



Regularizing fairness in optimal policy learning with distributional targets[★]

Anders Bredahl Kock^a, David Preinerstorfer^{b,*}

^a Department of Economics, University of Oxford and St. Hilda's College, 10 Manor Rd, OX1 3UQ, Oxford, UK

^b Institute for Statistics and Mathematics, WU Vienna University of Economics and Business, Welthandelsplatz 1, 1020, Vienna, Austria

ARTICLE INFO

JEL classification:

C18
C21
C44

Keywords:

Treatment allocation
Pure exploration
Unfairness penalties

ABSTRACT

A decision maker typically (i) incorporates training data to learn about the relative effectiveness of treatments, and (ii) chooses an implementation mechanism that implies an “optimal” predicted outcome distribution according to some target functional. Nevertheless, a fairness-aware decision maker may not be satisfied achieving said optimality at the cost of being “unfair” against a subgroup of the population, in the sense that the outcome distribution in that subgroup deviates too strongly from the overall optimal outcome distribution. We study a framework that allows the decision maker to regularize such deviations, while allowing for a wide range of target functionals and fairness measures to be employed. We establish regret and consistency guarantees for empirical success policies with (possibly) data-driven preference parameters, and provide numerical results. Furthermore, we briefly illustrate the methods in two empirical settings.

1. Introduction

Optimal policy choice, i.e., inferring good answers to the question “whom to assign to which treatment” from data, is of fundamental importance to decision makers. Much progress has been achieved in the recent literature: seminal contributions (emphasizing the nonparametric nature of the problem) were made in Manski (2004), Dehejia (2005), Stoye (2009, 2012), Bhattacharya and Dupas (2012), Manski and Tetenov (2016), a study of asymptotic efficiency properties was provided in Hirano and Porter (2009), whereas more recent contributions focused on optimal policy choice from observational data and include Kitagawa and Tetenov (2018, 2021), Athey and Wager (2021), or Mbakop and Tabord-Meehan (2021); see also the recent overview given in Hirano and Porter (2020) for further references.

A number of recent advances in that literature, in particular Kleinberg et al. (2018), cf. also the recent review of Cowgill and Tucker (2019), stress that in addition to assigning individuals with the goal to maximize expected welfare, decision makers should incorporate “fairness” considerations into optimal policy choice. The problem of a fairness-aware decision maker that we consider in the present article is the following: although the decision rule inferred, e.g., by applying one of the approaches suggested in the literature discussed above, may be optimal on an *aggregate* level, there may be *subgroups* of the population defined by protected characteristics, e.g., gender, ethnicity, or religion, in which the effect of the decision rule deviates markedly from that on the aggregate level. In

* We thank Michael Lechner, Liyang Sun, Alex Tetenov, Alonso Zuniga, and several colleagues and seminar/workshop participants at WU Vienna, TU Graz, University of Vienna, University of Geneva, UPF Barcelona, and Erasmus University Rotterdam for helpful feedback and suggestions. We are grateful to an Associate Editor and two referees for their comments, which helped us to improve the paper.

* Corresponding author.

E-mail addresses: anders.kock@economics.ox.ac.uk (A.B. Kock), david.preinerstorfer@wu.ac.at (D. Preinerstorfer).

other words, although the decision maker may have succeeded in choosing a policy that leads to optimal results on the population level, this global optimality is not resembled within all subgroups and hence the policy may be considered unfair.

This concern fits into the broader context of fair Machine Learning or fair Artificial Intelligence, where the general goal is to avoid that algorithms inferred from data discriminate against certain groups; cf. [Corbett-Davies and Goel \(2018\)](#) and [Pessach and Shmueli \(2022\)](#) for recent overviews on the subject. They identify the following key features: (i) incorporating fairness is complicated by the lack of a single accepted definition of fairness (cf., e.g., Section 3 and Appendix B.1 in [Pessach and Shmueli, 2022](#)), (ii) impossibility results show that different types of fairness cannot be achieved at the same time, and (iii) achieving (a particular type of) fairness typically reduces the efficiency of the policy.

Fairness concerns can be incorporated into the optimal policy choice algorithms studied in [Kitagawa and Tetenov \(2018\)](#) or [Athey and Wager \(2021\)](#) by adjusting subpopulation weights in the objective function. Nevertheless, such weights need to be chosen by the decision maker upfront, and are typically difficult to defend. Motivated by the lack of methods allowing decision makers to incorporate more general forms of fairness considerations, [Viviano and Bradic \(2023\)](#) studied methods that combine fairness and optimal policy choice in a more distinguished way: instead of trying to infer a policy that maximizes expected population welfare, they formulate a multiobjective optimization problem, in which each objective corresponds to the welfare in the subpopulation obtained by considering a given value of a protected characteristic. After characterizing the Pareto frontier, the authors suggest to choose an element from the frontier which maximizes a fairness criterion, which is allowed to be very general, and can be chosen by the user in dependence on the notion of fairness that fits the specific problem at hand. Statistical performance guarantees are provided. Another recent contribution is [Fang et al. \(2023\)](#), who do not consider fairness between subgroups defined by a sensitive characteristic, but aim at lifting the lower tail of the outcome distribution implied by a policy to a given level, therefore leading to “fairness” between rich and poor individuals. To achieve this, in addition to maximizing expected welfare, they introduce a lower-bound constraint on a quantile chosen by the user, and suggest an algorithm to solve the resulting optimization problem, for which asymptotic guarantees are established. Both articles discuss observational as well as experimental data.

Although decision makers can draw on these methods in situations where expected utility is the main target, there is currently no theory available that shows how one could incorporate fairness considerations in settings with more flexible target measures. Both [Viviano and Bradic \(2023\)](#) and [Fang et al. \(2023\)](#) target expected outcomes, and exploit linearity properties of that target in their constructions and proofs. For example, the characterization of the Pareto frontier given in Lemma 2.1 of [Viviano and Bradic \(2023\)](#), on which their approach crucially rests, exploits linearity – or at least convexity/concavity – of the target measure.¹ Important alternative targets in economic decision making are welfare measures other than expected utility, or measures incorporating poverty concerns (cf. [Cowell, 2011](#) or [Chakravarty, 2009](#) for book-length treatments of measures that have been developed and studied in the economics and statistics literature). The importance of flexibility in choosing the target measure was emphasized in, e.g., [Bitler et al. \(2006\)](#) and [Rostek \(2010\)](#).

The goal of the present article is to develop a flexible theory of optimal treatment assignment that allows the decision maker to incorporate fairness considerations in a setting where the target is not necessarily the expected outcome, but, e.g., a welfare measure such as the Gini-welfare or a quantile such as the median. More specifically, we develop a theory that allows the decision maker to work with the rich class of targets introduced in [Kock et al. \(2022\)](#), while incorporating covariate information and monitoring fairness properties of the policy recommendation. To arrive at a widely applicable theory concerning the target functional, we do not take a Pareto frontier based approach as in [Viviano and Bradic \(2023\)](#) or a constraint approach as in [Fang et al. \(2023\)](#) (which requires the somewhat inconvenient assumption that the constraint is actually feasible, which is not necessary in our approach), but we proceed as discussed informally in [Section 1.1](#) below, and refer the reader to later sections for more details.

In [Appendix A](#) we discuss in more detail how the present article relates to [Kitagawa and Tetenov \(2021\)](#) and [Kock et al. \(2023\)](#), where optimal policy choice problems with general targets are treated (without fairness considerations).²

1.1. Informal summary of our methods and results, and practical guidelines

1.1.1. Objective

In [Section 2](#), we first decompose $\langle \delta, \mathcal{F} \rangle$, the (unknown) population distribution implied by a decision rule δ , as

$$\langle \delta, \mathcal{F} \rangle = \sum_{z \in \mathcal{Z}} p_Z(z) \langle \delta, \mathcal{F} \rangle_z,$$

where $\langle \delta, \mathcal{F} \rangle_z$ denotes the (unknown) distribution implied by the decision rule δ in the subpopulation of subjects with protected characteristic equal to $z \in \mathcal{Z}$, and where $p_Z(z)$ denotes the (unknown) proportion of such subjects. While the primary goal of the decision maker is to choose δ such as to maximize $T(\langle \delta, \mathcal{F} \rangle)$, where T is some (potentially nonlinear and nonconvex) functional, the secondary goal of the decision maker is to monitor fairness properties of the policy.³ Specifically, we consider a situation where, as the

¹ It needs to be emphasized again that (in contrast to their mean target, which is fixed) the *fairness* measure used by [Viviano and Bradic \(2023\)](#) is very general and allows for great flexibility.

² The class of targets allowed for in [Kock et al. \(2023\)](#) is the one established in [Kock et al. \(2022\)](#) (cf. in particular their Appendices F and G, and covers many popular welfare, inequality, and poverty indices. It also covers targets such as quantiles, U-statistics, L-statistics, and more. Recently, there has been an increased interest in quantile targets in optimal policy choice, cf., e.g., the work of [Chen et al. \(2025\)](#) and [Guggenberger et al. \(2026\)](#).

³ Because the decision maker does not know $\langle \delta, \mathcal{F} \rangle$ or $\langle \delta, \mathcal{F} \rangle_z$, the decision maker has to estimate these quantities, which introduces a nontrivial statistical aspect that needs to be addressed.

secondary goal, the decision maker finds *parity* across sub-populations (aka groups) defined by a protected characteristic desirable. That is, the decision maker wants to choose δ such that

$$\langle \delta, \mathcal{F} \rangle_z \approx \langle \delta, \mathcal{F} \rangle \quad \text{for every } z \in \mathcal{Z},$$

in the sense that

$$\max_{z \in \mathcal{Z}} S(\langle \delta, \mathcal{F} \rangle_z, \langle \delta, \mathcal{F} \rangle) \approx 0,$$

where $S \geq 0$ is a semi-metric on the set of distribution functions, e.g., the Kolmogorov-Smirnov distance, allowing for some flexibility. If δ satisfies such a parity condition, this means that the outcome distributions implied by the decision rule δ do not differ much across the groups defined by the protected characteristic (with respect to the properties encoded in the fairness measure S).

To combine the two goals outlined above, i.e., maximization of T evaluated at the implied outcome distribution (δ, \mathcal{F}) and achieving some degree of parity formalized via S , we adapt a penalty-based approach. This idea has repeatedly been applied to incorporate fairness considerations into regression and classification problems, cf., e.g., Kamishima et al. (2011, 2012), Berk et al. (2017), or Bechavod and Ligett (2017). Specifically, we aim at studying statistical methods to choose a decision rule δ that maximizes the penalized objective

$$(1 - \lambda)T(\langle \delta, \mathcal{F} \rangle) - \lambda \max_{z \in \mathcal{Z}} S(\langle \delta, \mathcal{F} \rangle_z, \langle \delta, \mathcal{F} \rangle), \quad (1)$$

where $\lambda \in [0, 1]$ is a preference parameter, which has to be chosen by the decision maker (possibly in a data-driven way). Note that the larger λ , the more emphasis is put on fairness.

1.1.2. Policies

The policies we study are all of the *empirical success* type: In the case of discrete covariates, which is our main focus, we study a plug-in approach, where we first estimate \mathcal{F} through $\hat{\mathcal{F}}$, say, and then solve (numerically) the empirical version of the penalized objective in (1), cf. Section 3 for a detailed description of the policy and its regret and consistency properties. We also discuss variants of our policies when the preference parameter λ in (1) is chosen in a data-driven way. To this end, the decision maker (DM) may obtain optimal policies for a set of preference parameters $\Lambda \subseteq [0, 1]$ first (making use of the just-mentioned empirical success type policy), and then select one of these policies by balancing the fairness/efficiency trade-off of the policies. This can be done informally, e.g., by investigating plots of the estimated fairness/efficiency properties of the policies, or more formally, e.g., by fixing a “budget” the DM can afford spending on fairness (cf. Berk et al., 2017), and then selecting the policy which realizes the maximal degree of fairness affordable. The latter approach is described and studied in detail in Section 4. The former, more graphical approach, is illustrated in the applications in Section 7 in the context of the Pennsylvania reemployment bonus experiment (cf. Biliias, 2000) and the entrepreneurship program data from Lyons and Zhang (2017), which was also investigated in Viviano and Bradic (2023).

The case of non-discrete covariates is studied in Section 5, cf. Section 5.1 for a detailed description of the policies and a regret upper bound. In this context, we do not work with a plug-in for $\hat{\mathcal{F}}$ but estimate the objective function (introduced earlier) more directly, an approach that may be viewed as an extension of Kitagawa and Tetenov (2021) to our setting.

1.1.3. Practical guideline

In a concrete application, one can proceed as follows:

1. First, one has to fix a target functional T and a measure of unfairness S . Canonical choices are, e.g., the Gini-welfare measure or the expectation for the former and the Kolmogorov-Smirnov distance (possibly one-sided) for the latter. Many relevant target functionals are described in Chakravarty (2009) or Cowell (2011), to which we refer the practitioner who is unsure about which target functional fits the concrete application best. Once this choice has been made, the next step depends on the nature of the covariates.
2. If the covariates are discrete, one can use the policy in Section 3, otherwise one uses the policy in Section 5.1. How this is precisely carried out further depends on whether a concrete value for λ is available or not:
 - If a concrete value for the preference parameter λ is known, approximately solve the objective in Eq. (19) (for the discrete case) or in Eq. (41) (otherwise) with a numerical optimization procedure, and roll out the policy obtained by random assignment; e.g., by making use of the concrete assignment mechanism described in Footnote 6.
 - If λ is not specified, we recommend to solve the respective objective for a grid of λ values, and use graphical inspection and/or the procedure in Section 4 to select a suitable value of λ ; cf. the empirical examples in Section 7 for some guidance. Then, roll out the policy by random assignment (e.g., by making use of the concrete assignment mechanism described in Footnote 6).

2. Setting

2.1. Observational structure and two assumptions

We consider a DM, who must decide at which proportions $K \geq 2$ available treatments should be rolled out to a heterogeneous population. Before the “roll-out phase” begins, the DM needs to find out which proportions lead to the “optimal” (this will be made precise below) outcome distribution in the population. To make this decision, we assume that the DM is in possession of a training sample of n subjects that:

1. were independently drawn from the target population,
2. were assigned to one of the K treatments, possibly by making use of covariate information in the assignment decision,
3. and were subjected to an outcome measurement.

We adopt the potential outcomes framework, where to each subject $j = 1, \dots, n$, there corresponds a potential outcomes vector

$$Y_j = (Y_{1,j}, \dots, Y_{K,j})' \in \mathbb{R}^K.$$

That is, $Y_{i,j}$ is the potential outcome of subject j under treatment i . As usual, for each subject $j = 1, \dots, n$ in the training sample, the decision maker does not observe the whole potential outcome vector Y_j , but only observes the single coordinate

$$\tilde{Y}_j := Y_{D_j,j},$$

where $D_j \in \{1, \dots, K\}$ denotes the treatment subject j was assigned to, which the DM observes as well. In addition to \tilde{Y}_j and D_j , the DM observes two *categorical* covariates

$$X_j \in \mathcal{X}, \quad 0 < |\mathcal{X}| < \infty, \quad \text{and} \quad Z_j \in \mathcal{Z}, \quad 0 < |\mathcal{Z}| < \infty.$$

For now we focus on the case where both covariates X_j and Z_j can take on only finitely many values, because this simplifies the analysis and exposition considerably, and because this is an important case for practice. In [Section 5](#) we consider the case where \mathcal{X} is not finite.

Remark 1. At this stage, the reader may think of X_j as a covariate that can be used to assign individuals to treatments during the roll-out phase. That is, the proportions used during the roll-out phase can depend on these covariates, in the sense that different proportions are rolled out for different values of $x \in \mathcal{X}$. The covariate Z_j , on the other hand, should be thought of as a “protected characteristic” that *cannot* be used for assigning individuals in the roll-out phase, e.g., because it is simply not possible to measure this information during the roll-out phase or because it is not allowed to make different assignment decisions for different values of the protected characteristic.⁴ Nevertheless, this covariate is assumed to be available for all subjects in the training sample of the DM. [Tan et al. \(2022\)](#), who consider robust decision learning targeting the expected outcome, use an analogous setting concerning the treatment of sensitive characteristics.

Summarizing, each subject is characterized by an underlying (only partly observed) random vector

$$W_j := (Y_j', X_j, Z_j, D_j)' \in \mathbb{R}^K \times \mathcal{X} \times \mathcal{Z} \times \{1, \dots, K\}.$$

The *training sample* available to the DM is (a realization of)

$$\tilde{W}_j := (\tilde{Y}_j, X_j, Z_j, D_j)' \in \mathbb{R} \times \mathcal{X} \times \mathcal{Z} \times \{1, \dots, K\} \quad \text{for } j = 1, \dots, n. \tag{2}$$

All random variables are defined on an underlying probability space $(\Omega, \mathcal{A}, \mathbb{P})$ with corresponding expectation operator \mathbb{E} ; outer probability and outer expectation are denoted by \mathbb{P}^* and \mathbb{E}^* , respectively, which will help us to avoid imposing measurability conditions.

Without further mentioning in the theoretical results, we shall *throughout this article impose* an independent sampling condition:

Assumption 1. The random vectors $W_j, j = 1, \dots, n$, are i.i.d.

The independent sampling assumption is standard in the treatment assignment literature, and weakening it is not an objective of the present article. Also without further mentioning in the theoretical results, we *impose throughout this article* that the assignment mechanism is unconfounded conditional on the covariates observed, that is:

Assumption 2. Conditional on X_j and Z_j the potential outcome vector Y_j is independent of D_j ; i.e.,

$$Y_j \perp\!\!\!\perp D_j \mid (X_j, Z_j).$$

2.2. Decision rules and policies

Based on the training sample in (2), the DM needs to specify how to proceed in the roll-out phase. The assignment decision can only be based on the observed value of the unprotected characteristic $x \in \mathcal{X}$, cf. [Remark 1](#). Formally, and allowing for randomized decision rules, for each $x \in \mathcal{X}$ the DM therefore has to decide (based on the training sample) on the proportions at which to roll out the treatments $i = 1, \dots, K$. A *decision rule* thus specifies, for each covariate value $x \in \mathcal{X}$, a “probability vector”

$$\delta(x) = (\delta_1(x), \dots, \delta_K(x))' \in \mathcal{S}_K. \tag{3}$$

⁴ Note that the variables X_j and Z_j may be dependent, and that a decision depending on the former characteristic only may therefore also depend on Z_j . In this sense, the decision rules we allow for are certainly not “independent” of the information inherent in the protected characteristic in general. If such an “independence” condition is required, a different approach that tries to eliminate this dependence has to be taken, which is not the goal in the present article. In the terminology of [Frauen et al. \(2023\)](#), who consider optimal policy choice with fairness considerations for an expected utility target, our focus is oriented towards “value fairness” whereas the fairness notion discussed in this footnote corresponds to “action fairness”.

where $\mathcal{S}_K = \{\alpha = (\alpha_1, \dots, \alpha_K)' \in [0, 1]^K : \alpha_1 + \dots + \alpha_K = 1\}$ denotes the standard simplex in \mathbb{R}^K . This class of decision rules allows for (but does not impose) randomization.

The i th coordinate $\delta_i(x)$ of $\delta(x)$ corresponds to the probability of assigning a subject with covariate x to treatment i during the roll-out phase. The dependence of the proportions on x permits the DM to take advantage of the population’s heterogeneity. Note again (cf. Remark 1) that δ does not have z as an argument.

Every data-driven method that turns the training sample defined in Eq. (2) into a decision rule as just described shall be referred to as a *policy*; i.e., a policy is a function

$$\pi_n : \mathbb{R}^n \times \mathcal{X}^n \times \mathcal{Z}^n \times \{1, \dots, K\}^n \rightarrow \mathbb{S}_K, \tag{4}$$

where $\mathbb{S}_K = \mathcal{S}_K^{|\mathcal{X}|}$ denotes the set of decision rules, i.e., all functions from \mathcal{X} to the simplex \mathcal{S}_K . Throughout, we shall equip \mathbb{S}_K with the distance d_1 , which, for δ and $\tilde{\delta}$ in \mathbb{S}_K , is defined as

$$d_1(\delta, \tilde{\delta}) := \sum_{x \in \mathcal{X}} \|\delta(x) - \tilde{\delta}(x)\|_1, \tag{5}$$

where, for $p \geq 1$, the symbol $\|\cdot\|_p$ denotes the p -norm on a Euclidean space.

Remark 2. As in Kock et al. (2023), we could also consider policies that are restricted to lie in a non-empty and closed subset of \mathbb{S}_K (corresponding to budget, similarity, or compatibility constraints). The theory and algorithms developed in Sections 2 and 3 would go through upon fairly straightforward modification of the statements and proofs. To avoid notation overload, however, we do not pursue this level of generality here. In Section 5, where we consider non-finitely valued covariates, restrictions on the set of decision rules will be imposed.

2.3. Goal of the DM

We have not yet formally defined the actual *goal* of the DM. In order to decide which decision rules δ are preferable in a specific context, we

1. need to investigate the structure of the population cdf (cumulative distribution function) implied by rolling out a decision rule δ , and
2. single out a criterion on the basis of which to compare different population cdfs.

This is addressed in the following subsections: After introducing some notation in Section 2.3.1, we address the first item in the above enumeration in Section 2.3.2 and the second in Section 2.3.3. An illustrative toy example is studied in Section 6.1.

2.3.1. Some notation

We denote the joint probability mass function of (X_j, Z_j, D_j) by $p_{X,Z,D}(x, z, i) = \mathbb{P}(X_j = x, Z_j = z, D_j = i)$, write $p_{X,Z}$, p_X , and p_Z for the corresponding marginal probability mass functions of (X_j, Z_j) , X_j and Z_j , respectively, and set

$$p_{X|Z=z}(x) := \frac{p_{X,Z}(x, z)}{p_Z(z)} \quad \text{and} \quad p_{Z|X=x}(z) := \frac{p_{X,Z}(x, z)}{p_X(x)}. \tag{6}$$

On top of Assumptions 1 and 2, we shall assume (until Section 5) that the common support of (X_j, Z_j, D_j) is $\mathcal{X} \times \mathcal{Z} \times \{1, \dots, K\}$ so that, in particular, the quotients in (6) are well defined:

Assumption 3. We have $p_{X,Z,D}(x, z, i) > 0$ for every $x \in \mathcal{X}$, $z \in \mathcal{Z}$, and $i \in \{1, \dots, K\}$.

For $x \in \mathcal{X}$ and $z \in \mathcal{Z}$ we denote by $F^i(\cdot | x, z)$ the cumulative distribution function (cdf) of

$$Y_{i,j} | X_j = x, Z_j = z.$$

Furthermore, we denote by $F^i(\cdot | x)$ the cdf of $Y_{i,j} | X_j = x$, and by F^i the cdf of $Y_{i,j}$. We write $F = (F^1, \dots, F^K)'$, and analogously define for every x and z the vectors of conditional cdfs $F(\cdot | x)$, and $F(\cdot | x, z)$. For notational convenience, we shall collect the conditional cdfs of a random vector Y'_j given $X_j = x$ and $Z_j = z$ and the corresponding probability mass functions in an array

$$F := [F^i(\cdot | x, z), p_{X,Z}(x, z) : i = 1, \dots, K; x \in \mathcal{X}; z \in \mathcal{Z}]. \tag{7}$$

Given F as in (7), i.e., an array consisting of cdfs and a probability mass function over $\mathcal{X} \times \mathcal{Z}$, we shall sometimes write

$$(Y'_j, X_j, Z_j)' \sim F$$

to signify that F is the array corresponding⁵ to the distribution of $(Y'_j, X_j, Z_j)'$. Generically, we shall throughout denote by F the array corresponding to the distribution of $(Y'_j, X_j, Z_j)'$, which is *unknown* to the DM.

⁵ Note that F corresponds to but does not specify completely the joint distribution of $(Y'_j, X_j, Z_j)'$, because it does not specify the dependence between the potential outcomes, an aspect that is irrelevant for our results.

2.3.2. Distributions generated by rolling out a decision rule

We now consider the effect of rolling out a (fixed) decision rule $\delta \in \mathbb{S}_K$ to the whole population via randomized assignment. More specifically, consider a new subject independently drawn from the same population as the training sample in the sense that

$$(Y', X, Z)' \sim \mathcal{F}, \tag{8}$$

and assume that its assignment D (chosen by the DM), say, satisfies:

1. conditionally on X , it is distributed according to $\delta(X)$, and
2. conditionally on (X, Z) , it is independent of every outcome Y_i for $i = 1, \dots, K$;

which, for every $x \in \mathcal{X}, y \in \mathbb{R}, z \in \mathcal{Z}$ and treatment i , we summarize in the property⁶

$$P(D = i, Y_i \leq y \mid X = x, Z = z) = \delta_i(x)P(Y_i \leq y \mid X = x, Z = z). \tag{9}$$

Denoting Y_D by \tilde{Y} , it is elementary to verify that the cdf of \tilde{Y} then equals

$$\mathbb{P}(\tilde{Y} \leq \cdot) = \sum_{x \in \mathcal{X}} \sum_{i=1}^K \delta_i(x) p_X(x) F^i(\cdot \mid x) =: \langle \delta, \mathcal{F} \rangle. \tag{10}$$

For every $z \in \mathcal{Z}$, we can analogously consider the cdf of \tilde{Y} conditional on $Z = z$, i.e., the cdf in the *subpopulation* of subjects with protected characteristic equal to z , which is given by

$$\mathbb{P}(\tilde{Y} \leq \cdot \mid Z = z) = \sum_{x \in \mathcal{X}} \sum_{i=1}^K \delta_i(x) p_{X \mid Z=z}(x) F^i(\cdot \mid x, z) =: \langle \delta, \mathcal{F} \rangle_z, \tag{11}$$

and for which we thus have $\langle \delta, \mathcal{F} \rangle = \sum_{z \in \mathcal{Z}} p_Z(z) \langle \delta, \mathcal{F} \rangle_z$.

2.3.3. Objective: Target functional penalized for unfairness

In the previous subsection we illustrated how the choice of δ implies a population cdf (10), which is further decomposed as a weighted average of subpopulation cdfs (11). Which types of population cdfs are desirable to the DM depends strongly on the context. Typically, “desirability” will be measured through a user-specified real-valued functional T , say, of the cdf generated; such as its mean, a quantile such as the median, or a trimmed mean, possibly augmented by a poverty or inequality measure.

Given the functional T —the concrete form of which we leave unspecified for now to ease the presentation—we consider the problem where the DM wants to trade off two potentially conflicting objectives described next:

1. The DM wants to choose $\delta \in \mathbb{S}_K$ to maximize $T(\langle \delta, \mathcal{F} \rangle)$, i.e., the functional T evaluated at the outcome distribution $\langle \delta, \mathcal{F} \rangle$.
2. The DM wants to roll out a δ that is fair in the sense that the subpopulation cdfs $\langle \delta, \mathcal{F} \rangle_z$ are all “close” to the population cdf $\langle \delta, \mathcal{F} \rangle$ singled out by the first objective. If this condition is violated, there are subpopulations defined by the protected characteristic, the cdf of which substantially deviates from the “global” outcome cdf $\langle \delta, \mathcal{F} \rangle$, which the DM wants to avoid.

While the first objective is governed by a functional T , we need to make the second part of the DM’s objective more precise. To this end, because the similarity of cdfs can be measured in different ways, we introduce another functional S that:

- (a) maps pairs of cdfs to $[0, \infty)$,
- (b) equals 0 if the two input cdfs coincide,
- (c) is intended to measure the “similarity” between two cdfs, high values corresponding to a strong degree of dissimilarity.

For concreteness, one may think of S as the Kolmogorov-Smirnov distance (or more generally, any pseudo-distance) on the set of cdfs, or the absolute difference in the functional T evaluated at two input cdfs. One could also introduce some “directionality effects” by working, e.g., with a one-sided Kolmogorov-Smirnov pseudo-distance (for cdfs F and G consider $S(F, G) = \sup_{x \in \mathbb{R}} \max(F(x) - G(x), 0)$ instead of $S(F, G) = \sup_{x \in \mathbb{R}} |F(x) - G(x)|$). These are important examples, but the theory developed below applies more generally.

The remaining question now is how the DM could balance the two optimization objectives. In the present article, we consider rules $\delta^\lambda \in \mathbb{S}_K$, say, that maximize, for a chosen penalty $\lambda \in [0, 1]$, the *penalized* objective

$$\Omega_{\lambda, \mathcal{F}}(\delta) := (1 - \lambda)T(\langle \delta, \mathcal{F} \rangle) - \lambda \max_{z \in \mathcal{Z}} S(\langle \delta, \mathcal{F} \rangle_z, \langle \delta, \mathcal{F} \rangle). \tag{12}$$

The objective $\Omega_{\lambda, \mathcal{F}}(\delta)$ in (12) is composed of two terms: The first summand depends on the functional T evaluated at the cdf generated by the decision rule $\delta \in \mathbb{S}_K$. The DM’s primary goal is to maximize this quantity. The second term adds the maximal dissimilarity between a subpopulation cdf $\langle \delta, \mathcal{F} \rangle_z$ and the population cdf $\langle \delta, \mathcal{F} \rangle$. Note that if S is chosen suitably (e.g., when S is a pseudo-metric on the set of cdfs), and

$$\max_{z \in \mathcal{Z}} S(\langle \delta, \mathcal{F} \rangle_z, \langle \delta, \mathcal{F} \rangle)$$

is large, this corresponds to a situation where the cdf generated by rolling out δ leads to very different cdfs between subpopulations defined by the protected characteristics and the population cdf. In this sense, rolling out the treatments according to δ may be considered unfair.

