Minimax regret treatment rules with finite samples when a quantile is the object of interest*

Patrik Guggenberger
Department of Economics
Pennsylvania State University

Nihal Mehta
Department of Economics
Pennsylvania State University

Nikita Pavlov
Department of Economics
Pennsylvania State University

First Version: July, 2023; Revised: June 17, 2024

Abstract

Consider the setup in which a policymaker is informed about the population by a finite sample and based on that sample has to decide whether or not to apply a certain treatment to the population. We work out finite sample minimax regret treatment rules under various sampling schemes when outcomes are restricted onto the unit interval. In contrast to Stoye (2009) where the focus is on maximization of expected utility the focus here is instead on a particular quantile of the outcome distribution. We find that in the case where the sample consists of a fixed number of untreated and a fixed number of treated units, any treatment rule is minimax regret optimal. The same is true in the case of random treatment assignment in the sample with any assignment probability and in the case of testing an innovation when the known quantile of the untreated population equals 1/2. However if the known quantile exceeds 1/2 then never treating is the unique optimal rule and if it is smaller than 1/2 always treating is optimal. We also consider the case where a covariate is included.

*We would like to thank Antonio Galvao, seminar participants at (in chronological order) UC Davis, UC Riverside, Colorado Boulder, SUNY Stony Brook, McGill, München, Konstanz, Frankfurt, Heidelberg, and conference participants at the Rochester conference in Econometrics for helpful comments.

Keywords: Finite sample theory, minimax regret, quantile, statistical decision theory, treatment assignments, treatment choice

JEL codes: C44

1 Introduction

Consider the setup in which a policymaker is informed about the population by a finite sample and based on that sample has to decide whether or not to apply a certain treatment to the population or whether to potentially randomize treatment assignment. This paper is adding to a short but growing literature on finding treatment rules, i.e. measurable mappings from the sample to the unit interval, that have finite-sample optimality properties.

In most of the literature, the focus is on the expected outcome under the chosen treatment rule.[1] Given that typically there is no treatment rule that is uniformly best over all possible joint distributions for (Y_0, Y_1), where Y_0 and Y_1 denote random variables for outcomes without and with treatment, respectively, one has to resort to other criteria for optimality. One option is to consider a prior over the space of joint distributions and maximize expected outcome for this particular prior. Another option is to focus on admissible treatment rules but that criterion typically does not single out an individual treatment rule, see Manski and Tetenov (2023) and Montiel Olea, Qiu, and Stoye (2023). Alternatively, one might consider finding a treatment rule that maximizes minimal expected outcome where the minimum is taken over all joint distributions of (Y_0, Y_1). However, if there exists a distribution that assigns the minimal possible values in the shared domains of Y_0 and Y_1 with probability one then any treatment rule is going to be optimal according to this criterion and therefore also the "max-min" approach is not pinning down a unique rule. For that reason, instead, the so-called "minimax regret" criterion is often adopted that determines treatment rules that minimize the maximal regret where regret is defined as the difference between the largest expected outcome that could be achieved for any treatment rule and the expected outcome under the treatment rule under consideration, and the maximum is taken with respect to all possible distributions for (Y_0, Y_1). Stoye (2009) derives minimax regret rules in finite samples for the case of two treatments under various sampling schemes, namely matched pairs, random

[1] See e.g. Manski (2004), Manski and Tetenov (2007), Stoye (2007, 2009, 2012), Tetenov (2012), Masten (2023), Montiel Olea, Qiu, and Stoye (2023), Yata (2023), Chen and Guggenberger (2024), Kitagawa, Lee, and Qiu (2024), and additional references in these papers. Hirano and Porter (2009), Kitagawa and Tetenov (2018), and Christensen, Moon, and Schorfheide (2023) and many other references therein also use the minimax regret criterion but consider an asymptotic, rather than a finite sample, framework. This literature is inspired by the classical work of Wald (1950). In a recent paper Manski and Tetenov (2023) study several potential features of the state-dependent distribution of loss (rather than just its expectation) that a decision rule generates across potential samples.

sampling, and testing an innovation, and furthermore provides near-uniqueness results.

In this project we are interested in a setup where rather than expected outcome, the policymaker is concerned about the α-quantile of the outcome (for a given $\alpha \in [0,1]$). For example, one might be interested in median income (case $\alpha = 1/2$) or a minimal education achievement for school kids, "no child left behind" (case $\alpha = 0$). See Manski (1988) who studies the "quantile utility" model whose predictions, unlike the expected utility model, are invariant under ordinal transformations of utility. Subsequently, Rostek (2010) axiomatizes quantile preference and, recently, Manski and Tetenov (2023) discuss various deviations from mean loss. Also see Chambers (2009).[2]

As the main contribution, in this paper, we derive minimax regret treatment rules δ in finite samples when an α-quantile is the focus of interest and outcomes Y_0 and Y_1 take values in the unit interval. Somewhat surprisingly, we show that under various sampling schemes *all* treatment rules are minimax regret, namely in the case i) when the sample consists of a fixed number of treated and a fixed number of untreated units (that is, unbalanced panels are allowed for) and in the case ii) under random assignment with arbitrary treatment assignment probability in the sample equal to $p \in (0,1)$. In both cases i) and ii) maximal regret equals 1 for any treatment rule.

On the other hand, in the case iii) "testing an innovation", that is, the case where only data on Y_1 is observed and the quantile $q_{s,\alpha}(Y_0)$ of Y_0 is known, where the subscript s denotes the joint distribution of (Y_0, Y_1), if $q_{s,\alpha}(Y_0) > 1/2$ then $\delta \equiv 0$ is the unique minimax rule, if $q_{s,\alpha}(Y_0) < 1/2$ then $\delta \equiv 1$ is a minimax rule, and finally, if $q_{s,\alpha}(Y_0) = 1/2$ then any treatment rule δ is minimax regret; in each case, for minimax treatment rules δ we obtain the formula $\min\{q_{s,\alpha}(Y_0), 1 - q_{s,\alpha}(Y_0)\}$ for the maximal regret.

Obviously, like for the case where expected outcomes are the focus, also here the max-min criterion is not informative. Namely, as long as the joint distribution of (Y_0, Y_1) can be chosen as $q_{s,\alpha}(Y_0) = q_{s,\alpha}(Y_1) = 0$, that is, the worst possible outcome under the given restriction that $Y_0, Y_1 \in [0,1]$, any treatment rule would be max-min optimal. In contrast though, to the case where expected outcomes are the focus, for quantiles also minimax regret is not informative for sampling designs i) and ii) and also for iii) when $q_{s,\alpha}(Y_0) = 1/2$.

The typical strategy in the extant literature for determining minimax regret rules when the focus is on expected outcomes is via a Nash equilibrium approach in a fictitious zero sum game in which the policymaker plays against an antagonistic nature whose payoff equals regret. To establish that a particular treatment rule δ is minimax regret one attempts to

[2]Related (but in a non-finite-sample setup) Qi, Cui, Liu, and Pang (2019) and Qi, Pang, and Liu (2023) consider optimal decision rules based on the conditional value at risk (CVaR) measure. Quantile preferences have been attracting growing interest in the literature, e.g. De Castro, Galvao, and Ota (2023) consider a model in which an economic agent maximizes the discounted value of a stream of future α-quantile utilities.

guess a "state of nature" s , allowed to be mixed strategy (over a finite number of states), called a *least favorable distribution*, for which the pair (δ, s) constitutes a Nash equilibrium. Existence of such a pair (δ, s) implies that δ is indeed a minimax regret rule, see for example Berger (1985), Stoye (2009), and Chen and Guggenberger (2024). Often, in a first step, one restricts outcomes to be Bernoulli and finds a minimax rule in this simplified setup and then, in a second step, uses the so-called coarsening approach to tackle the general case, see e.g. Cucconi (1968), Gupta and Hande (1992), and Schlag (2003, 2006). When nature picks a state of the world trying to inflict high regret it faces the trade-off that on the one hand a high differential between expected outcomes with and without treatment is needed but on the other hand the more different the distributions of Y_0 and Y_1 are the easier the policymaker can tell them apart using the sample information.

The proof structure for the main results in this paper differs from the one just described. Namely for cases i) and ii) we show that for any treatment rule δ one can find a state of nature $s = s_\delta$ such that s inflicts the highest possible regret, namely 1. That insight is sufficient to establish that all rules are minimax regret and that maximal regret equals 1 for all treatment rules. It is noteworthy and remarkable that, despite the trade-off just described, nature is powerful enough to inflict maximal regret on the policymaker. Even more surprising the conclusion under i) and ii) continue to be true even if nature is restricted to only Bernoulli distributions. The proof for part iii) relies on the main insight that if $\delta \neq 0$ then there exists a state of nature such that regret equals the known α-quantile of Y_0. For that to be true it is sufficient for nature to have discrete distributions, supported on $N+1$, points at its disposal.

We start off with the case with no covariates and then allow for a discrete-valued covariate in the model. In the latter case, a treatment rule maps a sample onto treatment probabilities for each of the K possible outcomes of the covariate. Not surprisingly, given the results without a covariate, we find that in cases i), ii), and case iii) with known quantile equal to 1/2, again, each treatment rule is minimax regret, while in case iii) when the known quantile is different from 1/2, no-data rules are minimax regret.

We include a small finite sample simulation study in the case of "testing an innovation" where we simulate regret of various treatment rules, namely the empirical success rule and several no-data rules. The study corroborates our theoretical findings about minimax regret treatment rules and the formulas for maximal regret that we derive.

The remainder of the paper is organized as follows. Section 2 introduces the theoretical setup and contains results when no covariate is included. Namely, Subsection 2.1 derives minimax regret treatment rules under various sampling schemes and also provides a brief discussion of minimax regret treatment rules when certain restrictions are imposed on the states of nature. Subsection 2.2 simulates the performance of various treatment rules in finite samples in the case of testing an innovation. Finally, Section 3 derives minimax regret treatment rules in the case where a covariate is included in the model.

2 Theoretical Setup Without Covariates

A policymaker has to decide whether or not to assign treatment to the population after being informed about the population by a finite sample. The setup is very similar to the one in Stoye (2009) with one key modification. Namely, rather than focussing on mean outcomes, here we are concerned with the α-quantile of the outcome distribution.

As in Stoye (2009) we assume there is an uncountable population whose members each have deterministic response functions both in the case without treatment and under treatment, but for the policymaker, who cannot distinguish between members of the population, (potential) outcomes become random variables Y_0 and Y_1. Throughout the paper potential outcomes for untreated/treated individuals are restricted to the unit interval[3]

$$Y_0, Y_1 \in S := [0, 1], \tag{2.1}$$

where S is assumed known to the policymaker. At first, different members of the population are all identical to the decision maker. Later, we will consider the case where a covariate is included and treatment assignment can be made conditional on the realization of the covariate.

By $\mathbb{S}$ we denote the set of "states of the world" s where a s denotes a possible joint probability distribution over the potential outcomes for Y_0 and Y_1. Unless otherwise stated $\mathbb{S}$ is unrestricted and contains all possible joint distributions for Y_0 and Y_1 on S^2. Upon observing a sample $w_N = (t_i, Y_{t_i,i})_{i=1,\ldots,N}$ of size N the task for the decision maker is to choose

$$\delta(w_N) \in [0, 1], \tag{2.2}$$

which denotes the probability with which treatment is assigned. Namely, which treatment the policymaker assigns is determined as an independent draw of the Bernoulli random variable $B = B(\delta(w_N)) \in \{0, 1\}$ that equals 1 with probability $\delta(w_N)$ and is assumed to be independent of all other random objects. Therefore, the setup here allows for randomized treatment rules.

Let $\alpha \in [0, 1]$. Given a state of the world s and a statistical treatment rule δ the objective

[3] We also briefly discuss the case where $Y_0, Y_1 \in (0, 1)$.

function for the decision maker is

$$u(\delta, s) = q_{s,\alpha}(Y_{B(\delta(w))}), \tag{2.3}$$

where $Y_{B(\delta(w))}$ denotes random outcomes generated when the treatment rule δ is used, and $q_{s,\alpha}(Y_{B(\delta(w))})$ denotes the α-quantile of $Y_{B(\delta(w))}$ when the state of the world is $s \in \mathbb{S}$. In particular, when the treatment rule $\delta(w)$ equals 0 (or 1) then with probability 1 $Y_{B(\delta(w))}$ equals Y_0 (or Y_1).

By definition, an α-quantile of a scalar valued random variable X with domain D is any number $q \in D$ that satisfies

$$P(X \leq q) \geq \alpha \text{ and } P(X \geq q) \geq 1 - \alpha. \tag{2.4}$$

Clearly then, any α-quantile of Y_0, Y_1 and $Y_{B(\delta(w))}$ is an element of [0,1]. In general, this definition does not lead to a unique α-quantile. The definition allows for the case where a quantile has a non-zero point mass. To be explicit in cases where there is non-uniqueness, we make the following definition that hinges on a choice $r \in [0, 1]$.

Definition α-quantile: Let Q denote the set of all α-quantiles q. For $\alpha \in (0, 1)$ we define the α-quantile as

$$r \sup Q + (1 - r) \inf Q. \tag{2.5}$$

When $\alpha = 0$ we use $\sup Q$ (that is, we use $r = 1$) and when $\alpha = 1$ we use $\inf Q$ (that is, we use $r = 0$).

From now on, by $q_{s,\alpha}(Y_{B(\delta(w))})$ we denote the α-quantile of $Y_{B(\delta(w))}$ when the state of the world is s. For simplicity of notation, unless needed for clarity, we do not index that expression by r.

By definition, regret of the treatment rule δ for a given distribution s of (Y_0, Y_1) equals

$$R(\delta, s) = \sup_{d \in \mathbb{D}} u(d, s) - u(\delta, s). \tag{2.6}$$

In words, regret equals the difference between the highest α-quantile that could have been achieved for any treatment rule $d \in \mathbb{D}$ and the α-quantile obtained for the particular treatment rule δ for the given state of nature s. If $q_{s,\alpha}(Y_1) \geq q_{s,\alpha}(Y_0)$ ($q_{s,\alpha}(Y_1) < q_{s,\alpha}(Y_0)$) and $d^* = 1$ ($d^* = 0$) is an element of $\mathbb{D}$ then $\sup_{d \in \mathbb{D}} u(d, s)$ is taken on by d^* and equals $q_{s,\alpha}(Y_1)$ ($q_{s,\alpha}(Y_0)$), see Lemma 3(ii) in the Appendix. If restrictions are imposed on $\mathbb{D}$ then $\sup_{d \in \mathbb{D}} u(d, s)$ may not be taken on by any element in $\mathbb{D}$.

In this paper, we focus on minimax regret treatment rules. By definition, if it exists, such a rule satisfies

$$\delta^* \in \arg\min_{\delta \in \mathbb{D}} \sup_{s \in \mathbb{S}} R(\delta, s). \tag{2.7}$$

In contrast, a maximin treatment rule, if its exists, satisfies

$$\delta^+ \in \arg\max_{\delta \in \mathbb{D}} \inf_{s \in \mathbb{S}} u(\delta, s). \tag{2.8}$$

As discussed already elsewhere (see e.g. Manski (2004), Stoye (2009)) the maximin criterion may lead to the uninformative result that *all* $\delta \in \mathbb{D}$ are maximizers. That occurs e.g. if for a particular $s^+ \in \mathbb{S}$, $u(\delta, s^+)$ does not depend on δ and takes on its smallest possible value, $u(\delta, s^+) = \inf_{s \in \mathbb{S}} u(\delta, s)$. That situation also occurs in our setup where the objective is to maximize the quantile of the outcome distribution, namely when $s^+ \in \mathbb{S}$ is chosen such that $q_{s,\alpha}(Y_0) = q_{s,\alpha}(Y_1) = 0$. In contrast to the setup where expected outcomes are the objective, as we will establish next, it turns out that for quantiles, depending on the particular sampling design, also the minimax regret criterion may be uninformative.

