Minimax Regret Treatment Choice with Covariates or with Limited Validity of Experiments

Jörg Stoye*

Cornell University

April 25, 2011

Abstract

This paper continues the investigation of minimax regret treatment choice initiated by Manski (2004). Consider a decision maker who must assign treatment to future subjects after observing outcomes experienced in a sample. A certain scoring rule is known to achieve minimax regret in simple versions of this decision problem. I investigate its sensitivity to perturbations of the decision environment in realistic directions. They are: (i) Treatment outcomes may be influenced by a covariate whose effect on outcome distributions is bounded (in one of numerous probability metrics). This is interesting because introduction of a covariate with unrestricted effects leads to a pathological result. (ii) The experiment may have limited validity, for example because of selective noncompliance or because the sampling universe is a potentially selective subset of the treatment population. Thus, even large samples may generate misleading signals. These problems are formalized via a "bounds" approach that turns the problem into one of partial identification.

In both scenarios, small but positive perturbations leave the minimax regret decision rule unchanged. Thus, minimax regret analysis is not knife-edge dependent on ignoring certain aspects of realistic decision problems. Indeed, it recommends to entirely disregard covariates whose effect is believed to be positive but small, as well as small enough amounts of missing data or selective attrition. All findings are finite sample results derived by game theoretic analysis.

Keywords: Finite sample theory, statistical decision theory, minimax regret, treatment response, treatment choice.

*Paper prepared for the CeMMAP conference on Identification and Decisions in honor of Charles Manski's 60^{th} birthday. For useful comments, I thank audiences at that conference as well as at Arizona, Collegio Carlo Alberto, Duke, NYU, Texas-Austin, UCLA, UCSD, Virginia, and Yale, and especially a guest editor and two anonymous referees. I also thank audiences at the Cowles Foundation and ISIPTA 07 (Prague) for comments on older work superseded by this paper. Part of the paper was written while visiting the Cowles Foundation, whose hospitality is gratefully acknowledged. Address: Jörg Stoye, Department of Economics, Cornell University, Uris Hall, Ithaca, NY 14853, stoye@cornell.edu.

1 Introduction

One recent focus of Charles Manski's research is to apply statistical decision theory, and in particular the minimax regret criterion, to problems of treatment choice (Manski (2000, 2004, 2006, 2007a, 2007b, 2009); Brock and Manski (2011)). This has spawned a small but active literature[1] with numerous results, not all of which were positive. In particular, Stoye's (2009a) findings on covariates – to be elaborated below – appear discouraging.

This paper provides new results that may be considered better news. I reconsider a stylized model of sample-based treatment assignment, variations of which were previously investigated by Canner (1970) and in a number of recent papers (Hirano and Porter (2009), Manski (2004), Schlag (2006), Stoye (2009a), Tetenov (2009a)). Minimax regret treatment rules for this problem are known and, by and large, intuitively reasonable. It is also known, however, that an apparently modest modification of the decision problem, namely the introduction of an observable covariate, induces a pathological result (Stoye (2009a)). Minimax regret then recommends to treat subjects with different covariate values as if they were sampled from completely unrelated populations; in the extreme, this leads to a "no-data rule" that completely ignores sample information. This finding raises at least two questions: Can we redeem minimax regret when there are covariates? More generally, are the minimax regret recommendations previously discovered overly dependent on certain simplifications made in the stylized problem?

I investigate this question by analyzing two modifications of the problem. First, a covariate is introduced, but its effect on potential outcome distributions is bounded in some distance metric (a menu of such metrics is offered). Second, the assumption that the experiment under consideration has perfect internal as well as external validity is relaxed. Internal validity is weakened by allowing for selective noncompliance or misclassification. External validity is weakened by allowing sampling populations to differ from treatment populations. The second class of modifications turn the decision problem into one of partial identification, that is, the average treatment effect's sign and thus the identity of the optimal treatment need not be identified.

The extensions have some notable features in common. The perturbations of the decision problem can be scaled from zero (i.e. no perturbation) to very large, and a sufficiently large perturbation will induce some pathology, mostly in the form of a no-data rule. Yet both cases also lead to very strong "local robustness" results: For any sample size N, there exists a positive (albeit of order $O(N^{-1/2})$) size of the perturbation such that the finite sample minimax regret treatment rule is completely unchanged. Thus, while there are many directions in which the treatment choice problem could be generalized,

[1] Other than references cited below, see Brock (2006) and Stoye (2007). Minimax regret was also independently reconsidered by Bergemann and Schlag (2007, 2008), Eozenou et al. (2006), and Chamberlain (2000).

minimax regret is locally insensitive to generalization in some salient such directions. The result is especially welcome in the case of covariates, where existing results seem to spell serious trouble.

