi/(\eta e)})^2}{(\sqrt{n}K)^{1+\eta/(2+\eta)}} \sum_{r \in \mathcal{R}} \frac{1}{r^{1+\eta/(2+\eta)}} \leq \frac{(1 + K\sqrt{16\pi/(\eta e)})^2}{(\sqrt{n}K)^{1+\eta/(2+\eta)}} \frac{2(1+\eta)}{\eta},$$

and also by

$$\frac{2K^{(1-\eta)/2}}{n^{(1+\eta)/4}} \sum_{r \in \mathcal{R}} \frac{1}{r^{(1+\eta)/2}}.$$

Proof. Parts 1 and 2 of the lemma follow from an optional sampling theorem. A suitable result in our context is Theorem 1 in Bélisle and Melfi (2008), cf. also their Section 5.3. The conditions in their Theorem 1 are easy to verify using the natural filtration generated by the i.i.d. sequence $(Y_t)_{t \in \mathbb{N}}$.

To prove Part 3, for every $s = 1, \dots, n$ set $\check{F}_{i,s}(\cdot) := \frac{1}{s} \sum_{r=1}^s \mathbf{1} \{ Y_{i,\tau_{i,r}} \leq \cdot \}$, and note that for ω such that $S_i(n) = s$ it holds that $\hat{F}_{i,n,n} = \check{F}_{i,s}$ by construction of f_t . Therefore, the left-hand side in the inequality in Part 3 coincides with

$$\mathbb{E} \max_{i=1}^K \left[\sum_{s=\lfloor n/K \rfloor}^R \|\check{F}_{i,s} - F^i\|_\infty \mathbf{1} \{ S_i(n) = s \} \right] \leq \mathbb{E} \max_{i=1}^K \max_{s \in \{\lfloor n/K \rfloor, \dots, R\}} \|\check{F}_{i,s} - F^i\|_\infty,$$

where the inequality follows from $\sum_{s=\lfloor n/K \rfloor}^R \mathbb{1}\{S_i(n) = s\} \leq 1$. Part 2 of this lemma implies that $\check{M}_i := \max_{s \in \{\lfloor n/K \rfloor, \dots, R\}} \|\check{F}_{i,s} - F^i\|_\infty$ for $i \in \mathcal{I}$ are independent random variables. Furthermore, by Part 1 of this lemma and since $(Y_{i,t})$ is an i.i.d. sequence, the random variables \check{M}_i for $i = 1, \dots, K$ and M_i for $i = 1, \dots, K$ have the same joint distribution for

$$M_i := \max_{s \in \{\lfloor n/K \rfloor, \dots, R\}} \|F_{i,s} - F^i\|_\infty, \quad \text{where} \quad F_{i,s}(\cdot) := \frac{1}{s} \sum_{r=1}^s \mathbb{1}\{Y_{i,i+(r-1)K} \leq \cdot\}.$$

We conclude that the random variables $\max_{i=1}^K \check{M}_i$ and $\max_{i=1}^K M_i$ have the same distribution and proceed by upper bounding $\mathbb{E} \max_{i=1}^K M_i$. To this end, note that by Lemma A.2 (applied with $k = e$) for any $x > 0$ the probability $\mathbb{P}(M_i > x)$ is not greater than

$$\mathbb{P}\left(\max_{s \in \{\lfloor n/K \rfloor, \dots, R\}} s \|F_{i,s} - F^i\|_\infty > \lfloor n/K \rfloor x\right) \leq \left(1 + 2\sqrt{2\pi}\right) \exp\left(-\frac{2x^2 \lfloor n/K \rfloor^2}{(1 + e^{-1})R}\right).$$

By definition $R = \underline{r} + \lceil \frac{n-rK}{2} \rceil$, thus

$$\frac{\lfloor n/K \rfloor^2}{R} \geq \frac{(\frac{n}{K} - 1)^2}{\underline{r} + \frac{n-rK}{2} + 1} = \frac{n^2(\frac{1}{K} - \frac{1}{n})^2}{n(\frac{1}{2} + \frac{\underline{r}}{n} - \frac{rK}{2n} + \frac{1}{n})} = \frac{n(1-a)^2}{f_{n,K}(\underline{r})K^2},$$

where $f_{n,K}(u) := \frac{1}{2} + \frac{u}{n} - \frac{uK}{2n} + \frac{1}{n}$ for $2 \leq K \leq n$, and $a = K/n$. Clearly $f_{n,K}(\cdot)$ is non-increasing, and $\frac{1}{K} \leq f_{n,K}(u) \leq 1$ for $u \in [1, \lfloor \frac{n}{K} \rfloor]$. Since $\underline{r} \in [1, \lfloor \frac{n}{K} \rfloor]$, we conclude

$$\mathbb{P}(M_i > x) \leq \left(1 + 2\sqrt{2\pi}\right) \exp\left(-\frac{2x^2 n(1-a)^2}{(1 + e^{-1})f_{n,K}(\underline{r})K^2}\right) \text{ for any } x > 0.$$