⁶ For concreteness, let U be a uniformly distributed random variable on $(0, 1)$ that is independent of the training sample and of (Y', X, Z) . Set $\gamma_0(x) := 0$ and $\gamma_i(x) = \sum_{j=1}^i \delta_j(x)$ for $i = 1, \dots, K$ and $x \in \mathcal{X}$, and assign the new subject to treatment i if $U \in [\gamma_{i-1}(X), \gamma_i(X))$. Then (9) is satisfied (note that it also remains satisfied conditionally on the training sample in case δ depends on the training sample).

Remark 3. Note that we have chosen the penalty term $\max_{z \in \mathcal{Z}} S(\langle \delta, F \rangle_z, \langle \delta, F \rangle)$ not to incorporate the “sizes” $p_Z(z)$ of the protected groups. Alternatively, one could also have used a penalty term of the form $\max_{z \in \mathcal{Z}} p_Z(z) S(\langle \delta, F \rangle_z, \langle \delta, F \rangle)$, which, however, has the disadvantage of down-weighting small marginalized groups.

The degree at which unfairness is penalized depends on the *preference parameter* λ , which has to be chosen by the DM. Choosing $\lambda \approx 0$ corresponds to a situation where fairness is essentially not taken into account, and where the DM is mostly interested in determining which treatment combination maximizes T . On the other hand, choosing $\lambda \approx 1$ corresponds to a situation where unfairness is severely penalized.

Remark 4. Typically, the DM will determine solutions to (suitable approximations of) the optimization problem (13) for different values of λ , and will then determine the decision rule to eventually roll out by comparing aspects of the corresponding path of solutions δ^λ . That is, the DM will study aspects of (estimates of) the functions

$$\lambda \mapsto T(\langle \delta^\lambda, F \rangle) \quad \text{and} \quad \lambda \mapsto S(\langle \delta^\lambda, F \rangle_z, \langle \delta^\lambda, F \rangle),$$

cf. Section 7 for some illustration in the context of two empirical examples. Based on these functions, the DM can decide which λ to choose and which corresponding δ^λ to roll out. The decision rule a DM applies proceeding in this way incorporates a *data-driven* choice of λ . Because this choice may be highly complicated or subjective, we shall establish regret guarantees *irrespective* of what the data-driven preference parameter choice of the DM is.

Remark 5. Alternatively to the objective in (12), one could consider an objective function that maximizes $T(\langle \delta, F \rangle)$ subject to the constraint that δ satisfies

$$\max_{z \in \mathcal{Z}} S(\langle \delta, F \rangle_z, \langle \delta, F \rangle) \leq \varepsilon.$$

As mentioned in the introduction, an advantage of working with the penalized criterion (12) over this constrained maximization problem is that a solution to the latter may fail to exist for a given ε , because no δ may satisfy the restriction.

Proposition 1 below shows that under weak assumptions the optimization problem

$$\max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, F}(\delta) \tag{13}$$

permits a (possibly non-unique) solution for every $\lambda \in [0, 1]$. Note that the optimization problem in (13) cannot directly be solved by the DM, because F is unknown. Therefore, the DM first has to estimate F and will then solve an approximation to that optimization problem. We address this aspect further below.

A detailed example that illustrates the concepts introduced so far can be found in Section 6.

2.4. Main technical assumption concerning the functionals T and S and some preliminary observations

To establish theoretical guarantees, we impose throughout the assumption that all cdfs (and conditional cdfs) we are working with are supported on a common interval $[a, b]$, where $a < b$ are real numbers (a cdf F is supported on $[a, b]$ if $F(a-) = 0$ and $F(b) = 1$). We denote the set of cdfs supported on $[a, b]$ by $D_{cdf}([a, b])$. Furthermore, to allow the researcher to incorporate known structural properties of the cdfs, we shall introduce a “parameter space” $\mathcal{D} \subseteq D_{cdf}([a, b])$, in which the conditional cdfs appearing in F given in (7) are assumed to lie in. The following assumption on the target functional T and the functional measuring fairness S in relation to \mathcal{D} is crucial and will be assumed to hold throughout the article.

Assumption 4. The functionals

$$T : D_{cdf}([a, b]) \rightarrow \mathbb{R} \quad \text{and} \quad S : D_{cdf}([a, b]) \times D_{cdf}([a, b]) \rightarrow [0, \infty) \tag{14}$$

and the non-empty convex set $\mathcal{D} \subseteq D_{cdf}([a, b])$ satisfy

$$|T(F) - T(G)| \leq \|F - G\|_\infty \quad \text{and} \quad |S(F, \tilde{F}) - S(G, \tilde{G})| \leq \|F - G\|_\infty + \|\tilde{F} - \tilde{G}\|_\infty \tag{15}$$

for every $(F, \tilde{F}) \in \mathcal{D} \times \mathcal{D}$ and every $(G, \tilde{G}) \in D_{cdf}([a, b]) \times D_{cdf}([a, b])$. Furthermore, it holds that $S(F, F) = 0$ for every $F \in D_{cdf}([a, b])$.

Assumption 4 imposes an asymmetric Lipschitz-type condition on the functionals T and S that will allow us, e.g., to establish concentration results for plug-in estimators based on empirical cdfs through the Dvoretzky-Kiefer-Wolfowitz-Massart (DKWM) inequality, cf. Massart (1990).

It is not difficult to verify that the functionals discussed in the example presented in Section 6.1 (i.e., where T is a normalized Gini-welfare measure and S the Kolmogorov-Smirnov distance) satisfy Assumption 4.⁷ Note that functionals that satisfy a variant of (15), but where constants appear in front of the norms in the upper bounds can be brought into the framework of Assumption 4 by normalization. We could also work with a version of Assumption 4 with constants appearing in the upper bounds, but this would be notationally more cumbersome. These constants would then enter the regret upper bounds established below multiplicatively, but would not enter the policies introduced (hence their knowledge is not required by the DM to carry out the policy). Appendices F and G in Kock et al. (2022) verify the latter version of the above assumption for various poverty, inequality and welfare measures, as well as for functionals such as quantiles, U-statistics, L-statistics, and more.

⁷ Use the reverse triangle inequality for the Kolmogorov-Smirnov distance to verify the condition on S_{KS} there, and the discussion after Lemma E.9 in Kock et al. (2022) to verify the condition on T_{Gini} .

2.4.1. Some consequences

An important first question is whether the optimization problem in (13) permits a solution and if this is the case, how it depends on λ . For an array F as in (7), the following result establishes, among other useful properties, that the objective $\Omega_{\lambda,F}$ defined in Eq. (12) indeed has a maximizer (which need not be unique, cf. the example discussed in Section 6.1). For an array F as in (7) and \mathcal{H} a set of cdfs, we write $F \subset \mathcal{H}$ if all cdfs contained in the array F are elements of \mathcal{H} , i.e.,

$$F \subset \mathcal{H} \Leftrightarrow F^i(\cdot | x, z) \in \mathcal{H} \text{ for every } i = 1, \dots, K, x \in \mathcal{X}, \text{ and } z \in \mathcal{Z}.$$

For a definition of upper hemicontinuity of a set-valued function (correspondence) we refer to Section 17.2 in Aliprantis and Border (2006).

Proposition 1. *The following statements hold for $F \subset \mathcal{D}$ (cf. Assumption 4):*

1. $\arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda,F}(\delta)$ is a non-empty and compact subset of \mathbb{S}_K for every $\lambda \in [0, 1]$;
2. the function $\lambda \mapsto \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda,F}(\delta)$ is continuous and convex on $[0, 1]$;
3. the correspondence $\lambda \mapsto \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda,F}(\delta)$ is upper hemicontinuous on $[0, 1]$;
4. for every array G , say, as in (7) that satisfies $G \subset D_{cdf}([a, b])$ (but not necessarily $G \subset \mathcal{D}$), it holds that $\sup_{\delta \in \mathbb{S}_K} \Omega_{\lambda,G}(\delta)$ is finite.

As a consequence of Proposition 1, if F entering the objective $\Omega_{\lambda,F}(\delta)$ were known, the DM could simply attempt to solve the optimization problem $\max_{\delta \in \mathbb{S}_K} \Omega_{\lambda,F}(\delta)$ and would then roll out an (approximate) solution to the population (possibly after comparing solutions for different values of the preference parameter λ). However, note that this is *not* feasible, because the conditional cdfs entering the problem are not available to the DM. Nevertheless, the DM can use the training sample to estimate those quantities and use a plug-in approach to empirically determine an (approximate) best policy, which we shall detail in Section 3.

2.5. Regret

The goal of the DM is to use the training sample $\tilde{W}_j = (\tilde{Y}_j, X_j, Z_j, D_j)$ for $j = 1, \dots, n$, for which $(Y'_j, X_j, Z_j)' \sim F$, unknown, to learn an element of

$$\arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda,F}(\delta), \quad \text{for } \lambda \in [0, 1],$$

a non-empty set (given the assumptions imposed throughout) as just established in Proposition 1. Given a policy π_n , i.e., an algorithm to single out a decision rule in \mathbb{S}_K based on the training sample, the regret we shall work with is

$$r(\pi_n; \lambda, F) := \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda,F}(\delta) - \Omega_{\lambda,F}(\pi_n(\tilde{W}_j, j = 1, \dots, n)). \tag{16}$$

The regret measures the quality of the recommendation the DM makes based on the sample of subjects $j = 1, \dots, n$ in terms of the value of the objective function.

If the DM chooses from a set of policies π_n^λ indexed by $\lambda \in \Lambda \subseteq [0, 1]$ and the preference parameter $\hat{\lambda}_n \in \Lambda$ is chosen in a *data-driven* way (i.e., as a function of \tilde{W}_j for $j = 1, \dots, n$), the relevant regret notion becomes

$$r(\pi_n^{\hat{\lambda}_n}; \hat{\lambda}_n, F) = \max_{\delta \in \mathbb{S}_K} \Omega_{\hat{\lambda}_n,F}(\delta) - \Omega_{\hat{\lambda}_n,F}(\pi_n^{\hat{\lambda}_n}(\tilde{W}_j, j = 1, \dots, n)), \tag{17}$$

the behavior of which we study in such situations. That is, given the DM has decided to use the concrete value $\hat{\lambda}_n$, we ask how good the performance of the corresponding policy $\pi_n^{\hat{\lambda}_n}$ is relative to the best one for the selected value of the preference parameter. An alternative regret notion, which is interesting if $\hat{\lambda}_n$ estimates a certain target preference parameter λ^* (cf. the situation described in Section 4) is briefly discussed in Appendix B.

3. Empirical success policies

We now consider *empirical success policies*. Such policies depend on two quantities: First, the preference parameter $\lambda \in [0, 1]$, which defines the target, and a tuning parameter $\varepsilon > 0$ regulating the optimization accuracy. Given these two parameters, the DM proceeds in the following way (for some quantities, we abstain from signifying their dependence on n notationally to keep the expressions easy to read):

1. For every $i \in \{1, \dots, K\}$, $x \in \mathcal{X}$, and $z \in \mathcal{Z}$, determine $\hat{F}^i(\cdot | x, z)$, the empirical cdf based on all observations j in

$$\mathcal{M}_{x,z}^i := \{j = 1, \dots, n : D_j = i, X_j = x, Z_j = z\}, \tag{18}$$

i.e.,

$$\hat{F}^i(\cdot | x, z) := |\mathcal{M}_{x,z}^i|^{-1} \sum_{j \in \mathcal{M}_{x,z}^i} \mathbb{1}(Y_{D_j,j} \leq \cdot) = |\mathcal{M}_{x,z}^i|^{-1} \sum_{j \in \mathcal{M}_{x,z}^i} \mathbb{1}(Y_{i,j} \leq \cdot),$$

which we set equal to, e.g., the cdf corresponding to point mass at b (the upper endpoint of the support of the cdfs one is working with) in case $\mathcal{M}_{x,z}^i$ is empty.

2. For every $x \in \mathcal{X}$ and $z \in \mathcal{Z}$, obtain estimates

$$\hat{p}_X(x) := \frac{|\{j : X_j = x\}|}{n}, \hat{p}_Z(z) := \frac{|\{j : Z_j = z\}|}{n}, \text{ and } \hat{p}_X(x, z) := \frac{|\{j : X_j = x, Z_j = z\}|}{n},$$

and set $\hat{p}_{Z|X=x}(z) := \hat{p}_{X,Z}(x, z) / \hat{p}_X(x)$, which we set 0 if $\hat{p}_X(x) = 0$.

3. Set $\hat{F}_n := [(\hat{F}^i(\cdot | x, z), \hat{p}_{X,Z}(x, z)) : i = 1, \dots, K; x \in \mathcal{X}; z \in \mathcal{Z}]$.

4. Choose

$$\pi_n^{\varepsilon, \lambda}(\bar{W}_j, j = 1, \dots, n) \in \left\{ \gamma \in \mathbb{S}_K : \sup_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \hat{F}_n}(\delta) - \Omega_{\lambda, \hat{F}_n}(\gamma) \leq \varepsilon \right\}. \tag{19}$$

Some technical remarks are discussed in [Appendix C.1](#).

Remark 6 (Data-driven preference parameters). A data-driven preference parameter is any function $\hat{\lambda}_n$, say, that maps the training sample to $[0, 1]$, i.e.,

$$\hat{\lambda}_n : \mathbb{R}^n \times \mathcal{X}^n \times \mathcal{Z}^n \times \{1, \dots, K\}^n \rightarrow [0, 1].$$

To extract a decision rule, the function $\hat{\lambda}_n$ is applied to the data to obtain a value in $[0, 1]$, which is then used as the input in the above description of the policy to obtain the decision rule

$$\pi_n^{\varepsilon, \hat{\lambda}_n}(\bar{W}_j, j=1, \dots, n) (\bar{W}_j, j = 1, \dots, n);$$

we shall (with some abuse of notation) drop the dependence on the training sample and simply write $\pi_n^{\varepsilon, \hat{\lambda}_n}$ if no confusion can arise. For example, one could proceed in three steps:

1. Fix a finite set of preference parameters $\Lambda \subseteq [0, 1]$, which one intends to consider, and obtain (approximate) solutions $\pi_n^{\varepsilon, \lambda}$ as in (19) for every $\lambda \in \Lambda$.
2. Compare the solutions according to some criterion (cf. [Section 4](#) for a specific suggestion) and choose $\hat{\lambda}_n$ accordingly.
3. Roll out $\pi_n^{\varepsilon, \hat{\lambda}_n}$.

3.1. Regret bounds

We now establish an upper bound on the expected regret of any empirical success policy with data-dependent preference parameter. In addition to the bound on expected regret, the result also contains high-probability bounds on the regret.

Theorem 1. Let $(Y'_j, X_j, Z_j)' \sim \mathcal{F} \subset \mathcal{D}$ for $j = 1, \dots, n$ and abbreviate

$$\eta(\mathcal{F}) := \max_{x, z} \sum_{i=1}^K \frac{1}{\sqrt{\mathbb{P}(D_j = i | X_j = x, Z_j = z)}} \times \sum_{\bar{z} \in \mathcal{Z}} \frac{1}{\sqrt{p_Z(\bar{z})}}. \tag{20}$$

Then, for any data-driven preference parameter $\hat{\lambda}_n : \mathbb{R}^n \times \mathcal{X}^n \times \mathcal{Z}^n \rightarrow [0, 1]$, the expected regret

$$\mathbb{E}^* \left(r(\pi_n^{\varepsilon, \hat{\lambda}_n}; \hat{\lambda}_n, \mathcal{F}) \right) \leq 22\eta(\mathcal{F}) \sqrt{\frac{\log(K)|\mathcal{X}|}{n}} + 4\sqrt{\frac{|\mathcal{Z}|}{n}} + \varepsilon; \tag{21}$$

furthermore, for every $\rho > 0$ it holds that

$$\mathbb{P}^* \left(r(\pi_n^{\varepsilon, \hat{\lambda}_n}; \hat{\lambda}_n, \mathcal{F}) \geq \rho + \varepsilon \right) \leq 4|\mathcal{Z}||\mathcal{X}|K^2 \max_{x, z, i} e^{-\frac{nq_{x, z, i}\rho^2}{512(|\mathcal{X}||\mathcal{Z}|)^2}}, \tag{22}$$

and where $q_{x, z, i} := p_{X, Z, D}(x, z, i)$.

It is important to emphasize that the upper bound just given decays at the parametric rate $1/\sqrt{n}$ in the size n of the training sample, given that one chooses the optimization procedure in (19) in such a way that the optimization error is of the order $1/\sqrt{n}$. By similar methods to the ones in [Kock et al. \(2023\)](#), this could formally be achieved, e.g., by optimizing the empirical objective function $\delta \mapsto \Omega_{\lambda, \hat{F}_n}(\delta)$ over a fine enough grid in \mathbb{S}_K (exploiting that the true underlying objective function $\delta \mapsto \Omega_{\lambda, \mathcal{F}}(\delta)$ is Lipschitz continuous, cf. [Lemma 4](#)). We have decided not to specify the optimization method explicitly in this article. In practice, running an optimization algorithm such as a Nelder-Mead search (possibly multiple times at randomly chosen initial values) with a choice of accuracy parameters that one can afford in a particular application (in terms of runtime) is more convenient than optimizing over a grid and is what we do in our numerical and empirical results. The question which optimization heuristic should be used in terms of efficiency is interesting and deserves further investigation, but goes beyond the scope of the present article.

Next, we comment on the quantity $\eta(\mathcal{F})$ that enters the upper bound in [Theorem 1](#). This quantity is composed of two factors. The first,

$$\max_{x, z} \sum_{i=1}^K \frac{1}{\sqrt{\mathbb{P}(D_j = i | X_j = x, Z_j = z)}}, \tag{23}$$

is related to what is typically referred to as an overlap condition. It is particularly large if there exists a pair x and z and a treatment i , say, that is only assigned with a very low probability whenever $X_j = x$ and $Z_j = z$; this then leads to less accurate estimation of the cdf $\hat{F}^i(\cdot | x, z)$. On the other hand, if

$$\mathbb{P}(D_j = i | X_j = x, Z_j = z) \approx 1/K,$$

i.e., if the assignment is balanced, then (23) is approximately bounded from above by $K^{3/2}$. The second factor

$$\sum_{\bar{z} \in \mathcal{Z}} \frac{1}{\sqrt{p_Z(\bar{z})}}$$

is smallest if the protected subgroups are of equal size, in which case it equals $|\mathcal{Z}|^{3/2}$ and becomes large if a protected subgroup is rather small. This is not surprising, because the objective (deliberately) does not down-weight unfairness against small groups (cf. the discussion in Remark 3), and therefore needs accurate estimates for the cdfs also within small groups. Hence, if there are small groups, the problem of determining the optimal policy becomes more difficult, which is reflected in the larger upper bound.

We note that the theoretical guarantees given in Theorem 1 are informative if the cardinalities of \mathcal{X} and \mathcal{Z} are small relative to the size of the training sample. For a discussion of the setting where \mathcal{X} is not finite, we refer the reader to Section 5. In Appendix C.2 we discuss how the value function $\lambda \mapsto \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta)$ can be estimated by interpolation, which may be of some independent interest.

3.2. Consistency

As the next theoretical result in this article, we give conditions under which the d_1 -distance of the policy $\pi^{\varepsilon_n, \hat{\lambda}_n}$ to the (non-empty) random set of maximizers⁸

$$\arg \max_{\delta \in \mathbb{S}_K} \Omega_{\hat{\lambda}_n, \mathcal{F}}(\delta) \tag{24}$$

converges to zero in probability as $n \rightarrow \infty$; i.e., we give conditions under which the policy $\pi^{\varepsilon_n, \hat{\lambda}_n}$ is consistent. Furthermore, we shall also quantify the rate of convergence to zero of the probability that the decision rule deviates from its target in (24) by more than a given distance $\rho > 0$.

As in the previous section, we allow for data-driven preference parameters $\hat{\lambda}_n$, so that the set of maximizers in the previous display is data-dependent (unless $\hat{\lambda}_n$ is constant). In the context of establishing consistency, this introduces some technical difficulties that—in contrast to Theorem 1—need to be traded off with additional assumptions on the data-driven preference parameter choice. In Appendix C.3 we exhibit the subtleties that arise for consistency results in the presence of data-driven preference parameters $\hat{\lambda}_n$, which explains why additional restrictions on $\hat{\lambda}_n$ need to be imposed for a consistency result.

3.2.1. Consistency of the policy

Recall that all random variables are defined on an underlying probability space $(\Omega, \mathcal{A}, \mathbb{P})$. We denote convergence in \mathbb{P} -probability by “ $\xrightarrow{\mathbb{P}}$ ” and convergence in outer \mathbb{P} -probability by “ $\xrightarrow{\mathbb{P}^*}$ ”. Furthermore, for $\rho > 0$ and $\lambda \in [0, 1]$ we denote

$$\mathbb{M}_{\mathcal{F}, \lambda}^\rho = \{ \gamma \in \mathbb{S}_K : d_1(\gamma, \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta)) \geq \rho \}. \tag{25}$$

For every $\rho > 0$ we denote

$$c_{\mathcal{F}}(\rho, \lambda) := \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta) - \max_{\delta \in \mathbb{M}_{\mathcal{F}, \lambda}^\rho} \Omega_{\lambda, \mathcal{F}}(\delta), \tag{26}$$

if $\mathbb{M}_{\mathcal{F}, \lambda}^\rho \neq \emptyset$ and we set $c_{\mathcal{F}}(\rho, \lambda) = 0$ in case $\mathbb{M}_{\mathcal{F}, \lambda}^\rho$ is empty. Proposition 1 shows that the quantities appearing in (25) and (26) are well defined (noting that $\mathbb{M}_{\mathcal{F}, \lambda}^\rho$ is a compact set). Based on Lemma 2, the following result exhibits a general upper bound on the probability that the d_1 -distance between $\pi^{\varepsilon_n, \hat{\lambda}_n}$ and $\arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta)$ exceeds ρ , which then leads to consistency statements under further conditions on the quantities appearing in that upper bound.

Theorem 2. Let $(Y'_j, X_j, Z_j)' \sim \mathcal{F} \subset \mathcal{D}$ for $j = 1, \dots, n$. Let a sequence of preference parameters

$$\hat{\lambda}_n : \mathbb{R}^n \times \mathcal{X}^n \times \mathcal{Z}^n \rightarrow \Lambda \subseteq [0, 1]$$

be given, where $\Lambda \neq \emptyset$ does not depend on n . Then, for every $\rho > 0$, the function $\lambda \mapsto c_{\mathcal{F}}(\rho, \lambda)$ is Borel measurable, and for every positive real number ζ , it holds that

$$\begin{aligned} \mathbb{P}^* \left(\pi^{\varepsilon_n, \hat{\lambda}_n} \in \mathbb{M}_{\mathcal{F}, \hat{\lambda}_n}^\rho \right) &\leq 2|\mathcal{X}|(K^2 + 1) \sum_{z \in \mathcal{Z}} \max_{x, j} e^{-\frac{n d_{x, z, j} \zeta^2}{8|\mathcal{X}|^2}} + 2|\mathcal{Z}| e^{-\frac{n \zeta^2}{2|\mathcal{Z}|^2}} \\ &\quad + \mathbb{P}^*(0 < c_{\mathcal{F}}(\rho, \hat{\lambda}_n) < \varepsilon_n + 6\zeta), \end{aligned} \tag{27}$$

⁸ Here, as usual, the distance of a point x in a metric space (X, d) to a non-empty subset $A \subseteq X$ is defined as $\inf_{z \in A} d(x, z)$.

for $q_{x,z,i} := p_{X,Z,D}(x, z, i)$; and where the event in the latter probability is measurable if $\hat{\lambda}_n$ is measurable. Therefore, if there exists a sequence $\zeta_n > 0$ such that for every $\rho > 0$ we have

$$n\zeta_n^2 \rightarrow \infty \quad \text{and} \quad \mathbb{P}^*(0 < c_{\mathcal{F}}(\rho, \hat{\lambda}_n) < \varepsilon_n + 6\zeta_n) \rightarrow 0,$$

then the policy $\pi_n^{\varepsilon_n, \hat{\lambda}_n}$ is consistent, i.e.,

$$d_1\left(\pi_n^{\varepsilon_n, \hat{\lambda}_n}, \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\hat{\lambda}_n, \mathcal{F}}(\delta)\right) \xrightarrow{\mathbb{P}^*} 0, \quad \text{as } n \rightarrow \infty. \tag{28}$$

In particular, the policy $\pi_n^{\varepsilon_n, \hat{\lambda}_n}$ is consistent if Λ consists of finitely many elements and $\varepsilon_n \rightarrow 0$.

We note that the upper bound in (27) holds under minimal assumptions on the preference parameter. While the first two summands in that upper bound converge to 0 as $n \rightarrow \infty$ (due to Assumption 3), there is no guarantee that the third summand converges to 0, in general. This is only true under additional conditions, e.g., finiteness of Λ . Note that the finiteness condition does not present a serious practical limitation, as one can typically only consider a finite grid of λ values.

The final question that remains is how to choose the preference parameter λ . Before we address this question, we show that a policy that is consistent for $\arg \max_{\delta \in \mathbb{S}_K} \Omega_{\hat{\lambda}_n, \mathcal{F}}(\delta)$ and is based on a data-driven preference parameter $\hat{\lambda}_n$ that converges to a non-random $\lambda^* \in [0, 1]$, is also consistent for $\arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda^*, \mathcal{F}}(\delta)$; i.e., the policy is consistent for the target in which $\hat{\lambda}_n$ is replaced by the limiting value λ^* . A similar statement is established if the preference parameter approaches a random variable with finite support. Such a result is interesting, in case an *infeasible* “oracle” preference parameter $\lambda^* = \lambda^*(\mathcal{F})$ depending on the unknown \mathcal{F} is targeted by a consistent data-driven preference parameter selection mechanism, as will be done in Section 4.

Proposition 2. Let $(Y'_j, X_j, Z_j)' \sim \mathcal{F} \subset \mathcal{D}$ for $j = 1, \dots, n$. Let a sequence of preference parameters $\hat{\lambda}_n : \mathbb{R}^n \times \mathcal{X}^n \times \mathcal{Z}^n \rightarrow [0, 1]$ be given that converges in (outer) probability to a (non-random) $\lambda^* \in [0, 1]$ and let π_n be a sequence of policies. Then

$$d_1\left(\pi_n, \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\hat{\lambda}_n, \mathcal{F}}(\delta)\right) \xrightarrow{\mathbb{P}^*} 0 \quad \Rightarrow \quad d_1\left(\pi_n, \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda^*, \mathcal{F}}(\delta)\right) \xrightarrow{\mathbb{P}^*} 0. \tag{29}$$

If $\hat{\lambda}_n \in \Lambda$, a finite subset of $[0, 1]$ that does not depend on n , and if $\hat{\lambda}_n - \lambda_n^* \xrightarrow{\mathbb{P}^*} 0$ for a sequence random variables $\lambda_n^* \in \Lambda$, then (29) holds with λ^* replaced by λ_n^* .

Proposition 2 follows from upper-hemicontinuity of $\lambda \mapsto \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta)$ that was established in Proposition 1. Proposition 2 establishes that if one can single out a preferable theoretical preference parameter λ^* that itself is infeasible but consistently estimable, and if the policy π_n is consistent for the *random* set of maximizers $\arg \max_{\delta \in \mathbb{S}_K} \Omega_{\hat{\lambda}_n, \mathcal{F}}(\delta)$, then the policy is also consistent for the *non-random* set of maximizers $\arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda^*, \mathcal{F}}(\delta)$.