2.1 Minimax Regret Treatment Rules

We consider three different sample designs, namely:

(i) *Fixed number of untreated/treated units* with $N_0 \in \{0, ..., N\}$ iid observations of Y_0 and $N_1 := N - N_0$ iid observations of Y_1 for some $N \in \mathbb{N} \cup \{0\}$. Here, a treatment rule $\delta \in \mathbb{D}$ is any mapping $[0,1]^N \to [0,1]$. The rule δ assigns a treatment probability $\delta(w) \in [0,1]$ after observing a sample $w = (y(0)', y(1)')' \in [0,1]^N$ of N_0 untreated and N_1 treated units $y(0) \in [0,1]^{N_0}$ and $y(1) \in [0,1]^{N_1}$, respectively.

(ii) *Random assignment* with $N \in \mathbb{N} \cup \{0\}$ iid observations, where in the sample, the treatment probability equals $p \in (0,1)$.[4] A treatment rule $\delta \in \mathbb{D}$ is any mapping $\{0,1\}^N \times [0,1]^N \to [0,1]$. The rule δ assigns a treatment probability $\delta(t, y) \in [0,1]$ after observing a sample $w = (t, y)$ of treatment statuses $t = (t_1, ..., t_N)$ and realizations $y = (y_1, ..., y_N)$ where the i-th component of y, y_i, is an independent realization of Y_{t_i}.

(iii) *Testing an innovation*, is the case where aspects of the distribution of Y_0 are known to the policymaker; in particular, we assume that the α-quantile of Y_0 is known (but nature can pick arbitrary distributions for Y_0 subject to that restriction). That is, in this case the set $\mathbb{S}$ consists of all joint distributions s for (Y_0, Y_1) with the restriction that the α-quantile of the marginal for Y_0 equals a certain value $q_\alpha(Y_0)$. In this case, $N \in \mathbb{N} \cup \{0\}$ iid observations of Y_1 are observed. A treatment rule $\delta \in \mathbb{D}$ is any mapping $[0,1]^N \to [0,1]$. The rule δ

[4] Note that $p \in \{0,1\}$ leads back to design (i) with N_0 or N_1 equal to N.

assigns a treatment probability $\delta(w) \in [0,1]$ after observing a sample $w = (y_{1,1}, ..., y_{1,N})$ of N independent realizations $y_{1,i}$, $i = 1, ..., N$ of Y_1.

We write w_N for the observed sample w when we want to emphasize the sample size N.

The designs above nest the ones in Stoye (2009). In contrast to Stoye (2009), in design (i) we allow for arbitrary numbers of treated and untreated units rather than $N/2$ units each as in "matched pairs" and in design (ii) the treatment probability is any fixed number $p \in (0,1)$ rather than necessarily $p = .5$.

The following statement provides the analogue to Proposition 1 in Stoye (2009) and is the main contribution of this paper.

Proposition 1 (*i*) *Consider the case where the sample consists of a* "fixed number of treated/ untreated units". *If* ($\alpha \in (0,1)$ *and* $r \in [0,1]$) *or* ($\alpha = 0$ *and* $r = 1$) *then any treatment rule* $\delta \in \mathbb{D}$ *is minimax regret and* $\max_{s \in \mathbb{S}} R(\delta, s) = 1$ *for any* $\delta \in \mathbb{D}$. *If* ($\alpha = 1$ *and* $r = 0$) *then exactly those treatment rules* $\delta \in \mathbb{D}$, *that are not equal to 1 wp1 and not equal to 0 wp1, are minimax regret and satisfy* $\max_{s \in \mathbb{S}} R(\delta, s) = 0$.

(*ii*) *Consider the case where the sample is generated via* "random assignment". *Then the same statement as in part* (*i*) *holds.*

(*iii*) *In the case of* "testing an innovation" *let* ($\alpha \in (0,1)$ *and* $r \in [0,1]$) *or* ($\alpha = 0$ *and* $r = 1$). *If the known* α*-quantile of* Y_0, $q_\alpha(Y_0)$, *equals* $1/2$ *then any treatment rule* $\delta \in \mathbb{D}$ *is minimax regret; if instead* $q_\alpha(Y_0) > 1/2$ *then* $\delta^0 \equiv 0$ *is the unique minimax regret rule; if* $q_\alpha(Y_0) < 1/2$ *then* $\delta^1 \equiv 1$ *is a minimax regret rule; in each case, for minimax regret treatment rules* δ *we obtain* $\max_{s \in \mathbb{S}} R(\delta, s) = \min\{q_\alpha(Y_0), 1 - q_\alpha(Y_0)\}$. *If* ($\alpha = 1$ *and* $r = 0$) *then exactly those treatment rules* $\delta \in \mathbb{D}$, *that are not equal to 1 wp1 and not equal to 0 wp1, are minimax regret and satisfy* $\max_{s \in \mathbb{S}} R(\delta, s) = 0$.

Comments. 1. Proposition 1 establishes that the minimax regret criterion when applied to α-quantiles does not favor data-driven rules. In fact, in the *"testing an innovation"* case data-driven rules are strictly dominated by $\delta^0 \equiv 0$ or $\delta^1 \equiv 1$, unless $q_\alpha(Y_0) = 1/2$ in which case all treatment rules are minimax. This is in stark contrast to the results in Stoye (2009) about minimax regret treatment rules when the focus is on mean outcomes. Namely Stoye (2009) shows that e.g. in the case of binary outcomes where the sample is obtained as a matched pair, the treatment that has more successes in the sample should be chosen with probability one.

2. For some intuition of the result in (i)-(ii) consider the simplest case where the sample size N is 0. In that case, a treatment rule δ is simply an element in $[0,1]$ that denotes the probability of assigning treatment. For given δ, the objective is to find distributions for

Y_0, Y_1 such that their α-quantiles have maximal distance (that is distance 1) and such that the α-quantile of $Y_{B(\delta)}$ is zero. If $\delta > 0$ assume nature picks Y_0 and Y_1 as Bernoulli

$$P(Y_1 = 0) = 1, \; P(Y_0 = 0) = \alpha - \varepsilon \text{ for some } \varepsilon > 0 \tag{2.9}$$

and, consequently, $P(Y_0 = 1) = 1 - (\alpha - \varepsilon)$. Thus $q_{s,\alpha}(Y_1) = 0$ and $q_{s,\alpha}(Y_0) = 1$ and

$$P(Y_{B(\delta)} = 0) = \delta + (1 - \delta)(\alpha - \varepsilon) \tag{2.10}$$

Thus, $P(Y_{B(\delta)} = 0) > \alpha$ iff $\delta(1 - \alpha) > \varepsilon(1 - \delta)$ which holds for ε small enough. Thus

$$\max_{s \in \mathbb{S}} R(\delta, s) = 1 \tag{2.11}$$

for all $\delta \in (0, 1]$. If instead $\delta^0 \equiv 0$ then nature chooses $P(Y_1 = 1) = 1$ and $P(Y_0 = 0) = 1$ which leads to regret of 1. Surprisingly, the intuition of that no-data example generalizes to cases with arbitrary sample size.

3. In part (iii), in the case $q_\alpha(Y_0) < 1/2$ we could not rule out that there are other minimax regret rules besides $\delta^1 \equiv 1$ when $N > 0$. When $N = 0$, it is obvious that δ^1 is the unique minimax rule. A minimax regret treatment rule, if it is not unique, may be inadmissible. In the case of "testing an innovation" $\delta^0 \equiv 0$ (when $q_\alpha(Y_0) > 1/2$) is admissible, but it is unclear a priori whether that is true for δ^1 (when $q_\alpha(Y_0) < 1/2$) and $\delta^{.5} \equiv 1/2$ (when $q_\alpha(Y_0) = 1/2$).

4. In case $Y_0, Y_1 \in (0, 1)$ rather than $Y_0, Y_1 \in [0, 1]$, $\max_{s \in \mathbb{S}} R(\delta, s)$ does not generally exist. E.g. in cases (i) and (ii) nature choosing only distributions with support on $[\varepsilon, 1 - \varepsilon]$ for some small $\varepsilon > 0$ it can generate regret of $1 - 2\varepsilon$ and therefore $\sup_{s \in \mathbb{S}} R(\delta, s) = 1$ for the unrestricted space of distributions for $Y_0, Y_1 \in (0, 1)$. Therefore, if minimax rules are defined with $\sup_{s \in \mathbb{S}} R(\delta, s)$ (as we do in (2.7)) rather than $\max_{s \in \mathbb{S}} R(\delta, s)$ the results in Proposition 1(i)-(ii) continue to hold.

Restrictions on nature's action space

In what follows we impose various restrictions on nature's action space and study the implications on the results obtained in Proposition 1. For simplicity assume $\alpha \in (0, 1)$.

Corollary 1 *(a) The results in Proposition 1(i)-(ii) remain valid when rather than $Y_0, Y_1 \in [0, 1]$ the setup is altered to $Y_0, Y_1 \in \{0, 1\}$. Similarly, the result in (iii) remains valid if the distributions of Y_0 and Y_1 are restricted to be discrete and supported on at most $N + 1$ points in $[0, 1]$.*

(*b*) *If* $\mathbb{S}$ *equals the set of distributions whose marginals are all continuous* (*with respect to Lebesgue measure*) *then Proposition 1 continues to hold.*

Comments. 1. Part (a) of Corollary 1 is a direct corollary from the proof of Proposition 1. In the proof for parts (i)-(ii) only Bernoulli distributions are used for nature while in part (iii) only discrete distributions are used.

2. Part (b) of Corollary 1 considers the case where nature is restricted to continuous distributions. We have seen in Proposition 1 that in cases (i)-(ii) any treatment rule $\delta \in \mathbb{D}$ is minimax and $\max_{s\in\mathbb{S}} R(\delta, s) = 1$. Because without pointmasses it is impossible for a random variable to have α-quantile equal to 0 it follows that $R(\delta, s)$ is always strictly smaller than 1 for any pair (δ, s) when s is continuous with respect to Lebesgue measure. The main construction in the proof of Proposition 1 still goes through when one considers a sequence of continuous distributions that converge to the Bernoulli distributions that are used in that proof. As a technical detail it is important to use $\sup_{s\in\mathbb{S}} R(\delta, s)$ rather than $\max_{s\in\mathbb{S}} R(\delta, s)$ in the definition of regret, because $\max_{s\in\mathbb{S}} R(\delta, s)$ would not exist in the case considered here. One can establish that $\sup_{s\in\mathbb{S}} R(\delta, s) = 1$ for any treatment rule $\delta \in \mathbb{D}$.

3. In the case of *"testing an innovation"* a restriction on nature's action space occurs if one assumes that the entire distribution of Y_0 is known, not just its α-quantile. Namely, in that case the set $\mathbb{S}$ consists of all possible distributions s for $Y_1 \in [0, 1]$ (while the distribution of Y_0 is given). The analysis of that problem is more difficult. Denote by F_{Y_0} and $q_\alpha(Y_0)$ the cdf and α-quantile of Y_0, respectively.

Take $r = 0$ in (2.5) and assume $F_{Y_0}(q_\alpha(Y_0)) > \alpha$. Consider the case $N = 0$ in which case a treatment rule $\delta \in [0, 1]$ denotes the treatment probability. Under these assumptions the following statements hold.

Denote by $q(\delta)$ the smallest $q \in [0, q_\alpha(Y_0)]$ such that

$$\delta + (1 - \delta)F_{Y_0}(q) \geq \alpha. \tag{2.12}$$

Then, when $q_\alpha(Y_0) < 1/2$, $\delta^1 \equiv 1$ is the only minimax regret rule and $\max_{s\in\mathbb{S}} R(\delta^1, s) = q_\alpha(Y_0)$. When $q_\alpha(Y_0) = 1/2$ any $\delta \in [0, 1]$ is minimax regret with maximal regret equal to $q_\alpha(Y_0)$. Finally, when $q_\alpha(Y_0) > 1/2$, $\delta \in [0, 1]$ is minimax regret iff

$$q_\alpha(Y_0) - q(\delta) \leq 1 - q_\alpha(Y_0). \tag{2.13}$$

Given that $q_\alpha(Y_0) - q(\delta)$ is a weakly increasing function in δ, if we denote by $\delta^* \in [0, 1]$ an intersection of $q_\alpha(Y_0) - q(\delta)$ and $1 - q_\alpha(Y_0)$ then any $\delta \in [0, \delta^*]$ is minimax regret and for those δ we have $\max_{s\in\mathbb{S}} R(\delta, s) = 1 - q_\alpha(Y_0)$.

In the Appendix, we give a proof of the statements made above. The results for the case $q_\alpha(Y_0) > 1/2$ are partly in contrast to Proposition 1(iii) where there is a *unique* minimax regret rule. However, maximal regret is the same here as in Proposition 1(iii). We have not yet generalized the results to arbitrary sample sizes $N > 0$ and to the case $F_{Y_0}(q_\alpha(Y_0)) = \alpha$.

2.2 Finite sample simulation

For the case of "testing an innovation" Proposition 1(iii) establishes that the minimax regret criterion when applied to α-quantiles does not favor data-driven rules. In this section we conduct a simulation experiment to juxtapose the pointwise (in s) regret of the data-driven empirical success rule δ^{ES}, defined by

$$\delta^{ES}(w) = I(q_\alpha(Y_1, w) > q_\alpha(Y_0)) + .5I(q_\alpha(Y_1, w) = q_\alpha(Y_0)), \tag{2.14}$$

where $q_\alpha(Y_1, w)$ denotes the α-sample quantile of Y_1 for the sample w, to the regret of the minimax regret rule $\delta^1 \equiv 1$ (in the case when $q_\alpha(Y_0) \leq 1/2$), $\delta^{.5} \equiv 1/2$ (in the case when $q_\alpha(Y_0) = 1/2$), and the minimax regret rule $\delta^0 \equiv 0$ (in the case when $q_\alpha(Y_0) \geq 1/2$).

Obviously, when simulating regret we cannot possibly include all states of nature $s \in \mathbb{S}$. For the simulation experiment, we create a "sufficiently rich" subset of distributions for (Y_0, Y_1). Namely, for some $n, w \in \mathbb{N}$, we consider the set of states of nature $\mathbb{S}^E = \mathbb{S}^E(n, w)$ that consists of all discrete distributions s for Y_1 supported on a grid

$$\{0, 1/n, 2/n, ..., n/n\} \tag{2.15}$$

with probabilities

$$P_s(Y_1 = j/n) \tag{2.16}$$

being of the form i/w, $i = 0, ..., w$; while for the distribution of Y_0 we consider two choices, namely

I) $Y_0 \in \{0, q_\alpha(Y_0)\}$ with $P(Y_0 = 0) = \alpha - \varepsilon$ and $P(Y_0 = q_\alpha(Y_0)) = 1 - (\alpha - \varepsilon)$ for $\varepsilon = .000001$ and

II) Y_0 being continuously distributed on [0,1] with density $f(x)$ equal to $\alpha/q_\alpha(Y_0)$ for $x \leq q_\alpha(Y_0)$ and equal to $(1-\alpha)/(1-q_\alpha(Y_0))$ otherwise. Note that in both I) and II) the α-quantile of Y_0, $q_{s,\alpha}(Y_0)$, equals $q_\alpha(Y_0)$. We report results for all choices of $\alpha \in \{.1, .5, .9\}$, $q_\alpha(Y_0) \in \{.1, .5, .9\}$, sample size $N = 30$, and $(n, w) = (6, 12)$. The latter leads to S^E having cardinality 18564.