Lemma A.1 with $k = e$, $D = (1 + \sqrt{8\pi})$ and $\sigma^2 = \frac{(1+e^{-1})f_{n,K}(\underline{r})K^2}{4(1-a)^2n}$ yields

$$\mathbb{E}e^{tM_i} \leq \alpha_1 e^{\alpha_2 t^2} \text{ for any } t > 0, \quad (65)$$

with $\alpha_1 := \lceil 1 + D\sqrt{2\pi} \rceil$ and $\alpha_2 := (\frac{1}{2} + \frac{1}{2e})\sigma^2$. Hence, by Jensen's inequality

$$\exp\left(t\mathbb{E} \max_{i=1}^K M_i\right) \leq \sum_{i=1}^K \mathbb{E}e^{tM_i} \leq K\alpha_1 e^{\alpha_2 t^2}, \text{ for any } t > 0,$$

and we obtain

$$\mathbb{E} \max_{i=1}^K M_i \leq \frac{\log(K\alpha_1)}{t} + \alpha_2 t, \text{ for any } t > 0.$$

Setting $t = \sqrt{\log(K\alpha_1)/\alpha_2}$ now yields

$$\mathbb{E} \max_{i=1}^K M_i \leq 2\sqrt{\log(K\alpha_1)\alpha_2}, \quad (66)$$

which upon inserting α_1, α_2, D and σ yields

$$\mathbb{E} \max_{i=1}^K M_i \leq \frac{1+e}{\sqrt{2}e(1-a)} \sqrt{\log(K [1 + (1 + \sqrt{8\pi})\sqrt{2\pi}])} \frac{K}{\sqrt{n}} \sqrt{f_{n,K}(r)} = A_{n,K}.$$

To show $\mathbb{E} \max_{i=1}^K M_i \leq B_{n,K}$, note that for all $t > 0$, $\mathbb{E} e^{tM_i} \leq \sum_{s=\lfloor n/K \rfloor}^R \mathbb{E} e^{t\|F_{i,s} - F^i\|_\infty}$ which applying the bound in Equation (37) is further upper bounded by

$$\sum_{s=\lfloor n/K \rfloor}^R (1 + \sqrt{8\pi}) e^{(1+\frac{1}{e})\frac{1}{8s}t^2} \leq (R - \lfloor n/K \rfloor + 1) (1 + \sqrt{8\pi}) e^{(\frac{1}{8} + \frac{1}{8e})\frac{1}{\lfloor \frac{n}{K} \rfloor}t^2}.$$

Applying the argument that led from (65) to (66) with $\alpha_1 = (R - \lfloor n/K \rfloor + 1) (1 + \sqrt{8\pi})$ and $\alpha_2 = (\frac{1}{8} + \frac{1}{8e})\frac{1}{\lfloor \frac{n}{K} \rfloor}$ we obtain (using $\lfloor n/K \rfloor \geq (n/K) - 1$ that

$$\mathbb{E} \max_{i=1}^K M_i \leq 2\sqrt{\log(K (R - \lfloor n/K \rfloor + 1) (1 + \sqrt{8\pi}))} \left(\frac{1}{8} + \frac{1}{8e}\right)^{0.5} \frac{1}{\lfloor \frac{n}{K} \rfloor^{0.5}} \leq B_{n,K}.$$