4. Choosing the preference parameter

Theorem 1 shows that the expected regret of the policies considered decays at the parametric rate regardless of how $\hat{\lambda}_n$ is chosen. Furthermore, we have seen in Theorem 2 that whenever the preference parameter is selected from a finite set of candidate values, the corresponding policy is consistent in the sense that

$$d_1\left(\pi_n^{\varepsilon_n, \hat{\lambda}_n}, \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\hat{\lambda}_n, \mathcal{F}}(\delta)\right) \xrightarrow{\mathbb{P}^*} 0, \quad \text{as } n \rightarrow \infty,$$

and Proposition 2 even shows that if $\hat{\lambda}_n$ “converges” in (outer) probability to λ_n^* (in the sense that their distance converges to 0 in (outer) probability), then

$$d_1\left(\pi_n^{\varepsilon_n, \hat{\lambda}_n}, \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda_n^*, \mathcal{F}}(\delta)\right) \xrightarrow{\mathbb{P}^*} 0, \quad \text{as } n \rightarrow \infty.$$

Those guarantees are not in place if Λ is not finite, but this is not a strong restriction: In practice, a DM will initially fix a *finite* set of preference parameters Λ , say, from which an element is then selected in a data-driven way, possibly by comparing properties of the inferred policies $\pi_n^{\varepsilon_n, \lambda}$ for $\lambda \in \Lambda$. In this section, we assume that Λ is finite and consider in more detail the question of which preference parameter λ^* in Λ should be used. A first answer to this question was already briefly touched upon in Remark 4, where it was recommended that the DM study aspects of (estimates of) the functions

$$\lambda \mapsto \mathbb{T}\left(\langle \pi_n^{\varepsilon_n, \lambda}, \mathcal{F} \rangle\right) \quad \text{and} \quad \lambda \mapsto \mathbb{S}\left(\langle \pi_n^{\varepsilon_n, \lambda}, \mathcal{F} \rangle_z, \langle \pi_n^{\varepsilon_n, \lambda}, \mathcal{F} \rangle\right).$$

The empirical examples in Section 7 will illustrate this approach.

To provide another answer to this question, we adapt the “price of fairness” criterion in Berk et al. (2017) to our context. Suppose that the DM has obtained the policies $\pi_n^{\varepsilon_n, \lambda}$ for every $\lambda \in \Lambda$, which we shall assume to contain 0 in what follows, and now wants to decide which one among those to roll out. Recall that if the DM decides to roll out the policy $\pi_n^{\varepsilon_n, \lambda}$ for a given $\lambda \in \Lambda$, this results in the *population* cdf $\langle \pi_n^{\varepsilon_n, \lambda}, \mathcal{F} \rangle$, which is unknown to the DM, because \mathcal{F} is unknown. Ignore this “estimation step” for a moment, and define for every $\lambda \in \Lambda$ the difference

$$\Delta_n(\lambda, \mathcal{F}) := \mathbb{T}(\langle \pi_n^{\varepsilon_n, 0}, \mathcal{F} \rangle) - \mathbb{T}(\langle \pi_n^{\varepsilon_n, \lambda}, \mathcal{F} \rangle),$$

which depends on the data due to its dependence on the policies. The difference $\Delta_n(\lambda, F)$ measures, for a given λ , the difference in the target functional T applied to the policy $\pi_n^{\varepsilon_n^0}$ (that does not penalize unfairness) and the policy $\pi_n^{\varepsilon_n^\lambda}$ (that penalizes unfairness with penalty λ). Note that if $\Delta_n(\lambda, F) > 0$, then penalizing unfairness leads to a decrease/loss in the target functional.⁹ Therefore, one may interpret $\Delta(\lambda, F)$ as the “price” the DM is paying for penalizing unfairness by rolling out $\pi_n^{\varepsilon_n^\lambda}$ instead of $\pi_n^{\varepsilon_n^0}$. We now assume that there is a maximal “budget” the DM can afford to spend in penalizing for unfairness, i.e., a maximal decrease in the target functional T that can be tolerated. Call this budget $\beta > 0$. Given that the DM intends to select the preference parameter λ from Λ , we thus define the targeted preference parameter as

$$\lambda_n^*(\beta, F) := \max\{\lambda \in \Lambda : \Delta_n(\lambda, F) \leq \beta\}; \tag{30}$$

note that $\lambda^*(\beta, F)$ exists because Λ is finite and contains 0. The so-defined “oracle” target $\lambda_n^*(\beta, F)$ according to budget β is the maximal unfairness penalty that satisfies the budget constraint. Because F is unknown, we have to estimate $\lambda_n^*(\beta, F)$.

Taking care of the estimation step is more difficult than it may perhaps seem at first sight. To see this, note that the function h , say, that maps $z = (z_0, z_1, \dots, z_M) \in \{0\} \times \mathbb{R}^M$ ($M \in \mathbb{N}$) to $\max\{i : z_i \leq \beta\}$ is not everywhere continuous; the set of discontinuity points of h is given by

$$D(h) := \{z \in \{0\} \times \mathbb{R}^M : z_{h(z)} = \beta\}. \tag{31}$$

Therefore, replacing F by \hat{F}_n in the definition of the preference parameter in the penultimate display in a plug-in fashion, and even though $\Delta_n(\lambda, \hat{F}_n) \approx \Delta_n(\lambda, F_n)$ for every $\lambda \in \Lambda$ actually holds under weak assumptions (cf. the proof of [Theorem 3](#)), we can in general not easily conclude (e.g., by a suitable version of the continuous mapping theorem) that the plug-in preference parameter that replaces F by \hat{F}_n in (30) consistently estimates its target without further assumptions.

What can be achieved—without additional conditions beyond [Assumptions 3](#) and [4](#) and a finiteness condition on Λ —is to consistently underestimate the oracle and to consistently satisfy the budget constraint. This is achieved by replacing F by \hat{F}_n in (30) and by introducing some “slack”, i.e., we use the estimator

$$\hat{\lambda}_n^*(\beta, \hat{F}_n) = \max\{\lambda \in \Lambda : \Delta_n(\lambda, \hat{F}_n) \leq \beta(1 - c_n)\} \quad \text{with} \quad c_n := \sqrt{\log(n)/n}. \tag{32}$$

where

$$\Delta_n(\lambda, \hat{F}_n) := T((\pi_n^{\varepsilon_n^0}, \hat{F}_n)) - T((\pi_n^{\varepsilon_n^\lambda}, \hat{F}_n)).$$

The precise result is as follows.

Theorem 3. Let $(Y_j', X_j, Z_j)' \sim F \sqsubset \mathcal{D}$ for $j = 1, \dots, n$. Let $\Lambda = \{0, \lambda_1, \dots, \lambda_M\} \subseteq [0, 1]$ with $0 < \lambda_1 < \lambda_2 < \dots < \lambda_M \leq 1$ for some fixed $M \in \mathbb{N}$. Fix a budget $\beta > 0$. If $\varepsilon_n \rightarrow 0$, it follows that

$$\mathbb{P}^*(\hat{\lambda}_n^*(\beta, \hat{F}_n) > \lambda_n^*(\beta, F)) \rightarrow 0 \quad \text{and} \quad \mathbb{P}^*(\Delta_n(\hat{\lambda}_n^*(\beta, \hat{F}_n), F) > \beta) \rightarrow 0, \tag{33}$$

i.e., $\hat{\lambda}_n^*(\beta, \hat{F}_n)$ consistently underestimates $\lambda_n^*(\beta, F)$ and consistently satisfies the budget constraint. Under the additional condition that for some $\alpha > 0$ it holds that

$$\mathbb{P}^*(\beta - \Delta_n(\lambda_n^*(\beta, F), F) \leq [\alpha + \beta]c_n) \rightarrow 0, \tag{34}$$

it even holds that

$$\hat{\lambda}_n^*(\beta, \hat{F}_n) - \lambda_n^*(\beta, F) \xrightarrow{\mathbb{P}^*} 0,$$

from which it then also follows that

$$d_1\left(\pi_n^{\varepsilon_n \hat{\lambda}_n^*(\beta, \hat{F}_n)}, \arg \max_{\delta \in \mathcal{S}_K} \Omega_{\lambda_n^*(\beta, F), F}(\delta)\right) \xrightarrow{\mathbb{P}^*} 0. \tag{35}$$

[Theorem 3](#) provides conditions under which the data-driven preference parameter $\hat{\lambda}_n^*(\beta, \hat{F}_n)$ converges in outer probability to its target, and that selecting the policy from a set of empirical success policies indexed over Λ based on $\hat{\lambda}_n^*(\beta, \hat{F}_n)$ is consistent for the target that is defined by spending a maximal budget to penalize for unfairness.

The main condition needed for consistency is the one in (34). It requires that the quantity $\Delta_n(\lambda_n^*(\beta, F), F)$, which never exceeds β by definition, takes its values in an $(\alpha + \beta)c_n$ neighborhood of β with outer probability converging to 0, thus avoiding the discontinuity problems implied by the function h as described around (31). That $\Delta_n(\lambda_n^*(\beta, F), F)$ concentrates at the budget constraint β so that this condition is violated may arise in special cases but is not too likely given that Λ consists only of finitely many values chosen by the user. Hence, the extra condition needed for consistency does not seem to be a very strong restriction.

If the extra condition (34) for consistency is not satisfied, the data-driven preference parameter is “robust” in the sense that it consistently satisfies the budget constraint and is at least consistently not larger than the oracle, which is defined as the maximal penalty for unfairness that satisfies the budget constraint.

That we can (only) consistently underestimate the oracle in general is conceptually in line with general findings in [Sun \(2021\)](#), where impossibility results in the related setting of empirical welfare maximization under (estimated) constraints are studied, and ways to take the corresponding estimation error into account are suggested.

⁹ Note that $\Delta_n(\lambda, F)$ can in principle also be negative, due to the policy being based on estimates of F .

5. Generalization to non-discrete covariates

Throughout this article, we consider the case where \mathcal{Z} is finite. So far, we focused on a setting where \mathcal{X} is finite as well. In case \mathcal{X} contains infinitely many elements, one could simply “discretize” \mathcal{X} , and then apply the policy developed and studied for the case of finitely supported covariates in Section 3. Granted the assumptions used to establish theoretical guarantees in earlier sections of this article are satisfied for the *discretized* covariates, the theoretical guarantees (expected regret bounds, consistency, etc.) carry over immediately. One disadvantage of this approach is that the performance guarantees developed, e.g., in Theorem 1, get weaker as the discretization gets finer. Yet, the finer the discretization, the more reasonable is, e.g., the conditional independence Assumption 2, conditioning on the discretized covariates, which restricts the scope of such a “direct” application of our results.

An alternative approach, which is not based on discretization, but can be viewed as a generalization of the approach of Kitagawa and Tetenov (2021) to our context, is discussed in the present section. The non-finiteness of the number of elements of \mathcal{X} necessarily requires some notational and conceptual modifications, which are only maintained locally in this section and the corresponding proofs in Appendix G. The objects studied (e.g., the training sample, policies, or the distributions generated in the roll-out phase) have the same conceptual meaning as before, and hence their interpretation does not require additional discussion on top of that already given earlier. In this section, we maintain throughout the following *list of assumptions and notational conventions, without further mentioning them in the theoretical results*:

1. We assume that \mathcal{Z} is finite, but we do not assume that \mathcal{X} is finite (although this is not ruled out formally).
2. We impose Assumptions 1, 2, and 4 throughout, but no longer impose Assumption 3 (which is replaced, to some extent, by Assumption 5 below).
3. We assume that $\mathbf{W}_j = (Y'_j, X_j, Z_j, D_j)'$ takes its values in $\mathbb{R}^K \times \mathcal{X} \times \mathcal{Z} \times \{1, \dots, K\}$ equipped with the product sigma algebra, where finite sets (in particular \mathcal{Z}) are equipped with the power set (set of all subsets), \mathbb{R}^K is equipped with the Borel sigma algebra, and the non-empty set \mathcal{X} is equipped with some sigma algebra.
4. We denote a regular conditional distribution of $Y_{i,j}$ given (X_j, Z_j) by $K_{x,z}^i(\cdot)$.¹⁰ For $x \in \mathcal{X}$ and $z \in \mathcal{Z}$ the conditional cdf $F^i(\cdot | x, z) := K_{x,z}^i((-\infty, \cdot])$. Furthermore, we denote by F^i the cdf of $Y_{i,j}$. We collect the conditional cdfs of a random vector Y'_j given $X_j = x$ and $Z_j = z$ and $\mathbb{P}_{X,Z}$, the joint distribution of X_j and Z_j , in an array

$$F := [(F^i(\cdot | x, z), \mathbb{P}_{X,Z}) : i = 1, \dots, K; x \in \mathcal{X}; z \in \mathcal{Z}]. \tag{36}$$

The symbol \mathbb{P}_X denotes the distribution of X_j , \mathbb{P}_Z denotes the distribution of Z_j with corresponding probability mass function p_Z , which we assume to be strictly positive at every element of \mathcal{Z} , and $\mathbb{P}_{X|Z=z}(\cdot) = \mathbb{P}_{X,Z}(\cdot, \{z\})/p_Z(z)$ is a regular conditional distribution of X_j given Z_j evaluated at $z \in \mathcal{Z}$ (recall that \mathcal{Z} is assumed finite).

5. A decision rule δ is a *measurable* function from \mathcal{X} to \mathcal{S}_K , the simplex in \mathbb{R}^K . We consider situations, where the decision rules are chosen from a non-empty set Π of measurable functions from \mathcal{X} to \mathcal{S}_K .
6. For every decision rule δ and every $z \in \mathcal{Z}$, we define (analogously to (11)) the cdf

$$\langle \delta, F \rangle_z(y) := \int \sum_{i=1}^K \delta_i(x) F^i(y | x, z) d\mathbb{P}_{X|Z=z}(x), \tag{37}$$

and set (analogously to (10))

$$\langle \delta, F \rangle := \sum_{z \in \mathcal{Z}} \langle \delta, F \rangle_z p_Z(z).$$

Analogously to Section 2.3.2, the cdfs $\langle \delta, F \rangle_z$ and $\langle \delta, F \rangle$ are the conditional and unconditional cdfs obtained by rolling out the decision rule δ .¹¹

7. We write $F \sqsubset_{\Pi} \mathcal{H}$ if for every $\delta \in \Pi$ and every $z \in \mathcal{Z}$ the cdf $\langle \delta, F \rangle_z$ is an element of the closure of \mathcal{H} (w.r.t. the supremum metric).

In addition, we impose the following “overlap” condition throughout this section (in place of working with Assumption 3):

Assumption 5 (Overlap condition). Let

$$e_i(x, z) := \mathbb{P}(D_j = i | X_j = x, Z_j = z) \quad \text{for } i = 1, \dots, K,$$

define a regular conditional distribution of D_j given X_j and Z_j , and assume that (it can be chosen such that) $0 < e_i(x, z)$ for every $i = 1, \dots, K$, every $x \in \mathcal{X}$, and every $z \in \mathcal{Z}$, and such that $\mathbb{E}(e_i^{-1}(X_j, Z_j)) < \infty$ for every $i = 1, \dots, K$.

Granted all the assumptions above, we introduce the class of (non-negative) measurable functions

$$\mathcal{G}_{\Pi} := \{ f_{\delta,c,z} : \delta \in \Pi, c \in \mathbb{R}, z \in \mathcal{Z} \},$$

¹⁰ Because $Y_{i,j}$ is real-valued, the existence of a regular conditional distribution is guaranteed, cf. Definition A.7 and Theorem A.37 in Liese and Miescke (2008).

¹¹ More precisely, consider a new draw from the underlying population as in (8), and suppose that its assignment, D , say, satisfies Eq. (9) for every $i, x \in \mathcal{X}, y \in \mathbb{R}$ and $z \in \mathcal{Z}$, e.g., because the mechanism as discussed in Footnote 6 was implemented by the DM. Then, for every $z \in \mathcal{Z}$, $\langle \delta, F \rangle_z$ is the cdf of Y_D given $Z = z$, and the cdf of Y_D is $\langle \delta, F \rangle$.

where $f_{\delta,c,z} : \mathbb{R} \times \mathcal{X} \times \mathcal{Z} \times \{1, \dots, K\} \rightarrow \mathbb{R}$ is defined as (recall that $p_Z(z) > 0$ is assumed for every $z \in \mathcal{Z}$)

$$f_{\delta,c,z}(y, x, z^*, d) := \delta_d(x) \times \mathbb{1}(y \leq c) \times \frac{\mathbb{1}(z^* = z)}{e_d(x, z)p_Z(z)}.$$

That the expectation of $f_{\delta,c,z}(Y_{D_j,j}, X_j, Z_j, D_j)$ equals $\langle \delta, \mathcal{F} \rangle_z(c)$ is shown in Lemma 12 in Appendix G, delivering (recall \tilde{W}_j from (2)) the unbiased estimator $n^{-1} \sum_{j=1}^n f_{\delta,c,z}(\tilde{W}_j)$ of $\langle \delta, \mathcal{F} \rangle_z(c)$, cf. also the discussion in Section 3 of Kitagawa and Tetenov (2021) for a similar approach. We abbreviate,

$$\|P_n - P\|_{\mathcal{G}_\Pi} := \sup_{f \in \mathcal{G}_\Pi} |n^{-1} \sum_{i=1}^n f(\tilde{W}_i) - \mathbb{E}(f(\tilde{W}_i))|.$$

Note that suitably “small”, i.e., “parametric”, sets of decision rules Π (and hence \mathcal{G}_Π) allow one to establish, for some finite constant $c(\mathcal{G}_\Pi, P) > 0$, the bound

$$\mathbb{E}^*(\|P_n - P\|_{\mathcal{G}_\Pi}) \leq c(\mathcal{G}_\Pi, P)/\sqrt{n} < \infty. \tag{38}$$

For example, if one additionally assumes that $0 < \epsilon < e_i(x, z)$ for every $i = 1, \dots, K$, every $x \in \mathcal{X}$, and every $z \in \mathcal{Z}$, the class of (non-negative) functions \mathcal{G}_Π is bounded from above by $\max_{z \in \mathcal{Z}} 1/(\epsilon \times p_Z(z))$, so that Theorem 2.14.1 in van der Vaart and Wellner (1996) delivers (38) under (i) finiteness of a uniform entropy integral (satisfied in particular for VC classes) and (ii) suitable measurability conditions and some structural conditions of the underlying probability space.

We abstain from formulating expected regret bounds in the present setting in terms of other possible sufficient conditions for (38), but develop an upper bound on the regret itself that depends on $\|P_n - P\|_{\mathcal{G}_\Pi}$. This approach has the benefit of also being informative if (38) is not satisfied, e.g., because \mathcal{G}_Π is nonparametric, and a different rate of decay holds for the quantity to the left in (38). We refer the reader to, e.g., van der Vaart and Wellner (1996) and Giné and Nickl (2015) for tools to establish (38) under low-level assumptions on Π .

The goal of the decision maker is to choose δ such as to approximate

$$\sup_{\delta \in \Pi} \Omega_{F,\lambda}(\delta),$$

with the penalized objective $\Omega_{F,\lambda}(\delta)$ as defined in (12) (Lemma 13 in Appendix G shows that, under the maintained assumptions, the supremum in the previous display is finite). Similar to the definition in (4), a policy is a function $\pi_n : \mathbb{R}^n \times \mathcal{X}^n \times \mathcal{Z}^n \times \{1, \dots, K\}^n \rightarrow \Pi$, and its regret is defined as

$$r(\pi_n; \lambda, F) := \sup_{\delta \in \Pi} \Omega_{\lambda,F}(\delta) - \Omega_{\lambda,F}(\pi_n(\tilde{W}_j, j = 1, \dots, n)), \tag{39}$$

where we replace λ by $\hat{\lambda}_n$ if the preference parameter is selected in a data-driven way.

5.1. Empirical success type policy

For a given preference parameter $\lambda \in [0, 1]$ and tuning parameter $\epsilon > 0$, the policy $\pi_n^{\epsilon,\lambda}$ we shall consider here is a suitably adapted version of the policy discussed in Section 3. Conceptually, we do no longer work with a plug-in version of the objective (simply replacing F by an estimate), but we somewhat more carefully estimate the quantities through which the objective actually depends on directly. Here, we exploit that the expectation of $f_{\delta,c,z}(Y_{D_j,j}, X_j, Z_j, D_j)$ equals $\langle \delta, \mathcal{F} \rangle_z(c)$, which is in line with the approach in Kitagawa and Tetenov (2021).

We work with the “empirical target” $\hat{\Omega}_\lambda(\delta)$, say, which, for $\delta \in \Pi$ is defined as (cf. Assumption 4)

$$\hat{\Omega}_\lambda(\delta) := (1 - \lambda)\mathbb{T}(\widehat{\langle \delta, \mathcal{F} \rangle}) - \lambda \max_{z \in \mathcal{Z}} \mathbb{S}(\widehat{\langle \delta, \mathcal{F} \rangle}_z, \widehat{\langle \delta, \mathcal{F} \rangle}),$$

for estimates $\widehat{\langle \delta, \mathcal{F} \rangle}_z$ and $\widehat{\langle \delta, \mathcal{F} \rangle}$ defined in (40) below.

To obtain $\widehat{\langle \delta, \mathcal{F} \rangle}_z$, we project the function $c \mapsto n^{-1} \sum_{j=1}^n f_{\delta,c,z}(Y_{D_j,j}, X_j, Z_j, D_j)$, which (although being non-negative, non-decreasing, càdlàg, and pointwise unbiased for $\langle \delta, \mathcal{F} \rangle_z$) is not guaranteed to be a cdf, cf. also the discussion in Section 3 of Kitagawa and Tetenov (2021), onto the set of cdfs; here, the projection of a non-decreasing, non-negative, càdlàg function $G : \mathbb{R} \rightarrow \mathbb{R}$ onto $D_{cdf}([a, b])$ (equipped with the Kolmogorov-Smirnov distance) is denoted by $M_{a,b}G$, i.e., $M_{a,b}G(x)$ equals 0 for $x < a$, equals $\min(G(x), 1)$ for $x \in [a, b]$, and equals 1 for $x \geq b$.

1. For every $z \in \mathcal{Z}$ and every $\delta \in \mathbb{S}_K$ define the following cdfs in $D_{cdf}([a, b])$

$$\widehat{\langle \delta, \mathcal{F} \rangle}_z := M_{a,b} \left[n^{-1} \sum_{j=1}^n f_{\delta,\cdot,z}(Y_{D_j,j}, X_j, Z_j, D_j) \right], \text{ and } \widehat{\langle \delta, \mathcal{F} \rangle} := \sum_{z \in \mathcal{Z}} p_Z(z) \widehat{\langle \delta, \mathcal{F} \rangle}_z. \tag{40}$$

2. Choose

$$\pi_n^{\epsilon,\lambda}(\tilde{W}_j, j = 1, \dots, n) \in \left\{ \gamma \in \Pi : \sup_{\delta \in \Pi} \hat{\Omega}_\lambda(\delta) - \hat{\Omega}_\lambda(\gamma) \leq \epsilon \right\}. \tag{41}$$

Lemma 13 in Appendix G shows that the supremum defining the set in (41) is finite.

If the probabilities p_Z and the propensities $e_i(x, z)$ are unknown, a DM needs to replace them by estimates $\hat{p}_Z(z)$ and $\hat{e}_{i,n}(x, z)$ (the former defined as in Section 3, while we keep the latter abstract), respectively. In this situation, set

$$\tilde{f}_{\delta,c,z}(y, x, z^*, d) := \delta_d(x) \times \mathbb{1}(y \leq c) \times \frac{\mathbb{1}(z^* = z)}{\hat{e}_d(x, z) \hat{p}_Z(z)},$$

which we leave undefined if the denominator equals 0. Then, we consider the adapted policy $\tilde{\pi}_n^{\epsilon, \lambda}$, say, that differs from $\pi_n^{\epsilon, \lambda}$ in that it is not based on the quantities defined in (40) when computing the empirical target $\hat{\Omega}_\lambda(\delta)$, but replaces them by

$$\langle \widehat{\delta}, \widetilde{F} \rangle_z := M_{a,b} \left[n^{-1} \sum_{j=1}^n \tilde{f}_{\delta, \cdot, z}(Y_{D_j, j}, X_j, Z_j, D_j) \right], \text{ and } \langle \widehat{\delta}, F \rangle := \sum_{z \in \mathcal{Z}} \hat{p}_Z(z) \langle \widehat{\delta}, \widetilde{F} \rangle_z. \tag{42}$$

5.2. Performance bound

We now provide an upper bound on the regret of the policies just defined.

Theorem 4. Let $(Y'_j, X_j, Z_j)' \sim F$ for $j = 1, \dots, n$, and assume that $F \sqsubset_{\Pi} \mathcal{D}$. Then, for any data-driven preference parameter $\hat{\lambda}_n : \mathbb{R}^n \times \mathcal{X}^n \times \mathcal{Z}^n \rightarrow [0, 1]$, we have

$$r(\tilde{\pi}_n^{\epsilon, \hat{\lambda}_n}; \hat{\lambda}_n, F) \leq 6 \times \|P_n - P\|_{\mathcal{G}_{\Pi}} + \epsilon, \tag{43}$$

and $r(\tilde{\pi}_n^{\epsilon, \hat{\lambda}_n}; \hat{\lambda}_n, F)$ is bounded from above by the upper bound in (43) plus

$$\|\hat{p}_Z - p_Z\| + 6 \times \max_{z \in \mathcal{Z}} n^{-1} \sum_{j=1}^n \left| \hat{p}_Z^{-1}(z) \hat{e}_{D_j}^{-1}(X_j, z) - p_Z^{-1}(z) e_{D_j}^{-1}(X_j, z) \right|, \tag{44}$$

which is to be interpreted as ∞ if one of the denominator vanishes.

Because $\mathbb{E}\|\hat{p}_Z - p_Z\| = O(n^{-1/2})$, e.g., cf. Lemma 9, even if the propensity scores are unknown but (correctly) specified through a suitably regular parametric model, the expectation of the term in (44) will typically be of order $O(n^{-1/2})$. Together with the behavior of the (outer) expectation of $\|P_n - P\|_{\mathcal{G}_{\Pi}}$ being of order $O(n^{-1/2})$ if Π is sufficiently constrained (cf. the discussion around (38)), this then leads to a $O(n^{-1/2})$ upper bound on the expected regret of the policy in the context studied in this section provided $\epsilon = O(n^{-1/2})$.

It is also apparent from Theorem 4 that “large” nonparametric classes of policies or propensities will lead to slower rates of convergence to 0 of the expectation of the upper bound on the regret developed in Theorem 4.

6. An example and numerical results

6.1. A toy example

In this section, we provide a toy example that illustrates the setting outlined in Section 2. We shall re-investigate the example considered here in the numerical results below. Consider the case where T is the Gini-welfare functional, which is a popular welfare measure in the economics literature; cf. Sen (1976), and for results in an optimal policy choice context without fairness considerations see Kitagawa and Tetenov (2021). For a cdf F , the Gini-welfare functional is defined (assuming that the Lebesgue-Stieltjes integrals involved in its definition exist) as

$$W(F) := \int x dF(x) - \frac{1}{2} \int \int |x - y| dF(x) dF(y). \tag{45}$$

This welfare measure trades-off the expected outcome $\mu(F) = \int x dF(x)$, say, of F with $\sigma(F) = \frac{1}{2} \int \int |x - y| dF(x) dF(y)$, say, the “inequality” inherent in the distribution F , measured as its mean absolute difference (divided by 2). Hence, targeting the Gini-welfare measure, instead of simply the expected outcome, the decision maker not only aims at maximizing the expected outcome in the population considered, but also incorporates inequality concerns. We furthermore consider the similarity measure S_{KS} , say, which maps two cdfs F and G to the Kolmogorov-Smirnov distance between them, i.e.,

$$S_{KS}(F, G) := \|F - G\|_{\infty} := \sup_{x \in \mathbb{R}} |F(x) - G(x)|.$$

We shall here consider the case where $F \in \mathcal{D} = D_{cdf}([0, 1])$. Denoting by $T_{Gini} = W/2$ the Gini-welfare functional normalized by 2, the penalized objective we shall work with is defined as

$$\Omega_{\lambda, F}(\delta) = (1 - \lambda) T_{Gini}(\langle \delta, F \rangle) - \lambda \max_{z \in \mathcal{Z}} \|\langle \delta, F \rangle_z - \langle \delta, F \rangle\|_{\infty};$$

the normalization guarantees that Assumption 4 holds with $\mathcal{D} = D_{cdf}([0, 1])$.