For the given choices of n, w, and N and for each choice of α and $q_\alpha(Y_0)$, for each state of nature $s \in \mathbb{S}^E(n,w)$ and one (of the two possible) choice of distribution for Y_0 we simulate regret for the four treatment rules δ^{ES}, δ^1, $\delta^{.5}$, and δ^0 by generating $R = 100K$ samples of size N by drawing iid observations of the distribution of Y_1. We use $r = 0$ when simulating α-quantiles of the outcome distribution under the various treatment rules. We analytically calculate α-quantiles for Y_1 for $Y_{B(\delta^{.5})}$, and likewise use the true α-quantile $q_\alpha(Y_0)$ of Y_0 when calculating regret. For a given treatment rule δ and a given state of nature s, regret is calculated as $R(\delta,s) = \max\{q_\alpha(Y_0), q_{s,\alpha}(Y_1)\} - q_{s,\alpha}(Y_{B(\delta(w))})$.

We compare the treatment rules along several dimensions, namely

a) mean regret over all 18564 states of nature $s \in \mathbb{S}^E(n,w)$,

b) maximal regret over all states of nature $s \in \mathbb{S}^E(n,w)$,

c) minimal regret over all states of nature $s \in \mathbb{S}^E(n,w)$, and

d) the proportion of $s \in \mathbb{S}^E(n,w)$ for which regret for the empirical success rule δ^{ES} is smaller than regret for each one of its three competitors.

Just for clarity, in a) for each treatment rule we sum up its regret over all 18564 states of nature $s \in \mathbb{S}^E(n,w)$ and then report that sum divided by 18564. For a given $s \in \mathbb{S}^E(n,w)$, in our simulations, we interpret regret of δ^{ES} as smaller than regret of another rule, say δ^0, if the simulated regret of δ^{ES} is smaller than the simulated regret of the rule δ^0 minus a threshold of ξ. If instead the simulated regret of δ^{ES} falls into the interval [(simulated regret of δ^0)$-\xi$,(simulated regret of δ^0)$+\xi$] we record regrets of the two rules as equal for that state of nature. We take $\xi = .00000001$ below. Similarly, when programming the empirical success rule in (2.14) the event "$q_\alpha(Y_1,w) = q_\alpha(Y_0)$" is implemented as the simulated α-quantile of Y_1 falling into the interval $[q_\alpha(Y_0) - \xi, q_\alpha(Y_0) + \xi]$.

All reported results are rounded to the second digit after the comma and so a reported zero could be as large as .004.

In the tables below we do not include results for c) because those results turn out to be equal to zero for all treatment rules and all designs except for $\delta^{.5}$ when $q_\alpha(Y_0) \in \{.1, .9\}$ in which case minimal regret equals .067 for all choices of α and both choices of distributions for Y_0.

TABLE I provides results for maximal and mean regret over all 18564 states of nature $s \in S^E(6,12)$ for the four different treatment rules. Results for rules that are minimax regret in a given setting are reported in bold.

We first discuss results for maximal regret. The treatment rules that are known to be minimax regret for unrestricted $\mathbb{S}$ also have smallest maximal regret over $\mathbb{S}^E(n,w)$ (relative to the other treatment rules considered here) for all cases except for case II) with $\alpha = .9$, when $q_\alpha(Y_0) = .1$ (in which case δ^{ES} has smaller maximal regret than the optimal rule δ^1), when $q_\alpha(Y_0) = .5$ (in which case it is known that all four rules are optimal but in finite samples for $\mathbb{S}^E(n,w)$ again δ^{ES} does best), and finally when $q_\alpha(Y_0) = .1$ (in which case δ^{ES} has smaller maximal regret than the optimal rule δ^0). The explanation is of course that $\mathbb{S}^E(n,w)$ does not contain those states of nature that would inflict the highest regret on δ^{ES} (like the particular Bernoulli and discrete distributions for Y_0 and Y_1, respectively, that are used in the proof of Proposition 1(iii)). Furthermore, reported finite sample maximal regret over $\mathbb{S}^E(n,w)$ for the optimal rules matches the theoretical value $\min\{q_\alpha(Y_0), 1-q_\alpha(Y_0)\}$ reported in Proposition 1(iii) except for the case II) with $\alpha = .9$ when $q_\alpha(Y_0) = .5$ where the optimal rule δ^{ES} has maximal regret smaller than .5 (which occurs, again, because $\mathbb{S}^E(n,w)$ is not rich enough). An open question from Proposition 1(iii) is whether for $q_\alpha(Y_0) < .5$ other minimax regret rules besides δ^1 might exist. The simulations for $q_\alpha(Y_0) = .1$ are compatible with the possibility that when $\alpha = .5$ or $.9$ also δ^{ES} might be minimax regret.

We next discuss results for mean regret. In most cases where $q_\alpha(Y_0) \in \{.1, .9\}$ in which case either δ^0 or δ^1 are minimax regret, their mean performance is also best (or very close to best) among the four treatment rules. On the other hand, when $q_\alpha(Y_0) = .5$ (in which case all four treatment rules are minimax regret according to Proposition 1(iii)) we see huge difference in mean performance across the four treatment rules for a given case I) or II) and α, but also huge differences in performance for a given treatment rule and α across cases I) and II). With respect to the former point, in case I) when $\alpha = .1$ mean regret for the four rules are in the interval [0,.43]. With regards to the latter point, e.g. for δ^{ES} when $\alpha = .1$ mean regret equals .43 and .06, in cases I) and II) respectively. Differences across cases I) and II) are also often huge for other quantiles. E.g. again for δ^{ES} when $q_\alpha(Y_0) = .9$ and $\alpha = .1$ mean regret equals .9 and 0, in cases I) and II) respectively. (Recall that we round results to the second digit. When $q_\alpha(Y_0) = .9$ and $\alpha = .1$ there are not many distributions for Y_1 in $\mathbb{S}^E(n,w)$ that have an α-quantile that exceeds .9.)

We next discuss the results for exercise d) contained in TABLE II where we report the proportion of the 18564 states of nature $s \in S^E(6,12)$ for which the regret of δ^{ES} is smaller than the regret of δ where consider all δ from the set of no-data rules $\{\delta^1, \delta^{.5}, \delta^0\}$. The results indicate that for many states s, choices of $q_\alpha(Y_0)$, α, and the distribution for Y_0, $R(\delta^{ES}, s)$ and $R(\delta, s)$ are very close which leads to numerically unstable results. To deal with the instability we introduce the buffer ξ as explained above and $Prop(R(\delta^{ES}, s) < R(\delta, s))$ and $Prop(R(\delta^{ES}, s) \leq R(\delta, s))$ in TABLE II represent the proportion of states s for which $R(\delta^{ES}, s) < R(\delta, s) - \xi$ and $R(\delta^{ES}, s) < R(\delta, s) + \xi$, respectively. Those two proportions can be drastically different, suggesting that in many scenarios regret for δ^{ES} and δ are (virtually) identical. According to the measure $Prop(R(\delta^{ES}, s) \leq R(\delta, s))$, maybe somewhat surprisingly given the results from TABLE I, δ^{ES} is to be preferred over the other three rules in all scenarios considered in Case II) (except the case $\alpha = q_\alpha(Y_0) = .9$ for δ^0) and in the majority of scenarios in Case I) (except compared to δ^0 when $q_\alpha(Y_0) = .5, \alpha = .1$ and except for most cases with $q_\alpha(Y_0) = .9$ and all three rules). Quite often $Prop(R(\delta^{ES}, s) \leq R(\delta, s))$ is reported as higher than 95%, but the improvement in regret for many states of nature is minuscule. For example compared to $\delta^{.5}$ in Case I) with $\alpha = q_\alpha(Y_0) = .1$ only in 33.3% of the cases $Prop(R(\delta^{ES}, s) < R(\delta, s))$ while for 100% of the cases $Prop(R(\delta^{ES}, s) \leq R(\delta, s))$ implying that the regret of δ^{ES} in 66.4% of the cases is at most ξ smaller than the regret of $\delta^{.5}$.

3 Treatment choice with a covariate

Next, as in Stoye (2009) we next allow for a discrete covariate $X \in \mathcal{X} = \{x_1, ..., x_K\}$ that is observed both in the sample and in the treatment data. Outcomes $Y_{t,x}$ now carry a double subindex to indicate the treatment status $t \in \{0, 1\}$ and the value of the covariate $x \in \mathcal{X}$. It is assumed that x_k occurs with positive probability for each $k = 1, ..., K$. Denote by F_X the distribution of X. A state of the world $s \in \mathbb{S}$ now represents a joint distribution for $(Y_{t,x})_{t\in\{0,1\},\ x\in\mathcal{X}}$.

A sample $w = w_N$ now consists of realizations (t_i, x_i, y_i) for $i = 1, ..., N$ of $(T, X, Y_{T,X})$, where y_i is a realization of Y_{t_i,x_i} and we consider again sampling designs (i)-(iii) from Section 2.1. It is assumed that x_i, $i = 1, ..., n$, are iid and F_X is independent of s and T.

In design (i), "fixed number of treated/untreated units", there are N_0 observations $(0, x_i, y_i)$ with y_i being iid draws of Y_{0,x_i}, $i = 1, ..., N_0$ and $N_1 = N - N_0$ observations $(1, x_i, y_i)$ with y_i being iid draws of Y_{1,x_i}, $i = N_0 + 1, ..., N$. Note that there are not typically equally many treated and untreated observations for each covariate x_k, $k = 1, ..., K$. In fact, there may be zero observations altogether in the sample for a given x_k.

In design (ii), "random sampling", the observations (t_i, x_i, y_i) for $i = 1, ..., N$ are iid realizations of $(T, X, Y_{T,X})$ with $P(T_i = 1) = p \in (0, 1)$ with T, X, and s being independent of each other.

In design (iii), "testing an innovation", $(1, x_i, y_i)$ for $i = 1, ..., N$ are observed where y_i is an independent realization of Y_{1,x_i}, for $i = 1, ..., N$.

In design (iii), we assume the policymaker knows the α-quantile, denoted by $q_\alpha(Y_{0,X})$, of $Y_{0,X}$ and nature can choose from joint distributions $s \in \mathbb{S}$ for $(Y_{t,x})_{t\in\{0,1\},\ x\in\mathcal{X}}$ such that the α-quantile of $Y_{0,X}$ equals $q_\alpha(Y_{0,X})$.

A treatment rule δ maps a sample w onto treatment probabilities for each x_k for $k = 1, ..., K$, that is $\delta(w) \in [0, 1]^K$, where the k-th component of $\delta(w)$ indicates the treatment probability for individuals with covariate x_k for $k = 1, ..., K$. Denote the k-th component of

$\delta(w)$ by $\delta_{x_k}(w)$ for $k = 1, ..., K$. With some abuse of notation, for each design (i)-(iii) we denote the set of all treatment rules by the same symbols $\mathbb{D}$ (even though it means something different for different designs).

The object of interest is

$$u(\delta, s) = q_{s,F_X,\alpha}(Y_{B(\delta_X(w)),X}) \tag{3.1}$$

the α-quantile of the outcome distribution. With that definition, regret is then defined formally analogously to the setup without a covariate, namely, $R(\delta, s) = \sup_{d\in\mathbb{D}} u(d, s) - u(\delta, s)$. Alternatively, one could focus on the α-quantile of the outcome distribution for a particular covariate only, x_1 say, $u_{x_1}(\delta, s) = q_{s,F_X,\alpha}(Y_{B(\delta_{x_1}(w)),x_1})$ and obtain analogous results to the ones in Corollary 2 below.

If the space of probability distributions for nature $\mathbb{S}$ is unrestricted and thus equals the space of all distributions for $(Y_{t,x})_{t\in\{0,1\},\ x\in\mathcal{X}}$ one can show in designs (i)-(ii) (as an implication of the proof of Proposition 1) that for every treatment rule δ maximal risk over $\mathbb{S}$ equals 1, $\max_{s\in\mathbb{S}} R(\delta, s) = 1$. This result then immediately implies that every treatment rule is minimax regret. The corollary (to Proposition 1) that follows gives a stronger result; it shows that maximal risk continues to be 1 even if certain restrictions are imposed on $\mathbb{S}$. For simplicity of the presentation assume $\alpha \in (0, 1)$.

Corollary 2 *(i)-(ii) In the case of the designs "fixed number of treated/untreated units" and "random sampling" for any treatment rule $\delta \in \mathbb{D}$, $\max_{s\in\mathbb{S}} R(\delta, s) = 1$ if $\mathbb{S}$ includes as a subset all joint distributions for $(Y_{t,x})_{t\in\{0,1\},\ x\in\mathcal{X}}$ whose marginals are independent Bernoulli distributions. Therefore,* any *$\delta \in \mathbb{D}$ is minimax regret.*

(iii) In the case of "testing an innovation" assume the policymaker knows the α-quantile, denoted by $q_\alpha(Y_{0,X})$, of $Y_{0,X}$. If $q_\alpha(Y_{0,X}) = 1/2$ then any $\delta \in \mathbb{D}$ is minimax regret; if $q_\alpha(Y_{0,X}) > 1/2$ then $\delta^0 = 0$ is the unique minimax regret rule, and if $q_\alpha(Y_{0,X}) < 1/2$ then $\delta^1 = 1$ is a minimax regret rule. In each case, $max_{s\in\mathbb{S}} R(\delta, s) = \min\{q_\alpha(Y_{0,X}), 1 - q_\alpha(Y_{0,X})\}$.

Comments: 1. Various variants could be considered in design (iii). E.g. one could instead assume that the joint distribution $(Y_{0,x})_{x\in\mathcal{X}}$ is known (in which case nature only chooses a joint distribution for $(Y_{1,x})_{x\in\mathcal{X}}$) or one could assume that the policymaker knows the vector $(q_\alpha(Y_{0,x_1}), ..., q_\alpha(Y_{0,x_K}))$ of α-quantiles of all the marginal distributions $(Y_{0,x})_{x\in\mathcal{X}}$ and nature can choose from joint distributions $s \in \mathbb{S}$ for $(Y_{t,x})_{t\in\{0,1\},\ x\in\mathcal{X}}$ such that all the marginals $(Y_{0,x})_{x\in\mathcal{X}}$ have the required α-quantiles.

2. One could also consider alternative sampling designs (i)-(iii), where in the sampling stage, rather than being randomly assigned, the values of the covariate X are assigned

deterministically, as is done for treatment status in design (i). This could be referred to as "stratified sampling."

For brevity we do not explicitly deal with these variations.

4 Conclusion

We derive minimax regret treatment rules in finite samples when an α-quantile of the outcome distribution is the focus of interest. We establish that when the sample i) consists of a fixed number of untreated/treated units or ii) is generated via random treatment assignment then *all* treatment rules are minimax regret and therefore the minimax regret criterion is not helpful in singling out a recommended treatment rule. Given that the same shortcoming applies to the max-min criterion, an important open question concerns finding a meaningful criterion in this setup based on which an optimal treatment rule should be chosen. The idea from Montiel Olea, Qiu, and Stoye (2023) to look for rules that randomize "the least" in a situation where there are multiple minimax regret rules would not lead to a unique rule in our setup because in cases i) and ii) both δ^0 and δ^1 are minimax regret and never randomize.