For Part 4, pick a $\delta^* \in \arg \max_{\delta \in \mathcal{M}_K^n} \mathbb{T}(\langle \delta, \mathbf{F} \rangle)$. Let $i^* \in \mathcal{I}$ be the index such that $\delta^* \in \mathcal{M}_{A_{i^*}, K}^n$. For $\omega \in \Omega \setminus \mathcal{G}$ we certainly have $\mathcal{M}_{K, \bar{r}}^n \not\ni \delta^*$. Hence, for such ω , there exists an elimination round $r_*(\omega) \in \mathcal{R}$, say, where δ^* is *removed* from $\mathcal{M}_{A_{i^*}, K, r_*(\omega)}^n$. That is, there exists an index $l_*(\omega) \neq i^*$, say, a $\delta_*(\omega) \in \mathcal{M}_{A_{l_*(\omega)}, K, r_*(\omega)}^n$, and an index $t(\omega)$ such that

$$\mathbb{T}(\langle \delta_*(\omega), \hat{\mathbf{F}}_{t(\omega), n} \rangle) - \mathbb{T}(\langle \delta^*, \hat{\mathbf{F}}_{t(\omega), n} \rangle) > u_\eta(r_*(\omega), n).$$

Note that $\delta_*(\omega)$ places all its weight on indices in $A_{l_*(\omega)}$, and δ^* places all its weight on indices in A_{i^*} . Furthermore, by definition of the policy $\tilde{\pi}$, as long as a set $A \in \{A_1, \dots, A_m\}$ of treatments has not been *eliminated*, each treatment in that set is assigned once per round. Therefore, the cdfs in coordinates of $\hat{\mathbf{F}}_{t(\omega), n}$ with indices in $A_{l_*(\omega)} \cup A_{i^*}$ are all based on $r_*(\omega)$ observations. Thus, we may equivalently replace $\hat{\mathbf{F}}_{t(\omega), n}$ by $\check{\mathbf{F}}_{r_*(\omega)} = (\check{F}_{1, r_*(\omega)}, \dots, \check{F}_{m, r_*(\omega)})$ in the previous display, where $\check{F}_{i,s}(\cdot) := \frac{1}{s} \sum_{r=1}^s \mathbb{1}\{Y_{i, \tau_{i,r}} \leq \cdot\}$ was defined in the Proof of Part 3 above. Hence, we proceed with

$$\begin{aligned} \Omega \setminus \mathcal{G} &\subseteq \bigcup_{r \in \mathcal{R}} \bigcup_{l \neq i^*} \bigcup_{\delta_* \in \mathcal{M}_{A_l, K}^n} \left\{ \mathbb{T}(\langle \delta_*, \check{\mathbf{F}}_r \rangle) - \mathbb{T}(\langle \delta^*, \check{\mathbf{F}}_r \rangle) > u_\eta(r, n) \right\} \\ &\subseteq \bigcup_{r \in \mathcal{R}} \bigcup_{l \neq i^*} \bigcup_{\delta_* \in \mathcal{M}_{A_l, K}^n} \left\{ \mathbb{T}(\langle \delta_*, \check{\mathbf{F}}_r \rangle) - \mathbb{T}(\langle \delta_*, \mathbf{F} \rangle) + \mathbb{T}(\langle \delta^*, \mathbf{F} \rangle) - \mathbb{T}(\langle \delta^*, \check{\mathbf{F}}_r \rangle) > u_\eta(r, n) \right\} \\ &\subseteq \bigcup_{r \in \mathcal{R}} \bigcup_{l \neq i^*} \left\{ \max_{i \in A_l} \|\check{F}_{i,r} - F^i\|_\infty + \max_{i \in A_{i^*}} \|\check{F}_{i,r} - F^i\| > \frac{u_\eta(r, n)}{C} \right\} \\ &= \bigcup_{r \in \mathcal{R}} \left\{ \max_{i \in \{1, \dots, m\} \setminus A_{i^*}} \|\check{F}_{i,r} - F^i\|_\infty + \max_{i \in A_{i^*}} \|\check{F}_{i,r} - F^i\|_\infty > \frac{u_\eta(r, n)}{C} \right\}. \end{aligned}$$

Abbreviating $A_{i^*}^c = \{1, \dots, m\} \setminus A_{i^*}$, we thus have

$$\mathbb{P}(\Omega \setminus \mathcal{G}) \leq \sum_{r \in \mathcal{R}} \mathbb{P} \left(\max_{i \in A_{i^*}^c} \|\check{F}_{i,r} - F^i\|_\infty + \max_{i \in A_{i^*}} \|\check{F}_{i,r} - F^i\|_\infty > \frac{u_\eta(r, n)}{C} \right). \quad (67)$$