Having fixed the functionals defining the target, we now consider the case where there are only $K = 2$ treatments, $\mathcal{X} = \{0\}$ and $\mathcal{Z} = \{0, 1\}$. That is, for simplicity of discussion there is only one covariate value, but there are two possible values for the protected characteristic. To specify an array F as in (7), we consider the cdfs

$$F^1(y | 0, 0) = F^2(y | 0, 1) = \sqrt{y} =: G(y) \text{ and } F^2(y | 0, 0) = F^1(y | 0, 1) = y^2 =: H(y), \tag{46}$$




Fig. 1. The figure plots the cdf $\langle \delta, \mathcal{F} \rangle$ (left panel) and the cdf $\langle \delta, \mathcal{F} \rangle_0$ (right panel) for different values of $\delta \in \{0, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1\}$ and for $p = 3/4$. Lower values of δ lead to stochastically larger cdfs. Note that $\langle \delta, \mathcal{F} \rangle_1$ can also be read off from the right panel due to symmetry, but the dependence on δ is now “reverted”, cf. Eq. (47).

for $y \in [0, 1]$ (0 for $y < 0$ and 1 for $y > 1$), and the probability mass function

$$p_{X,Z}(0, 0) =: p > 1/2 \quad \text{and} \quad p_{X,Z}(0, 1) = (1 - p),$$

i.e., the subpopulation corresponding to $Z = 0$ constitutes the “majority” whereas the subpopulation corresponding to $Z = 1$ is the “minority.” It holds that $\mu(G) = 1/3 < 2/3 = \mu(H)$ and $\sigma(G) = 1/6 > \sigma(H) = 2/15$, from which it follows that

$$\mathbb{T}_{Gini}(G) = (\mu(G) - \sigma(G))/2 = 1/12 \quad \text{and} \quad \mathbb{T}_{Gini}(H) = (\mu(H) - \sigma(H))/2 = 4/15.$$

Hence, the cdf H has larger Gini-welfare than the cdf G , and, from this point-of-view, is preferable. As a consequence, for the subpopulation defined by $Z = 0$ the second treatment (the cdf of which is H) is better than the first treatment (the cdf of which is G) in terms of \mathbb{T}_{Gini} ; the contrary is the case for the subpopulation defined by $Z = 1$. Note that in this example

$$F^1(\cdot | 0) = pG(\cdot) + (1 - p)H(\cdot) \quad \text{and} \quad F^2(\cdot | 0) = pH(\cdot) + (1 - p)G(\cdot).$$

The cdf generated by rolling out the decision rule $\delta = (\delta, 1 - \delta)'$ for a $\delta \in [0, 1]$ equals¹²

$$\langle \delta, \mathcal{F} \rangle = [\delta p + (1 - \delta)(1 - p)]G + [\delta(1 - p) + (1 - \delta)p]H,$$

with corresponding expectation $\mu(\langle \delta, \mathcal{F} \rangle) = (\delta + (1 - 2\delta)p + 1)/3$ and

$$\sigma(\langle \delta, \mathcal{F} \rangle) = \frac{1}{210} (\delta(20 - 27\delta) - 27(1 - 2\delta)^2 p^2 + 2(\delta(54\delta - 47) + 10)p + 35),$$

so that

$$\mathbb{T}_{Gini}(\langle \delta, \mathcal{F} \rangle) = \frac{1}{420} (50\delta + 27\delta^2(1 - 2p)^2 - 2\delta p(54p + 23) + p(27p + 50) + 35).$$

The subpopulation cdfs are given by

$$\langle \delta, \mathcal{F} \rangle_0 = \delta G + (1 - \delta)H \quad \text{and} \quad \langle \delta, \mathcal{F} \rangle_1 = \delta H + (1 - \delta)G. \tag{47}$$

Fig. 1 illustrates the cdfs $\langle \delta, \mathcal{F} \rangle$ and $\langle \delta, \mathcal{F} \rangle_0$ in dependence on δ in case $p = 3/4$.

A simple computation now shows that the penalty term equals (recalling that $p > 1/2$)

$$\max_{z \in Z} \|\langle \delta, \mathcal{F} \rangle_z - \langle \delta, \mathcal{F} \rangle\|_\infty = p|2\delta - 1| \max_{y \in (0,1)} y(1 - y^3) = p \frac{3}{4 \times 2^{2/3}} |2\delta - 1|.$$

To be explicit, the objective $\Omega_{\lambda, \mathcal{F}}((\delta, 1 - \delta)')$ hence equals

$$\frac{1}{420} (1 - \lambda) (50\delta + 27\delta^2(1 - 2p)^2 - 2\delta p(54p + 23) + p(27p + 50) + 35) - \frac{3\lambda p |2\delta - 1|}{4 \times 2^{2/3}}, \tag{48}$$

and is plotted in the left panel in Fig. 2 (for $p = 3/4$).

Fix a $\lambda \in [0, 1)$. From (48) we then see that the function $\delta \mapsto \Omega_{\lambda, \mathcal{F}}((\delta, 1 - \delta)')$ when restricted to $[0, 1/2]$ and also when restricted to $[1/2, 1]$, respectively, coincides with a quadratic polynomial with positive leading coefficient. It thus follows that $\arg \max_{\delta \in [0,1]} \Omega_{\lambda, \mathcal{F}}((\delta, 1 - \delta)) \subseteq \{0, 1/2, 1\}$. Using $p > 1/2$ together with (48), we see that $\Omega_{\lambda, \mathcal{F}}((0, 1)') > \Omega_{\lambda, \mathcal{F}}((1, 0)')$. It follows that either the argmax is $\{0\}$, or $\{1/2\}$, or their union $\{0, 1/2\}$, the latter occurring whenever $\Omega_{\lambda, \mathcal{F}}((0, 1)') = \Omega_{\lambda, \mathcal{F}}((1/2, 1/2)')$, which is the

¹² Note that in case \mathcal{X} a singleton set and $K = 2$, a decision rule δ is uniquely characterized by the probability at which the first treatment is assigned, which we denote by δ throughout this example.




Fig. 2. Objective function $\Omega_{\lambda, \mathcal{F}}(\delta)$ for $p = 3/4$ (left panel), and the corresponding maximum value of the objective function for $p = 3/4$ in dependence on λ (right panel).

case precisely if $\lambda = c(p) := 1 - \frac{630 \sqrt[3]{2} p}{2p(54p+315 \sqrt[3]{2+100})-127}$. For $\lambda = 1$ it is easy to see that the argmax is $\{1/2\}$. In general, the argmax is given by

$$\arg \max_{\delta \in [0,1]} \Omega_{\lambda, \mathcal{F}}((\delta, 1 - \delta)') = \begin{cases} 0 & 0 \leq \lambda < c(p) \\ \{0, 1/2\} & \lambda = c(p) \\ 1/2 & c(p) < \lambda \leq 1, \end{cases} \tag{49}$$

cf. Fig. 1 for the corresponding cdfs when $p = 3/4$, with a maximum value of

$$\max_{\delta \in [0,1]} \Omega_{\lambda, \mathcal{F}}((\delta, 1 - \delta)') = \begin{cases} \frac{p(-5(20+63 \sqrt[3]{2})\lambda - 54(\lambda-1)p + 100) - 70(\lambda-1)}{840} & 0 \leq \lambda \leq c(p) \\ \frac{89}{560}(1 - \lambda) & c(p) < \lambda \leq 1, \end{cases}$$

which we depict for $p = 3/4$ in Fig. 2 (note that $c(3/4) \approx 0.123$).

From the structure of the argmax in (49) we see that there is a phase-transition phenomenon in the value of the preference parameter λ . Up to the threshold $c(p)$, the optimal policy always assigns the second treatment, which is the optimal treatment within the majority group, but is also the sub-optimal treatment within the minority group. However, if λ surpasses the threshold $c(p)$, i.e., if the degree of fairness increases sufficiently, the optimal policy changes to a 50 : 50 assignment rule.

Remark 7. In this example, the value function is convex (cf. also Proposition 1), but is not differentiable in λ , as there is a kink at $c(p)$, which is related to the argmax not being a singleton set at $c(p)$. Note also that the argmax does *not* continuously depend on the preference parameter λ , even in this simple example. Also note that for $\lambda > c(p)$ the optimal policy is *randomized*. Therefore, this example also illustrates that focusing on non-randomized decision rules would come with an efficiency loss.

6.2. Numerical results

We now simulate data in the setting just considered in Section 6.1, and estimate the empirical success policy (as introduced in Section 3) on a grid of preference parameters $\lambda \in \{0, 1/49, 2/49, \dots, 1\}$, where the optimization step in the policy is achieved via a Nelder-Mead algorithm (as it is implemented in R’s “constrOptim” routine), and for which the starting points are selected as the best among 50 randomly drawn points (according to a uniform distribution) from \mathbb{S}_2 , respectively. Recall that in this example \mathcal{X} consists only of one element, whereas there are two possible values in \mathcal{Z} . To be comparable to the theoretical results in Section 6.1 we fix $p = 3/4$, recall that p corresponds to the size of the majority group.

The treatment assignment mechanisms D_i studied in the numerical results (besides satisfying Assumption 2) were generated according to the 2 scenarios as follows (recall that $Z_i = 0$ corresponds to being a member of the majority, whereas $Z_i = 1$ corresponds to being a member of the minority):

- A1 the conditional distribution of D_i given Z_i satisfies $P(D_i = 1 \mid Z_i = 0) = 1/4$ and $P(D_i = 1 \mid Z_i = 1) = 3/4$; thus, being assigned to Treatment 1 in the majority group is less likely than being assigned to Treatment 2, whereas being assigned Treatment 1 in the minority group is more likely than being assigned to Treatment 2.
- A2 the conditional distribution of D_i given Z_i satisfies $P(D_i = 1 \mid Z_i = 0) = 3/4$ and $P(D_i = 1 \mid Z_i = 1) = 1/4$; thus, being assigned to Treatment 1 in the majority group is more likely than being assigned to Treatment 2, whereas being assigned Treatment 1 in the minority group is less likely than being assigned to Treatment 2.

Recall from Section 6.1 that in the majority group the second treatment is better, while the first treatment is better in the minority group. Therefore, the assignment mechanism A1 corresponds to a setting where the respective optimal treatment was assigned more likely within each subgroup, whereas assignment mechanism A2 corresponds to the opposite of that situation. In addition to considering the mechanisms A1 and A2, we consider different sample sizes $n \in \{100, 1000, 10000\}$, resulting in 6 different simulation

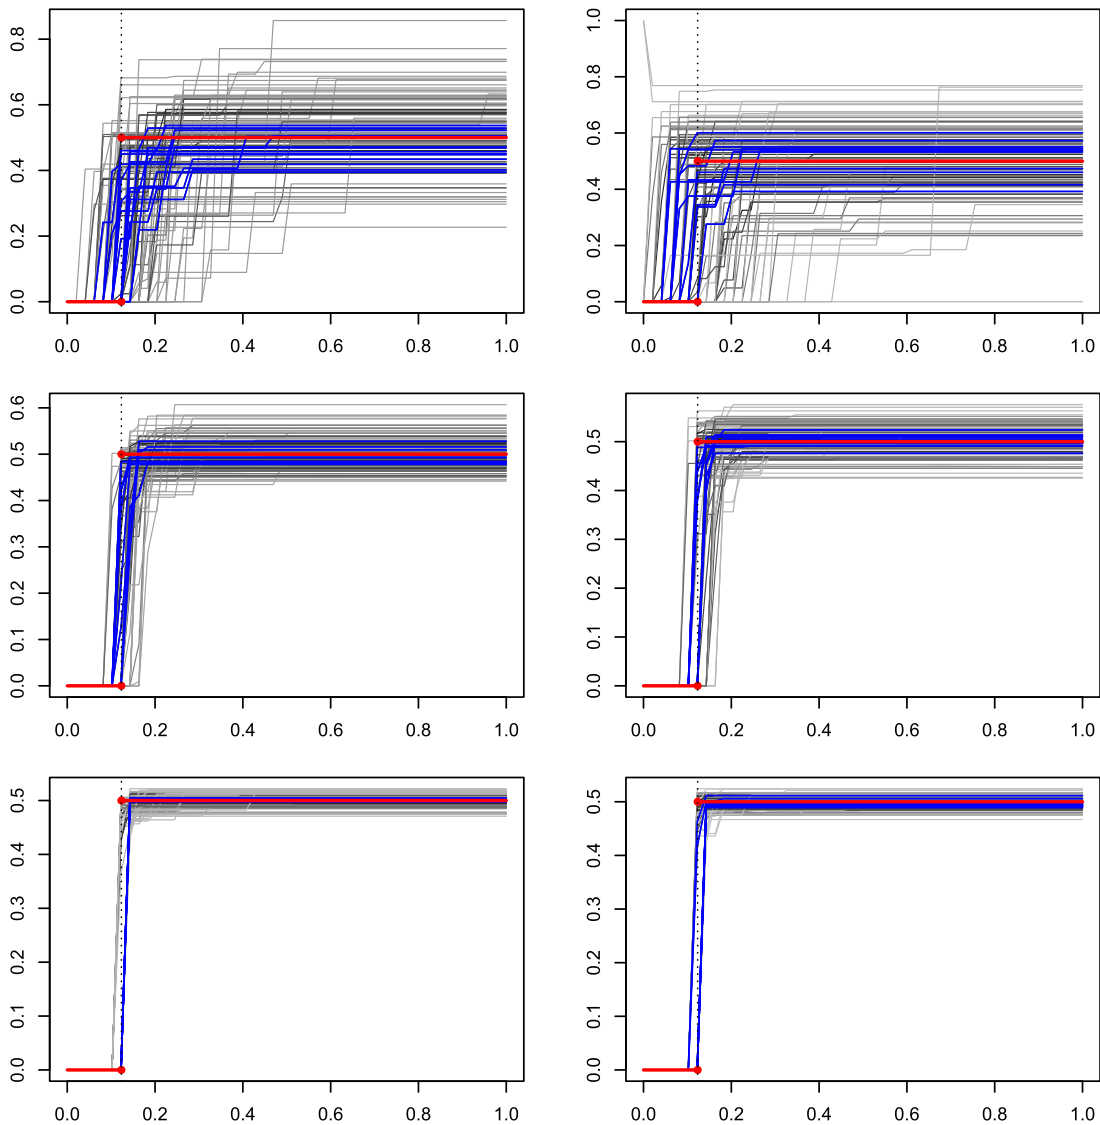


Fig. 3. Figures (rows 1, 2, and 3 are for differing sample sizes $n = 100, 1000, 10000$, respectively, columns for different assignment mechanisms A1 (left column) and A2 (right column), respectively) showing the inferred probability (y-axis) to assign Treatment 1 over all 100 replications, in dependence on the preference parameter λ (x-axis), with linear interpolation in between the points $\{0, 1/49, 2/49, \dots, 1\}$, at which the policy was actually estimated. The gray-scale coloring of the curves is chosen according to the depth of the curves, darker shades of gray corresponding to a stronger degree of typicality/centrality of the curve; see text for more explanation. The 10% of most central curves are highlighted in blue color, whereas the true argmax in dependence on λ is highlighted in red color. The vertical dashed line intersects the abscissa at $c(p) \approx 0.123$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

scenarios. In each of these scenarios we generate 100 datasets and infer the optimal empirical-success policy together with the value function for every $\lambda \in \{0, 1/49, 2/49, \dots, 1\}$. The results are shown in Figs. 3 and 4, which we shall explain and discuss in the following paragraphs.

For each of the 6 scenarios considered, Fig. 3 shows 100 curves of policies (corresponding to the 100 datasets generated) extracted for the whole range of preference parameters $\lambda \in [0, 1]$, where we interpolated linearly in between the grid-points $\{0, 1/49, 2/49, \dots, 1\}$. Note that because there are only two treatments, we can identify a policy with the corresponding probability to assign Treatment 1 (as we already did in Section 6.1). We decided to highlight the functions according to their “typicality”, which we measured through Tukey’s halfspace depth (cf. Tukey, 1975) as implemented in the R package **DepthProc** by Kosiorowski and Zawadzki (2022): the darker a gray curve is in Fig. 3, the higher its Tukey-depth (i.e., its “typicality”) among the 100 functions estimated; furthermore, the top 10% of curves in terms of depth are colored in blue. Additionally, Fig. 3 also shows the true argmax correspondence in red




Fig. 4. Figures (rows 1, 2, and 3 are for differing sample sizes $n = 100, 1000, 10000$, respectively, columns for different assignment mechanisms A1 (left column) and A2 (right column), respectively) showing the empirical value function (y-axis) over all 100 replications, in dependence on the preference parameter λ (x-axis), with linear interpolation in between the points $\{0, 1/49, 2/49, \dots, 1\}$, at which the policy was actually estimated. The gray-scale coloring of the curves is chosen according to the depth of the curves, darker shades of gray corresponding to a stronger degree of typicality/centrality of the curve; see text for more explanation. The 10% of most central curves are highlighted in blue color, whereas the true value function in dependence on λ is highlighted in red color. The vertical dashed line intersects the abscissa at $c(p) \approx 0.123$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

and the cut-off $c(3/4)$ is highlighted as a vertical dashed line (cf. Section 6.1 for additional explanation). Fig. 3 clearly illustrates the effect of a higher sample size: the larger the sample size, the closer are the estimated policy functions to the true argmax. Recalling from Section 6.1 that there is a phase-transition phenomenon concerning the optimal policy present at $c(p) = 0.123$, we note that this phenomenon is reflected quite accurately in the estimated policies for higher sample sizes, whereas for $n = 100$ there is a lot of variation concerning the value of λ at which the estimated policy changes, which is even more pronounced in scenario A2 (i.e., the second column in Fig. 3). In this context, we also note that the threshold $c(p) = 0.123$ is not an element of the grid $\{0, 1/49, 2/49, \dots, 1\}$ of λ values for which the policy actually has been estimated. Although using 0.123 as an additional grid point would have potentially improved the policy, this would give an overly optimistic impression, because the threshold depends on \mathcal{F} , which is unknown in practice. Therefore, using 0.123 as an additional grid point would have distorted the results in favor of the method studied. To increase the accuracy, one could nevertheless work with a finer grid, at higher computational costs.




Fig. 5. Average regret for sample sizes $n = 100, 1000, 10000$, respectively, assignment mechanisms A1 (left column) and A2 (right column), in dependence on the preference parameter λ (x-axis), with linear interpolation in between the points $\{0, 1/49, 2/49, \dots, 1\}$, at which the policy was actually estimated. The vertical dashed line intersects the abscissa at $c(p) \approx 0.123$.

In Fig. 4, we show the corresponding estimated value functions, with an analogous coloring scheme as in Fig. 3 and where the red curve depicts the true value function. Interestingly, the figure shows that (in the example considered) there appears to be a downward bias present in the estimation of the value function, which diminishes as sample size grows. In general, a larger sample size leads to more accurate results.

For completeness, we also plot the average regret in all 6 scenarios and for all 50 values of the preference parameter in Fig. 5. As expected from the theoretical results, the regret gets smaller as sample size increases. Furthermore, larger values of the preference parameter correspond to relatively larger values of the average regret.

7. Empirical illustrations

To showcase the methods discussed in this article in an empirical context, we shall now apply them (more precisely the policies in Section 3) to two datasets:

- data on the Pennsylvania reemployment bonus experiment underlying the analysis in Biliás (2000) (the dataset can be obtained, e.g., via the replication material for the article Chernozhukov et al. (2018) on the publisher’s website¹³);
- the entrepreneurship program dataset from Lyons and Zhang (2017) (available in the replication material to that article via the publisher’s website¹⁴), which builds the basis of the empirical illustration in Viviano and Bradic (2023).

We use the same target T (Gini-welfare) and similarity measure S (Kolmogorov-Smirnov distance) as in Section 6.2 (after rescaling the outcome to $[0, 1]$ in the first dataset). Furthermore, we used the same optimization heuristic (based on the same grid and starting value selection procedure). For both datasets we summarize the numerical computations in figures showing the following quantities in dependence on the preference parameter $\lambda \in [0, 1]$ based on linear interpolation in between the grid points

$$\max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \hat{F}_n}(\delta), \quad T(\langle \pi_n^\lambda, \hat{F}_n \rangle), \quad \text{and} \quad \max_{z \in Z} S(\langle \pi_n^\lambda, \hat{F}_n \rangle_z, \langle \pi_n^\lambda, \hat{F}_n \rangle); \tag{50}$$

i.e., the (empirical) value function, the Gini-welfare corresponding to the predicted outcome distribution $\langle \pi_n^\lambda, \hat{F}_n \rangle$ when implementing the policy π_n^λ , which (nearly) maximizes the empirical objective function $\delta \mapsto \Omega_{\lambda, \hat{F}_n}(\delta)$, and the (empirical) similarity measure at the policy π_n^λ .

¹³ URL <https://doi.org/10.1111/ectj.12097> retrieved on 25.01.2024.

¹⁴ URL <https://www.openicpsr.org/openicpsr/project/113492/version/V1/view> retrieved on 26.08.2023.




Fig. 6. Results for the Pennsylvania reemployment bonus experiment data. From left to right, the three plots show (in dependence on the preference parameter λ , which is plotted on the x -axes): the empirical value function, the predicted Gini-welfare when the optimal policy is implemented, and the corresponding similarity measure. See Eq. (50) for a precise description of the functions that are plotted. The dashed vertical line corresponds to $\lambda = 1/49$. The dotted vertical line corresponds to $\lambda = 18/49$. The dashed horizontal lines (figure in the middle) indicate the efficiency of the policy corresponding to $\lambda = 0$ minus $\beta = 0.05$ and minus $\beta(1 - c_n) \approx 0.0048$, respectively. The crossed region between those lines corresponds to the slack. The colored region above the crossed region contains the predicted Gini-welfare of the “affordable” policies; cf. the description in the text.

7.1. Pennsylvania reemployment bonus experiment

In this dataset the treatment is a bonus for unemployment insurance claimants granted if they become employed within a given period of time. We here focus on treatment group 4 as the treatment vs. a control group, cf. [Biliias \(2000\)](#) for details. In particular, it follows that $K = 2$. The total sample size then is $n = 5099$. The outcome variable measures the duration (in weeks) of the first spell of unemployment. We incorporated the following covariates into our analysis:

1. a categorical variable (with 3 categories) indicating the number of dependents of the claimant being 0, 1, or ≥ 2 .
2. a categorical variable (with 3 categories) indicating whether the claimant’s age was < 35 , between 35 and 54, or higher;
3. a categorical variable (with 3 categories) indicating whether the claimant was working in the sector of durable manufacturing, the sector of nondurable manufacturing, or in another sector.

After combining them into a single factor (and dropping one category for which there was no observation) this resulted in a factor with 26 levels, i.e., $|\mathcal{X}| = 26$.

Furthermore, we worked with the following sensitive characteristics, after combining them into a single factor with 8 levels, i.e., $|\mathcal{Z}| = 8$:

1. a dummy variable indicating whether or not the claimant is female;
2. a categorical variable (4 categories) indicating whether the claimant is Black, Hispanic, White, or other;

[Fig. 6](#) summarizes the results. The figure to the right, which shows the KS-based similarity measure (low values corresponding to higher similarity, i.e., higher fairness), shows that the difference between the distributions in the subpopulations (defined by the sensitive characteristics in the above enumeration) decreases sharply when passing from $\lambda = 0$ to $\lambda = 1/49$, whereas it remains essentially constant for $\lambda \geq 2/49$. Although there is also an initial drop (of about 0.0024) in Gini-welfare in the total population, which is shown in the center of [Fig. 6](#), the Gini-welfare keeps decreasing at a still notable degree also for $\lambda \geq 2/49$. If a DM can afford a decrease in total Gini-welfare of 0.0024, the empirical results therefore suggest to use the policy corresponding to $\lambda = 1/49$, as this already leads to a major reduction of unfairness. Higher levels of the preference parameter λ seem to mainly decrease welfare while not increasing fairness much further.

In addition to the “qualitative” or inspection-based approach to choosing the preference parameter just outlined, [Fig. 6](#) (center) also illustrates the more quantitative and budget-based approach laid out in [Section 4](#) for the example of $\beta = 0.005$. That is, when the DM deems an efficiency loss of 0.005 spent on increasing the degree of fairness affordable. In such a situation, one would, according to [\(32\)](#), introduce the slack $c_n = \sqrt{\log(n)/n} \approx 0.041$, and then choose the policy corresponding to the largest preference parameter that realizes an efficiency not smaller than the efficiency realized at the policy corresponding to $\lambda = 0$ (which equals about 0.072 in this example) minus $\beta(1 - c_n) \approx 0.0048$, resulting in 0.0667 (as opposed to just subtracting β itself from 0.072, which would result in 0.0665). This is depicted in [Fig. 6](#), where the effect of introducing some slack is also illustrated. In this case, the figure shows that (according to this procedure) the DM would use the preference parameter 18/49.

The actual assignment probabilities to the treatment category in dependence on the covariates can be seen in [Fig. 7](#), where we focus on smaller levels of λ , as this enhances readability. Given that introducing a small preference parameter of $\lambda = 1/49$ already substantially increases the fairness in the output distributions, it is interesting to investigate that region of the plot in more detail:

- in the non-manufacturing sector (i.e., figure to the left in [Fig. 7](#)), such a small change of preference parameter is related to a strong decrease in the treatment assignment probabilities in the “0 dependents & middle age” (solid blue) category, and to a




Fig. 7. Treatment assignment probabilities (y-axis) for all groups are plotted against the preference parameter λ (x-axis) for the Pennsylvania reemployment bonus experiment data. The three figures correspond to non-manufacturing (left), durable manufacturing (middle), and nondurable manufacturing (right); the colors indicate age (green indicating < 35 , blue indicating an age between 35 and 54, and black indicating a higher age than 54); the line types solid, dashed and dotted indicate the number of dependents 0, 1 and ≥ 2 , respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 8. Results for the entrepreneurship program data studied in Lyons and Zhang (2017). From left to right, the three plots show (in dependence on the preference parameter λ , which is plotted on the x-axes): the empirical value function, the predicted Gini-welfare when the optimal policy is implemented, and the corresponding similarity measure. See Eq. (50) for a precise description of the functions that are plotted. The dashed vertical line corresponds to $\lambda = 11/49$. The dotted vertical line corresponds to $\lambda = 21/49$. The dashed horizontal lines indicate the efficiency of the policy corresponding to $\lambda = 0$ minus $\beta = 0.02$ and minus $\beta(1 - c_n) \approx 0.017$, respectively. The crossed region in between those lines corresponds to the slack. The colored region above the crossed region contains the predicted Gini-welfare of the “affordable” policies; cf. the description in the text.

strong increase in the treatment assignment probabilities in the “0 dependents & young age” (solid green) and the “2 dependents & middle age” (dotted blue) categories.

- in the durable manufacturing sector (i.e., middle figure in Fig. 7) there is no comparable change in the assignment probabilities for very small λ .
- in the nondurable manufacturing sector (i.e., figure to the right in Fig. 7) a small change of preference parameter is related to a strong increase in the assignment probabilities in the “0 dependents & middle age” category.

7.2. Entrepreneurship program

The treatment variable here is an entrepreneurship training and incubation program for students and the outcome variable is subsequent entrepreneurial activity, which is a dummy variable, cf. Lyons and Zhang (2017) and Section 6 of Viviano and Bradic (2023) for more details. Also in this dataset $K = 2$. The total sample size (after eliminating observations with missing values) is $n = 335$. Following “Case 2” in Viviano and Bradic (2023) we incorporated the following covariates into our analysis:

1. a (rounded) score assigned to the candidate by the interviewer taking values in 1, 2, ..., 10 (cf. Footnote 3 in Lyons and Zhang, 2017);

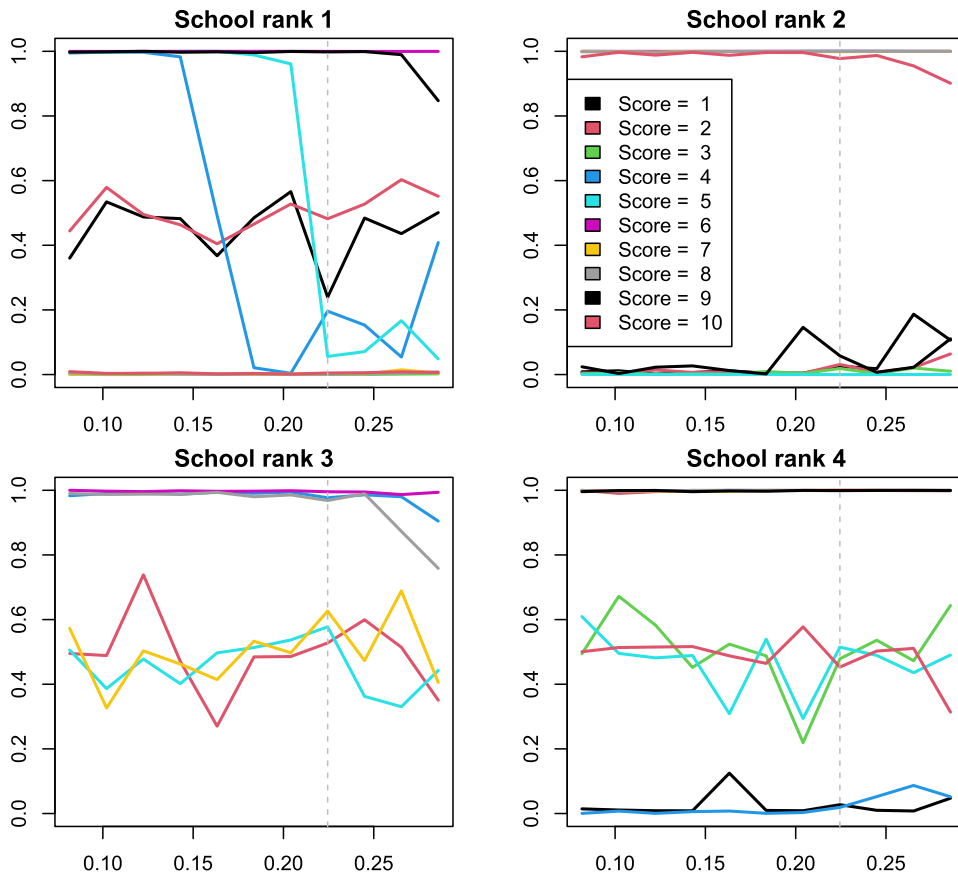


Fig. 9. Treatment assignment probabilities (y-axis) for all groups are plotted against the preference parameter λ (x-axis) for the entrepreneurship program data studied in Lyons and Zhang (2017). The four figures correspond to the school ranks and the colors indicate the score as shown in the legend. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

2. the school rank, taking values in 1, 2, 3, 4.