We also establish that when iii) the sample consists of only realizations from the treated population while the α-quantile of the untreated population is known, never treating is the unique minimax rule if the known quantile exceeds .5, while always treating is a minimax rule if that quantile is strictly smaller than .5, and finally, if the quantile equals .5, then any treatment rule is minimax regret. It follows that based on the minimax regret criterion, rules that take the data into consideration are never strictly preferred.

An interesting but quite difficult extension that we are currently investigating concerns applying the minimax regret criterion in finite samples to conditional value at risk $S_{s,\alpha}(Y) := \alpha^{-1}E_s[Y1(Y \leq q_{s,\alpha}(Y))]$, for the random variable Y (or, in its more general form, $S_{s,\alpha}(Y) := \sup_{\gamma\in R}\{\gamma - \alpha^{-1}E_s[\gamma - Y]_+\}$, see Qi, Pang, and Liu (2023)) rather than to the α-quantile. Here, $1(\cdot)$ denotes the indicator function and $[Y]_+ = \max\{Y, 0\}$ denotes the positive part.

5 Appendix

Proof of Proposition 1. We use the notation $P_s(A)$ to denote probability of the event A when nature chooses the state of the world $s \in \mathbb{S}$ and $w = w_N$ denotes the sample of size N.

Note that for a state of the world $s \in \mathbb{S}$ and treatment rule $\delta \in \mathbb{D}$

$$\begin{aligned}
&P_s(Y_{B(\delta(w))} \leq q) \\
&= P_s(B(\delta(w)) = 1 \ \& \ Y_1 \leq q) + P_s(B(\delta(w)) = 0 \ \& \ Y_0 \leq q) \\
&= P_s(B(\delta(w)) = 1)P_s(Y_1 \leq q) + P_s(B(\delta(w)) = 0)P_s(Y_0 \leq q) \qquad (5.1)
\end{aligned}$$

and analogously for $P_s(Y_{B(\delta(w))} \geq q)$, where the second equality uses independence. By definition an α-quantile q of $Y_{B(\delta(w))}$ satisfies $P_s(Y_{B(\delta(w))} \leq q) \geq \alpha$ and $P_s(Y_{B(\delta(w))} \geq q) \geq 1 - \alpha$.

The proofs of (i)-(ii) proceed by showing that for any treatment rule $\delta \in \mathbb{D}$ there exists a state of the world $s_\delta \in \mathbb{S}$ for which $R(\delta, s_\delta) = 1$. Because it is also true that for any treatment rule $\delta \in \mathbb{D}$, $\max_{s \in \mathbb{S}} R(\delta, s) \leq 1$ it follows that any treatment rule is minimax.

Lemma 3(ii) below establishes the (unsurprising) result that for any arbitrary $s \in \mathbb{S}$

$$\max_{d \in \mathbb{D}} u(d, s) = \max_{d \in \mathbb{D}} q_{s,\alpha}(Y_{B(d(w))}) = \max\{q_{s,\alpha}(Y_0), q_{s,\alpha}(Y_1)\}. \qquad (5.2)$$

This result implies a simplified formula for $R(\delta, s)$ that we will use from now on.

In the proof of (i)-(ii) that follows, for any given treatment rule $\delta \in \mathbb{D}$ we will construct a $s_\delta \in \mathbb{S}$ for which $\max\{q_{s_\delta,\alpha}(Y_0), q_{s_\delta,\alpha}(Y_1)\} = 1$ and $u(\delta, s_\delta) = 0$. Namely, throughout the proof of (i)-(ii) for an arbitrary treatment rule $\delta \in \mathbb{D}$ we construct a $s_\delta \in \mathbb{S}$ whose independent marginals for Y_0 and Y_1 are Bernoulli (supported on $\{0, 1\}$) and for which $R(\delta, s_\delta) = 1$. Under these restrictions, the distribution s_δ for (Y_0, Y_1) is then fully defined by

$$a_0 := P_{s_\delta}(Y_0 = 0) \text{ and } a_1 := P_{s_\delta}(Y_1 = 0). \qquad (5.3)$$

The α-quantiles of Y_0, Y_1, and $Y_{B(\delta(w))}$ in the constructions below will be unique when $\alpha \in (0, 1)$. Therefore, the particular value of $r \in [0, 1]$ has no impact on the results. As in (5.1) and with s_δ just defined

$$\begin{aligned}
&P_{s_\delta}(Y_{B(\delta(w))} = 0) \\
&= \textstyle\sum_{u \in \{0,1\}^{N_0}} \sum_{v \in \{0,1\}^{N_1}} a_0^{N_0 - \overline{u}}(1 - a_0)^{\overline{u}} a_1^{N_1 - \overline{v}}(1 - a_1)^{\overline{v}} [\delta((u', v')')a_1 + (1 - \delta((u', v')'))a_0] \\
&= \textstyle\sum_{i=0}^{N_0} \sum_{j=0}^{N_1} a_0^{N_0 - i}(1 - a_0)^i a_1^{N_1 - j}(1 - a_1)^j [\delta^{i,j} a_1 + (\binom{N_0}{i}\binom{N_1}{j} - \delta^{i,j})a_0], \qquad (5.4)
\end{aligned}$$

where for a vector $u \in \{0, 1\}^{N_0}$ we denote by $\overline{u} \in \mathbb{N}$ the sum of its components (that is the number of 1's in the vector) and similar for other vectors, and

$$\delta^{i,j} := \sum_{\substack{u\in\{0,1\}^{N_0}, v\in\{0,1\}^{N_1}\\ \overline{u}=i \text{ and } \overline{v}=j}} \delta((u',v')'). \tag{5.5}$$

Note that in the second line of (5.4), the term $a_0^{N_0-\overline{u}}(1-a_0)^{\overline{u}} a_1^{N_1-\overline{v}}(1-a_1)^{\overline{v}}$ is the probability to observe particular vectors u and v while the term $\delta((u',v')')a_1 + (1-\delta((u',v')'))a_0$ is the probability that an outcome 0 is reached when the sample consists of u and v.

Part (i). Assume $\alpha \in (0,1)$ first. When in the definition of s_δ in (5.3) we take $(a_0, a_1) = (1, \alpha)$ we obtain

$$\begin{aligned} P_{s_\delta}(Y_{B(\delta(w))} = 0) &= \sum_{j=0}^{N_1} \alpha^{N_1-j}(1-\alpha)^j [\delta^{0,j}\alpha + (\binom{N_1}{j} - \delta^{0,j}))] \\ &= \sum_{j=0}^{N_1} [\alpha^{N_1-j}(1-\alpha)^j(\alpha-1)]\delta^{0,j} + \sum_{j=0}^{N_1} \alpha^{N_1-j}(1-\alpha)^j \binom{N_1}{j}. \end{aligned} \tag{5.6}$$

Given that the coefficient $\alpha^{N_1-j}(1-\alpha)^j(\alpha-1)$ in front of $\delta^{0,j}$ is negative given $\alpha \in (0,1)$ it follows that $P_{s_\delta}(Y_{B(\delta(w))} = 0)$ is strictly decreasing in each $\delta^{0,j}$. Therefore, we obtain

$$P_{s_\delta}(Y_{B(\delta(w))} = 0) \geq \sum_{j=0}^{N_1} \alpha^{N_1-j}(1-\alpha)^j \alpha \binom{N_1}{j} = \alpha, \tag{5.7}$$

where the right hand bound follows from (5.6) by replacing $\delta^{0,j}$ by its maximal value $\binom{N_1}{j}$.

Consider first the case where at least one of the $\delta^{0,j}$ (for $j = 0, ..., N_1$) is smaller than $\binom{N_1}{j}$. In that case it follows that $P_{s_\delta}(Y_{B(\delta(w))} = 0) > \alpha$. Furthermore, notice that $P_{s_\delta}(Y_{B(\delta(w))} = 0)$ is continuous in a_1 and therefore $P_{s_\delta}(Y_{B(\delta(w))} = 0) > \alpha$ when instead of $(a_0, a_1) = (1, \alpha)$, s_δ is taken as $(a_0, a_1) = (1, \alpha - \varepsilon)$ for some small enough $\varepsilon > 0$. But then for that choice of s_δ, $u(\delta, s_\delta) = 0$ and $\max\{q_{s_\delta,\alpha}(Y_0), q_{s_\delta,\alpha}(Y_1)\} = q_{s_\delta,\alpha}(Y_1) = 1$. Therefore $R(\delta, s_\delta) = 1$ as desired.

Next, consider the case where $\delta^{0,j} = \binom{N_1}{j}$ for all $j = 0, ..., N_1$ (which implies that $\delta((0^{N_0\prime}, v')') = 1$ for all $v \in \{0,1\}^{N_1}$, where 0^{N_0} denotes an N_0-dimensional column vector of zeros). In that case, from (5.4) we obtain with $(a_0, a_1) = (\alpha, 1)$

$$\begin{aligned}
&P_{s_\delta}(Y_{B(\delta(w))} = 0) \\
&= \textstyle\sum_{i=0}^{N_0} \alpha^{N_0-i}(1-\alpha)^i[\delta^{i,0} + (\binom{N_0}{i} - \delta^{i,0}))\alpha] \\
&= \alpha^{N_0} + \textstyle\sum_{i=1}^{N_0} \alpha^{N_0-i}(1-\alpha)^i[\delta^{i,0}(1-\alpha) + \binom{N_0}{i}\alpha] \\
&\geq \alpha^{N_0} + \alpha \textstyle\sum_{i=1}^{N_0} \alpha^{N_0-i}(1-\alpha)^i \binom{N_0}{i} \\
&= \alpha^{N_0} + \alpha(1-\alpha^{N_0}) \\
&> \alpha, \qquad (5.8)
\end{aligned}$$

where the second equality uses $\delta^{0,0} = \binom{N_1}{0} = 1$, the first inequality follows from setting $\delta^{i,0} = 0$ for $i = 1, ..., N_0$, and the second inequality follows from $\alpha \in (0, 1)$. The remainder of the proof is as above, namely, by continuity $P_{s_\delta}(Y_{B(\delta(w))} = 0) > \alpha$ for a s_δ with $(a_0, a_1) = (\alpha - \varepsilon, 1)$ for a small enough $\varepsilon > 0$. For such s_δ regret equals 1.

Now consider the case $\alpha = 0$. Take any δ and define s_δ by setting $a_0 = P_{s_\delta}(Y_0 = 0) = 0$ and $a_1 = P_{s_\delta}(Y_1 = 0) = \varepsilon > 0$. Then, $q_{s_\delta,0}(Y_1) = 0$ and any any $q \in [0, 1]$ satisfies the condition in (2.4) for $X = Y_0$. Thus, with $r = 1$ the 0-quantile of Y_0 equals 1. Thus, in order for regret to equal 1 under s_δ it is enough to establish that the 0-quantile of $Y_{B(\delta(w))}$ equals 0. The latter follows if $P_{s_\delta}(Y_{B(\delta(w))} = 0) > 0$. From (5.4)

$$P_{s_\delta}(Y_{B(\delta(w))} = 0) = \varepsilon \textstyle\sum_{j=0}^{N_1} \varepsilon^{N_1-j}(1-\varepsilon)^j \delta^{N_0,j} \qquad (5.9)$$

with $\delta^{N_0,j} = \sum_{v \in \{0,1\}^{N_1},\, \overline{v}=j} \delta((1, ..., 1, v')')$. If $P_{s_\delta}(Y_{B(\delta(w))} = 0) > 0$ the proof is complete. If not, and $P_{s_\delta}(Y_{B(\delta(w))} = 0) = 0$, it must be that $\delta((1, ..., 1, v')') = 0$ for all $v \in \{0, 1\}^{N_1}$. If $\delta((1, ..., 1, v')') = 0$ for all $v \in \{0, 1\}^{N_1}$, define s'_δ by setting $a_0 = P_{s'_\delta}(Y_0 = 0) = \varepsilon$ and $a_1 = P_{s'_\delta}(Y_1 = 0) = 0$. Then, by (5.4), with $\delta^{i,N_1} = \sum_{u \in \{0,1\}^{N_0}, \overline{u}=i} \delta((u', 1, ..., 1)') \leq \binom{N_0}{i}$

$$P_{s'_\delta}(Y_{B(\delta(w))} = 0) = \varepsilon \textstyle\sum_{i=0}^{N_0} \varepsilon^{N_0-i}(1-\varepsilon)^i (\binom{N_0}{i} - \delta^{i,N_1}) \qquad (5.10)$$

But $1 - \delta((1, ..., 1, v')') = 1$ for all $v \in \{0, 1\}^{N_1}$ implies $\binom{N_0}{i} - \delta^{i,N_1} = 1$ when $i = N_0$. Thus, by (5.10), $P_{s'_\delta}(Y_{B(\delta(w))} = 0) > 0$. But, the latter implies that regret equals 1 under s'_δ.

Finally, consider the case ($\alpha = 1$ and $r = 0$). Consider a treatment rule δ that is not equal to 1 wp1 and not equal to 0 wp1. We will show that $\max_{s \in \mathbb{S}} R(\delta, s) = 0$. To see this, fix $s \in \mathbb{S}$. By the definition of an α-quantile in (2.4) and the choice $r = 0$ note that for any

$q < q_{s,1}(Y_0)$ it must be that $P_s(Y_0 \le q) < 1$ and likewise for any $q < q_{s,1}(Y_1)$ it must be that $P_s(Y_1 \le q) < 1$. Therefore, using (5.1), for any $q < \max\{q_{s,1}(Y_0), q_{s,1}(Y_1)\}$ we have

$$P_s(Y_{B(\delta(w))} \le q) = P_s(B(\delta(w)) = 1)P_s(Y_1 \le q) + P_s(B(\delta(w)) = 0)P_s(Y_0 \le q) < 1$$

which implies $q < q_{s,1}(Y_{B(\delta(w))})$. By Lemma 3 $q_{s,1}(Y_{B(\delta(w))}) \le \max\{q_{s,1}(Y_0), q_{s,1}(Y_1)\}$ and thus $q_{s,1}(Y_{B(\delta(w))}) = \max\{q_{s,1}(Y_0), q_{s,1}(Y_1)\}$. But that implies that $R(\delta, s) = 0$ and because $s \in \mathbb{S}$ that proves the desired result.