From Part 2 of this lemma we already know that, for every $r \in \mathbb{N}$, the random variables $\|\check{F}_{i,r} - F^i\|_\infty$ for $i \in \mathcal{I}$ are independent. Part 1 furthermore shows that the same statement holds if $\check{F}_{i,r}$ is replaced by $F_{i,r}$, and that (for every $r \in \mathbb{N}$ and every $i = 1, \dots, K$) the distributions of $\|\check{F}_{i,r} - F^i\|_\infty$ and $\|F_{i,r} - F^i\|_\infty$ coincide. Consequently, for every $r \in \mathbb{N}$ the random variables

$$\max_{i \in A_{i^*}^c} \|\check{F}_{i,r} - F^i\|_\infty + \max_{i \in A_{i^*}} \|\check{F}_{i,r} - F^i\|_\infty \quad \text{and} \quad \max_{i \in A_{i^*}^c} \|F_{i,r} - F^i\|_\infty + \max_{i \in A_{i^*}} \|F_{i,r} - F^i\|_\infty,$$

have the same distribution. Applying Corollary A.4 with $k = 2/\eta$ to each summand in the upper bound of (67), after replacing the cdfs $\check{F}_{i,r}$ by the cdfs $F_{i,r}$, we get

$$\mathbb{P}(\Omega \setminus \mathcal{G}) \leq \left(1 + K\sqrt{16\pi/(\eta e)}\right)^2 \sum_{r \in \mathcal{R}} e^{-\frac{u_\eta^2(r, n)r}{(1+\eta/2)C^2}},$$

which, using that

$$u_\eta(r, n) = C\sqrt{\frac{(1+\eta/2)(1+\eta/(2+\eta))}{r}}[0.5 \log(n) + \log(rK)],$$

coincides with

$$\frac{(1 + K\sqrt{16\pi/(\eta e)})^2}{(\sqrt{n}K)^{1+\eta/(2+\eta)}} \sum_{r \in \mathcal{R}} \frac{1}{r^{1+\eta/(2+\eta)}} \leq \frac{(1 + K\sqrt{16\pi/(\eta e)})^2}{(\sqrt{n}K)^{1+\eta/(2+\eta)}} \frac{2(1+\eta)}{\eta},$$

where we used $\sum_{r=1}^\infty \frac{1}{r^{1+b}} \leq 1 + \int_1^\infty \frac{1}{x^{1+b}} dx$ for all $b > 0$.

After replacing the cdfs $\check{F}_{i,r}$ by the cdfs $F_{i,r}$ in (67), we can use a union bound and the DKWM inequality to obtain

$$\begin{aligned} \mathbb{P}(\Omega \setminus \mathcal{G}) &\leq \sum_{r \in \mathcal{R}} \sum_{i \in A_{i^*}^c} \mathbb{P} \left(\|F_{i,r} - F^i\|_\infty > \frac{u_\eta(r, n)}{2C} \right) + \sum_{r \in \mathcal{R}} \sum_{i \in A_{i^*}} \mathbb{P} \left(\|F_{i,r} - F^i\|_\infty > \frac{u_\eta(r, n)}{2C} \right) \\ &\leq 2K \sum_{r \in \mathcal{R}} e^{-\frac{u_\eta^2(r, n)r}{2C^2}} = \frac{2K^{(1-\eta)/2}}{n^{(1+\eta)/4}} \sum_{r \in \mathcal{R}} \frac{1}{r^{(1+\eta)/2}}. \end{aligned}$$

This establishes Part 4. □

References

Athey, S. and S. Wager (2017). Efficient policy learning. *arXiv preprint arXiv:1702.02896*.