In principle, this would result in 40 levels after combining the two variables into a single factor. However, there is no data available for 4 of the 40 levels, hence $|\mathcal{X}| = 36$ in this application. The sensitive characteristic we worked with in this application is a dummy variable indicating whether the student is female, so that $|\mathcal{Z}| = 2$. Fig. 8 contains the results of the computations. The similarity measure reveals three substantial drops. Each of them could correspond to a possible choice of preference parameter. Unless budget constraints do not allow for this possibility, the DM could implement the policy corresponding to $\lambda = 11/49$: the KS-based similarity measure is essentially 0 for the policy corresponding to this preference parameter and hence cannot reduce substantially for larger values of the preference parameter. This leads to a “loss” in the population Gini-welfare of about 0.0124 relative to the policy corresponding to $\lambda = 0$.

Fig. 8 (center) also illustrates the more quantitative and budget-based approach lined out in Section 4 for the example of $\beta = 0.02$. That is, when the DM deems an efficiency loss of 0.02 spent on increasing the degree of fairness affordable. Here, the slack $c_n = \sqrt{\log(n)/n} \approx 0.132$, and one would choose the policy corresponding to the largest preference parameter that realizes an efficiency not smaller than the efficiency realized at the policy corresponding to $\lambda = 0$ (which equals about 0.128 in this example) minus $\beta(1 - c_n) \approx 0.017$, resulting in 0.111 (as opposed to just subtracting β from 0.128, which would result in 0.108). This is depicted in Fig. 8, where the effect of introducing some slack is also illustrated. In this example, the DM would conclude to use the preference parameter 21/49. Because using $\lambda = 21/49$ instead of $\lambda = 11/49$ (as we concluded in the previous paragraph) does not come with any notable decrease in unfairness (but decreases efficiency), a DM should perhaps stick to $\lambda = 11/49$, even though a higher preference parameter could be afforded. This application hence also shows that even when using a budget-based approach, inspecting the fairness and efficiency implications of the policies graphically is advisable for practice.

The actual assignment probabilities to the treatment in dependence on the covariates can be seen in Fig. 9, where we focus on levels of λ close to 11/49 as this enhances readability and corresponds to the interesting region of preference parameters according to the above discussion. One can observe that apart from schools ranked first and fourth, there are no fundamental changes in the assignment probabilities when λ increases from 10/49 to 11/49. In schools ranked first, one can observe that the assignment

probabilities of students with a score of 5/10 and 9/10 drop, whereas the assignment probabilities of students with a score of 4/10 increase to some degree. In schools ranked fourth the assignment probabilities of students with scores 3/10 and 5/10 increase.

8. Conclusion

We develop tools that allow a DM to choose a policy targeting a general functional (e.g., the Gini-welfare) of the population cdf, while taking into account the objective of treating protected subgroups fairly. The magnitude by which unfairness is penalized can be regulated by the DM via a preference parameter. We study ways to select this preference parameter, and show that the policy’s excess regret converges to 0 at the parametric rate. Furthermore, going beyond regret guarantees, we also establish the consistency of the policy under conditions that are satisfied in particular under the assumption that the preference parameter is selected (in an arbitrary way) from a finite set of preference parameters. We also introduce and discuss theoretical properties of a policy that applies in settings with non-discrete covariates, and illustrate our results numerically and in the context of two empirical applications.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper. The first author is an Associate Editor for Journal of Econometrics and was not involved in the editorial review or the decision to publish this article.

Appendix A. Relation to Kitagawa and Tetenov (2021) and Kock et al. (2023)

Kitagawa and Tetenov (2021) study treatment choice problems where the target functional is an element of the class of “rank-dependent” social welfare functions, cf. their Equation 1. This class includes, in particular, the Gini-welfare measure. Their results extend optimal policy design in the spirit of Kitagawa and Tetenov (2018) to a more flexible class of targets, whereas the rest of the literature typically concerns expected outcome as the target measure. Kitagawa and Tetenov (2021) do not investigate how fairness regarding subgroups of the population can be achieved, or how fairness and efficiency properties can be traded-off against each other. Subpopulation distributions and their similarity are central to our approach, but play no role in their context. We also mention here that their approach to estimating outcome distribution functions provides an important guideline in our Section 5, where we show how the assumption of discrete covariates, which is imposed in most parts of the present article, can be avoided.

The observational setting we consider in the present article is related to that in Kock et al. (2023): we also consider a pure-exploration problem in which the decision maker can target a *general* functional of the outcome distribution. In contrast to Kock et al. (2023) we

1. explicitly incorporate covariates into the decision making process to exploit observed heterogeneity in the population;
2. consider a more general objective that allows the DM to recommend a treatment (combination) that penalizes unfairness;
3. monitor how different degrees of fairness awareness affect the inferred decision, and investigate how to choose the relevant preference parameter in practice.

The first item provides a practically relevant extension to Kock et al. (2023). The second and third constitute the main contribution of the present article: we study the question how a DM can incorporate fairness considerations into the decision making process, and we give practical guidelines how the fairness-optimality trade-off can be monitored and chosen by the DM. In particular, our theory allows for data-dependent choices of the preference parameter that determines how strongly the fairness penalty enters the objective function. In contrast to Kock et al. (2023) we do not consider sequential procedures, because we here want to leave the sampling mechanism imposed on the data available to the DM unspecified.

Appendix B. Alternative regret notion: anchoring at a target preference parameter

In situations where $\hat{\lambda}_n$ estimates an underlying target λ_n^* (cf. Section 4 for a concrete example), one could *alternatively* consider anchoring the regret at the target λ_n^* itself instead of $\hat{\lambda}_n$, i.e.,

$$r(\boldsymbol{\pi}_n^{\hat{\lambda}_n}; \lambda_n^*, \mathcal{F}) = \max_{\delta \in \mathcal{S}_K} \Omega_{\lambda_n^*, \mathcal{F}}(\delta) - \Omega_{\lambda_n^*, \mathcal{F}}(\boldsymbol{\pi}_n^{\hat{\lambda}_n}).$$

This regret notion evaluates the policy $\boldsymbol{\pi}_n^{\hat{\lambda}_n}$ relative to the best one for the unknown value λ_n^* . It is straightforward to verify (cf. the proof of Lemma 4) that for a constant $C > 0$ depending on \mathcal{S} , \mathcal{T} , and \mathcal{D} it holds that

$$\left| r(\boldsymbol{\pi}_n^{\hat{\lambda}_n}; \hat{\lambda}_n, \mathcal{F}) - r(\boldsymbol{\pi}_n^{\hat{\lambda}_n}; \lambda_n^*, \mathcal{F}) \right| \leq 2 \max_{\delta \in \mathcal{S}_K} |\Omega_{\hat{\lambda}_n, \mathcal{F}}(\delta) - \Omega_{\lambda_n^*, \mathcal{F}}(\delta)| \leq C |\hat{\lambda}_n - \lambda_n^*|,$$

so that

$$\mathbb{E}^*(r(\boldsymbol{\pi}_n^{\hat{\lambda}_n}; \lambda_n^*, \mathcal{F})) \leq \mathbb{E}^*(r(\boldsymbol{\pi}_n^{\hat{\lambda}_n}; \hat{\lambda}_n, \mathcal{F})) + C \mathbb{E}^* |\hat{\lambda}_n - \lambda_n^*|.$$

This quantifies, to some extent, the (expected behavior) that for $\hat{\lambda}_n \approx \lambda_n^*$ the two regret notions behave similarly, whereas they may behave quite differently in case $\hat{\lambda}_n \not\approx \lambda_n^*$.

Section 4 investigates consistency properties for a specific choice of λ_n^* and a corresponding estimate $\hat{\lambda}_n$, under which $\hat{\lambda}_n \approx \lambda_n^*$ and hence $r(\pi_n^{\hat{\lambda}_n}; \hat{\lambda}_n, \mathcal{F}) \approx r(\pi_n^{\lambda_n^*}; \lambda_n^*, \mathcal{F})$ can be expected to hold. Nevertheless, these results do not allow one to obtain regret guarantees for $r(\pi_n^{\hat{\lambda}_n}; \lambda_n^*, \mathcal{F})$ from those of $r(\pi_n^{\lambda_n^*}; \hat{\lambda}_n, \mathcal{F})$ based on the above upper bound for two reasons:

1. we do not study rates of convergence, so that the guarantees obtained for the regret we work with, i.e., $r(\pi_n^{\hat{\lambda}_n}; \hat{\lambda}_n, \mathcal{F})$, cannot be immediately shown to carry over to $r(\pi_n^{\lambda_n^*}; \lambda_n^*, \mathcal{F})$;
2. in the example studied in Section 4, $\hat{\lambda}_n$ can only be expected to consistently underestimate the target λ_n^* , and the results in Sun (2021) suggest that this cannot be avoided, in general. Hence, $\mathbb{E}^*|\hat{\lambda}_n - \lambda_n^*|$ may not converge to 0 and may actually dominate the behavior of $r(\pi_n^{\lambda_n^*}; \lambda_n^*, \mathcal{F})$.

Appendix C. Supplementary results and discussions for Section 3

In Appendix C we maintain (without further mentioning) all assumptions and notational conventions that are maintained in Section 3.

C.1. Technical remarks concerning the empirical success policy

We here summarize some technical remarks concerning the policy stated in the beginning of Section 3.

Remark 8. The policy stated in the beginning of Section 3 allows for the possibility that the empirical version

$$\sup_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \hat{F}_n}(\delta) \tag{C.1}$$

of the maximization problem $\max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta)$ is only solved up to a given level of accuracy ε , cf. (19). While the problem $\max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta)$ is guaranteed to permit a solution by Proposition 1, there is no such guarantee for the problem in (C.1). The level of accuracy ε can be chosen by the user, and can be made arbitrarily small, balancing accuracy with computational (e.g., runtime) considerations. For the policy stated in the beginning of Section 3 to be feasible, in the sense that the set in (19) is non-empty, it is important that the supremum that is approximated is actually finite, which is guaranteed by the last item in Proposition 1.

Remark 9. Note that the policy $\pi_n^{\varepsilon, \lambda}$ stated in the beginning of Section 3 is not required to be a measurable selection of the set in (19) and hence may not be measurable. Our theoretical results take care of this possibility.

Remark 10. If decision rules are to be obtained for multiple values of λ , Step 4 in the policy stated in the beginning of Section 3 (but not the other steps) need to be repeated for all preference parameters one considers. Because the maximizer of the objective $\max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta)$ may not be unique (cf. the example discussed in Section 6.1 for an illustration), there is no guarantee that specific approximate solutions (as in (19)) to $\sup_{\delta \in \mathbb{S}_K} \Omega_{\lambda_1, \hat{F}_n}(\delta)$ and $\sup_{\delta \in \mathbb{S}_K} \Omega_{\lambda_2, \hat{F}_n}(\delta)$ are “close” (relative to ε) even if $\lambda_1 \approx \lambda_2$ and $\hat{F}_n \approx \mathcal{F}$. What is true is that the corresponding values of the optimization problem’s value will be similar (due to a continuity property established in Lemma 4 and results entailing $\hat{F}_n \approx \mathcal{F}$), i.e., $\sup_{\delta \in \mathbb{S}_K} \Omega_{\lambda_1, \hat{F}_n}(\delta) \approx \sup_{\delta \in \mathbb{S}_K} \Omega_{\lambda_2, \hat{F}_n}(\delta)$.

C.2. Value function estimation by linear interpolation

We here show how the value of the decision problem under consideration in Section 3 as a function of the preference parameter, i.e., the value function

$$\lambda \mapsto \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta), \tag{C.2}$$

which we know to be convex from Proposition 1, can be uniformly consistently estimated by suitable interpolation methods (for pointwise estimation see Remark 11 below). An estimate of the value function allows the decision maker to investigate the effect of different preference parameters on the maximal value of the objective (cf. the figures in Sections 6 and 7 and their discussion).

Due to numerical restrictions, it will typically only be possible to determine the empirical success policy $\pi^{\varepsilon_n, \lambda}$ over a finite grid of preference parameters λ , from which one then infers the value function (C.2). We now show that if the empirical success policy $\pi^{\varepsilon_n, \lambda}$ is determined for every $\lambda \in G_m := \{0, \frac{1}{m}, \frac{2}{m}, \dots, 1\}$ with $m = m_n \rightarrow \infty$, and if $\varepsilon_n \rightarrow 0$, then the value function is uniformly consistently estimated in (outer) probability by linearly interpolating the values of the empirical success policies determined for all preference parameters in the grid. The result follows from a finite-sample probability bound on

$$\sup_{\delta \in \mathbb{S}_K, \lambda \in [0,1]} \left| \Omega_{\lambda, \mathcal{F}}(\delta) - \Omega_{\lambda, \hat{F}_n}(\delta) \right|$$

that is also used in the proof of Theorem 1 (and which is given in Lemma 10 in the Appendix). To define the interpolation method properly, denote the set of continuous functions from $[0, 1] \rightarrow \mathbb{R}$ by $C([0, 1])$, let $m \in \mathbb{N}$, and let the “linear interpolation operator”

$$\text{LIP}_m : \mathbb{R}^{m+1} \rightarrow C([0, 1])$$

associate to every $v = (v_0, \dots, v_m)' \in \mathbb{R}^{m+1}$ the piecewise linear function

$$\lambda \mapsto \text{LIP}_m(\lambda; v) := v_i + \frac{v_{i+1} - v_i}{(i+1)/m - i/m}(\lambda - i/m), \text{ with } 0 \leq i < m \text{ so that } \lambda \in [i/m, (i+1)/m],$$

i.e., it holds that $\text{LIP}_m(i/m; v) = v_i$ for every $i = 0, \dots, m$, with linear interpolation in between the grid points i/m for $i = 0, \dots, m$.

Theorem 5. Let $(Y_j', X_j, Z_j)' \sim \mathcal{F} \subset \mathcal{D}$ for $j = 1, \dots, m$, $m_n \in \mathbb{N}$, and $\varepsilon_n > 0$. Then, there exists a real number $\alpha(\mathcal{F}) > 0$ such that

$$\sup_{\lambda \in [0,1]} \left| \text{LIP}_{m_n} \left(\lambda; \left[\Omega_{i/m_n, \hat{\mathcal{F}}_n}(\boldsymbol{\pi}_n^{\varepsilon_n, i/m_n}) \right]_{i=0}^{m_n} \right) - \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta) \right| \tag{C.3}$$

is bounded from above by

$$2\alpha(\mathcal{F})/m_n + \sup_{\delta \in \mathbb{S}_K, \lambda \in [0,1]} |\Omega_{\lambda, \mathcal{F}}(\delta) - \Omega_{\lambda, \hat{\mathcal{F}}_n}(\delta)| + \varepsilon_n,$$

where for every $\rho > 0$ it holds that

$$\mathbb{P}^* \left(\sup_{\delta \in \mathbb{S}_K, \lambda \in [0,1]} |\Omega_{\lambda, \mathcal{F}}(\delta) - \Omega_{\lambda, \hat{\mathcal{F}}_n}(\delta)| \geq \rho \right) \leq 4|\mathcal{Z}||\mathcal{X}|K^2 \max_{x,z,j} e^{-\frac{nq_{x,z,j}\rho^2}{128(|\mathcal{X}||\mathcal{Z}|)^2}}, \tag{C.4}$$

for $q_{x,z,i} := p_{X,Z,D}(x, z, i)$. In particular, the expression in (C.3) converges to 0 in outer probability if $m_n \rightarrow \infty$ and $\varepsilon_n \rightarrow 0$.

To prove Theorem 5 we need an auxiliary lemma:

Lemma 1. If $f : [0, 1] \rightarrow \mathbb{R}$ is Lipschitz continuous with constant $c > 0$ and $m \in \mathbb{N}$, then

$$\sup_{\lambda \in [0,1]} |\text{LIP}_m(\lambda; [f(i/m)]_{i=0}^m) - f(\lambda)| \leq 2c/m.$$

Proof. We give the proof for the case $c = 1$ (which one can then apply to $\tilde{f} := f/c$ to obtain the general statement). Fix $\lambda \in [0, 1]$ and let $i \in \{0, \dots, m-1\}$ be such that $\lambda \in [i/m, (i+1)/m]$. By definition,

$$\text{LIP}_m(\lambda; [f(i/m)]_{i=0}^m) = f(i/m) + m(f((i+1)/m) - f(i/m))(\lambda - i/m),$$

from which we obtain that $|\text{LIP}_m(\lambda; [f(i/m)]_{i=0}^m) - f(\lambda)|$ is bounded from above by

$$|f(i/m) - f(\lambda)| + |f((i+1)/m) - f(i/m)| \leq 2/m.$$

□

Proof of Theorem 5: Proposition 1 shows that the value function $\lambda \mapsto \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta)$ is continuous and convex on $[0, 1]$. With the same argument as in the proof of that proposition one can show that the value function is continuous and convex on, e.g., $[-1, 2]$. It follows that the value function is Lipschitz continuous on $[0, 1]$ with finite constant $\alpha(\mathcal{F}) > 0$, say. From Lemma 1 it hence follows that

$$\sup_{\lambda \in [0,1]} \left| \text{LIP}_{m_n} \left(\lambda; \left[\max_{\delta \in \mathbb{S}_K} \Omega_{i/m_n, \mathcal{F}}(\delta) \right]_{i=0}^{m_n} \right) - \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta) \right| \leq 2\alpha(\mathcal{F})/m_n. \tag{C.5}$$

Next, we note that (by the definition of LIP_m) the expression

$$\sup_{\lambda \in [0,1]} \left| \text{LIP}_{m_n} \left(\lambda; \left[\max_{\delta \in \mathbb{S}_K} \Omega_{i/m_n, \mathcal{F}}(\delta) \right]_{i=0}^{m_n} \right) - \text{LIP}_{m_n} \left(\lambda; \left[\Omega_{i/m_n, \hat{\mathcal{F}}_n}(\boldsymbol{\pi}_n^{\varepsilon_n, i/m_n}) \right]_{i=0}^{m_n} \right) \right| \tag{C.6}$$

coincides with

$$\max_{i=0, \dots, m} \left| \max_{\delta \in \mathbb{S}_K} \Omega_{i/m_n, \mathcal{F}}(\delta) - \Omega_{i/m_n, \hat{\mathcal{F}}_n}(\boldsymbol{\pi}_n^{\varepsilon_n, i/m_n}) \right|.$$

Fix an index $i \in \{0, \dots, m\}$ and upper bound $|\max_{\delta \in \mathbb{S}_K} \Omega_{i/m_n, \mathcal{F}}(\delta) - \Omega_{i/m_n, \hat{\mathcal{F}}_n}(\boldsymbol{\pi}_n^{\varepsilon_n, i/m_n})|$ by¹⁵

$$\left| \sup_{\delta \in \mathbb{S}_K} \Omega_{i/m_n, \mathcal{F}}(\delta) - \sup_{\delta \in \mathbb{S}_K} \Omega_{i/m_n, \hat{\mathcal{F}}_n}(\delta) \right| + \left| \sup_{\delta \in \mathbb{S}_K} \Omega_{i/m_n, \hat{\mathcal{F}}_n}(\delta) - \Omega_{i/m_n, \hat{\mathcal{F}}_n}(\boldsymbol{\pi}_n^{\varepsilon_n, i/m_n}) \right|,$$

which, by definition of the policy, is upper bounded by

$$\sup_{\delta \in \mathbb{S}_K, \lambda \in [0,1]} |\Omega_{\lambda, \mathcal{F}}(\delta) - \Omega_{\lambda, \hat{\mathcal{F}}_n}(\delta)| + \varepsilon_n.$$

Combining the just-obtained bound with (C.5), we can upper bound

$$\sup_{\lambda \in [0,1]} \left| \text{LIP}_{m_n} \left(\lambda; \left[\Omega_{i/m_n, \hat{\mathcal{F}}_n}(\boldsymbol{\pi}_n^{\varepsilon_n, i/m_n}) \right]_{i=0}^{m_n} \right) - \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta) \right|$$

by

$$2\alpha(\mathcal{F})/m_n + \sup_{\delta \in \mathbb{S}_K, \lambda \in [0,1]} |\Omega_{\lambda, \mathcal{F}}(\delta) - \Omega_{\lambda, \hat{\mathcal{F}}_n}(\delta)| + \varepsilon_n.$$

Eq. (C.4) is established as (F.19) in Lemma 10. The final statement in the theorem is now obvious. □

¹⁵ Recall from the last item in Proposition 1 that the supremum inserted is finite.

Remark 11. In a situation where one is not interested in uniform estimation of the value function, but is only interested in estimating the value function at a given (possibly data-dependent) preference parameter $\hat{\lambda}_n \in [0, 1]$ (e.g., the one that is intended to be used), one trivially obtains a “pointwise” approximation result by directly applying the upper bound given in Eq. (C.4), noting that

$$\sup_{\delta \in \mathbb{S}_K} \left| \Omega_{\hat{\lambda}_n, \mathcal{F}}(\delta) - \Omega_{\hat{\lambda}_n, \hat{\mathcal{F}}_n}(\delta) \right| \geq \rho$$

implies

$$\sup_{\delta \in \mathbb{S}_K, \lambda \in [0, 1]} \left| \Omega_{\lambda, \mathcal{F}}(\delta) - \Omega_{\lambda, \hat{\mathcal{F}}_n}(\delta) \right| \geq \rho.$$

C.3. An auxiliary result on M-estimators

The strategy we employ to establish consistency of the policy in Section 3.2 is based on the standard consistency proof for M-estimators. We need to take care of the possibility, however, that our (population) objective function evaluated at the data-driven tuning parameter, i.e., $\delta \mapsto \Omega_{\hat{\lambda}_n, \mathcal{F}}(\delta)$,

- (i) may not permit a unique optimizer (a problem that is fairly standard), and that
- (ii) the corresponding $\arg \max_{\delta \in \mathbb{S}_K} \Omega_{\hat{\lambda}_n, \mathcal{F}}(\delta)$ depends on the data-driven tuning parameter $\hat{\lambda}_n$ (a problem which substantially complicates the situation).

To exhibit the problem and to connect it to classical situations that do not suffer from (i) and (ii), cf., e.g., Chapter 5 of van der Vaart (1998), we start with the following auxiliary lemma. The statement is not formulated within the framework laid out in Section 2, but is put in a more general context. In that result θ plays the role of a target parameter that one wishes to estimate consistently by making use of the sequence of estimators $\hat{\theta}_n$, whereas γ plays the role of a tuning parameter that defines the objective function and which is selected in a data-driven way through $\hat{\gamma}_n$. Recall that all random variables are assumed to be defined on an underlying probability space $(\Omega, \mathcal{A}, \mathbb{P})$. As usual, we shall notationally suppress the dependence of a function on $\omega \in \Omega$; e.g., for a function $f : A \times \Omega \rightarrow \mathbb{R}$, say, we write $f(a)$ instead of $f(a, \omega)$ if there is no reason for confusion.

Lemma 2. Let (Θ, d) be a compact metric space (Θ non-empty), $\Gamma \subseteq \mathbb{R}$ be non-empty, and let the functions

$$Q_n : \Theta \times \Gamma \times \Omega \rightarrow \mathbb{R}, \quad n \in \mathbb{N}, \quad \text{and} \quad Q : \Theta \times \Gamma \rightarrow \mathbb{R}$$

be such that $Q(\cdot, \gamma) : \Theta \rightarrow \mathbb{R}$ is continuous for every $\gamma \in \Gamma$. For $\rho > 0$ and $\gamma \in \Gamma$ define

$$\mathbb{M}_\gamma^\rho := \{\theta' \in \Theta : d(\theta', \arg \max_{\theta \in \Theta} Q(\theta, \gamma)) \geq \rho\},$$

and let $\hat{\gamma}_n : \Omega \rightarrow \Gamma$ and $\hat{\theta}_n : \Omega \rightarrow \Theta$. Then, the following holds:

1. For every $\rho > 0$, it holds that $\hat{\theta}_n \in \mathbb{M}_{\hat{\gamma}_n}^\rho$ implies $\max_{\theta \in \Theta} Q(\theta, \hat{\gamma}_n) - \max_{\theta \in \mathbb{M}_{\hat{\gamma}_n}^\rho} Q(\theta, \hat{\gamma}_n) > 0$ and

$$\frac{2 \sup_{\theta \in \Theta} |Q_n(\theta, \hat{\gamma}_n) - Q(\theta, \hat{\gamma}_n)| + (\sup_{\theta \in \Theta} Q_n(\theta, \hat{\gamma}_n) - Q_n(\hat{\theta}_n, \hat{\gamma}_n))}{\max_{\theta \in \Theta} Q(\theta, \hat{\gamma}_n) - \max_{\theta \in \mathbb{M}_{\hat{\gamma}_n}^\rho} Q(\theta, \hat{\gamma}_n)} \geq 1. \tag{C.7}$$

2. For every $\rho > 0$, it holds that if $\mathbb{M}_\gamma^\rho \neq \emptyset$ for all $\gamma \in \Gamma$, if Γ is compact, and if $\gamma \mapsto \max_{\theta \in \Theta} Q(\theta, \gamma) - \max_{\theta \in \mathbb{M}_\gamma^\rho} Q(\theta, \gamma)$ is continuous (e.g., because $Q : \Theta \times \Gamma \rightarrow \mathbb{R}$ and $\gamma \mapsto \max_{\theta \in \mathbb{M}_\gamma^\rho} Q(\theta, \gamma)$ are continuous), then

$$\inf_{\gamma \in \Gamma} \left[\max_{\theta \in \Theta} Q(\theta, \gamma) - \max_{\theta \in \mathbb{M}_\gamma^\rho} Q(\theta, \gamma) \right] > 0. \tag{C.8}$$

3. The relation (C.8) for every $\rho > 0$, together with

$$\sup_{\theta \in \Theta} |Q_n(\theta, \hat{\gamma}_n) - Q(\theta, \hat{\gamma}_n)| \quad \text{and} \quad \sup_{\theta \in \Theta} Q_n(\theta, \hat{\gamma}_n) - Q_n(\hat{\theta}_n, \hat{\gamma}_n)$$

both converging (as $n \rightarrow \infty$) to zero in outer probability, implies consistency of $\hat{\theta}_n$, in the sense that for every $\rho > 0$ it holds that

$$\mathbb{P}^*(d(\hat{\theta}_n, \arg \max_{\theta \in \Theta} Q(\theta, \gamma)) \geq \rho) = \mathbb{P}^*(\hat{\theta}_n \in \mathbb{M}_{\hat{\gamma}_n}^\rho) \rightarrow 0.$$

Proof. Let $\rho > 0$ be given. Suppose that $\omega \in \Omega$ satisfies $\hat{\theta}_n(\omega) \in \mathbb{M}_{\hat{\gamma}_n(\omega)}^\rho$ (we carefully signify the dependence on ω of all respective quantities in this proof). Then $\mathbb{M}_{\hat{\gamma}_n(\omega)}^\rho$ cannot be empty. The set $\arg \max_{\theta \in \Theta} Q(\theta, \hat{\gamma}_n(\omega))$ is non-empty and closed due to compactness of Θ and continuity of $\theta \mapsto Q(\theta, \hat{\gamma}_n(\omega))$. Therefore, $\theta' \mapsto d(\theta', \arg \max_{\theta \in \Theta} Q(\theta, \hat{\gamma}_n(\omega)))$ is Lipschitz-continuous on Θ . It follows that $\mathbb{M}_{\hat{\gamma}_n(\omega)}^\rho$ is a closed, and hence compact, non-empty subset of Θ . Therefore, $\max_{\theta \in \mathbb{M}_{\hat{\gamma}_n(\omega)}^\rho} Q(\theta, \hat{\gamma}_n(\omega))$ exists and, using the definition of $\mathbb{M}_{\hat{\gamma}_n(\omega)}^\rho$ together with $\rho > 0$, satisfies

$$\infty > \max_{\theta \in \Theta} Q(\theta, \hat{\gamma}_n(\omega)) - \max_{\theta \in \mathbb{M}_{\hat{\gamma}_n(\omega)}^\rho} Q(\theta, \hat{\gamma}_n(\omega)) =: \Delta(\rho, \omega) > 0,$$

which establishes the first statement in Part 1. From $\hat{\theta}_n(\omega) \in \mathbb{M}_{\hat{\gamma}(\omega)}^\rho$ we therefore obtain

$$\max_{\theta \in \Theta} Q(\theta, \hat{\gamma}_n(\omega)) - Q(\hat{\theta}_n(\omega), \hat{\gamma}_n(\omega)) \geq \Delta(\rho, \omega). \tag{C.9}$$

We now note that the inequality in (C.7) is trivially satisfied if the supremum $\sup_{\theta \in \Theta} Q_n(\theta, \hat{\gamma}_n(\omega)) = \infty$, and we can therefore assume henceforth that this supremum is a real number (it obviously cannot equal $-\infty$). Adding and subtracting $\sup_{\theta \in \Theta} Q_n(\theta, \hat{\gamma}_n(\omega))$ and $Q_n(\hat{\theta}_n(\omega), \hat{\gamma}_n(\omega))$, respectively, and using

$$\left| \max_{\theta \in \Theta} Q(\theta, \hat{\gamma}_n(\omega)) - \sup_{\theta \in \Theta} Q_n(\theta, \hat{\gamma}_n(\omega)) \right| \leq \sup_{\theta \in \Theta} \left| Q(\theta, \hat{\gamma}_n(\omega)) - Q_n(\theta, \hat{\gamma}_n(\omega)) \right|,$$

we can upper bound the left-hand side in (C.9) by

$$2 \sup_{\theta \in \Theta} \left| Q(\theta, \hat{\gamma}_n(\omega)) - Q_n(\theta, \hat{\gamma}_n(\omega)) \right| + \left(\sup_{\theta \in \Theta} Q_n(\theta, \hat{\gamma}_n(\omega)) - Q_n(\hat{\theta}_n(\omega), \hat{\gamma}_n(\omega)) \right).$$

Using $\Delta(\rho, \omega) > 0$ we now obtain the inequality in (C.7).