Part (ii). The proof of (ii) is similar to case (i). Again we start off with $\alpha \in (0, 1)$ and define s_δ as in (5.3). We obtain

$$\begin{aligned}
&P_{s_\delta}(Y_{B(\delta(w))} = 0) \\
&= \sum_{\substack{t=(t_1,...,t_N)' \in \{0,1\}^N, \\ y=(y_{t_1},...,y_{t_N})' \in \{0,1\}^N}} p^{\bar{t}}(1-p)^{N-\bar{t}} a_0^{N_{0ty}}(1-a_0)^{N-\bar{t}-N_{0ty}} a_1^{N_{1ty}}(1-a_1)^{\bar{t}-N_{1ty}} \\
&\times [\delta((t,y))a_1 + (1-\delta((t,y)))a_0] \\
&= \sum_{\substack{T=0,...,N, \\ i=0,...,N-T, \\ j=0,...,T}} p^T(1-p)^{N-T} a_0^i (1-a_0)^{N-T-i} a_1^j (1-a_1)^{T-j} \\
&\times [\delta^{T,i,j} a_1 + (\binom{N}{T}\binom{N-T}{i}\binom{T}{j} - \delta^{T,i,j})a_0] \\
&= \sum_{\substack{T=0,...,N, \\ i=0,...,N-T, \\ j=0,...,T}} p^T(1-p)^{N-T} a_0^i (1-a_0)^{N-T-i} a_1^j (1-a_1)^{T-j} \\
&\times [(a_1 - a_0)\delta^{T,i,j} + \binom{N}{T}\binom{N-T}{i}\binom{T}{j} a_0],
\end{aligned} \tag{5.11}$$

where for given vectors $t \in \{0,1\}^N$ and $y \in \{0,1\}^N$, N_{0ty} denotes the number of y_{t_i}, $i = 1, ..., N$, for which $y_{t_i} = 0$ and $t_i = 0$, N_{1ty} is the number of y_{t_i}, $i = 1, ..., N$, for which $y_{t_i} = 0$ and $t_i = 1$, where again $\bar{t} \in \mathbb{N}$ for a vector $t \in \{0,1\}^N$ denotes the sum of its components, and

$$\delta^{T,i,j} := \sum_{\substack{t=(t_1,...,t_N)' \in \{0,1\}^N, \\ y=(y_{t_1},...,y_{t_N})' \in \{0,1\}^N, \\ \bar{t}=T,\ N_{0ty}=i,\ N_{1ty}=j}} \delta((t,y)). \tag{5.12}$$

Note that $N_{0ty} \le N - \bar{t}$ and $N_{1ty} \le \bar{t}$. In the second line of (5.11), the factor $p^{\bar{t}}(1 - p)^{N-\bar{t}}$ captures the probability of observing exactly $\bar{t}$ treatments and $N - \bar{t}$ controls in the sample, while the factor $a_0^{N_{0ty}}(1-a_0)^{N-\bar{t}-N_{0ty}}$ captures the probability that among the $N - \bar{t}$ observations from the control group exactly N_{0ty} zeros are observed, while, finally, the factor $a_1^{N_{1ty}}(1-a_1)^{\bar{t}-N_{1ty}}$ captures the probability that among the $\bar{t}$ observations from treated individuals exactly N_{1ty} zeros are observed.

Evaluating $P_{s_\delta}(Y_{B(\delta(w))}=0)$ when s_δ takes $(a_0,a_1)=(1,\alpha)$ we obtain

$$P_{s_\delta}(Y_{B(\delta(w))}=0)=\sum_{\substack{T=0,...,N,\\ j=0,...,T}} p^T(1-p)^{N-T}\alpha^j(1-\alpha)^{T-j}[(\alpha-1)\delta^{T,N-T,j}+\binom{N}{T}\binom{T}{j}]. \quad (5.13)$$

Because $(\alpha-1)<0$ and $\alpha>0$, $P_{s_\delta}(Y_{B(\delta(w))}=0)$ is strictly decreasing in $\delta^{T,N-T,j}$ for all $T=0,...,N$ and $j=0,...,T$ and is minimized when $\delta^{T,N-T,j}$ takes on its maximal value $\binom{N}{T}\binom{T}{j}$. Therefore

$$\begin{aligned}
P_{s_\delta}(Y_{B(\delta(w))}=0) &\geq \sum_{\substack{T=0,...,N,\\ j=0,...,T}} p^T(1-p)^{N-T}\alpha^j(1-\alpha)^{T-j}\alpha\binom{N}{T}\binom{T}{j}\\
&=\alpha\sum_{T=0,...,N}\binom{N}{T}p^T(1-p)^{N-T}\sum_{j=0,...,T}\binom{T}{j}\alpha^j(1-\alpha)^{T-j}\\
&=\alpha. \qquad (5.14)
\end{aligned}$$

If $\delta^{T,N-T,j}<\binom{N}{T}\binom{T}{j}$ for any $T=0,...,N$ and $j=0,...,T$ then by an argument as in part (ii) it follows that an s_δ with $(a_0,a_1)=(1,\alpha-\varepsilon)$ for $\varepsilon>0$ small enough leads to regret of 1. This holds because for that s_δ we have $q_{s_\delta,\alpha}(Y_0)=0$, $q_{s_\delta,\alpha}(Y_1)=1$, and $q_{s_\delta,\alpha}(Y_{B(\delta(w))})=0$.

If on the other hand, $\delta^{T,N-T,j}=\binom{N}{T}\binom{T}{j}$ for all $T=0,...,N$ and $j=0,...,T$, then, as in part (ii) one can show that s_δ defined by $(a_0,a_1)=(\alpha-\varepsilon,1)$ for a small enough $\varepsilon>0$ leads to regret of 1. Namely, for s_δ with $(a_0,a_1)=(\alpha,1)$ we get from (5.11)

$$\begin{aligned}
&P_{s_\delta}(Y_{B(\delta(w))}=0)\\
&=\sum_{T=0,...,N}p^T(1-p)^{N-T}\sum_{i=0,...,N-T}\alpha^i(1-\alpha)^{N-T-i}[(1-\alpha)\delta^{T,i,T}+\binom{N}{T}\binom{N-T}{i}\alpha]\\
&=\sum_{T=0,...,N}p^T(1-p)^{N-T}\\
&\times\{\alpha^{N-T}\binom{N}{T}+\sum_{i=0,...,N-T-1}\alpha^i(1-\alpha)^{N-T-i}[(1-\alpha)\delta^{T,i,T}+\binom{N}{T}\binom{N-T}{i}\alpha]\}\\
&\geq\sum_{T=0,...,N}\binom{N}{T}p^T(1-p)^{N-T}\{\alpha^{N-T}+\alpha\sum_{i=0,...,N-T-1}\binom{N-T}{i}\alpha^i(1-\alpha)^{N-T-i}\}\\
&=\sum_{T=0,...,N}\binom{N}{T}p^T(1-p)^{N-T}\{\alpha^{N-T}+\alpha(1-\alpha^{N-T})\}\\
&>\alpha\sum_{T=0,...,N}\binom{N}{T}p^T(1-p)^{N-T}\\
&=\alpha, \qquad (5.15)
\end{aligned}$$

where the second equality simply takes the last summand with index $i=N-T$ out of the second summation sign and uses $\delta^{T,N-T,T}=\binom{N}{T}$, the inequality is obtained by setting

$\delta^{T,i,T} = 0$ for all $T = 0, ..., N$ and $i = 0, ..., N - T - 1$, and the strict inequality uses $\alpha^{N-T} + \alpha(1 - \alpha^{N-T}) > \alpha$. By an argument as in part (i), (5.15) implies that for s'_δ with $(a_0, a_1) = (\alpha - \varepsilon, 1)$ with sufficiently small $\varepsilon > 0$ regret of 1 can be obtained because under s'_δ we have $q_{s'_\delta,\alpha}(Y_0) = 1$, $q_{s'_\delta,\alpha}(Y_1) = 0$, and $q_{s'_\delta,\alpha}(Y_{B(\delta(w))}) = 0$.

Next, assume $\alpha = 0$. Define s_δ such that $a_0 = P_{s_\delta}(Y_0 = 0) = 0$ and $a_1 = P_{s_\delta}(Y_1 = 0) = \varepsilon > 0$. That and $r = 1$ imply that $q_{s_\delta,\alpha}(Y_0) = 1$ and $q_{s_\delta,\alpha}(Y_1) = 0$. Then, from (5.11)

$$P_{s_\delta}(Y_{B(\delta(w))} = 0) = \sum_{\substack{t=(t_1,...,t_N)'\in\{0,1\}^N,\\ y=(y_{t_1},...,y_{t_N})'\in\{0,1\}^N,\\ y_{t_j}=1 \text{ if } t_j=0 \text{ for } j=1,...,N}} p^{\bar{t}}(1-p)^{N-\bar{t}}\varepsilon^{N_{1ty}}(1-\varepsilon)^{\bar{t}-N_{1ty}}\delta((t,y))\varepsilon. \quad (5.16)$$

Thus, $P_{s_\delta}(Y_{B(\delta(w))} = 0) > 0$ (which implies $R(\delta, s_\delta) = 1$), unless $\delta((t, y)) = 0$ for all vectors $t, y \in \{0, 1\}^N$ such that $y_{t_j} = 1$ if $t_j = 0$ for $j = 1, ..., N$. For those δ consider instead, s'_δ such that $a_0 = P_{s'_\delta}(Y_0 = 0) = \varepsilon > 0$ and $a_1 = P_{s'_\delta}(Y_1 = 0) = 0$. Then

$$P_{s'_\delta}(Y_{B(\delta(w))} = 0) = \sum_{\substack{t=(t_1,...,t_N)'\in\{0,1\}^N,\\ y=(y_{t_1},...,y_{t_N})'\in\{0,1\}^N\\ y_{t_j}=1 \text{ if } t_j=1 \text{ for } j=1,...,N}} p^{\bar{t}}(1-p)^{N-\bar{t}}\varepsilon^{N_{0ty}}(1-\varepsilon)^{N-\bar{t}-N_{0ty}}(1-\delta((t,y)))\varepsilon \quad (5.17)$$

exceeds zero, because $1 - \delta((0^{N\prime}, 1^{N\prime})') = 1$ and the vectors $t = 0^N$ and $y = 1^N$ appear in the sum in (5.17). Thus $R(\delta, s'_\delta) = 1$.

Finally, the case ($\alpha = 1$ and $r = 0$) is handled exactly as in part (i).

Part (iii). We first show the following lemma.

Lemma 1 *Let ($\alpha \in (0, 1)$ and $r \in [0, 1]$) or ($\alpha = 0$ and $r = 1$). If $\delta \neq 0$ then there exists a $s_\delta \in \mathbb{S}$ such that $R(\delta, s_\delta) = q_\alpha(Y_0)$.*

Proof of Lemma 1. The case $q_\alpha(Y_0) = 0$ is trivial. Simply take s_δ such that $P_{s_\delta}(Y_0 = 0) = P_{s_\delta}(Y_1 = 0) = 1$. Thus, from now on we assume $q_\alpha(Y_0) > 0$.

Assume first $\alpha \in (0, 1)$. Because $\delta \neq 0$ there exists a vector

$$\widetilde{w} = (y_{1,1}, ..., y_{1,N}) \text{ such that } \delta(\widetilde{w}) > 0, \quad (5.18)$$

where $y_{1,i} \in [0, 1]$ for $i = 1, ..., N$. (When $N = 0$ then we don't need to define $\widetilde{w}$).

Definition of distributions of Y_0 and Y_1 : We now define a $s_\delta \in \mathbb{S}$ for which $R(\delta, s_\delta) = q_\alpha(Y_0)$ and under $s_\delta \in \mathbb{S}$, Y_0 and Y_1 have independent discrete marginals. Namely, Y_0 is defined by $P_{s_\delta}(Y_0 = 0) = \alpha - \varepsilon$ for some small $\varepsilon > 0$ (and obviously $\varepsilon < \alpha$) to be further restricted below and $P_{s_\delta}(Y_0 = q_\alpha(Y_0)) = 1 - (\alpha - \varepsilon)$. Thus $q_{s_\delta,\alpha}(Y_0) = q_\alpha(Y_0)$.

When $N = 0$ set $P_{s_\delta}(Y_1 = 0) = 1$ and note that $P_{s_\delta}(Y_{B(\delta)} = 0) = \delta + (1-\delta)(\alpha - \varepsilon)$ which exceeds α for ε small enough (using $\alpha < 1$). Thus $q_{s_\delta,\alpha}(Y_1) = 0$ and $q_{s_\delta,\alpha}(Y_{B(\delta)}) = 0$ which implies the desired result $R(\delta, s_\delta) = q_\alpha(Y_0)$.

When $N \geq 1$ the discrete distribution of Y_1 under s_δ is defined by

$$P_{s_\delta}(Y_1 = 0) = \frac{\alpha+1}{2} + \frac{1}{N}[1 - \frac{\alpha+1}{2}]\sum_{j=1}^{N} I(y_{1,j} = 0),$$
$$P_{s_\delta}(Y_1 = y_{1,i}) = \frac{1}{N}[1 - \frac{\alpha+1}{2}]\sum_{j=1}^{N} I(y_{1,j} = y_{1,i}) \text{ for all } i = 1, ..., N \text{ s.t. } y_{1,i} \neq 0, \qquad (5.19)$$

where the $y_{1,i}$ $i = 1, ..., N$ appear as the components of $\widetilde{w}$ in (5.18). Note that all probabilities in (5.19) are properly defined, that is, they are all in the interval $(0, 1]$, and sum up to 1. Just for clarity $\sum_{j=1}^{N} I(y_{1,j} = 0)$ equals the number of zero components in the vector $\widetilde{w}$. Note that several of the $y_{1,i}$ for $i = 1, ..., N$ may be identical and therefore Y_1 maybe supported on strictly fewer than $N + 1$ points.

With these definitions we obtain $q_{s_\delta,\alpha}(Y_1) = 0$ (because $.5(\alpha + 1) > \alpha$) and

$$\begin{aligned}
&P_{s_\delta}(Y_{B(\delta(w))} = 0) \\
&= \sum_{\substack{\widehat{w}=(w_1,...,w_N) \\ w_i \in \{0, y_{1,1},...,y_{1,N}\},\ i=1,...,N}} \prod_{i=1}^{N} P_{s_\delta}(Y_1 = w_i)[\delta(\widehat{w})P_{s_\delta}(Y_1 = 0) + (1 - \delta(\widehat{w}))P_{s_\delta}(Y_0 = 0)] \\
&= \{\sum_{\substack{\widehat{w}=(w_1,...,w_N) \\ w_i \in \{0, y_{1,1},...,y_{1,N}\},\ i=1,...,N}} \prod_{i=1}^{N} P_{s_\delta}(Y_1 = w_i)\delta(\widehat{w})[P_{s_\delta}(Y_1 = 0) - P_{s_\delta}(Y_0 = 0)]\} + P_{s_\delta}(Y_0 = 0) \\
&\geq \{\prod_{i=1}^{N} P_{s_\delta}(Y_1 = y_{1,i})\delta(\widetilde{w})[\frac{\alpha+1}{2} - (\alpha - \varepsilon)]\} + \alpha - \varepsilon, \qquad (5.20)
\end{aligned}$$

where the inequality is obtained by setting $\delta(\widehat{w}) = 0$ except for when $\widehat{w} = \widetilde{w}$, using that $P_{s_\delta}(Y_1 = 0) - P_{s_\delta}(Y_0 = 0) \geq \frac{\alpha+1}{2} - (\alpha - \varepsilon) > 0$. We claim that $P_{s_\delta}(Y_{B(\delta(w))} = 0) > \alpha$ (and thus $q_{s,\alpha}(Y_{B(\delta(w))}) = 0$) for sufficiently small $\varepsilon > 0$. But the claim is equivalent to

$$\prod_{i=1}^{N} P_{s_\delta}(Y_1 = y_{1,i})\delta(\widetilde{w})[\frac{1-\alpha}{2} + \varepsilon)] > \varepsilon \qquad (5.21)$$

which clearly holds true for $\varepsilon > 0$ sufficiently small (using $\alpha < 1$).

When $\alpha = 0$ (in which case we take $r = 1$) we can use the same proof structure except we use s_δ such that $P_{s_\delta}(Y_0 = q_\alpha(Y_0)) = 1$ (which implies $q_{s_\delta,\alpha}(Y_0) = q_\alpha(Y_0)$) and the distribution for Y_1 defined in (5.19) with $\alpha = 0$. □

Note throughout that for $s \in \mathbb{S}$ we have $q_\alpha(Y_0) = q_{s,\alpha}(Y_0)$.