All findings are finite sample results derived by analysis of a fictitious game between the decision maker and a malicious Nature. In contrast, Manski (2004) used large deviations inequalities to find finite sample bounds on regret. These bounds can be very slack; see the web appendix of Stoye (2009a) for numerical illustrations. Hirano and Porter (2009) show how to conduct asymptotic analysis of treatment choice problems through local experiments. They explicitly exThis approach is complementary to the discovery of global finite sample results and motivates one of the experiments analyzed below.

The remainder of this paper is structured as follows. Section 2 sets up the decision problem and recapitulates a rather general "benchmark result" that amalgamates previous work. It is the sensitivity of this result that will be examined. Section 3 does so for covariates, section 4 for limitations of the validity of experiments. Section 5 concludes. While intuitions for most results are given in the text, all technical arguments are collected in an appendix.

2 Setting the Stage

2.1 The Decision Problem

The decision problem is mostly as in Manski (2004), and notation mostly follows Stoye (2009a). A decision maker must assign one of two treatments $T \in \{0,1\}$ to members j of a treatment population J. Each member of the treatment population has a response function $y^j(t) : \{0,1\} \to \mathcal{Y}$ that maps treatments onto outcomes.[2] The population is a probability space (J, Σ_J, P) and is "large" in the sense that J is uncountable and $P(j) = 0$ for all j. The decision maker cannot distinguish between members of J, hence from her point of view, assigning treatment t induces a random variable Y_t (the *potential outcome*) with distribution $P(y^j(t))$. We will focus on the distribution $P(Y_0, Y_1)$ as the unknown quantity. Specifically, a *state of the world* s will be identified with $P(Y_0, Y_1)$, which is partially characterized by a couplet $(\mu_0, \mu_1) = \mathbb{E}(Y_0, Y_1)$. The set $\mathcal{S}$ collects feasible states of the world. Notation will later be extended to accommodate modifications of the problem.

The decision maker has access to data generated by a statistical experiment. Two such experiments will be considered.

[2] An important limitation of the setup is that the notation implies the "Stable Unit Treatment Value Assumption" (SUTVA), i.e. outcomes experienced by one individual may not depend on treatment received by others, as would be the case with externalities or general equilibrium effects. Manski (2010) shows how to extend partial identification analysis of treatment effects to the case where this fails. Combining these findings with statistical decision theory is an interesting question for further research.

Binomial experiment: This experiment precisely models sampling of binary random variables. Thus, $\mathcal{Y}^B = \{0,1\}$ and $\mathcal{S}^B = \Delta\{0,1\}^2$, the set of distributions over $\{0,1\}^2$. The sample space is $\Omega^B = (\{0,1\} \times \{0,1\})^N$ with typical element $\omega^B = (t_n, y_n)_{n=1}^N$. I mostly take N to be known, although the generalization to N being a random variable with known distribution is conceptually simple and will be established along the way to proposition 6(ii). Conditional on a realization t_n, y_n is distributed Bernoulli with parameter μ_{t_n}. The sampling distribution of T, in turn, depends on the *sample design*. The following designs will be considered:

- **Matched pairs:** N is even, and $\{t_1, t_2, \ldots\} = \{0, 1, 0, 1, \ldots\}$.
- **Independent Randomization:** N is odd or even, and $\{t_1, t_2, \ldots\}$ are i.i.d. realizations of Bernoulli variables with parameter 1/2 (i.e., independent tosses of fair coins).
- **Constrained Randomization:** N is odd or even, and $\{t_1, t_2, \ldots\}$ is equally likely to be $\{0, 1, 0, 1, \ldots\}$ or $\{1, 0, 1, 0, \ldots\}$.
- **Free Treatment Assignment:** N is odd or even, and within-sample treatment assignment is a choice variable: The decision maker can choose the distribution of $(t_n)_{n=1}^N$ from the set $\Delta\{0,1\}^N$ of distributions over $\{0,1\}^N$.[3]

The binomial experiment is more general than one might at first think. Substantively, the real restriction is that a priori bounds on treatment outcomes exist, are known, and coincide across treatments. Restricting them to lie in $[0, 1]$ is then a normalization. The additional restriction to binary outcomes can be justified by appealing to a "binary randomization" technique due to Schlag (2006).[4] Specifically, one can define treatment rules for non-binary outcomes by first replacing every observed outcome with one realization of a binary, mean-preserving spread of itself and then operating rules defined for binary outcomes. This leaves intact minimax regret values and can, therefore, be used to find minimax regret efficiency bounds as well as decision rules that attain them.