To verify the statement in Part 2, just note that under the continuity condition imposed the infimum in (C.8) is attained at $\gamma^* \in \Gamma$, say, because Γ is assumed to be compact (and non-empty). This proves the statement, cf. also the definition of $\mathbb{M}_{\gamma^*}^\rho$. To see that the statement in parenthesis is sufficient for the continuity assumption in Part 2, we can use Berge’s maximum theorem, e.g., Theorem 17.31 in Aliprantis and Border (2006).

Part 3 is an immediate consequence of (the imposed assumptions and) Part 1. \square

Remark 12. Note that the quotient on the left-hand side of (C.7) relates the worst-case approximation error and the optimization error, i.e.,

$$\sup_{\theta \in \Theta} |Q_n(\theta, \hat{\gamma}_n) - Q(\theta, \hat{\gamma}_n)| \quad \text{and} \quad \sup_{\theta \in \Theta} Q_n(\theta, \hat{\gamma}_n) - Q_n(\hat{\theta}_n, \hat{\gamma}_n)$$

to the minimal change in the value of the objective function $Q(\cdot, \hat{\gamma}_n)$ when being ρ -separated from its set of maximizers, i.e., to

$$\max_{\theta \in \Theta} Q(\theta, \hat{\gamma}_n) - \max_{\theta \in \mathbb{M}_{\hat{\gamma}_n}^\rho} Q(\theta, \hat{\gamma}_n).$$

In principle, the optimization error can be controlled by the practitioner and is therefore often, but not always, ignored, i.e., assumed to be 0 when studying consistency of an estimator defined as the solution of an optimization problem. Under suitable assumptions (e.g., a uniform law of large numbers in case Q_n is an average), the approximation error typically converges to 0 in (outer) probability.

Remark 13. Part 3 of Lemma 2 shows that (C.8) together with convergence to zero of the approximation and optimization errors in (outer) probability is enough to ensure consistency of the estimator. The second part of that lemma shows that (C.8) follows from a set of conditions including a continuity assumption. The condition in (C.8), however, is not generally met in the problem we consider in the present article without imposing conditions on Γ , e.g., finiteness. This is illustrated in the context of the example discussed in Section 6.1 in the following remark, and implies that additional assumptions on $\hat{\lambda}_n$ are necessary to deduce a consistency statement through the approach outlined above.

Example 1. Reconsider the example discussed in Section 6.1. Given $\rho \in (0, 1/2)$, we also have to consider the maximum of the objective in (48) over the set $\mathbb{M}_{F, \lambda}^\rho$, say, of all decision rules $\eta = (\eta, 1 - \eta)'$, $\eta \in [0, 1]$, that are separated away from the argmax in d_1 -distance, i.e.,

$$\mathbb{M}_{F, \lambda}^\rho := \{ \eta \in \mathbb{S}_K : d_1(\eta, \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, F}(\delta)) \geq \rho \}.$$

Due to the specific structure of the argmax in that example, cf. (49), and continuity properties of the objective, is not difficult to see that in that example

$$0 \leq \left[\max_{\delta \in \mathbb{S}_K} \Omega_{\lambda_m, F}(\delta) - \max_{\delta \in \mathbb{M}_{F, \lambda_m}^\rho} \Omega_{\lambda_m, F}(\delta) \right] \leq \Omega_{\lambda_m, F}((0, 1)') - \Omega_{\lambda_m}((1/2, 1/2)') \rightarrow 0$$

along any sequence $\lambda_m \in [0, c(p))$ that converges to $c(p)$, showing that (C.8) is violated in this example.

Hence, the condition in (C.8) cannot be expected to hold in our framework when selecting tuning parameters from $[0, 1]$, and therefore obtaining consistency by controlling the approximation and estimation error as outlined above requires further conditions on $\hat{\lambda}_n$ that were not necessary in the context of Theorem 1.

Appendix D. Notation and some preparatory results

Recall that for an array F as in Eq. (7), i.e.,

$$F = [F^i(\cdot | x, z), p_{X, Z}(x, z)] : i = 1, \dots, K; x \in \mathcal{X}; z \in \mathcal{Z}],$$

and a given (non-empty) set of cdfs $\mathcal{H} \subseteq D_{cdf}([a, b])$ we write

$$F \sqsubset \mathcal{H} \Leftrightarrow F^i(\cdot | x, z) \in \mathcal{H} \text{ for every } i = 1, \dots, K, x \in \mathcal{X}, \text{ and } z \in \mathcal{Z}.$$

For a vector $G = (G_1, \dots, G_l)'$ of functions $G : \mathbb{R} \rightarrow \mathbb{R}$ we write

$$\|G\|_\infty := \max_{i=1}^l \|G_i\|_\infty = \max_{i=1}^l \sup_{v \in \mathbb{R}} |G_i(v)|,$$

which may be infinite. Recall that the symbol $\|\cdot\|_p$ denotes the p -norm on a Euclidean space.

For two arrays F and \tilde{F} as in Eq. (7) that satisfy $F \subset D_{cdf}([a, b])$ and $\tilde{F} \subset D_{cdf}([a, b])$, we obtain (for $x \in \mathcal{X}$ and $z \in \mathcal{Z}$ fixed, and cf. Section 2.3.1 for the definition of $F(\cdot | x)$ and similar quantities)

$$\|F(\cdot | x) - \tilde{F}(\cdot | x)\|_\infty = \max_{i=1}^K \|F^i(\cdot | x) - \tilde{F}^i(\cdot | x)\|_\infty,$$

and

$$\|F(\cdot | x, z) - \tilde{F}(\cdot | x, z)\|_\infty = \max_{i=1}^K \|F^i(\cdot | x, z) - \tilde{F}^i(\cdot | x, z)\|_\infty;$$

furthermore, we associate to an array F as in Eq. (7) the quantities

$$p_X = (p_X(x))_{x \in \mathcal{X}} \quad \text{and} \quad p_{X,Z} = (p_{X,Z}(x, z))_{x \in \mathcal{X}, z \in \mathcal{Z}}$$

and write, e.g., \tilde{p}_X or \hat{p}_X , respectively for the same quantities but corresponding to \tilde{F} or \hat{F}_n . Similarly, we associate to F (and analogously to \tilde{F} or \hat{F}_n), for every fixed $z \in \mathcal{Z}$, the vectors $p_{X|Z=z} = (p_{X|Z=z}(x))_{x \in \mathcal{X}}$ and $p_Z = (p_Z(z))_{z \in \mathcal{Z}}$.

Recall that for $\delta \in \mathbb{S}_K$ and $z \in \mathcal{Z}$ we write

$$\langle \delta, F \rangle = \sum_{x \in \mathcal{X}} \sum_{i=1}^K \delta_i(x) p_X(x) F^i(\cdot | x) \quad \text{and} \quad \langle \delta, F \rangle_z = \sum_{x \in \mathcal{X}} \sum_{i=1}^K \delta_i(x) p_{X|Z=z}(x) F^i(\cdot | x, z),$$

and denote

$$u(F, \tilde{F}) := \sum_{x \in \mathcal{X}} \|p_X(x) F(\cdot | x) - \tilde{p}_X(x) \tilde{F}(\cdot | x)\|_\infty.$$

and

$$u_z(F, \tilde{F}) := \sum_{x \in \mathcal{X}} \|p_{X|Z=z}(x) F(\cdot | x, z) - \tilde{p}_{X|Z=z}(x) \tilde{F}(\cdot | x, z)\|_\infty. \tag{D.1}$$

Remark 14. Throughout, we shall use the symbols $\langle \gamma, F \rangle$ and $\langle \gamma, F \rangle_z$ for any $\gamma : \mathcal{X} \rightarrow \mathbb{R}^K$ (not necessarily a map to \mathbb{S}_K). Note, however, that for such γ , these quantities are no longer cdfs in general.

With this notation (recalling d_1 from (5)), we obtain the following result.

Lemma 3. For two arrays F and \tilde{F} as in Eq. (7) and δ and $\tilde{\delta}$ in \mathbb{S}_K , we have

$$\|\langle \delta, F \rangle - \langle \tilde{\delta}, \tilde{F} \rangle\|_\infty \leq d_1(\delta, \tilde{\delta}) + u(F, \tilde{F}), \tag{D.2}$$

and for every $z \in \mathcal{Z}$ it holds that

$$\|\langle \delta, F \rangle_z - \langle \tilde{\delta}, \tilde{F} \rangle_z\|_\infty \leq d_1(\delta, \tilde{\delta}) + u_z(F, \tilde{F}). \tag{D.3}$$

Furthermore, it holds that

$$u(F, \tilde{F}) \leq \sum_{z \in \mathcal{Z}} p_Z(z) u_z(F, \tilde{F}) + \|p_Z - \tilde{p}_Z\|_1 \leq \max_{z \in \mathcal{Z}} u_z(F, \tilde{F}) + \|p_Z - \tilde{p}_Z\|_1. \tag{D.4}$$

Proof. We first establish (D.2). Writing

$$\langle \tilde{\delta} - \delta, \tilde{F} \rangle := \sum_{x \in \mathcal{X}} \sum_{i=1}^K (\tilde{\delta}_i(x) - \delta_i(x)) \tilde{p}_X(x) \tilde{F}^i(\cdot | x),$$

the triangle inequality delivers

$$\|\langle \delta, F \rangle - \langle \tilde{\delta}, \tilde{F} \rangle\|_\infty \leq \|\langle \delta, F \rangle - \langle \delta, \tilde{F} \rangle\|_\infty + \|\langle \tilde{\delta} - \delta, \tilde{F} \rangle\|_\infty \leq \|\langle \delta, F \rangle - \langle \delta, \tilde{F} \rangle\|_\infty + d_1(\delta, \tilde{\delta}).$$

From $\delta(x) \in \mathcal{S}_K$ for every $x \in \mathcal{X}$, we obtain

$$\|\langle \delta, F \rangle - \langle \delta, \tilde{F} \rangle\|_\infty \leq \sum_{x \in \mathcal{X}} \sum_{i=1}^K \delta_i(x) \|p_X(x) F^i(\cdot | x) - \tilde{p}_X(x) \tilde{F}^i(\cdot | x)\|_\infty \leq u(F, \tilde{F}),$$

which delivers (D.2). The inequality in (D.3) is obtained analogously. Concerning the inequalities in (D.4), we note that

$$p_X(x) F(\cdot | x) = \sum_{z \in \mathcal{Z}} p_{X,Z}(x, z) F(\cdot | x, z)$$

allows us to write

$$u(F, \tilde{F}) = \sum_{x \in \mathcal{X}} \left\| \sum_{z \in \mathcal{Z}} p_{X,Z}(x, z) \{p_{X|Z=z}(x) F(\cdot | x, z) - \tilde{p}_{X|Z=z}(x) \tilde{F}(\cdot | x, z)\} - \sum_{z \in \mathcal{Z}} \tilde{p}_{X|Z=z}(x) (\tilde{p}_Z(z) - p_Z(z)) \tilde{F}(\cdot | x, z) \right\|_\infty,$$

which (using the triangle inequality) we can upper bound by

$$\sum_{z \in \mathcal{Z}} p_Z(z) u_z(\mathcal{F}, \tilde{\mathcal{F}}) + \sum_{z \in \mathcal{Z}} \left[\sum_{x \in \mathcal{X}} \tilde{p}_{X|Z=z}(x) \right] |p_Z(z) - \tilde{p}_Z(z)|,$$

from which two inequalities in (D.4) follow. \square

Appendix E. Auxiliary results and proofs for Section 2

For the following result, we recall from Eq. (12) the definition of the objective function

$$\Omega_{\lambda, \mathcal{F}}(\delta) = (1 - \lambda) \mathbb{T}(\langle \delta, \mathcal{F} \rangle) - \lambda \max_{z \in \mathcal{Z}} \mathbb{S}(\langle \delta, \mathcal{F} \rangle_z, \langle \delta, \mathcal{F} \rangle).$$

Lemma 4. Fix $\lambda \in [0, 1]$, let $\mathcal{F} \subset \mathcal{D}$ and $\tilde{\mathcal{F}} \subset D_{cdf}([a, b])$ be arrays as in Eq. (7), and let δ and $\tilde{\delta}$ be elements of \mathbb{S}_K . Then,

$$|\Omega_{\lambda, \mathcal{F}}(\delta) - \Omega_{\lambda, \tilde{\mathcal{F}}}(\tilde{\delta})| \leq (1 + \lambda) \left[d_1(\delta, \tilde{\delta}) + \max_{z \in \mathcal{Z}} u_z(\mathcal{F}, \tilde{\mathcal{F}}) \right] + \|p_Z - \tilde{p}_Z\|_1. \tag{E.1}$$

In particular, the function $\delta \mapsto \Omega_{\lambda, \mathcal{F}}(\delta)$ is Lipschitz continuous on \mathbb{S}_K with constant $1 + \lambda$. Furthermore, the function $(\lambda, \delta) \mapsto \Omega_{\lambda, \mathcal{F}}(\delta)$ is Lipschitz continuous on $[0, 1] \times \mathbb{S}_K$.

Proof of Lemma 4. We upper bound

$$|\Omega_{\lambda, \mathcal{F}}(\delta) - \Omega_{\lambda, \tilde{\mathcal{F}}}(\tilde{\delta})| \leq (1 - \lambda)A + \lambda B,$$

where

$$A := |\mathbb{T}(\langle \delta, \mathcal{F} \rangle) - \mathbb{T}(\langle \tilde{\delta}, \tilde{\mathcal{F}} \rangle)| \leq d_1(\delta, \tilde{\delta}) + \max_{z \in \mathcal{Z}} u_z(\mathcal{F}, \tilde{\mathcal{F}}) + \|p_Z - \tilde{p}_Z\|_1,$$

the inequality follows from Assumption 4 (note that \mathcal{D} is assumed to be convex and $\mathcal{F} \subset \mathcal{D}$ holds by assumption, implying that $\langle \delta, \mathcal{F} \rangle \in \mathcal{D}$) and (D.2) together with (D.4); and

$$B := \max_{z \in \mathcal{Z}} \left| \mathbb{S}(\langle \delta, \mathcal{F} \rangle_z, \langle \delta, \mathcal{F} \rangle) - \mathbb{S}(\langle \tilde{\delta}, \tilde{\mathcal{F}} \rangle_z, \langle \tilde{\delta}, \tilde{\mathcal{F}} \rangle) \right| \\ \leq \max_{z \in \mathcal{Z}} \|\langle \delta, \mathcal{F} \rangle_z - \langle \tilde{\delta}, \tilde{\mathcal{F}} \rangle_z\|_\infty + \|\langle \delta, \mathcal{F} \rangle - \langle \tilde{\delta}, \tilde{\mathcal{F}} \rangle\|_\infty,$$

where we used Assumption 4. Using Eqs. (D.2)–(D.4), we obtain

$$B \leq 2 \left[d_1(\delta, \tilde{\delta}) + \max_{z \in \mathcal{Z}} u_z(\mathcal{F}, \tilde{\mathcal{F}}) \right] + \|p_Z - \tilde{p}_Z\|_1.$$

Combining the upper bounds on A and B just obtained, the inequality in (E.1) and the Lipschitz-continuity statement in the lemma follows. To see the last statement in the lemma, note that the upper bounds on A and B show that the functions

$$\delta \mapsto \mathbb{T}(\langle \delta, \mathcal{F} \rangle) \quad \text{and} \quad \delta \mapsto \max_{z \in \mathcal{Z}} \mathbb{S}(\langle \delta, \mathcal{F} \rangle_z, \langle \delta, \mathcal{F} \rangle)$$

are Lipschitz continuous on \mathbb{S}_K , from which the Lipschitz continuity of $(\lambda, \delta) \mapsto \Omega_{\lambda, \mathcal{F}}(\delta)$ follows. \square

Proof of Proposition 1. Lemma 4 shows that the function $(\lambda, \delta) \mapsto \Omega_{\lambda, \mathcal{F}}(\delta)$ is continuous on $[0, 1] \times \mathbb{S}_K$. Therefore, all assumptions for Berge’s maximum theorem (as in, e.g., Aliprantis and Border, 2006, Theorem 17.31) are satisfied, which establishes the first three statements in the proposition apart from the convexity statement, which is a consequence of Danskin’s theorem (cf., e.g., Proposition B.25 in Bertsekas, 1999).

To establish the last statement, let $\mathcal{F} \subset \mathcal{D}$, let $\delta \in \mathbb{S}_K$ and use Lemma 4 to obtain

$$\Omega_{\lambda, \mathcal{G}}(\delta) \leq |\Omega_{\lambda, \mathcal{G}}(\delta) - \Omega_{\lambda, \mathcal{F}}(\delta)| + \Omega_{\lambda, \mathcal{F}}(\delta) \leq 2 \max_{z \in \mathcal{Z}} u_z(\mathcal{F}, \mathcal{G}) + 2 + \Omega_{\lambda, \mathcal{F}}(\delta). \tag{E.2}$$

Taking the supremum over $\delta \in \mathbb{S}_K$ in both sides of (E.2) and using the fact already established in the first part of this proof then delivers the last statement of the proposition. \square

Appendix F. Auxiliary results and proofs for Section 3

Throughout this section, we shall sometimes abbreviate the training sample from (2) as follows

$$(\tilde{W}_j, j = 1, \dots, n) =: \tilde{W}. \tag{F.1}$$

Correspondingly, we can write

$$\pi_n^{\epsilon, \lambda}(\tilde{W}_j, j = 1, \dots, n) = \pi_n^{\epsilon, \lambda}(\tilde{W}).$$

Lemma 5. Let $\lambda \in [0, 1]$ and \mathcal{F} be an array as in Eq. (7) satisfying $\mathcal{F} \sqsubset \mathcal{D}$. Let $\pi_n^{\varepsilon, \lambda}$ be an empirical success policy as defined in Section 3. Then,

$$\begin{aligned} r(\pi_n^{\varepsilon, \lambda}; \lambda, \mathcal{F}) &\leq 2 \sup_{\delta \in \mathbb{S}_K} |\Omega_{\lambda, \mathcal{F}}(\delta) - \Omega_{\lambda, \hat{\mathcal{F}}_n}(\delta)| + \varepsilon \\ &\leq \varepsilon + 2(1 + \lambda) \max_{z \in \mathcal{Z}} u_z(\mathcal{F}, \hat{\mathcal{F}}_n) + 2\|\mathbf{p}_Z - \hat{\mathbf{p}}_Z\|_1. \end{aligned}$$

Proof. Using the last item in Proposition 1 (guaranteeing that all suprema below are finite) and (19), we write

$$\begin{aligned} r(\pi_n^{\varepsilon, \lambda}; \lambda, \mathcal{F}) &= \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta) - \sup_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \hat{\mathcal{F}}_n}(\delta) + \sup_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \hat{\mathcal{F}}_n}(\delta) - \Omega_{\lambda, \mathcal{F}}(\pi_n^{\varepsilon, \lambda}(\bar{\mathbf{W}})) \\ &\leq \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta) - \sup_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \hat{\mathcal{F}}_n}(\delta) + \Omega_{\lambda, \hat{\mathcal{F}}_n}(\pi_n^{\varepsilon, \lambda}(\bar{\mathbf{W}})) - \Omega_{\lambda, \mathcal{F}}(\pi_n^{\varepsilon, \lambda}(\bar{\mathbf{W}})) + \varepsilon \\ &\leq 2 \sup_{\delta \in \mathbb{S}_K} |\Omega_{\lambda, \mathcal{F}}(\delta) - \Omega_{\lambda, \hat{\mathcal{F}}_n}(\delta)| + \varepsilon. \end{aligned}$$

This proves the first inequality; the second now follows from (E.1) in Lemma 4. \square

From (D.1) and Lemma 5 it follows that the quantities determining the behavior of the upper bound for the regret just developed in Lemma 5 are the random variables

$$\|p_{X|Z=z}(x)F(\cdot | x, z) - \hat{p}_{X|Z=z}(x)\hat{F}(\cdot | x, z)\|_\infty, \quad \text{for } x \in \mathcal{X}, z \in \mathcal{Z}, \tag{F.2}$$

where $\hat{p}_{X|Z=z}(x) = \hat{p}_{X,Z}(x, z)/\hat{p}_Z(z)$ whenever $\hat{p}_Z(z) > 0$ and $\hat{p}_{X|Z=z}(x) = 0$ else, and the random variable

$$\|\mathbf{p}_Z - \hat{\mathbf{p}}_Z\|_1.$$

We shall next develop suitable bounds for these quantities.

Let us recall here that for every $i \in \{1, \dots, K\}$, $x \in \mathcal{X}$ and $z \in \mathcal{Z}$ we denote by $\hat{F}^i(\cdot | x, z)$ the empirical cdf based on observations in

$$\mathcal{M}_{x,z}^i := \{j = 1, \dots, n : D_j = i, X_j = x, Z_j = z\}, \tag{F.3}$$

i.e.,

$$\hat{F}^i(\cdot | x, z) = |\mathcal{M}_{x,z}^i|^{-1} \sum_{j \in \mathcal{M}_{x,z}^i} \mathbb{1}(Y_{D_j, j} \leq \cdot) = |\mathcal{M}_{x,z}^i|^{-1} \sum_{j \in \mathcal{M}_{x,z}^i} \mathbb{1}(Y_{i, j} \leq \cdot),$$

which we set equal to the cdf corresponding to point mass at b (the upper endpoint of the support of the cdfs one is working with) in case $\mathcal{M}_{x,z}^i$ is empty.

Note that $\mathcal{M}_{x,z}^i$ is a random set, i.e., actually $\mathcal{M}_{x,z}^i = \mathcal{M}_{x,z}^i(D_j, X_j, Z_j; j = 1, \dots, n)$, which we suppress notationally. In the conditioning arguments in the following proof, given fixed numbers $(d_j, x_j, z_j) \in \{1, \dots, K\} \times \mathcal{X} \times \mathcal{Z}$ for $j = 1, \dots, n$ we shall analogously write

$$\mathcal{M}_{x,z}^i(d_j, x_j, z_j; j = 1, \dots, n) = \{j = 1, \dots, n : d_j = i, x_j = x, Z_j = z\}, \tag{F.4}$$

which is a non-random subset of $\{1, \dots, n\}$ (i.e., $\mathcal{M}_{x,z}^i(d_j, x_j, z_j; j = 1, \dots, n)$ is a possible realization of $\mathcal{M}_{x,z}^i$).

Lemma 6. Let $(Y'_j, X_j, Z_j)' \sim \mathcal{F}$ for $j = 1, \dots, n$, and let $x \in \mathcal{X}$ and $z \in \mathcal{Z}$. Then,

$$\mathbb{E}(\|\hat{F}(\cdot | x, z) - F(\cdot | x, z)\|_\infty) \leq 2.75 \times \sqrt{\frac{\log(K)}{n}} \sum_{j=1}^K (q_{x,z,j})^{-1/2}, \tag{F.5}$$

where $q_{x,z,i} := \mathbb{P}(X_1 = x, Z_1 = z, D_1 = i)$ for $i = 1, \dots, K$; furthermore, for every $\rho > 0$ it holds that

$$\mathbb{P}(\|\hat{F}(\cdot | x, z) - F(\cdot | x, z)\|_\infty \geq \rho) \leq 2K \sum_{i=1}^K e^{-\frac{n\rho^2 q_{x,z,i}}{2}}. \tag{F.6}$$

Proof. We first establish (F.5). Fix $x \in \mathcal{X}$ and $z \in \mathcal{Z}$. Recall that $q_{x,z,i} > 0$ for every $i = 1, \dots, K$ due to Assumption 3. In this proof we abbreviate

$$M_i := \|\hat{F}^i(\cdot | x, z) - F^i(\cdot | x, z)\|_\infty,$$

and we define the random variable

$$\beta := \min_{i=1}^K |\mathcal{M}_{x,z}^i|.$$

Define the set $S \subseteq (\{1, \dots, K\} \times \mathcal{X} \times \mathcal{Z})^n$ via

$$S := \left\{ (\bar{\mathbf{d}}, \bar{\mathbf{x}}, \bar{\mathbf{z}}) \in (\{1, \dots, K\} \times \mathcal{X} \times \mathcal{Z})^n : \beta(\bar{\mathbf{d}}, \bar{\mathbf{x}}, \bar{\mathbf{z}}) > 0 \right\},$$

where we denoted, cf. the discussion surrounding (F.4),

$$\beta(\bar{d}, \bar{x}, \bar{z}) := \min_{i=1}^K \left| \mathcal{M}_{x,z}^i(\bar{d}_j, \bar{x}_j, \bar{z}_j; j = 1, \dots, n) \right|.$$

Use the law of total probability, $\mathbb{P}(\beta = 0) = \mathbb{P}(\mathbf{D}, \mathbf{X}, \mathbf{Z} \notin S)$, and $M_i \leq 1$, to upper bound

$$\mathbb{E}(\|\hat{F}(\cdot | x, z) - F(\cdot | x, z)\|_\infty) = \mathbb{E}(\max_{i=1}^K M_i)$$

by the sum

$$\mathbb{P}(\beta = 0) + \sum_{(\bar{d}, \bar{x}, \bar{z}) \in S} \mathbb{E}_{\bar{d}, \bar{x}, \bar{z}} \left(\max_{i=1}^K M_i \right) \mathbb{P}(\mathbf{D} = \bar{d}, \mathbf{X} = \bar{x}, \mathbf{Z} = \bar{z}), \tag{F.7}$$

where $\mathbb{E}_{\bar{d}, \bar{x}, \bar{z}}$ denotes the expectation operator conditional on the event $\{\mathbf{D} = \bar{d}, \mathbf{X} = \bar{x}, \mathbf{Z} = \bar{z}\}$, where $\mathbf{X} := (X_1, \dots, X_n)'$, $\mathbf{Z} := (Z_1, \dots, Z_n)'$ and $\mathbf{D} := (D_1, \dots, D_n)'$.