Consider first the case $q_\alpha(Y_0) > 1/2$. For $\delta^0 \equiv 0$, we will show next that

$$\max_{s \in \mathbb{S}} R(\delta^0, s) = 1 - q_\alpha(Y_0). \qquad (5.22)$$

The statement in (5.22) follows easily because for any $s \in \mathbb{S}$ for which $q_{s,\alpha}(Y_1) = 1$ we obtain $R(\delta^0, s) = 1 - q_\alpha(Y_0)$. Furthermore, because always

$$q_{s,\alpha}(Y_{B(\delta(w))}) \in [\min\{q_{s,\alpha}(Y_0), q_{s,\alpha}(Y_1)\}, \max\{q_{s,\alpha}(Y_0), q_{s,\alpha}(Y_1)\}] \tag{5.23}$$

(as proven in Lemma 3(i) below) $1 - q_{s,\alpha}(Y_0)$ is the maximal possible regret for any $s \in \mathbb{S}$ for which $q_{s,\alpha}(Y_1) > q_{s,\alpha}(Y_0)$. For any $s \in \mathbb{S}$ such that $q_{s,\alpha}(Y_1) \leq q_{s,\alpha}(Y_0)$ regret is zero. That proves (5.22).

Because when $q_\alpha(Y_0) > 1/2$, we have $1 - q_{s,\alpha}(Y_0) < q_{s,\alpha}(Y_0)$, (5.22) combined with the lemma imply the desired result when $q_\alpha(Y_0) > 1/2$.

Next consider the case $q_\alpha(Y_0) = 1/2$. Statement (5.22) (that also holds when $q_\alpha(Y_0) = 1/2$), Lemma 1, together with the fact that regret is bounded by 1/2 (by (5.23) when $q_{s,\alpha}(Y_0) = 1/2$), establish that all treatment rules are minimax regret.

Finally, consider the case $q_\alpha(Y_0) < 1/2$. We have $\max_{s\in\mathbb{S}} R(\delta^*, s) = q_{s,\alpha}(Y_0)$ for $\delta^* = 1$. Lemma 1 establishes that for $\delta \neq 0$ $\max_{s\in\mathbb{S}} R(\delta, s) \geq q_{s,\alpha}(Y_0)$ while for $\delta^0 \equiv 0$, $\max_{s\in\mathbb{S}} R(\delta^0, s) = 1 - q_{s,\alpha}(Y_0) > q_{s,\alpha}(Y_0)$. It follows that $\delta^* = 1$ is a minimax rule.

Finally, the case where $\alpha = 1$ and $r = 0$ is proven as in part (i). □

Proof of Corollary 1(a) This statement follows as a corollary from the proof of Proposition 1 because that proof only used Bernoulli distributions for nature for parts (i)-(ii) and discrete distributions in part (iii).

(b) (i) Consider an arbitrary treatment rule δ. Consider a state of nature $s = s_{\alpha,\varepsilon,n}$ for $n \in \mathbb{N}$ such that Y_0 has density n on $[0, 1/n]$ and Y_1 has density $f = f_{\alpha,\varepsilon,n}$ which equals $n(\alpha - \varepsilon)$ on $[0, 1/n]$ and $n(1 - \alpha + \varepsilon)$ on $[1 - 1/n, 1]$ (and zero elsewhere) for some small $\varepsilon \in [0, \alpha/2]$ to be specified more precisely below. Note that these continuous distributions are chosen in order to approximate as $n \to \infty$ the Bernoulli distributions used in the proof of Proposition 1(i). With these definitions

$$q_{s,\alpha}(Y_0) = \alpha/n \text{ and } q_{s,\alpha}(Y_1) = 1 - 1/n + \varepsilon/(n(1 - \alpha + \varepsilon)). \tag{5.24}$$

Consider an arbitrary $q \in (0, 1)$. Then for all n such that $1/n < q < 1 - 1/n$

$$\begin{aligned}
&P_s(Y_{B(\delta(w))} \leq q) \\
&= \textstyle\int_0^{1/n} n ... \int_0^{1/n} n \int_0^1 f(y_{11}) ... \int_0^1 f(y_{1N_1}) [\delta(y)(\alpha - \varepsilon) + (1 - \delta(y))] dy_{1N_1} ... dy_{11} dy_{0N_0} ... dy_{01} \\
&\geq (\alpha - \varepsilon) \textstyle\int_0^{1/n} n ... \int_0^{1/n} n \int_0^1 f(y_{11}) ... \int_0^1 f(y_{1N_1}) dy_{1N_1} ... dy_{11} dy_{0N_0} ... dy_{01} \\
&= (\alpha - \varepsilon),
\end{aligned} \tag{5.25}$$

where $y := (y_{01}, ..., y_{0N_0}, y_{11}, ..., y_{1N_1})'$ and the inequality comes from picking $\delta^1 \equiv 1$ given that $\alpha - \varepsilon - 1 < 0$.

Note that if $P_s(\delta(w) = 1) < 1$ for one choice of $\varepsilon \in [0, \alpha/2]$ then it holds for any $\varepsilon \in [0, \alpha/2]$.

Case 1: $P_s(\delta(w) = 1) < 1$. By the same steps as in (5.25) with $\varepsilon = 0$ we obtain $P_s(Y_{B(\delta(w))} \leq q) > \alpha$ and, by continuity in ε, then also for $\varepsilon > 0$ sufficiently small. It follows that the α-quantile of $Y_{B(\delta(w))}$ is nonbigger than q.

Case 2: $P_s(\delta(w) = 1) = 1$. In that case, consider a state of nature $\widetilde{s} = \widetilde{s}_{\alpha,\varepsilon,n}$ for $n \in \mathbb{N}$ such that Y_1 has density n on $[0, 1/n]$ and Y_0 has density $\widetilde{f} = \widetilde{f}_{\alpha,\varepsilon,n}$ which equals $n(\alpha - \varepsilon)$ on $[0, 1/n]$ and $n(1 - \alpha + \varepsilon)$ on $[1 - 1/n, 1]$ (and zero elsewhere). Thus,

$$q_{\widetilde{s},\alpha}(Y_1) = \alpha/n \text{ and } q_{\widetilde{s},\alpha}(Y_0) = 1 - 1/n + \varepsilon/(n(1 - \alpha + \varepsilon)) \tag{5.26}$$

and, analogously to (5.25),

$$\begin{aligned}
&P_{\widetilde{s}}(Y_{B(\delta(w))} \leq q) \\
&= \int_0^1 \widetilde{f}(y_{01}) ... \int_0^1 \widetilde{f}(y_{0N_0}) \int_0^{1/n} n ... \int_0^{1/n} n[\delta(y) + (1 - \delta(y))(\alpha - \varepsilon)] dy_{1N_1} ... dy_{11} dy_{0N_0} ... dy_{01} \\
&= \int_0^{1/n} \widetilde{f}(y_{01}) ... \int_0^{1/n} \widetilde{f}(y_{0N_0}) \int_0^{1/n} n ... \int_0^{1/n} n dy_{1N_1} ... dy_{11} dy_{0N_0} ... dy_{01} \\
&+ \int ... \int_{A_n} \widetilde{f}(y_{01}) ... \widetilde{f}(y_{0N_0}) \int_0^{1/n} n ... \int_0^{1/n} n[\delta(y) + (1 - \delta(y))(\alpha - \varepsilon)] dy_{1N_1} ... dy_{11} dy_{0N_0} ... dy_{01} \\
&\geq (\alpha - \varepsilon)^{N_0} + (\alpha - \varepsilon) \int ... \int_{A_n} \widetilde{f}(y_{01}) ... \widetilde{f}(y_{0N_0}) dy_{0N_0} ... dy_{01} \\
&= (\alpha - \varepsilon)^{N_0} + (\alpha - \varepsilon)(1 - (\alpha - \varepsilon)^{N_0}),
\end{aligned} \tag{5.27}$$

where $A_n := [0,1]^{N_0} \backslash [0, 1/n]^{N_0}$, the second equality uses that $\delta(w) = 1$ a.s. on $[0, 1/n]^N$ (under s and therefore also under $\widetilde{s}$), and the inequality is obtained by taking $\delta(y) = 0$ on the domain of integration. If ε was equal to zero, the final expression in (5.27) would be strictly larger than α. By continuity in ε, that is then also true for sufficiently small $\varepsilon > 0$. It follows that the α-quantile of $Y_{B(\delta(w))}$ is nonbigger than q.

As q can be chosen arbitrarily small in that construction it follows that both in case 1 and case 2 regret arbitrarily close to 1 can be inflicted by nature given (5.24) and (5.26).

Parts (ii) and (iii) follow from similar constructions. □

Proof of the statements in Comment 3 below Corollary 1. Note first that $q(\delta)$ indeed exists because cdfs are continuous from the left. E.g. $q(0) = q_\alpha(Y_0)$ and $q(1) = 0$. Also note that for any $q \in [0, 1]$ the function $g(\delta) := \delta + (1 - \delta) F_{Y_0}(q)$ is weakly increasing in δ. It follows that $q(\delta)$ is weakly decreasing in δ. Therefore, $q_\alpha(Y_0) - q(\delta)$ is weakly increasing in δ.

If $\delta < 1$ then regret of $1 - q_\alpha(Y_0)$ can be achieved by nature by defining the distribution of Y_1 as $P_s(Y_1 = 0) = \alpha - \varepsilon$ and $P_s(Y_1 = 1) = 1 - (\alpha - \varepsilon)$ for some small $\varepsilon > 0$ (to be specified further below). Namely, with that definition, for $q_\alpha(Y_0) < 1$ we have

$$P_s(Y_{B(\delta)} \leq q_\alpha(Y_0)) = \delta(\alpha - \varepsilon) + (1 - \delta)F_{Y_0}(q_\alpha(Y_0)) \tag{5.28}$$

which exceeds α for ε small enough because $F_{Y_0}(q_\alpha(Y_0)) > \alpha$ and $\delta < 1$. And of course $P_s(Y_{B(\delta)} \leq q_\alpha(Y_0)) \geq \alpha$ when $q_\alpha(Y_0) = 1$. Because the α-quantile of $Y_{B(\delta)}$ can clearly not be smaller than $q_\alpha(Y_0)$ it follows that it is equal to $q_\alpha(Y_0)$.

But then, in the case $q_\alpha(Y_0) < 1/2$, it follows that for any $\delta \in [0, 1)$, $\max_{s\in\mathbb{S}} R(\delta, s) \geq 1 - q_\alpha(Y_0) > q_\alpha(Y_0)$. Given that for $\delta^1 \equiv 1$ maximal regret equals $q_\alpha(Y_0)$ it is the unique minimax regret rule. When $q_\alpha(Y_0) = 1/2$ regret cannot exceed $1/2$ and given regret of $1 - q_\alpha(Y_0)$ can always be achieved by the argument above (for $\delta < 1$ and for $q_\alpha(Y_0) = 1/2$ also for $\delta = 1$) it follows that all rules are minimax regret.

Furthermore, for any $\delta \in [0, 1]$ regret of $q_\alpha(Y_0) - q(\delta)$ can be achieved by nature by defining s, the distribution of Y_1, by $P_s(Y_1 = 0) = 1$. Then $P_s(Y_{B(\delta)} \leq q(\delta)) = \delta + (1 - \delta)F_{Y_0}(q(\delta))$ which by (2.12) and $r = 0$ implies $q_{s,\alpha}(Y_{B(\delta)}) = q(\delta)$. By construction, no bigger regret than $q_\alpha(Y_0) - q(\delta)$ can be achieved by nature for setups where $q_\alpha(Y_1) \leq q_\alpha(Y_0)$. That proves the claim in the case where $q_\alpha(Y_0) > 1/2$. $\square$

Proof of Corollary 2. Part (i). Let δ be an arbitrary treatment rule. We will construct a state of nature s_δ that yields regret equal to 1. Consider the case where all marginals under s_δ are independent Bernoulli distributions and define

$$a_{0x_k} = P_{s_\delta}(Y_{0,x_k} = 0) \text{ and } a_{1x_k} = P_{s_\delta}(Y_{1,x_k} = 0) \tag{5.29}$$

for $k = 1, ..., K$ (and, of course, $1 - a_{0x_k} = P_{s_\delta}(Y_{0,x_k} = 1)$ and $1 - a_{1x_k} = P_{s_\delta}(Y_{1,x_k} = 1)$). Notationwise, we index probabilities by those distributions that they depend on. Then,

$$P_{s_\delta,F_X}(Y_{B(\delta_X(w)),X} = 0) = \textstyle\sum_{k=1}^{K} P_{F_X}(X = x_k) P_{s_\delta,F_X}(Y_{B(\delta_{x_k}(w)),x_k} = 0) \tag{5.30}$$

and

$$\begin{aligned} &P_{s_\delta,F_X}(Y_{B(\delta_{x_k}(w)),x_k} = 0) \\ &= \sum_{\substack{\widetilde{x}=(\widetilde{x}_1,\ldots,\widetilde{x}_N)\in\mathcal{X}^N \\ y=(y_{t_1},\ldots,y_{t_N})\in\{0,1\}^N}} \textstyle\prod_{l=1}^{N_0} P_{F_X}(X = \widetilde{x}_l) a_{0\widetilde{x}_l}^{1-y_{t_l}} (1 - a_{0\widetilde{x}_l})^{y_{t_l}} \prod_{l'=N_0+1}^{N} P_{F_X}(X = \widetilde{x}_{l'}) a_{1\widetilde{x}_{l'}}^{1-y_{t_{l'}}} (1 - a_{1\widetilde{x}_{l'}})^{y_{t_{l'}}} \\ &\times [\delta_{x_k}(\widetilde{w}(\widetilde{x}, y)) a_{1x_k} + (1 - \delta_{x_k}(\widetilde{w}(\widetilde{x}, y))) a_{0x_k}], \end{aligned} \tag{5.31}$$

where in the last line $\widetilde{w}(\widetilde{x},y)$ denotes the sample associated with $\widetilde{x}$ and y. Note that expressions like $P_{s_\delta,F_X}(Y_{B(\delta_{x_k}(w)),x_k}=0)$ are still indexed by F_X because the distribution of w depends on F_X. In the case $(a_{0x_k},a_{1x_k})=(1,\alpha)$ for all $k=1,...,K$ (5.31) becomes

$$\begin{aligned}
&P_{s_\delta,F_X}(Y_{B(\delta_{x_k}(w)),x_k}=0)\\
&=\sum_{\substack{\widetilde{x}=(\widetilde{x}_1,...,\widetilde{x}_N)\in\mathcal{X}^N\\ y=(0,...,0,y_{t_{N_0+1}},...,y_{t_N})\in\{0\}^{N_0}\times\{0,1\}^{N_1}}}\prod_{l=1}^{N_0}P_{F_X}(X=\widetilde{x}_l)\prod_{l'=N_0+1}^{N}P_{F_X}(X=\widetilde{x}_{l'})\alpha^{1-y_{t_{l'}}}(1-\alpha)^{y_{t_{l'}}}\\
&\times[(\alpha-1)\delta_{x_k}(\widetilde{w}(\widetilde{x},y))+1]\\
&\geq[\sum_{\substack{\widetilde{x}=(\widetilde{x}_1,...,\widetilde{x}_N)\in\mathcal{X}^N\\ (y_{t_{N_0+1}},...,y_{t_N})\in\{0,1\}^{N-N_0}}}\prod_{l=1}^{N_0}P_{F_X}(X=\widetilde{x}_l)\prod_{l'=N_0+1}^{N}P_{F_X}(X=\widetilde{x}_{l'})\alpha^{1-y_{t_{l'}}}(1-\alpha)^{y_{t_{l'}}}]\alpha\\
&=\alpha,
\end{aligned}\tag{5.32}$$

where the inequality follows from taking $\delta_{x_k}(\widetilde{w}(\widetilde{x},y))=1$.