- Audibert, J.-Y., S. Bubeck, and R. Munos (2010). Best arm identification in multi-armed bandits. In *COLT-23th Conference on Learning Theory-2010*, pp. 13–p.
- Barrett, G. F. and S. G. Donald (2009). Statistical inference with generalized gini indices of inequality, poverty, and welfare. *Journal of Business & Economic Statistics* 27(1), 1–17.
- Bélisle, C. and V. Melfi (2008). Independence after adaptive allocation. *Statistics & Probability Letters* 78(3), 214–224.
- Bhattacharya, D. and P. Dupas (2012). Inferring welfare maximizing treatment assignment under budget constraints. *Journal of Econometrics* 167(1), 168–196.
- Biewen, M. (2002). Bootstrap inference for inequality, mobility and poverty measurement. *Journal of Econometrics* 108(2), 317–342.
- Bitler, M. P., J. B. Gelbach, and H. W. Hoynes (2006). What mean impacts miss: Distributional effects of welfare reform experiments. *American Economic Review* 96(4), 988–1012.
- Bomze, I. M., S. Gollowitzer, and E. A. Yildirim (2014). Rounding on the standard simplex: Regular grids for global optimization. *Journal of Global Optimization* 59(2-3), 243–258.
- Bubeck, S., R. Munos, and G. Stoltz (2009). Pure exploration in multi-armed bandits problems. In *International Conference on Algorithmic Learning Theory*, pp. 23–37. Springer.
- Cambini, A. and L. Martein (2009). *Generalized convexity and optimization*, Volume 616 of *Lecture Notes in Economics and Mathematical Systems*. Springer: Berlin.
- Carpentier, A. and A. Locatelli (2016). Tight (lower) bounds for the fixed budget best arm identification bandit problem. In *Conference on Learning Theory*, pp. 590–604.
- Cassel, A., S. Mannor, and A. Zeevi (2018). A general approach to multi-armed bandits under risk criteria. *arXiv preprint arXiv:1806.01380*.
- Davidson, R. and E. Flachaire (2007). Asymptotic and bootstrap inference for inequality and poverty measures. *Journal of Econometrics* 141(1), 141–166.
- Dehejia, R. H. (2005). Program evaluation as a decision problem. *Journal of Econometrics* 125(1), 141–173.

- Dufour, J.-M., E. Flachaire, and L. Khalaf (2019). Permutation tests for comparing inequality measures. *Journal of Business & Economic Statistics* 37(3), 457–470.
- Foster, J., J. Greer, and E. Thorbecke (1984). A class of decomposable poverty measures. *Econometrica* 52(3), 761–766.
- Hirano, K. and J. R. Porter (2009). Asymptotics for statistical treatment rules. *Econometrica* 77(5), 1683–1701.
- Jamieson, K., M. Malloy, R. Nowak, and S. Bubeck (2014). lil’ucb: An optimal exploration algorithm for multi-armed bandits. In *Conference on Learning Theory*, pp. 423–439.
- Karnin, Z., T. Koren, and O. Somekh (2013). Almost optimal exploration in multi-armed bandits. In *International Conference on Machine Learning*, pp. 1238–1246.
- Kasy, M. and A. Sautmann (2019). Adaptive treatment assignment in experiments for policy choice.
- Kaufmann, E., O. Cappé, and A. Garivier (2016). On the complexity of best-arm identification in multi-armed bandit models. *Journal of Machine Learning Research* 17(1), 1–42.
- Kitagawa, T. and A. Tetenov (2018). Who should be treated? Empirical welfare maximization methods for treatment choice. *Econometrica* 86(2), 591–616.
- Kitagawa, T. and A. Tetenov (2019). Equality-minded treatment choice. *Journal of Business & Economic Statistics*, 1–14.
- Kock, A. B., D. Preinerstorfer, and B. Veliyev (2020a). Functional sequential treatment allocation. *arXiv preprint arXiv:1812.09408*.
- Kock, A. B., D. Preinerstorfer, and B. Veliyev (2020b). Functional sequential treatment allocation with covariates. *arXiv preprint arXiv:2001.10996*.
- Kock, A. B. and M. Thyrgaard (2017). Optimal sequential treatment allocation. *arXiv preprint arXiv:1705.09952*.
- Manski, C. F. (2004). Statistical treatment rules for heterogeneous populations. *Econometrica* 72(4), 1221–1246.
- Manski, C. F. and A. Tetenov (2016). Sufficient trial size to inform clinical practice. *Proceedings of the National Academy of Sciences* 113(38), 10518–10523.

- Rothe, C. (2010). Nonparametric estimation of distributional policy effects. *Journal of 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.
- Shorack, G. R. and J. A. Wellner (2009). *Empirical processes with applications to statistics*. SIAM.
- Stoye, J. (2009). Minimax regret treatment choice with finite samples. *Journal of Econometrics* 151(1), 70–81.
- Stoye, J. (2010). Partial identification of spread parameters. *Quantitative Economics* 1(2), 323–357.
- Stoye, J. (2012). Minimax regret treatment choice with covariates or with limited validity of experiments. *Journal of Econometrics* 166(1), 138–156.
- Tran-Thanh, L. and J. Y. Yu (2014). Functional bandits. *arXiv preprint arXiv:1405.2432*.
- Tsybakov, A. B. (2009). *Introduction to Nonparametric Estimation*. New York: Springer.