Fix a $(\bar{d}, \bar{x}, \bar{z}) \in S$, and note that $\mathbb{P}(\mathbf{D} = \bar{d}, \mathbf{X} = \bar{x}, \mathbf{Z} = \bar{z}) > 0$ by Assumptions 1 and 3. On the event $\{\mathbf{D} = \bar{d}, \mathbf{X} = \bar{x}, \mathbf{Z} = \bar{z}\} \subseteq \Omega$, the random set $\mathcal{M}_{x,z}^i$ is constant and non-empty, i.e., there exist natural numbers $j_1 < \dots < j_{|\mathcal{M}_{x,z}^i|}$ for which $\mathcal{M}_{x,z}^i = \{j_1, \dots, j_{|\mathcal{M}_{x,z}^i|}\}$, and such that we can write

$$\hat{F}^i(\cdot | x, z) = |\mathcal{M}_{x,z}^i|^{-1} \sum_{j \in \mathcal{M}_{x,z}^i} \mathbb{1}(Y_{D_j, j} \leq \cdot) = |\mathcal{M}_{x,z}^i|^{-1} \sum_{l=1}^{|\mathcal{M}_{x,z}^i|} \mathbb{1}(Y_{i, j_l} \leq \cdot),$$

so that in particular

$$M_i = \|\mathcal{M}_{x,z}^i|^{-1} \sum_{l=1}^{|\mathcal{M}_{x,z}^i|} \mathbb{1}(Y_{i, j_l} \leq \cdot) - F^i(\cdot | x, z)\|_\infty.$$

Using the i.i.d. assumption, for any real numbers $c_1, \dots, c_{|\mathcal{M}_{x,z}^i|}$ it holds that the probability

$$\mathbb{P}\left(Y_{i, j_1} \leq c_1, \dots, Y_{i, |\mathcal{M}_{x,z}^i|} \leq c_{|\mathcal{M}_{x,z}^i|} \mid \mathbf{D} = \bar{d}, \mathbf{X} = \bar{x}, \mathbf{Z} = \bar{z}\right),$$

can be written as the product

$$\prod_{l=1}^{|\mathcal{M}_{x,z}^i|} \mathbb{P}\left(Y_{i, j_l} \leq c_l \mid D_{j_l} = i, X_{j_l} = x, Z_{j_l} = z\right) = \prod_{l=1}^{|\mathcal{M}_{x,z}^i|} F^i(c_l | x, z), \tag{F.8}$$

the equality following from the conditional independence Assumption 2. That is, conditional on the event $\{\mathbf{D} = \bar{d}, \mathbf{X} = \bar{x}, \mathbf{Z} = \bar{z}\}$, the empirical cdf $\hat{F}^i(\cdot | x, z)$ is based on $|\mathcal{M}_{x,z}^i|$ i.i.d. random variables with cdf $F^i(\cdot | x, z)$. Arguing analogously to the proof of Theorem 4.3 in Kock et al. (2023) (from their Equation C.2 to C.3), but replacing the expectation operator there by the conditional expectation $\mathbb{E}_{\bar{d}, \bar{x}, \bar{z}}$, we hence obtain

$$\mathbb{E}_{\bar{d}, \bar{x}, \bar{z}} \left(\max_{i=1}^K M_i \right) \leq 1.505 \times \sqrt{\frac{\log(K)}{\beta(\bar{d}, \bar{x}, \bar{z})}},$$

which together with (F.7) and upon denoting $w_{x,z}(i) := \mathbb{P}\left(\min_{j=1}^K |\mathcal{M}_{x,z}^j| = i\right)$, establishes

$$\mathbb{E}(\|\hat{F}(\cdot | x, z) - F(\cdot | x, z)\|_\infty) \leq w_{x,z}(0) + 1.505 \sqrt{\log(K)} \sum_{i=1}^n i^{-1/2} w_{x,z}(i). \tag{F.9}$$

To establish the inequality in (F.5), we use the bound

$$w_{x,z}(i) \leq \sum_{j=1}^K \mathbb{P}\left(|\mathcal{M}_{x,z}^j| = i\right),$$

and note that for every $j = 1, \dots, K$ the random variable $|\mathcal{M}_{x,z}^j|$ is Bernoulli distributed with parameter n and success probability $q_{x,z,j} := \mathbb{P}(X_1 = x, Z_1 = z, D_1 = j)$. From the upper bound in (F.9), together with the inequality

$$(1 - p)^m \leq e^{-pm} \leq (2epm)^{-1/2} \tag{F.10}$$

(for every $p \in (0, 1]$ and every $m \in \mathbb{N}$), it follows that $\mathbb{E}(\|\hat{F}(\cdot | x, z) - F(\cdot | x, z)\|_\infty)$ is bounded from above by (recalling Assumption 3)

$$\sum_{j=1}^K \left\{ (2eq_{x,z,j}n)^{-1/2} + 1.505 \sqrt{\log(K)} \sum_{i=1}^n i^{-1/2} \mathbb{P}\left(|\mathcal{M}_{x,z}^j| = i\right) \right\}.$$

To obtain (F.5) it remains to apply Lemma 11 (with $m = 1$ and $\alpha = 1/2$) to bound

$$\sum_{i=1}^n i^{-1/2} \mathbb{P}\left(|\mathcal{M}_{x,z}^j| = i\right) \leq \sqrt{\frac{2}{q_{x,z,j}n}},$$

and to collect terms (recalling that $K \geq 2$).

Next we establish (F.6). Note that for $\rho \geq 1$ the claimed inequality is trivially true and fix $\rho \in (0, 1)$. Then, arguing as above and using analogous notation, with $\mathbb{P}_{\bar{d}, \bar{x}, \bar{z}}$ the conditional probability corresponding to $\mathbb{E}_{\bar{d}, \bar{x}, \bar{z}}$, the law of iterated expectations shows that $\mathbb{P}(\|\hat{F}(\cdot | x, z) - F(\cdot | x, z)\|_\infty \geq \rho)$ is bounded from above by

$$\mathbb{P}(\beta = 0) + \sum_{(\bar{d}, \bar{x}, \bar{z}) \in S} \mathbb{P}_{\bar{d}, \bar{x}, \bar{z}} \left(\max_{i=1}^K M_i \geq \rho \right) \mathbb{P}(D = \bar{d}, X = \bar{x}, Z = \bar{z}), \tag{F.11}$$

which by a union bound and the DKWM-inequality, cf. Massart (1990), exploiting the conditional independence relation established around (F.8), can be upper bounded by

$$\begin{aligned} & \mathbb{P}(\beta = 0) + 2K \sum_{(\bar{d}, \bar{x}, \bar{z}) \in S} \mathbb{P}(D = \bar{d}, X = \bar{x}, Z = \bar{z}) e^{-2\beta(\bar{d}, \bar{x}, \bar{z})\rho^2} \\ &= w_{x,z}(0) + 2K \sum_{i=1}^n e^{-2i\rho^2} w_{x,z}(i) \\ &\leq 2K \sum_{j=1}^K \sum_{i=0}^n e^{-2i\rho^2} \mathbb{P}(|\mathcal{M}_{x,z}^j| = i) = 2K \sum_{j=1}^K (1 - q_{x,z,j}(1 - e^{-2\rho^2}))^n, \end{aligned}$$

the second equality following from $|\mathcal{M}_{x,z}^j|$ being binomially distributed with parameter n and success probability $q_{x,z,j}$ and its moment-generating function (mgf). The inequality in (F.6) follows from Lemma 7 below. \square

Lemma 7. For every $\rho \geq 0$ and $\rho \in [0, 1]$, it holds that

$$1 - \rho(1 - e^{-2\rho^2}) \leq e^{-\rho^2/2}.$$

Proof. Combine $1 - e^{-2\rho^2} \geq \rho^2/2$ with $1 - z \leq e^{-z}$ for every $z \in \mathbb{R}$. \square

Lemma 8. For every $z \in \mathcal{Z}$, it holds that

$$\mathbb{E}(\|\hat{p}_{X|Z=z} - p_{X|Z=z}\|_1) \leq 2\sqrt{\frac{|\mathcal{X}|}{p_Z(z)n}}; \tag{F.12}$$

furthermore, for every $\rho > 0$ it holds that

$$\mathbb{P}(\|\hat{p}_{X|Z=z} - p_{X|Z=z}\|_1 \geq \rho) \leq 2|\mathcal{X}| e^{-\frac{n p_Z(z) \rho^2}{2|\mathcal{X}|^2}}; \tag{F.13}$$

cf. the discussion after (F.2) for the definition of $\hat{p}_{X|Z=z}$.

Proof. Fix $z \in \mathcal{Z}$. We first claim that for every $x \in \mathcal{X}$ we have

$$\mathbb{E}(|\hat{p}_{X|Z=z}(x) - p_{X|Z=z}(x)|) \leq p_{X|Z=z}(x)(1 - p_Z(z))^n + \frac{\sqrt{2p_{X|Z=z}(x)(1 - p_{X|Z=z}(x))}}{\sqrt{p_Z(z)n}};$$

given this inequality, (F.12) follows by summing over $x \in \mathcal{X}$ using that for non-negative real numbers $p_i, i = 1, \dots, m$, adding up to 1 it holds that $\sum_{i=1}^m \sqrt{p_i(1 - p_i)} \leq \sqrt{m - 1}$ (by the Cauchy-Schwarz inequality), (F.10), and combining the resulting terms using $(2e)^{-1/2} + \sqrt{2}\sqrt{|\mathcal{X}| - 1} \leq 2\sqrt{|\mathcal{X}|}$. To prove the claim, fix $x \in \mathcal{X}$ and note that by the law of total probability and Jensen's inequality

$$\mathbb{E}(|\hat{p}_{X|Z=z}(x) - p_{X|Z=z}(x)|) \leq \sum_{z^* \in \mathcal{Z}^n} \sqrt{\mathbb{E}((\hat{p}_{X|Z=z}(x) - p_{X|Z=z}(x))^2 | Z = z^*)} \mathbb{P}(Z = z^*).$$

Conditional on $Z = z^* \in \mathcal{Z}^n$ for a z^* such that $\{i : z_i^* = z\} \neq \emptyset$, the estimator $\hat{p}_{X|Z=z}(x)$ is the relative frequency of successes out of $|\{i : z_i^* = z\}|$ independent Bernoulli trials with success probability $p_{X|Z=z}(x)$ (use an argument similar to the one leading up to (F.8)), and therefore satisfies

$$\mathbb{E}((\hat{p}_{X|Z=z}(x) - p_{X|Z=z}(x))^2 | Z = z^*) = \frac{p_{X|Z=z}(x)(1 - p_{X|Z=z}(x))}{|\{i : z_i^* = z\}|}.$$

Because $\hat{p}_{X|Z=z}(x) := 0$ if none of the coordinates of Z equals z , we have that $\mathbb{E}(|\hat{p}_{X|Z=z}(x) - p_{X|Z=z}(x)|)$ is bounded from above by

$$p_{X|Z=z}(x)(1 - p_Z(z))^n + \sqrt{p_{X|Z=z}(x)(1 - p_{X|Z=z}(x))} \sum_{l=1}^n \frac{\mathbb{P}(|\{i : Z_i = z\}| = l)}{\sqrt{l}}.$$

By Lemma 11 (with $m = 1$ and $\alpha = 1/2$) the latter sum is not larger than $\frac{\sqrt{2}}{\sqrt{p_Z(z)n}}$.

To prove the second statement in the lemma, we start with the following observation: fix $x \in \mathcal{X}$ and $\gamma \in (0, 1]$ and use the law of total probability with a similar conditioning argument as above and Hoeffding's inequality to bound $\mathbb{P}(|\hat{p}_{X|Z=z}(x) - p_{X|Z=z}(x)| \geq \gamma)$ by

$$2 \sum_{m=0}^n e^{-2\gamma^2 m} \mathbb{P}(|\{i : Z_i = z\}| = m) = 2 \left(1 - p_Z(z)(1 - e^{-2\gamma^2})\right)^n \leq 2e^{-np_Z(z)\gamma^2/2},$$

the equality following from the mgf of the Binomial distribution and the inequality from Lemma 7. Now, we consider two cases: if $|\mathcal{X}| = 1$, then (F.13) trivially holds for $\rho \geq 1$ and it holds by the bound just developed for all $\rho \in (0, 1)$. If $|\mathcal{X}| > 1$, (F.13) trivially holds for $\rho > 2$. Fix $\rho \in (0, 2]$. Then, the above observation (noting $\rho/|\mathcal{X}| \leq 1$) gives

$$\mathbb{P}(\|\hat{p}_{X|Z=z} - p_{X|Z=z}\|_1 \geq \rho) \leq \sum_{x \in \mathcal{X}} \mathbb{P}(|\hat{p}_{X|Z=z}(x) - p_{X|Z=z}(x)| \geq \rho/|\mathcal{X}|) \leq 2|\mathcal{X}|e^{-\frac{n p_Z(z) \rho^2}{2|\mathcal{X}|^2}}.$$

□

A completely analogous argument establishes the following result (the result actually follows immediately from Lemma 8 for the special case $\mathcal{Z} = 1$ and “inverting” the roles of X and Z). We abstain from spelling out the details.

Lemma 9. *It holds that*

$$\mathbb{E}(\|\hat{p}_Z - p_Z\|_1) \leq 2\sqrt{\frac{|\mathcal{Z}|}{n}}. \tag{F.14}$$

Furthermore, for every $\rho > 0$ it holds that

$$\mathbb{P}(\|\hat{p}_Z - p_Z\|_1 \geq \rho) \leq 2|\mathcal{Z}|e^{-\frac{n\rho^2}{2|\mathcal{Z}|^2}}. \tag{F.15}$$

Lemma 10. *Let $(Y'_j, X_j, Z_j)' \sim F$ for $j = 1, \dots, n$, and let $z \in \mathcal{Z}$. Then*

$$\mathbb{E}(u_z(F, \hat{F}_n)) \leq 5.5\kappa_z \sqrt{\frac{\log(K)|\mathcal{X}|}{p_Z(z)n}}, \quad \text{where } \kappa_z := \max_x \sum_{j=1}^K \sqrt{\frac{p_{X,Z}(x, z)}{q_{x,z,j}}}, \tag{F.16}$$

with $q_{x,z,j} := \mathbb{P}(X_1 = x, Z_1 = z, D_1 = j)$ for $x \in \mathcal{X}$ and $j = 1, \dots, K$, and for every $\rho > 0$ it holds that

$$\mathbb{P}(u_z(F, \hat{F}_n) \geq \rho) \leq 2|\mathcal{X}|(K^2 + 1) \max_{x,j} e^{-\frac{nq_{x,z,j}\rho^2}{8|\mathcal{X}|^2}}. \tag{F.17}$$

If $\mathcal{F} \sqsubset \mathcal{D}$, it follows with $\kappa := \max_z \kappa_z$ that

$$\mathbb{E}^*\left(\sup_{\delta \in \mathbb{S}_K, \lambda \in [0,1]} |\Omega_{\lambda, \mathcal{F}}(\delta) - \Omega_{\lambda, \hat{F}_n}(\delta)|\right) \leq 11\kappa \sqrt{\frac{\log(K)|\mathcal{X}|}{n}} \sum_{z \in \mathcal{Z}} \frac{1}{\sqrt{p_Z(z)}} + 2\sqrt{\frac{|\mathcal{Z}|}{n}}, \tag{F.18}$$

and that for every $\rho > 0$ we have

$$\mathbb{P}^*\left(\sup_{\delta \in \mathbb{S}_K, \lambda \in [0,1]} |\Omega_{\lambda, \mathcal{F}}(\delta) - \Omega_{\lambda, \hat{F}_n}(\delta)| \geq \rho\right) \leq 4|\mathcal{Z}||\mathcal{X}|K^2 \max_{x,z,j} e^{-\frac{nq_{x,z,j}\rho^2}{128(|\mathcal{X}||\mathcal{Z}|)^2}}. \tag{F.19}$$

Proof. We start with the inequality in (F.16). Recall that $p_Z(z)$ and the $q_{x,z,j}$ appearing in the upper bound are non-zero, because of Assumption 3. From the definition of u_z in Eq. (D.1) we obtain

$$u_z(F, \hat{F}_n) \leq \left[\sum_{x \in \mathcal{X}} p_{X|Z=z}(x) \|F(\cdot | x, z) - \hat{F}(\cdot | x, z)\|_\infty \right] + \|p_{X|Z=z} - \hat{p}_{X|Z=z}\|_1. \tag{F.20}$$

Lemmas 6 and 8 together hence imply, for every $z \in \mathcal{Z}$, the bound

$$\mathbb{E}(u_z(F, \hat{F}_n)) \leq T_1 + T_2, \tag{F.21}$$

where we set

$$T_1 := 2.75 \times \sqrt{\frac{\log(K)}{n}} \sum_{x \in \mathcal{X}} \sum_{j=1}^K p_{X|Z=z}(x) (q_{x,z,j})^{-1/2} \quad \text{and} \quad T_2 := 2\sqrt{\frac{|\mathcal{X}|}{p_Z(z)n}}.$$

Recalling $q_{x,z,j} = \mathbb{P}(X_1 = x, Z_1 = z, D_1 = j)$ and using the Cauchy-Schwarz inequality we get

$$\sum_{x \in \mathcal{X}} \sum_{j=1}^K p_{X|Z=z}(x) (q_{x,z,j})^{-1/2} \leq p_Z^{-1}(z) \kappa_z \sum_{x \in \mathcal{X}} \sqrt{p_{X,Z}(x, z)} \leq \kappa_z \sqrt{|\mathcal{X}|/p_Z(z)},$$

so that

$$T_1 \leq 2.75\kappa_z \sqrt{\frac{\log(K)|\mathcal{X}|}{p_Z(z)n}}.$$

From $\kappa_z \geq 1$ it follows that

$$T_1 + T_2 \leq 5.5\kappa_z \sqrt{\frac{\log(K)|\mathcal{X}|}{p_Z(z)n}},$$

which establishes (F.16).

To verify (F.17), we first use (F.20) to obtain

$$\mathbb{P}(u_z(\mathcal{F}, \hat{\mathcal{F}}_n) \geq \rho) \leq \sum_{x \in \mathcal{X}} \mathbb{P}(\|F(\cdot | x, z) - \hat{F}(\cdot | x, z)\|_\infty \geq \rho/2) + \mathbb{P}(\|p_{X|Z=z} - \hat{p}_{X|Z=z}\| \geq \rho/2),$$

which we then estimate using the bounds in Eq. (F.6) (Lemma 6) and (F.13) (Lemma 8) to arrive at (note that $q_{x,z,j} \leq p_Z(z)$)

$$2K \sum_{x \in \mathcal{X}} \sum_{j=1}^K e^{-\frac{n\rho^2 q_{x,z,j}}{8}} + 2|\mathcal{X}| e^{-\frac{np_Z(z)\rho^2}{8|\mathcal{X}|^2}} \leq 2|\mathcal{X}|(K^2 + 1) \max_{x,j} e^{-\frac{nq_{x,z,j}\rho^2}{8|\mathcal{X}|^2}}.$$

To verify (F.18), we recall from Lemma 4 that

$$\sup_{\delta \in \mathbb{S}_K, \lambda \in [0,1]} \left| \Omega_{\lambda, \mathcal{F}}(\delta) - \Omega_{\lambda, \hat{\mathcal{F}}_n}(\delta) \right| \leq 2 \max_{z \in \mathcal{Z}} u_z(\mathcal{F}, \hat{\mathcal{F}}_n) + \|p_Z - \hat{p}_Z\|_1.$$

Therefore, the already established inequality in (F.16) together with (F.14) in Lemma 9 establishes (F.18).

To prove (F.19) note that the previous display together with the already established relation (F.17) and (F.15) in Lemma 9 establishes the bound

$$\mathbb{P}^* \left(\sup_{\delta \in \mathbb{S}_K, \lambda \in [0,1]} \left| \Omega_{\lambda, \mathcal{F}}(\delta) - \Omega_{\lambda, \hat{\mathcal{F}}_n}(\delta) \right| \geq \rho \right) \leq 2|\mathcal{X}|(K^2 + 1) \sum_{z \in \mathcal{Z}} \max_{x,j} e^{-\frac{nq_{x,z,j}\rho^2}{128|\mathcal{X}|^2}} + 2|\mathcal{Z}| e^{-\frac{n\rho^2}{8|\mathcal{Z}|^2}},$$

which can further be upper bounded by

$$2|\mathcal{X}|(K^2 + 1)|\mathcal{Z}| \max_{x,z,j} e^{-\frac{nq_{x,z,j}\rho^2}{128|\mathcal{X}|^2}} + 2|\mathcal{Z}| e^{-\frac{n\rho^2}{8|\mathcal{Z}|^2}} \leq 4|\mathcal{Z}||\mathcal{X}|K^2 \max_{x,z,j} e^{-\frac{nq_{x,z,j}\rho^2}{128(|\mathcal{X}|+|\mathcal{Z}|)^2}}.$$

□

Proof of Theorem 1: Lemma 5 shows for every $\lambda \in [0, 1]$ that

$$r(\pi_n^{\varepsilon, \lambda}; \lambda, \mathcal{F}) \leq 2 \sup_{\delta \in \mathbb{S}_K} |\Omega_{\lambda, \mathcal{F}}(\delta) - \Omega_{\lambda, \hat{\mathcal{F}}_n}(\delta)| + \varepsilon.$$

We therefore obtain

$$\mathbb{E}^*(r(\pi_n^{\varepsilon, \hat{\lambda}_n}; \hat{\lambda}_n, \mathcal{F})) \leq \mathbb{E}^* \left(\sup_{\lambda \in [0,1]} r(\pi_n^{\varepsilon, \lambda}; \lambda, \mathcal{F}) \right) \leq 2\mathbb{E}^* \left(\sup_{\delta \in \mathbb{S}_K, \lambda \in [0,1]} |\Omega_{\lambda, \mathcal{F}}(\delta) - \Omega_{\lambda, \hat{\mathcal{F}}_n}(\delta)| \right) + \varepsilon,$$

which by (F.18) is bounded from above by

$$22\sqrt{\frac{\log(K)|\mathcal{X}|}{n}} \kappa \sum_{z \in \mathcal{Z}} \frac{1}{\sqrt{p_Z(z)}} + 4\sqrt{\frac{|\mathcal{Z}|}{n}} + \varepsilon. \tag{F.22}$$

Inserting the definition of κ delivers the bound in (21).

The inequality in (22), follows immediately from the first display in this proof together with (F.19). □

Proof of Theorem 2: Fix $\rho > 0$. Recall from (26) that for every $\lambda \in [0, 1]$ we defined

$$c_{\mathcal{F}}(\rho, \lambda) = \begin{cases} \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta) - \max_{\delta \in \mathbb{M}_{\mathcal{F}, \lambda}^\rho} \Omega_{\lambda, \mathcal{F}}(\delta) & \text{if } \mathbb{M}_{\mathcal{F}, \lambda}^\rho \neq \emptyset, \\ 0 & \text{else.} \end{cases} \tag{F.23}$$

In the following steps, we show that $\lambda \mapsto c_{\mathcal{F}}(\rho, \lambda)$ is Borel-measurable on $[0, 1]$; if not mentioned otherwise *all references to theorems, results, and pages in the following items 1–5 are to Aliprantis and Border (2006)*, the terminology of which we shall also use without repeating their definitions:

1. By our Lemma 4, the function $(\lambda, \delta) \mapsto \Omega_{\lambda, \mathcal{F}}(\delta)$ is continuous on $[0, 1] \times \mathbb{S}_K$. Hence, it is a Carathéodory Function (cf. Definition 4.50); the unit interval is throughout equipped with its Borel sigma algebra. As a consequence of Part 3 of the “Measurable Maximum Theorem” 18.19, the nonempty-valued correspondence $\lambda \mapsto \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta)$ is measurable on $[0, 1]$ (and hence weakly measurable by Lemma 18.2).
2. Due to Theorem 18.5 the distance function

$$(\lambda, \gamma) \mapsto d_1 \left(\gamma, \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta) \right)$$

associated with the nonempty-valued weakly measurable $\lambda \mapsto \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta)$ is a Carathéodory function.

3. Decomposing $\mathbb{M}_{\mathcal{F}, \lambda}^\rho$ as

$$\{ \gamma \in \mathbb{S}_K : d_1(\gamma, \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta)) > \rho \} \cup \{ \gamma \in \mathbb{S}_K : d_1(\gamma, \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, \mathcal{F}}(\delta)) - \rho = 0 \}$$

and using Lemmas 18.7 and 18.8 together with the union part of 18.4 shows that $\lambda \mapsto \mathbb{M}_{\mathcal{F}, \lambda}^\rho$ is a measurable correspondence.

4. The correspondence $\lambda \mapsto \check{\mathbb{M}}_{F,\lambda}^\rho$ defined via

$$\lambda \mapsto \begin{cases} \mathbb{M}_{F,\lambda}^\rho & \text{if } \mathbb{M}_{F,\lambda}^\rho \neq \emptyset \\ \mathbb{S}_K & \text{else.} \end{cases}$$

is weakly measurable, which is evident from Definition 18.1, and has non-empty and compact values (for the latter compare the proof of Lemma 2).

5. We can write

$$c_F(\rho, \lambda) = \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda,F}(\delta) - \max_{\delta \in \check{\mathbb{M}}_{F,\lambda}^\rho} \Omega_{\lambda,F}(\delta).$$

Part 1 in the ‘‘Measurable Maximum Theorem’’ 18.19 proves that $\lambda \mapsto \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda,F}(\delta)$ and $\lambda \mapsto \max_{\delta \in \check{\mathbb{M}}_{F,\lambda}^\rho} \Omega_{\lambda,F}(\delta)$ are measurable from which measurability of $\lambda \mapsto c_F(\rho, \lambda)$ follows.

To establish the inequality in (27), we use Part 1 of Lemma 2 with $(\Theta, d) = (\mathbb{S}_K, d_1)$, $\Gamma = \Lambda$, $Q_n = \Omega_{\cdot, \hat{F}_n}(\cdot)$, $Q = \Omega_{\cdot, F}(\cdot)$, $\hat{\theta}_n = \pi_n^{\varepsilon_n, \hat{\lambda}_n}$, and $\hat{\gamma}_n = \hat{\lambda}_n$. All assumptions in Part 1 of Lemma 2 are satisfied, the continuity condition regarding Q was verified in Lemma 4. Fix $\zeta > 0$, recall the definition of $\mathbb{M}_{F,\hat{\lambda}_n}^\rho$ from (25), and note that the second condition in

$$\varepsilon_n + 6\zeta \leq c_F(\rho, \hat{\lambda}_n) \quad \text{and} \quad \pi_n^{\varepsilon_n, \hat{\lambda}_n} \in \mathbb{M}_{F,\hat{\lambda}_n}^\rho \tag{F.24}$$

implies via the first part of Lemma 2 and the definition of $\pi_n^{\varepsilon_n, \hat{\lambda}_n}$ that

$$2 \sup_{\delta \in \mathbb{S}_K} |\Omega_{\hat{\lambda}_n, \hat{F}_n}(\delta) - \Omega_{\hat{\lambda}_n, F}(\delta)| + \varepsilon_n \geq c_F(\rho, \hat{\lambda}_n) \quad \text{and} \quad c_F(\rho, \hat{\lambda}_n) > 0,$$

so that an application of Lemma 4 then implies

$$4 \max_{z \in \mathcal{Z}} u_z(F, \hat{F}_n) + 2\|\mathbf{p}_Z - \hat{\mathbf{p}}_Z\|_1 + \varepsilon_n \geq c_F(\rho, \hat{\lambda}_n) \quad \text{and} \quad c_F(\rho, \hat{\lambda}_n) > 0.$$

Hence, the two conditions in (F.24) together imply

$$\frac{2}{3} \max_{z \in \mathcal{Z}} u_z(F, \hat{F}_n) + \frac{1}{3} \|\mathbf{p}_Z - \hat{\mathbf{p}}_Z\|_1 \geq \zeta \quad \text{and} \quad c_F(\rho, \hat{\lambda}_n) > 0.$$

We can thus upper bound

$$\begin{aligned} \mathbb{P}^* \left(\pi_n^{\varepsilon_n, \hat{\lambda}_n} \in \mathbb{M}_{F,\hat{\lambda}_n}^\rho \right) &\leq \mathbb{P} \left(\max_{z \in \mathcal{Z}} u_z(\hat{F}_n, F) \geq \zeta \right) + \mathbb{P} \left(\|\hat{\mathbf{p}}_Z - \mathbf{p}_Z\|_1 \geq \zeta \right) \\ &\quad + \mathbb{P}^* \left(0 < c_F(\rho, \hat{\lambda}_n) < \varepsilon_n + 6\zeta \right), \end{aligned} \tag{F.25}$$

noting that the first two events in the upper bound are measurable (due to the specific form of $u_z(\hat{F}_n, F)$), and that the event in the third probability is measurable whenever $\hat{\lambda}_n$ is so (as a composition of two measurable functions) from what has been shown above. Lemma 10 and a union bound shows that the first probability in the upper bound in (F.25) is not larger than

$$2|\mathcal{X}|(K^2 + 1) \sum_{z \in \mathcal{Z}} \max_{x,j} e^{-\frac{nq_{x,z,j}\zeta^2}{8|\mathcal{X}|^2}},$$

Lemma 9 shows that

$$\mathbb{P} \left(\|\hat{\mathbf{p}}_Z - \mathbf{p}_Z\|_1 \geq \zeta \right) \leq 2|\mathcal{Z}|e^{-\frac{n\zeta^2}{2|\mathcal{Z}|^2}},$$

which establishes (27). The consistency statement after (27) follows from (27). It remains to prove the last statement in the theorem. To this end, note that if $\Lambda = \{\lambda_1^*, \dots, \lambda_M^*\}$, for a fixed $M \in \mathbb{N}$, we can trivially bound

$$\mathbb{P}^* \left(0 < c_F(\rho, \hat{\lambda}_n) < \varepsilon_n + 6\zeta_n \right) \leq \sum_{i=1}^M \mathbb{1} \left(0 < c_F(\rho, \lambda_i^*) < \varepsilon_n + 6\zeta_n \right),$$

which converges to 0 for any $\zeta_n \rightarrow 0$ whenever $\varepsilon_n \rightarrow 0$. \square

Proof of Proposition 2: We start with (29). Suppose that

$$d_1 \left(\pi_n, \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\hat{\lambda}_n, F}(\delta) \right) \xrightarrow{\mathbb{P}^*} 0.$$

We first note that if $\arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda^*, F}(\delta) = \mathbb{S}_K$, there is nothing to show, and we are done. Hence, we assume from now on that $\arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda^*, F}(\delta) \neq \mathbb{S}_K$. Recall from (25) that

$$\mathbb{M}_{F,\lambda}^\rho := \{ \gamma \in \mathbb{S}_K : d_1(\gamma, \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, F}(\delta)) \geq \rho \},$$

and fix $\rho > 0$ small enough such that $\mathbb{M}_{F, \lambda^*}^\rho \neq \emptyset$ (such a $\rho > 0$ exists, because the closed set $\arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda^*, F}(\delta)$, cf. Proposition 1, is a strict subset of \mathbb{S}_K). Because the map

$$\gamma \mapsto d_1(\gamma, \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, F}(\delta))$$

is continuous (as the distance to a non-empty closed set, cf. also Part 1 of Proposition 1), the set $\mathbb{M}_{F, \lambda^*}^\rho$ is closed, so that its complement in \mathbb{S}_K is an open neighborhood of $\arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, F}(\delta)$. Proposition 1 established upper hemicontinuity, cf. Definition 17.2 in Aliprantis and Border (2006), of the correspondence

$$\lambda \mapsto \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, F}(\delta),$$

so that there exists an $\varepsilon = \varepsilon(\rho) > 0$ such that $\Lambda_\varepsilon := [\lambda^* - \varepsilon, \lambda^* + \varepsilon] \cap [0, 1] \subseteq [0, 1]$ and such that

$$\bigcup_{\lambda \in \Lambda_\varepsilon} \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda, F}(\delta) \subseteq \mathbb{S}_K \setminus \mathbb{M}_{F, \lambda^*}^\rho. \tag{F.26}$$

The (outer) probability of the event $\hat{\lambda}_n \notin \Lambda_\varepsilon$ goes to 0 as $n \rightarrow \infty$. Hence, we shall assume that $\hat{\lambda}_n \in \Lambda_\varepsilon$ in the following. Given a non-empty and compact subset A of a metric space (V, d_V) , we can associate to every $v \in V$ an element $\Pi_A v \in A$, say, that satisfies $d_V(v, A) = \inf_{w \in A} d_V(v, w) = d_V(v, \Pi_A v)$. Denoting

$$\pi_n^* := \Pi_{\arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda^*, F}(\delta)} \pi_n, \quad \hat{\pi}_n := \Pi_{\arg \max_{\delta \in \mathbb{S}_K} \Omega_{\hat{\lambda}_n, F}(\delta)} \pi_n,$$

as well as

$$\hat{\pi}_n^* := \Pi_{\arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda^*, F}(\delta)} \left[\Pi_{\arg \max_{\delta \in \mathbb{S}_K} \Omega_{\hat{\lambda}_n, F}(\delta)} \pi_n \right] = \Pi_{\arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda^*, F}(\delta)} \hat{\pi}_n,$$

we hence see that

$$d_1\left(\pi_n, \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\lambda^*, F}(\delta)\right) = d_1(\pi_n, \pi_n^*) \leq d_1(\pi_n, \hat{\pi}_n^*) \leq d_1(\pi_n, \hat{\pi}_n) + d_1(\hat{\pi}_n, \hat{\pi}_n^*).$$

The expression $d_1(\pi_n, \hat{\pi}_n)$ equals $d_1(\pi_n, \arg \max_{\delta \in \mathbb{S}_K} \Omega_{\hat{\lambda}_n, F}(\delta))$ by definition, and thus converges to 0 in (outer) probability by assumption. The expression $d_1(\hat{\pi}_n, \hat{\pi}_n^*)$ is bounded from above by ρ due to $\hat{\lambda}_n \in \Lambda_\varepsilon$ and Eq. (F.26). Because $\rho > 0$ was arbitrary, this proves the first statement in the theorem.