If for some $k\in\{1,...,K\}$ $\delta_{x_k}(\widetilde{w}(\widetilde{x},y))<1$ for some $\widetilde{w}(\widetilde{x},y)$ (with $\widetilde{x}=(\widetilde{x}_1,...,\widetilde{x}_N)\in\mathcal{X}^N$ and $y=(0,...,0,y_{t_{N_0+1}},...,y_{t_N})$ with $y_{t_j}\in\{0,1\}$ for $j=N_0+1,...,N$) it follows that $P_{s_\delta,F_X}(Y_{B(\delta_{x_k}(w)),x_k}=0)>\alpha$ and by (5.30) therefore also $P_{s_\delta,F_X}(Y_{B(\delta_X(w)),X}=0)>\alpha$. Notice that $P_{s_\delta,F_X}(Y_{B(\delta_X(w)),X}=0)$ is continuous in a_{1x_k} for $k=1,...,K$ and therefore $P_{s_\delta,F_X}(Y_{B(\delta_X(w)),X}=0)>\alpha$ when instead of $(a_{0x_k},a_{1x_k})=(1,\alpha)$ for $k=1,...,K$, s_δ is defined by $(a_{0x_k},a_{1x_k})=(1,\alpha-\varepsilon)$ for $k=1,...,K$ for some small enough $\varepsilon>0$. But then for that choice of s_δ, regret equals 1 (as can be seen by comparing to the quantile obtained for the treatment rule $\delta\equiv 1$) as desired.

If instead for all $k\in\{1,...,K\}$ $\delta_{x_k}(\widetilde{w}(\widetilde{x},y))=1$ for all $\widetilde{w}(\widetilde{x},y)$ (with $\widetilde{x}=(\widetilde{x}_1,...,\widetilde{x}_N)\in\mathcal{X}^N$ and $y=(0,...,0,y_{t_{N_0+1}},...,y_{t_N})$ with $y_{t_j}\in\{0,1\}$ for $j=N_0+1,...,N$) nature can cause regret of 1 by using distributions defined by $(a_{0x_k},a_{1x_k})=(\alpha-\varepsilon,1)$ for some small enough $\varepsilon>0$. Namely, note first that for s_δ defined by $(a_{0x_k},a_{1x_k})=(\alpha,1)$ for $k=1,...,K$ we obtain from (5.31)

$$\begin{aligned}
&P_{s_\delta,F_X}(Y_{B(\delta_{x_k}(w)),x_k}=0)\\
&=\sum_{\substack{\widetilde{x}=(\widetilde{x}_1,...,\widetilde{x}_N)\in\mathcal{X}^N\\ y=(y_{t_1},...,y_{t_{N_0}},0,...,0)\in\{0,1\}^{N_0}\times\{0\}^{N_1}}}\prod_{l=1}^{N_0}P_{F_X}(X=\widetilde{x}_l)\alpha^{1-y_{t_l}}(1-\alpha)^{y_{t_l}}\prod_{l'=N_0+1}^{N}P_{F_X}(X=\widetilde{x}_{l'})\\
&\times[\delta_{x_k}(\widetilde{w}(\widetilde{x},y))(1-\alpha)+\alpha]\\
&\geq\alpha^{N_0}\sum_{\widetilde{x}=(\widetilde{x}_1,...,\widetilde{x}_N)\in\mathcal{X}^N}\prod_{l=1}^{N}P_{F_X}(X=\widetilde{x}_l)\\
&+\alpha\sum_{\substack{\widetilde{x}=(\widetilde{x}_1,...,\widetilde{x}_N)\in\mathcal{X}^N\\ y=(y_{t_1},...,y_{t_{N_0}},0,...,0)\in\{0,1\}^{N_0}\times\{0\}^{N_1},y\neq 0^N}}\prod_{l=1}^{N_0}P_{F_X}(X=\widetilde{x}_l)\alpha^{1-y_{t_l}}(1-\alpha)^{y_{t_l}}(\prod_{l'=N_0+1}^{N}P_{F_X}(X=\widetilde{x}_{l'}))\\
&=\alpha^{N_0}+\alpha-\alpha[\sum_{\widetilde{x}=(\widetilde{x}_1,...,\widetilde{x}_N)\in\mathcal{X}^N}\prod_{l=1}^{N_0}(P_{F_X}(X=\widetilde{x}_l)\alpha)(\prod_{l'=N_0+1}^{N}P_{F_X}(X=\widetilde{x}_{l'}))]\\
&=\alpha^{N_0}+\alpha-\alpha^{N_0+1},
\end{aligned}\tag{5.33}$$

where the inequality follows from setting $\delta_{x_k}(\widetilde{w}(\widetilde{x}, y)) = 0$ for all elements in the sum except for those when $y = 0^N$ in which case we use $\delta_{x_k}(\widetilde{w}(\widetilde{x}, y)) = 1$. The second to last equality follows by first considering a sum over all $y = (y_{t_1}, ..., y_{t_{N_0}}, 0, ..., 0) \in \{0,1\}^{N_0} \times \{0\}^{N_1}$ in the second sum and then subtracting the summand associated with $y = 0^N$. The final expression is strictly larger than α which by an argument as above allows us to conclude that $P_{s_\delta, F_X}(Y_{B(\delta_{x_k}(w)), x_k} = 0) > \alpha$ when s_δ is defined by $(a_{0x_k}, a_{1x_k}) = (\alpha - \varepsilon, 1)$ for $k = 1, ..., K$ for some small enough $\varepsilon > 0$.

The proof of **part (ii)** is very similar to part (i) and omitted.

Part (iii). The key input of the proof is the following lemma.

Lemma 2 *For any treatment rule $\delta \in \mathbb{D}$ such that $\delta \neq 0$ there exists an $s_\delta \in \mathbb{S}$ such that $R(\delta, s_\delta) = q_\alpha(Y_{0,X})$.*

Proof of Lemma 2. We define an $s_\delta \in \mathbb{S}$ by letting all marginals (Y_{t,x_k}) be independent and discrete, namely for an $\varepsilon > 0$ (to be specified later) let $P_{s_\delta}(Y_{0,x_k} = 0) = \alpha - \varepsilon$ and $P_{s_\delta}(Y_{0,x_k} = q_\alpha(Y_{0,X})) = 1 - (\alpha - \varepsilon)$ for every $x_k \in \mathcal{X}$. Because $P_{s_\delta, F_X}(Y_{0,X} = 0) = \sum_{k=1}^K P_{F_X}(X = x_k) P_{s_\delta}(Y_{0,x_k} = 0) = \alpha - \varepsilon$, it follows that $q_{s, F_X, \alpha}(Y_{0,X}) = q_\alpha(Y_{0,X})$ as required.

For $N = 0$, complete the definition of s_δ by setting $P_{s_\delta}(Y_{1,x_k} = 0) = 1$ for every $x_k \in \mathcal{X}$. Wlog assume $\delta_{x_l} > 0$. Note that $P_{s_\delta}(Y_{B(\delta_{x_l}), x_l} = 0) = \delta_{x_l} + (1 - \delta_{x_l})(\alpha - \varepsilon)$ which, given that $\delta_{x_l} > 0$, is strictly larger than α when $\varepsilon > 0$ is chosen sufficiently small. For every $x_k \in \mathcal{X}$ such that $x_k \neq x_l$ we have $P_{s_\delta}(Y_{B(\delta_{x_k}), x_k} = 0) \geq \alpha - \varepsilon$. Therefore,

$$
\begin{aligned}
& P_{s_\delta, F_X}(Y_{B(\delta_X), X} = 0) \\
&= \sum_{k=1}^K P_{F_X}(X = x_k) P_{s_\delta}(Y_{B(\delta_{x_k}), x_k} = 0) \\
&\geq P_{F_X}(X = x_l)(\delta_{x_l} + (1 - \delta_{x_l})(\alpha - \varepsilon)) + \sum_{k=1, k \neq l}^K P_{F_X}(X = x_k)(\alpha - \varepsilon) \\
&= P_{F_X}(X = x_l)\delta_{x_l}(1 - (\alpha - \varepsilon)) + (\alpha - \varepsilon) \\
&> \alpha, \qquad (5.34)
\end{aligned}
$$

where the second equality follows by including the summand $P_{F_X}(X = x_l)(\alpha - \varepsilon)$ into the summation sign, and the inequality holds for $\varepsilon > 0$ chosen small enough.

Next consider the case $N > 0$. Because $\delta \neq 0$ there exists some vector $\widetilde{w} = ((\widetilde{x}_1, \widetilde{y}_1), ..., (\widetilde{x}_N, \widetilde{y}_N))$ with $(\widetilde{x}_i, \widetilde{y}_i) \in \mathcal{X} \times [0,1]$ for every $i = 1, ..., N$ and some covariate $x_l \in \mathcal{X}$ such that $\delta_{x_l}(\tilde{w}) > 0$. Similarly to the proof of Proposition 1(iii), for every $x_k \in \mathcal{X}$ define the distribution of Y_{1,x_k}

under s_δ as

$$P_{s_\delta}(Y_{1,x_k}=0)=\frac{\alpha+1}{2}+\frac{1}{N}[1-\frac{\alpha+1}{2}]\sum_{j=1}^{N}I(\widetilde{y}_j=0),$$
$$P_{s_\delta}(Y_{1,x_k}=\widetilde{y}_i)=\frac{1}{N}[1-\frac{\alpha+1}{2}]\sum_{j=1}^{N}I(\widetilde{y}_j=\widetilde{y}_i)\text{ for }i=1,...,N\text{ s.t. }\widetilde{y}_i\neq 0. \quad (5.35)$$

Because $P_{s_\delta,F_X}(Y_{1,X}=0)=\sum_{k=1}^{K}P_{F_X}(X=x_k)P_{s_\delta}(Y_{1,x_k}=0)$, it then follows that $q_{s,F_x,\alpha}(Y_{1,X})=0$. Now, focus on $P_{s_{\delta,F_X}}(Y_{B(\delta_{x_l}(w)),x_l}=0)$. Furthermore,

$$\begin{aligned}
&P_{s_{\delta,F_X}}(Y_{B(\delta_{x_l}(w)),x_l}=0)\\
&=\sum_{\substack{w=((a_1,b_1),...,(a_N,b_N));\\ a_i\in\mathcal{X};\ b_i\in\{0,\widetilde{y}_1,...,\widetilde{y}_N\},\\ i=1,...,N}}\prod_{i=1}^{N}P_{s_\delta,F_x}(Y_{1,a_i}=b_i,X=a_i)(\delta_{x_l}(w)P_{s_\delta}(Y_{1,x_l}=0)+(1-\delta_{x_l}(w))P_{s_\delta}(Y_{0,x_l}=0))\\
&=P_{s_\delta}(Y_{0,x_l}=0)+\sum_{\substack{w=((a_1,b_1),...,(a_N,b_N));\\ a_i\in\mathcal{X};\ b_i\in\{0,\widetilde{y}_1,...,\widetilde{y}_N\},\\ i=1,...,N}}\prod_{i=1}^{N}P_{s_\delta,F_x}(Y_{1,a_i}=b_i,X=a_i)(\delta_{x_l}(w)(P_{s_\delta}(Y_{1,x_l}=0)-P_{s_\delta}(Y_{0,x_l}=0))\\
&\geq\alpha-\varepsilon+\prod_{i=1}^{N}P_{s_\delta}(Y_{1,\widetilde{x}_i}=\widetilde{y}_i)P_{F_X}(X=\widetilde{x}_i)\delta_{x_l}(\tilde{w})(\frac{\alpha+1}{2}-(\alpha-\varepsilon))\\
&>\alpha,
\end{aligned}\quad (5.36)$$

where the first inequality follows by replacing $\delta_{x_l}(w)$ by zero for every $w\neq\tilde{w}$ and the last inequality holds for small enough $\varepsilon>0$. By the same derivations, for any other $x_k\neq x_l$ it follows that $P_{s_\delta,F_X}(Y_{B(\delta_{x_k}(w)),x_k}=0)\geq\alpha-\varepsilon$. Therefore, for sufficiently small $\varepsilon>0$ we have

$$\begin{aligned}
&P_{s_\delta,F_x}(Y_{B(\delta_X(w)),X}=0)\\
&=\sum_{k=1}^{K}P_{F_X}(X=x_k)P_{s_\delta}(Y_{B(\delta_{x_k}(w)),x_k}=0)\\
&\geq\alpha-\varepsilon+P_{F_X}(X=x_l)\prod_{i=1}^{N}P_{s_\delta}(Y_{1,\widetilde{x}_i}=\widetilde{y}_i)P_{F_X}(X=\widetilde{x}_i)\delta_{x_l}(\tilde{w})(\frac{\alpha+1}{2}-(\alpha-\varepsilon))\\
&>\alpha,
\end{aligned}\quad (5.37)$$

where for the inequality we use $P_{s_\delta}(Y_{B(\delta_{x_k}(w)),x_k}=0)\geq\alpha-\varepsilon$ and the second to last line in (5.36). Therefore, we have $q_{s,F_X,\alpha}(Y_{0,X})=q_\alpha(Y_{0,X})$ and $q_{s,F_X,\alpha}(Y_{1,X})=q_{s,F_X,\alpha}(Y_{B(\delta_X(w)),X})=0$. Hence, the statement of the lemma follows. □

Given Lemma 2 the remainder of the proof of part (iii) is exactly the same as the last part of the proof of Proposition 1(iii). The only detail to think about is establishing that

$$q_{s,F_X,\alpha}(Y_{B(\delta_X(w)),X}) \in [min\{q_{s,F_X,\alpha}(Y_{0,X}), q_{s,F_X,\alpha}(Y_{1,X})\}, max\{q_{s,F_X,\alpha}(Y_{0,X}), q_{s,F_X,\alpha}(Y_{1,X})\}]. \tag{5.38}$$

The proof of that statement is analogous to the proof of Lemma 3 and therefore omitted. □

The following is a technical lemma that is used in the proof of Proposition 1 above.