To prove the second statement, note that the finiteness of Λ and the assumed convergence in outer probability implies $\mathbb{P}^*(\hat{\lambda}_n \neq \lambda_n^*) \rightarrow 0$ as $n \rightarrow \infty$, which obviously proves the result. \square

Proof of Theorem 3: We first note that by Assumption 4 and Eqs. (D.2) and (D.4), we have

$$\begin{aligned} \max_{\lambda \in \Lambda} |\Delta_n(\lambda, F) - \Delta_n(\lambda, \hat{F}_n)| &\leq |\langle \pi_n^{\varepsilon_n, 0}, F \rangle - \langle \pi_n^{\varepsilon_n, 0}, \hat{F}_n \rangle| + \max_{\lambda \in \Lambda} |\langle \pi_n^{\varepsilon_n, \lambda}, F \rangle - \langle \pi_n^{\varepsilon_n, \lambda}, \hat{F}_n \rangle| \\ &\leq 2u(F, \hat{F}_n) \leq 2 \max_{z \in \mathcal{Z}} u_z(F, \hat{F}_n) + 2\|\rho_Z - \hat{\rho}_Z\|_1. \end{aligned}$$

It hence follows from Lemmas 10 and 9 and a union bound that for every $\rho > 0$ it holds that

$$\mathbb{P}^*\left(\max_{\lambda \in \Lambda} |\Delta_n(\lambda, F) - \Delta_n(\lambda, \hat{F}_n)| \geq \rho\right) \leq 2|\mathcal{X}|(K^2 + 1) \sum_{z \in \mathcal{Z}} \max_{x, j} e^{-\frac{nq_{x, z, j} \rho^2}{128|\mathcal{X}|^2}} + 2|\mathcal{Z}|e^{-\frac{n\rho^2}{32|\mathcal{Z}|^2}}.$$

In particular, recalling that $c_n = \sqrt{\log(n)/n}$, we obtain (recalling Assumption 3) for every $\alpha > 0$ that

$$\mathbb{P}^*\left(\max_{\lambda \in \Lambda} |\Delta_n(\lambda, F) - \Delta_n(\lambda, \hat{F}_n)| \geq \alpha c_n\right) \rightarrow 0. \tag{F.27}$$

To prove (33), just note that by definition it always holds that

$$\Delta_n(\hat{\lambda}_n^*(\beta, \hat{F}_n), \hat{F}_n) \leq \beta(1 - c_n),$$

and that $\hat{\lambda}_n^*(\beta, \hat{F}_n) > \lambda_n^*(\beta, F)$ implies $\Delta_n(\hat{\lambda}_n^*(\beta, \hat{F}_n), F) > \beta$. It hence follows that both sets of which outer probabilities are taken in (33) are contained in the set defined by the relation $\max_{\lambda \in \Lambda} |\Delta_n(\lambda, F) - \Delta_n(\lambda, \hat{F}_n)| \geq \beta c_n$. The outer probability of the latter event converges to 0 by (F.27). This verifies (33).

Next, we need to verify that under the additional condition that for some $\alpha > 0$ it holds that

$$\mathbb{P}^*(\beta - \Delta_n(\lambda_n^*(\beta, F), F) \leq [\alpha + \beta]c_n) \rightarrow 0,$$

we have

$$\hat{\lambda}_n^*(\beta, \hat{F}_n) - \lambda_n^*(\beta, F) \xrightarrow{\mathbb{P}^*} 0; \tag{F.28}$$

the convergence in (35) then directly follows from the second part of Proposition 2, noting that the consistency requirement to the left in (29) is satisfied as a consequence of Theorem 2. To show (F.28), because Λ is finite, we need to verify that

$$\mathbb{P}^*(\hat{\lambda}_n^*(\beta, \hat{F}_n) \neq \lambda_n^*(\beta, F)) \rightarrow 0.$$

Due to sub-additivity of outer probabilities and because we already know that $\mathbb{P}^*(\hat{\lambda}_n^*(\beta, \hat{F}_n) > \lambda_n^*(\beta, F)) \rightarrow 0$, it suffices to verify that

$$\mathbb{P}^*(\hat{\lambda}_n^*(\beta, \hat{F}_n) < \lambda_n^*(\beta, F)) \rightarrow 0.$$

To this end, due to (F.27), it suffices to show that for an $\alpha > 0$, which we also assume to satisfy (34), we have

$$\mathbb{P}^*\left(\hat{\lambda}_n^*(\beta, \hat{F}_n) < \lambda_n^*(\beta, F) \text{ and } \max_{\lambda \in \Lambda} |\Delta_n(\lambda, F) - \Delta_n(\lambda, \hat{F}_n)| < \alpha c_n\right) \rightarrow 0.$$

Note that $\hat{\lambda}_n^*(\beta, \hat{F}_n) < \lambda_n^*(\beta, F)$ implies $\Delta_n(\lambda_n^*(\beta, F), \hat{F}_n) > \beta(1 - c_n)$, which together with

$$\max_{\lambda \in \Lambda} |\Delta_n(\lambda, F) - \Delta_n(\lambda, \hat{F}_n)| < \alpha c_n$$

gives

$$\beta - \Delta_n(\lambda_n^*(\beta, F), F) \leq [\alpha + \beta]c_n,$$

the outer probability of which converges to 0 by assumption. \square

Lemma 11. *If the random variable B follows a Binomial distribution with parameters $n \in \mathbb{N}$ and $p \in (0, 1]$, then, for every $\alpha \in (0, 1]$ and every $m \in \mathbb{N} \cup \{0\}$,*

$$\mathbb{E}(B^{-\alpha m} \mathbb{1}\{B > 0\}) \leq \left[\frac{2^{m(m+1)/2} [1 - (1-p)^{n+m}]}{p^m \prod_{i=1}^m (n+i)} \right]^\alpha \leq \frac{2^{\alpha m(m+1)/2}}{(pn)^{\alpha m}}; \tag{F.29}$$

where for $m = 0$ we set $\prod_{i=1}^m (n+i) = 1$.

Proof. The second inequality in (F.29) is trivial. By Jensen’s inequality we only need to establish the first inequality in (F.29) for the special case where $\alpha = 1$. Furthermore, for $m = 0$ that inequality is easily seen to be an equality. It therefore remains to verify the first inequality in (F.29) for $\alpha = 1$ and for every $m \in \mathbb{N}$, $n \in \mathbb{N}$, and $p \in (0, 1]$. Fix such a triple m, n , and p . For every $k \in \mathbb{N}$, we shall write B_k for a binomially distributed random variable with parameters k and p , and note that

$$\begin{aligned} \mathbb{E}(B_n^{-m} \mathbb{1}\{B_n > 0\}) &\leq 2^m \mathbb{E}((B_n + 1)^{-m} \mathbb{1}\{B_n > 0\}) \\ &= 2^m \sum_{i=1}^n \frac{1}{(i+1)^m} \binom{n}{i} p^i (1-p)^{n-i}, \end{aligned}$$

which we can equivalently write (an analogous observation was used in Rempala (2004), cf. the last display on their p.262) as

$$\begin{aligned} &\frac{2^m}{p(n+1)} \sum_{i=1}^n \frac{1}{(i+1)^{m-1}} \binom{n+1}{i+1} p^{i+1} (1-p)^{n-i} \\ &= \frac{2^m}{p(n+1)} \sum_{i=2}^{n+1} \frac{1}{i^{m-1}} \binom{n+1}{i} p^i (1-p)^{n+1-i} \\ &\leq \frac{2^m}{p(n+1)} \mathbb{E}(B_{n+1}^{-(m-1)} \mathbb{1}\{B_{n+1} > 0\}). \end{aligned}$$

Summarizing, we have shown that

$$\mathbb{E}(B_n^{-m} \mathbb{1}\{B_n > 0\}) \leq \frac{2^m}{p(n+1)} \mathbb{E}(B_{n+1}^{-(m-1)} \mathbb{1}\{B_{n+1} > 0\}),$$

which we can iterate to obtain the first inequality in (F.29) from

$$\mathbb{E}(B_{n+m}^0 \mathbb{1}\{B_{n+m} > 0\}) = 1 - (1-p)^{n+m}.$$

\square

Appendix G. Proofs for Section 5

Throughout this section, all assumptions listed in Section 5 are maintained.

Lemma 12. $\mathbb{E}(f_{\delta,c,z}(Y_{D_j}, X_j, Z_j, D_j)) = \langle \delta, F \rangle_z(c)$ for every $\delta \in \Pi$, $c \in \mathbb{R}$, and $z \in \mathcal{Z}$.

Proof of Lemma 12: Fix $\delta \in \Pi$, $c \in \mathbb{R}$, and $z \in \mathcal{Z}$. We abbreviate $\tilde{Y}_j = Y_{D_j}$ throughout. Note that

$$0 \leq f_{\delta,c,z}(\tilde{Y}_j, X_j, Z_j, D_j) \leq \max_{i=1,\dots,K} \frac{1}{p_Z(z) e_i(X_j, Z_j)}.$$

Hence, as a consequence of the maintained assumptions that $p_Z(z) > 0$ and $\mathbb{E}(e_i^{-1}(X_j, Z_j)) < \infty$ for every $i = 1, \dots, K$, $f_{\delta,c,z}(\tilde{Y}_j, X_j, Z_j, D_j)$ is integrable. Thus, by the law of iterated expectations, we can write

$$\mathbb{E}(f_{\delta,c,z}(\tilde{Y}_j, X_j, Z_j, D_j)) = \sum_{i=1}^K \mathbb{E}[\mathbb{E}(\mathbb{1}(D_j = i) f_{\delta,c,z}(\tilde{Y}_j, X_j, Z_j, D_j) \mid X_j, Z_j)].$$

The inner expectation coincides (almost everywhere) with

$$\frac{\delta_i(X_j)}{p_Z(z)e_i(X_j, z)} \times \mathbb{1}(Z_j = z) \times \mathbb{E}(\mathbb{1}(Y_{i,j} \leq c, D_j = i) \mid X_j, Z_j),$$

which, making use of the conditional independence assumption, reduces to

$$\sum_{i=1}^K \frac{\delta_i(X_j)}{p_Z(z)} \times \mathbb{1}(Z_j = z) \times F^i(c \mid X_j, Z_j),$$

so that (upon taking the outer expectation) we arrive (cf. (37)) at

$$\mathbb{E}(f_{\delta,c,z}(Y_j, X_j, Z_j, D_j)) = \langle \delta, F \rangle_z(c).$$

□

Proof of Theorem 4: We first establish the bound for $r(\pi_n^{\varepsilon, \hat{\lambda}_n}; \hat{\lambda}_n, F)$: Using Lemma 13 below, we can write

$$r(\pi_n^{\varepsilon, \hat{\lambda}_n}; \hat{\lambda}_n, F) = \sup_{\delta \in \Pi} \Omega_{\hat{\lambda}_n, F}(\delta) - \sup_{\delta \in \Pi} \hat{\Omega}_{\hat{\lambda}_n}(\delta) + \sup_{\delta \in \Pi} \hat{\Omega}_{\hat{\lambda}_n}(\delta) - \Omega_{\hat{\lambda}_n, F}(\pi_n^{\varepsilon, \hat{\lambda}_n}),$$

which, by the definition of the policy, is bounded from above by

$$\sup_{\delta \in \Pi} \Omega_{\hat{\lambda}_n, F}(\delta) - \sup_{\delta \in \Pi} \hat{\Omega}_{\hat{\lambda}_n}(\delta) + \hat{\Omega}_{\hat{\lambda}_n}(\pi_n^{\varepsilon, \hat{\lambda}_n}) - \Omega_{\hat{\lambda}_n, F}(\pi_n^{\varepsilon, \hat{\lambda}_n}) + \varepsilon.$$

This is bounded from above by $2 \sup_{\delta \in \Pi} |\Omega_{\hat{\lambda}_n, F}(\delta) - \hat{\Omega}_{\hat{\lambda}_n}(\delta)| + \varepsilon$. It is easy to verify that the inequalities in Assumption 4 remain valid for F and \bar{F} in the closure of $\mathcal{D} \subseteq D_{cdf}([a, b])$ (equipped with the supremum metric); cf. Remark 2.3 and Footnote 4 in Kock et al. (2022). Making use of this observation and Assumption 4, together with the assumption $F \sqsubset_{\Pi} \mathcal{D}$ (implying also that $\langle \delta, F \rangle = \sum_{z \in \mathcal{Z}} p_Z(z) \langle \delta, F \rangle_z$ is an element of the closure of \mathcal{D} , because the latter is assumed to be convex), we can upper bound

$$2 \sup_{\delta \in \Pi} |\Omega_{\hat{\lambda}_n, F}(\delta) - \hat{\Omega}_{\hat{\lambda}_n}(\delta)| + \varepsilon \leq 6 \max_{z \in \mathcal{Z}} \sup_{\delta \in \Pi} \|\widehat{\langle \delta, F \rangle}_z - \langle \delta, F \rangle_z\|_{\infty} + \varepsilon. \tag{G.1}$$

Fix $\delta \in \Pi$, $c \in \mathbb{R}$, and $z \in \mathcal{Z}$, and note that (by the definition of $\widehat{\langle \delta, F \rangle}_z$ and $M_{a,b}$)

$$|\widehat{\langle \delta, F \rangle}_z(c) - \langle \delta, F \rangle_z(c)| \leq |n^{-1} \sum_{j=1}^n f_{\delta,c,z}(Y_{D_j,j}, X_j, Z_j, D_j) - \langle \delta, F \rangle_z(c)| \leq \|P_n - P\|_{\mathcal{G}_{\Pi}}, \tag{G.2}$$

where the last inequality made use of Lemma 12.

We now establish the bound for $r(\pi_n^{\varepsilon, \hat{\lambda}_n}; \hat{\lambda}_n, F)$: Clearly, we only need to consider the case where all denominators in (44) are nonzero. In this case, for later use, we first observe that

$$|f_{\delta,c,z}(y, x, z^*, d) - \tilde{f}_{\delta,c,z}(y, x, z^*, d)| \leq \left| \frac{1}{p_Z(z)e_d(x, z)} - \frac{1}{\hat{p}_Z(z)\hat{e}_d(x, z)} \right|, \tag{G.3}$$

from which it follows that, for every $\delta \in \Pi$ and every $z \in \mathcal{Z}$, we have

$$\left| \widehat{\langle \delta, F \rangle}_z - \langle \delta, F \rangle_z \right| \leq n^{-1} \sum_{j=1}^n \left| \frac{1}{p_Z(z)e_{D_j}(X_j, z)} - \frac{1}{\hat{p}_Z(z)\hat{e}_{D_j}(X_j, z)} \right|. \tag{G.4}$$

Similarly as in the steps leading up to (G.1) one obtains the bound

$$r(\pi_n^{\varepsilon, \hat{\lambda}_n}; \hat{\lambda}_n, F) \leq \|\hat{p}_Z - p_Z\|_1 + 6 \max_{z \in \mathcal{Z}} \sup_{\delta \in \Pi} \|\widehat{\langle \delta, F \rangle}_z - \langle \delta, F \rangle_z\|_{\infty} + \varepsilon.$$

Adding $\pm \widehat{\langle \delta, F \rangle}_z$ inside the supremum norm, the triangle inequality together with Eqs. (G.2) and (G.4) concludes the proof. □

Lemma 13. *The quantities $\sup_{\lambda \in [0,1]} \sup_{\delta \in \Pi} \Omega_{F,\lambda}(\delta)$ and $\sup_{\lambda \in [0,1]} \sup_{\delta \in \Pi} \hat{\Omega}_{\lambda}(\delta)$ are finite, regardless of whether the latter expression is based on $\widehat{\langle \delta, F \rangle}$ and $\widehat{\langle \delta, F \rangle}_z$ (defined in (40)) or on the quantities $\langle \delta, F \rangle$ and $\langle \delta, F \rangle_z$ (defined in (42)).*

Proof. Because S is non-negative, we have, for every $\lambda \in [0, 1]$ and every $\delta \in \Pi$, that

$$\Omega_{F,\lambda}(\delta) \leq |T(\langle \delta, F \rangle)| + \max_{z \in \mathcal{Z}} S(\langle \delta, F \rangle_z, \langle \delta, F \rangle);$$

and analogously

$$\hat{\Omega}_{\lambda}(\delta) \leq |T(\widehat{\langle \delta, F \rangle})| + \max_{z \in \mathcal{Z}} S(\widehat{\langle \delta, F \rangle}_z, \widehat{\langle \delta, F \rangle}),$$

with an analogous bound holding true when $\hat{\Omega}_{\lambda}(\delta)$ is based on $\widehat{\langle \delta, F \rangle}_z$ and $\widehat{\langle \delta, F \rangle}$. Because $\langle \delta, F \rangle$, $\langle \delta, F \rangle_z$, $\widehat{\langle \delta, F \rangle}$, $\widehat{\langle \delta, F \rangle}_z$, $\widehat{\langle \delta, F \rangle}$, and $\widehat{\langle \delta, F \rangle}_z$ are elements of $D_{cdf}([a, b])$, it suffices to verify that there exist real numbers C and C' , say, such that for any $G \in D_{cdf}([a, b])$ and $G_z \in D_{cdf}([a, b])$, $z \in \mathcal{Z}$, it holds that

$$|T(G)| \leq C \quad \text{and} \quad S(G_z, G) \leq C'.$$

To this end, fix a $H \in \mathcal{D}$. By Assumption 4,

$$|T(G)| \leq |T(H)| + \|H - G\|_\infty \leq |T(H)| + 1,$$

and similarly (using that $S(H, H) = 0$)

$$S(G_z, G) = S(G_z, G) - S(H, H) \leq \|G_z - H\|_\infty + \|G - H\|_\infty \leq 2.$$

□

References

- Aliprantis, C.D., Border, K.C., 2006. *Infinite Dimensional Analysis*. Springer.
- Athey, S., Wager, S., 2021. Policy learning with observational data. *Econometrica* 89 (1), 133–161.
- Bechavod, Y., Ligett, K., 2017. Penalizing unfairness in binary classification. [arXiv:1707.00044](https://arxiv.org/abs/1707.00044).
- Berk, R., Heidari, H., Jabbari, S., Joseph, M., Kearns, M., Morgenstern, J., Neel, S., Roth, A., 2017. A convex framework for fair regression. [arXiv:1706.02409](https://arxiv.org/abs/1706.02409).
- Bertsekas, D.P., 1999. *Nonlinear Programming*. Belmont, MA: Athena Scientific. 2nd edition.
- Bhattacharya, D., Dupas, P., 2012. Inferring welfare maximizing treatment assignment under budget constraints. *J. Econom.* 167 (1), 168–196.
- Biliyas, Y., 2000. Sequential testing of duration data: the case of the pennsylvania ‘reemployment bonus’ experiment. *J. Appl. Econom.* 15 (6), 575–594.
- Bitler, M.P., Gelbach, J.B., Hoynes, H.W., 2006. What mean impacts miss: distributional effects of welfare reform experiments. *Am. Econ. Rev.* 96 (4), 988–1012.
- Chakravarty, S.R., 2009. *Inequality, Polarization and Poverty*. Springer.
- Chen, Z., Chen, S., Qi, Z., Chen, X., Yang, Z., 2025. Quantile-optimal policy learning under unmeasured confounding [arXiv:2506.07140](https://arxiv.org/abs/2506.07140).
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., Robins, J., 2018. Double/debiased machine learning for treatment and structural parameters. *Econom. J.* 21 (1), C1–C68.
- Corbett-Davies, S., Goel, S., 2018. The measure and mismeasure of fairness: a critical review of fair machine learning. [arXiv:1808.00023](https://arxiv.org/abs/1808.00023).
- Cowell, F., 2011. *Measuring Inequality*. Oxford University Press.
- Cowgill, B., Tucker, C.E., 2019. Economics, fairness and algorithmic bias. *J. Econ. Perspect.*
- Dehejia, R.H., 2005. Program evaluation as a decision problem. *J. Econom.* 125 (1), 141–173.
- Fang, E.X., Wang, Z., Wang, L., 2023. Fairness-oriented learning for optimal individualized treatment rules. *J. Am. Stat. Assoc.* 118, 1733–1746.
- Frauen, D., Melnychuk, V., Feuerriegel, S., 2023. Fair off-policy learning from observational data. [arXiv:2303.08516](https://arxiv.org/abs/2303.08516).
- Giné, E., Nickl, R., 2015. *Mathematical Foundations of Infinite-Dimensional Statistical Models*. Cambridge University Press.
- Guggenberger, P., Mehta, N., Pavlov, N., 2026. Minimax regret treatment rules with finite samples when a quantile is the object of interest. [arXiv:2601.03428](https://arxiv.org/abs/2601.03428).
- Hirano, K., Porter, J.R., 2009. Asymptotics for statistical treatment rules. *Econometrica* 77 (5), 1683–1701.
- Hirano, K., Porter, J.R., 2020. Asymptotic analysis of statistical decision rules in econometrics. In: Durlauf, S.N., Hansen, L.P., Heckman, J.J., Matzkin, R.L. (Eds.), *Handbook of Econometrics, Volume 7A*. Elsevier. Vol. 7 of *Handbook of Econometrics*, pp. 283–354.
- Kamishima, T., Akaho, S., Asoh, H., Sakuma, J., 2012. Fairness-aware classifier with prejudice remover regularizer. In: Flach, P.A., De Bie, T., Cristianini, N. (Eds.), *Machine Learning and Knowledge Discovery in Databases*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 35–50.
- Kamishima, T., Akaho, S., Sakuma, J., 2011. Fairness-aware learning through regularization approach. In: 2011 IEEE 11th International Conference on Data Mining Workshops, pp. 643–650. <https://doi.org/10.1109/ICDMW.2011.83>
- Kitagawa, T., Tetenov, A., 2018. Who should be treated? Empirical welfare maximization methods for treatment choice. *Econometrica* 86 (2), 591–616.
- Kitagawa, T., Tetenov, A., 2021. Equality-minded treatment choice. *J. Bus. Econ. Stat.* 39, 561–574.
- Kleinberg, J., Ludwig, J., Mullainathan, S., Rambachan, A., 2018. Algorithmic fairness. *AEA Pap. Proc.* 108, 22–27. <https://doi.org/10.1257/pandp.20181018>
- Kock, A.B., Preinerstorfer, D., Veliyev, B., 2022. Functional sequential treatment allocation. *J. Am. Stat. Assoc.* 117, 1311–1323.
- Kock, A.B., Preinerstorfer, D., Veliyev, B., 2023. Treatment recommendation with distributional targets. *J. Econom.* 234, 624–646.
- Kosiorowski, D., Zawadzki, Z., 2022. DepthProc An R Package for Robust Exploration of Multidimensional Economic Phenomena.
- Liese, F., Miescke, K.J., 2008. *Statistical Decision Theory*. Springer.
- Lyons, E., Zhang, L., 2017. The impact of entrepreneurship programs on minorities. *Am. Econ. Rev.* 107 (5), 303–307.
- Manski, C.F., 2004. Statistical treatment rules for heterogeneous populations. *Econometrica* 72 (4), 1221–1246.
- Manski, C.F., Tetenov, A., 2016. Sufficient trial size to inform clinical practice. *Proc. Nat. Acad. Sci.* 113 (38), 10518–10523.
- Massart, P., 1990. The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality. *Ann. Probab.* 18 (3), 1269–1283.
- Mbakop, E., Tabord-Meehan, M., 2021. Model selection for treatment choice: penalized welfare maximization. *Econometrica* 89 (2), 825–848.
- Pessach, D., Shmueli, E., 2022. A review on fairness in machine learning. *ACM Comput. Surv.* 55 (3), 1–44.
- Rempala, G.A., 2004. Asymptotic factorial powers expansions for binomial and negative binomial reciprocals. *Proc. Am. Math. Soc.* 132 (1), 261–272.
- Rostek, M., 2010. Quantile maximization in decision theory. *Rev. Econ. Stud.* 77 (1), 339–371.
- Sen, A., 1976. Poverty: an ordinal approach to measurement. *Econometrica* 44, 219–231.
- Stoye, J., 2009. Minimax regret treatment choice with finite samples. *J. Econom.* 151 (1), 70–81.
- Stoye, J., 2012. Minimax regret treatments with covariates or with limited validity of experiments. *J. Econom.* 166 (1), 138–156.
- Sun, L., 2021. Empirical welfare maximization with constraints. [arXiv:2103.15298](https://arxiv.org/abs/2103.15298).
- Tan, X., Qi, Z., Seymour, C., Tang, L., 2022. Rise: robust individualized decision learning with sensitive variables. *Adv. Neural Inf. Process. Syst.* 35, 19484–19498.
- Tukey, J.W., 1975. Mathematics and the picturing of data. In: *Proceedings of the International Congress of Mathematicians (Vancouver, B.C., 1974)*, Vol. 2. *Canad. Math. Congr.*, Montreal, QC, pp. 523–531.
- van der Vaart, A., 1998. *Asymptotic Statistics*. Cambridge University Press.
- van der Vaart, A.W., Wellner, J., 1996. *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer.
- Viviano, D., Bradic, J., 2023. Fair policy targeting. *J. Am. Stat. Assoc.* 119, 730–743.