Lemma 3 *For any given* $r \in [0, 1]$ *and arbitrary* $s \in \mathbb{S}$ *we have* (*i*)

$$q_{s,\alpha}(Y_{B(\delta(w))}) \in [\min\{q_{s,\alpha}(Y_0), q_{s,\alpha}(Y_1)\}, \max\{q_{s,\alpha}(Y_0), q_{s,\alpha}(Y_1)\}]$$

and (*ii*)

$$\max_{\delta\in\mathbb{D}} u(\delta, s) = \max_{\delta\in\mathbb{D}} q_{s,\alpha}(Y_{B(\delta(w))}) = \max\{q_{s,\alpha}(Y_0), q_{s,\alpha}(Y_1)\}.$$

Proof of Lemma 3. (i) Denote by $\overline{q}_{0,\alpha}$, $\overline{q}_{1,\alpha}$, and $\overline{q}_{\delta,\alpha}$ the largest α-quantile (meaning employing $r = 1$ in definition (2.5)) of Y_0, Y_1, and $Y_{B(\delta(w))}$, respectively. Likewise, denote by $\underline{q}_{0,\alpha}$, $\underline{q}_{1,\alpha}$, and $\underline{q}_{\delta,\alpha}$ the smallest α-quantile (meaning employing $r = 0$ in definition (2.5)) of Y_0, Y_1, and $Y_{B(\delta(w))}$, respectively. As in (5.1)

$$P_s(Y_{B(\delta(w))} \geq q) = P_s(B(\delta(w)) = 1)P_s(Y_1 \geq q) + P_s(B(\delta(w)) = 0)P_s(Y_0 \geq q). \tag{5.39}$$

We show first that

$$\overline{q}_{\delta,\alpha} \in [\min\{\overline{q}_{0,\alpha}, \overline{q}_{1,\alpha}\}, \max\{\overline{q}_{0,\alpha}, \overline{q}_{1,\alpha}\}]. \tag{5.40}$$

Indeed, if $\overline{q}_{\delta,\alpha} > \max\{\overline{q}_{0,\alpha}, \overline{q}_{1,\alpha}\}$ then, from (5.39)

$$\begin{aligned} &P_s(Y_{B(\delta(w))} \geq \overline{q}_{\delta,\alpha}) \\ &= P_s(B(\delta(w)) = 1)P_s(Y_1 \geq \overline{q}_{\delta,\alpha}) + P_s(B(\delta(w)) = 0)P_s(Y_0 \geq \overline{q}_{\delta,\alpha}) \\ &< 1 - \alpha, \end{aligned} \tag{5.41}$$

a contradiction to the definition of an α-quantile. On the other hand, if $\overline{q}_{\delta,\alpha} < \min\{\overline{q}_{0,\alpha}, \overline{q}_{1,\alpha}\}$

then there exists $\varepsilon > 0$ such that $\overline{q}_{\delta,\alpha} + \varepsilon < \min\{\overline{q}_{0,\alpha}, \overline{q}_{1,\alpha}\}$. But then

$$
\begin{aligned}
& P_s(Y_{B(\delta(w))} \geq \overline{q}_{\delta,\alpha} + \varepsilon) \\
& = P_s(B(\delta(w)) = 1)P_s(Y_1 \geq \overline{q}_{\delta,\alpha} + \varepsilon) + P_s(B(\delta(w)) = 0)P_s(Y_0 \geq \overline{q}_{\delta,\alpha} + \varepsilon) \\
& \leq P_s(B(\delta(w)) = 1)P_s(Y_1 \geq \overline{q}_{1,\alpha}) + P_s(B(\delta(w)) = 0)P_s(Y_0 \geq \overline{q}_{0,\alpha}) \\
& = 1 - \alpha, \qquad (5.42)
\end{aligned}
$$

contradicting the fact that $\overline{q}_{\delta,\alpha}$ is the largest α-quantile of $Y_{B(\delta(w))}$.

By an analogous argument it follows that

$$
\underline{q}_{\delta,\alpha} \in [\min\{\underline{q}_{0,\alpha}, \underline{q}_{1,\alpha}\}, \max\{\underline{q}_{0,\alpha}, \underline{q}_{1,\alpha}\}]. \qquad (5.43)
$$

Define the interval $I = [\min\{\underline{q}_{0,\alpha}, \underline{q}_{1,\alpha}\}, \max\{\overline{q}_{0,\alpha}, \overline{q}_{1,\alpha}\}]$.

If $I = [\underline{q}_{i,\alpha}, \overline{q}_{j,\alpha}]$ for $i, j \in \{0, 1\}$, $i \neq j$ then

$$
q_{s,\alpha}(Y_{B(\delta(w))}) = r\overline{q}_{\delta,\alpha} + (1 - r)\underline{q}_{\delta,\alpha} \qquad (5.44)
$$

and thus by (5.40) and (5.43)

$$
\begin{aligned}
& q_{s,\alpha}(Y_{B(\delta(w))}) \leq r\overline{q}_{j,\alpha} + (1 - r)\underline{q}_{j,\alpha} = q_{s,\alpha}(Y_j) \leq \max\{q_{s,\alpha}(Y_0), q_{s,\alpha}(Y_1)\} \text{ and} \\
& q_{s,\alpha}(Y_{B(\delta(w))}) \geq r\overline{q}_{i,\alpha} + (1 - r)\underline{q}_{i,\alpha} = q_{s,\alpha}(Y_i) \leq \max\{q_{s,\alpha}(Y_0), q_{s,\alpha}(Y_1)\}. \qquad (5.45)
\end{aligned}
$$

For the remaining case we have $I = [\underline{q}_{i,\alpha}, \overline{q}_{i,\alpha}]$ for a $i \in \{0, 1\}$. By (5.40) and (5.43) it follows that $[\underline{q}_{\delta,\alpha}, \overline{q}_{\delta,\alpha}] \subset I$ and $[\underline{q}_{j,\alpha}, \overline{q}_{j,\alpha}] \subset [\underline{q}_{\delta,\alpha}, \overline{q}_{\delta,\alpha}]$.

If $P_s(B(\delta(w)) = i) = 1$ then $Y_{B(\delta(w))} = Y_i$ wp1 and the lemma trivially holds. If $P_s(B(\delta(w)) = i) < 1$, we must have $[\underline{q}_{j,\alpha}, \overline{q}_{j,\alpha}] = [\underline{q}_{\delta,\alpha}, \overline{q}_{\delta,\alpha}]$. If not, if e.g. $\overline{q}_{j,\alpha} < \overline{q}_{\delta,\alpha}$ then $P_s(Y_j \geq \overline{q}_{\delta,\alpha}) < 1 - \alpha$ and $P_s(Y_i \geq \overline{q}_{\delta,\alpha}) = 1 - \alpha$. But then, by (5.39) we get the contradiction

$$
P_s(Y_{B(\delta(w))} \geq \overline{q}_{\delta,\alpha}) = P_s(B(\delta(w)) = 1)P_s(Y_1 \geq \overline{q}_{\delta,\alpha}) + P_s(B(\delta(w)) = 0)P_s(Y_0 \geq \overline{q}_{\delta,\alpha}) < 1 - \alpha, \qquad (5.46)
$$

where the last inequality uses $P_s(B(\delta(w)) = i) < 1$. But if $[\underline{q}_{j,\alpha}, \overline{q}_{j,\alpha}] = [\underline{q}_{\delta,\alpha}, \overline{q}_{\delta,\alpha}]$ then $q_{s,\alpha}(Y_{B(\delta(w))}) = q_{s,\alpha}(Y_j)$ and the claim in the lemma holds.

(ii) Considering the two rules $\delta^1 \equiv 1$ or $\delta^0 \equiv 0$ (that is, the rules that pick 1 (or 0) with probability 1), it follows that $\max_{\delta \in \mathbb{D}} q_{s,\alpha}(Y_{B(\delta(w))}) \geq \max\{q_{s,\alpha}(Y_0), q_{s,\alpha}(Y_1)\}$. But given that $q_{s,\alpha}(Y_{B(\delta(w))}) \leq \max\{q_{s,\alpha}(Y_0), q_{s,\alpha}(Y_1)\}$ for any $\delta \in \mathbb{D}$ by part (ii), the claim follows. $\square$

References

Berger, J. (1985), Statistical Decision Theory and Bayesian Analysis, Second Edition, New York: SpringerVerlag.

Chen, H. and P. Guggenberger (2024), "A note on minimax regret rules with multiple treatments in finite samples," unpublished working paper.

Chambers, C. (2009), "An axiomatization of quantiles on the domain of distribution functions," *Mathematical Finance* 19, 335–342.

Christensen, T., R. Moon, and F. Schorfheide (2023), "Optimal Discrete Decisions when Payoffs are Partially Identified", unpublished working paper.

Cucconi, O. (1968), "Contributi all'Analisi Sequenziale nel Controllo di Accettazione per Variabili. *Atti dell' Associazione Italiana per il Controllo della Qualità* 6, 171–186.

De Castro, L., A. Galvao, and H. Ota (2023), "Quantile Approach to Intertemporal Consumption with Multiple Assets", unpublished working paper.

Gupta, S. and S. Hande (1992), "On some nonparametric selection procedures. In: Saleh, A.K.Md.E. (Ed.), Nonparametric Statistics and Related Topics. Elsevier.

Hirano, K. and J. Porter (2009), "Asymptotics for Statistical Treatment Rules," *Econometrica*, 77, 1683–1701.

Kitagawa, T., S. Lee, and C. Qiu (2024), "Treatment Choice with Nonlinear Regret," unpublished working paper.

Kitagawa, T. and A. Tetenov (2018), "Who Should be Treated? Empirical Welfare Maximization Methods for Treatment Choice," *Econometrica*, 86, 591–616.

Manski, C. (1988), "Ordinal UtilityModels of Decision Making Under Uncertainty," *Theory and Decision*, 25, 79–104.

_________(2004), "Statistical Treatment Rules for Heterogeneous Populations," *Econometrica*, 72, 221–246.

_________ and A. Tetenov (2007), "Admissible Treatment Rules for a Risk-averse Planner with Experimental Data on an Innovation," *Journal of Statistical Planning and Inference*, 137, 1998–2010.

_________(2023), "Statistical Decision Theory Respecting Stochastic Dominance," *The Japanese Economic Review*, 74, 447–469.

Masten, M. (2023), "Minimax-regret treatment rules with many treatments," *The Japanese Economic Review*, 74, 501–537.

Montiel Olea, J.L, C. Qiu, and J. Stoye (2023), "Decision Theory for Treatment Choice with Partial Identification," working paper, Cornell University.

Qi, Z., Y. Cui, Y. Liu, and J.-S. Pang (2019), "Estimation of Individualized Decision Rules Based on An Optimized Covariate-dependent Equivalent of Random Outcomes," SIAM Journal on Optimization 29 (3), 2337–2362

Qi, Z., J. Pang, and Y. Liu (2023), "On Robustness of Individualized Decision Rules," *Journal of the American Statistical Association,* 118 (543), 2143–2157.

Rostek, M.J. (2010), "Quantile Maximization in Decision Theory," *The Review of Economic Studies,* 77, 339–371.

Schlag, K. (2003), "How to minimize maximum regret under repeated decision-making," EUI working paper.

______ (2006), "ELEVEN - Tests needed for a recommendation," EUI working paper, ECO No. 2006/2.

Stoye, J. (2007), "Minimax regret treatment choice with incomplete data and many treatments," *Econometric Theory*, 23(1), 190–199.

______ (2009), "Minimax Regret Treatment Choice with Finite Samples," *Journal of Econometrics*, 151, 70–81.

______ (2012), "Minimax Regret Treatment Choice with Covariates or with Limited Validity of Experiments," *Journal of Econometrics*, 166, 138–156.

Tetenov, A. (2012), "Statistical Treatment Choice Based on Asymmetric Minimax Regret Criteria," *Journal of Econometrics*, 166, 157–165.

Wald, A. (1950), Statistical Decision Functions, New York: Wiley.

Yata, Kohei (2023), "Optimal Decision Rules Under Partial Identification," unpublished working paper.

TABLE I: Maximal and mean regret over all 18564 states of nature $s \in S^E(6,12)$ for four different treatment rules

Case I)	δ^{ES}	δ^1	$\delta^{.5}$	δ^0	δ^{ES}	δ^1	$\delta^{.5}$	δ^0	δ^{ES}	δ^1	$\delta^{.5}$	δ^0
α	.1	.1	.1	.1	.5	.5	.5	.5	.9	.9	.9	.9
$q_\alpha(Y_0) = .1$												
max	.9	**.1**	.9	.9	.1	**.1**	.9	.9	.1	**.1**	.9	.9
mean	.08	**.04**	.14	.1	.01	**.01**	.37	.37	0	**0**	.08	.75
$q_\alpha(Y_0) = .5$												
max	**.5**	**.5**	**.5**	**.5**	**.5**	**.5**	**.5**	**.5**	**.5**	**.5**	**.5**	**.5**
mean	**.43**	**.35**	**.36**	**0**	**.13**	**.11**	**.15**	**.08**	**.01**	**0**	**.07**	**.35**
$q_\alpha(Y_0) = .9$												
max	.9	.9	.9	**.1**	.9	.9	.9	**.1**	.9	.9	.9	**.1**
mean	.9	.75	.75	**0**	.8	.44	.37	**0**	.33	.1	.11	**.04**
Case II)	δ^{ES}	δ^1	$\delta^{.5}$	δ^0	δ^{ES}	δ^1	$\delta^{.5}$	δ^0	δ^{ES}	δ^1	$\delta^{.5}$	δ^0
α	.1	.1	.1	.1	.5	.5	.5	.5	.9	.9	.9	.9
$q_\alpha(Y_0) = .1$												
max	.84	**.1**	.9	.9	.1	**.1**	.9	.9	0	**.1**	.9	.9
mean	.05	**.04**	.14	.1	.01	**.01**	.37	.37	0	**0**	0.08	.75
$q_\alpha(Y_0) = .5$												
max	**.5**	**.5**	**.5**	**.5**	**.5**	**.5**	**.5**	**.5**	**.02**	**.5**	**.5**	**.5**
mean	**.06**	**.35**	**.36**	**0**	**.08**	**.11**	**.15**	**.08**	**0**	**0**	**0.07**	**.35**
$q_\alpha(Y_0) = .9$												
max	.9	.9	.9	**.1**	.9	.9	.9	**.1**	.03	.9	.9	**.1**
mean	0	.75	.75	**0**	0.04	.44	.37	**0**	0.01	.1	.11	**.04**

TABLE II: Proportion (in %) of states $s \in S^E(6,12)$ for which regret for the empirical success rule δ^{ES} is smaller than regret of non-data rules $\delta \in \{\delta^1, \delta^{.5}, \delta^0\}$

Case I)	δ^1	$\delta^{.5}$	δ^0	δ^1	$\delta^{.5}$	δ^0	δ^1	$\delta^{.5}$	δ^0
α	.1	.1	.1	.5	.5	.5	.9	.9	.9
$q_\alpha(Y_0) = .1$									
$Prop(R(\delta^{ES}, s) < R(\delta, s))$	0	33.3	33.3	0	92.5	92.5	0	33.3	100
$Prop(R(\delta^{ES}, s) \leq R(\delta, s))$	76.5	100	56.9	97.5	97.5	95.0	100	100	100
$q_\alpha(Y_0) = .5$									
$Prop(R(\delta^{ES}, s) < R(\delta, s))$	0	.5	.5	3.0	30.8	30.8	0	31.6	91.7
$Prop(R(\delta^{ES}, s) \leq R(\delta, s))$	61.0	62.7	8.3	90.7	76.2	58.6	99.5	99.5	97.8
$q_\alpha(Y_0) = .9$									
$Prop(R(\delta^{ES}, s) < R(\delta, s))$	0	0	0	0	2.5	2.5	0	16.2	43.1
$Prop(R(\delta^{ES}, s) \leq R(\delta, s))$	43.3	43.4	0	21.0	11.7	2.5	66.7	66.7	43.1
Case II)	δ^1	$\delta^{.5}$	δ^0	δ^1	$\delta^{.5}$	δ^0	δ^1	$\delta^{.5}$	δ^0
α	.1	.1	.1	.5	.5	.5	.9	.9	.9
$q_\alpha(Y_0) = .1$									
$Prop(R(\delta^{ES}, s) < R(\delta, s))$	16.2	73.0	73.0	.3	95.3	95.3	0	33.4	100
$Prop(R(\delta^{ES}, s) \leq R(\delta, s))$	76.5	100	73.0	100	97.5	95.3	100	100	100
$q_\alpha(Y_0) = .5$									
$Prop(R(\delta^{ES}, s) < R(\delta, s))$	71.2	73.4	71.6	13.6	40.7	39.8	2.2	33.8	91.7
$Prop(R(\delta^{ES}, s) \leq R(\delta, s))$	98.3	100	79.6	100	86.6	65.3	100	100	97.8
$q_\alpha(Y_0) = .9$									
$Prop(R(\delta^{ES}, s) < R(\delta, s))$	99.5	99.6	99.6	84.0	82.1	75.6	56.9	73.0	43.1
$Prop(R(\delta^{ES}, s) \leq R(\delta, s))$	100	100	99.6	100	95.7	75.6	100	100	43.1