Gaussian experiment: In this experiment, $\mathcal{Y}^G = [0,1]$, thus $\mathcal{S}^G = \Delta\,[0,1]^2$, the set of (Borel-measurable) distributions on $[0,1]^2$. The sample space is $\Omega^G = \mathbb{R}$, and the decision maker observes

[3] This excludes sequential sample designs, i.e. within-sample treatment assignment cannot depend on lagged (in terms of n) realizations. The according extension of this section's benchmark problem is actually easy; in particular, nothing is to be gained by sequential design. Making the case is not worth the effort for this paper's purpose, however, because the same extension of the more complex problems considered later is difficult.

[4] More precisely, the technique was first applied to statistical treatment choice by Schlag (2006); it had been independently discovered for related problems by Cucconi (1968), Gupta and Hande (1992), and Schlag (2003). See Stoye (2009a) for an elaboration in this paper's notation.

a single realization of a signal ω^G distributed according to $N(\mu_1 - \mu_0, \sigma^2)$, where $\sigma^2 > 0$ is known.[5] Intuitively, the signal estimates the welfare contrast $\mu_1 - \mu_0$. While analysis of this experiment will nominally generate finite sample results, its motivation lies in asymptotic approximation. Two examples are as follows:

- Consider a sequence of binomial experiments with independent randomization and $N \to \infty$ as well as a corresponding sequence of true states that imply $(\mu_0, \mu_1) = (\mu + h_0/\sqrt{N}, \mu + h_1/\sqrt{N})$ for some $\mu \in (0,1)$ and $(h_0, h_1) \in \mathbb{R}^2$. With a leap of faith that is justified by proposition 1 below, assume that the signal actually used by the decision maker is equivalent to a single realization of a variable z distributed binomially with parameters $((1 + h_1/\sqrt{N} - h_0/\sqrt{N})/2, N)$. Then $\sqrt{N}\,(z/N - 1/2) \xrightarrow{d} N((h_1 - h_0)/2, 1/4) = N(h_1 - h_0, 1)$, thus the experiments are increasingly well approximated by the Gaussian one with $\sigma^2 = 1$ and a true state of the world that implies $(\mu_0, \mu_1) = (h_0, h_1)$.

- Consider the binomial experiment but with "testing an innovation," that is, with known μ_0. Then there is only one reasonable sampling scheme, namely to assign all subjects to treatment 1. Under this scheme, it does not require any leap of faith to see that the signal can be summarized as a binomial random variable z with parameters (μ_1, N). Fix $\mu_0 \in (0,1)$ and consider a sequence of such experiments where $\mu_1 = \mu_0 + h_1/\sqrt{N}$ for some $h_1 \in \mathbb{R}$. Then $\sqrt{N}\,(z/N - \mu_0) \xrightarrow{d} N(h_1, (\mu_0(1-\mu_0)))$, thus the Gaussian approximation applies again.

Both examples use localizations that are natural because they keep the decision problem nontrivial in the limit, which is why they were suggested by Hirano and Porter (2009). More generally, Hirano and Porter (2009) show that the Gaussian experiment is the appropriate limit experiment for many experiments, most of which will not be amenable to finite sample analysis. Indeed, one important example is "testing an innovation" as just described, which was analyzed by Manski (2004), Manski and Tetenov (2007), Stoye (2009a), and Tetenov (2009b). An extension of these analyses along this paper's lines appears elusive (and not for lack of trying on this author's part), but they are covered by the Gaussian limit experiment.

Assume that if s were known, the decision maker would maximize expected outcomes, thus she would assign all subjects to $T = 1$ if $\mu_1 > \mu_0$, to $T = 0$ if $\mu_1 < \mu_0$, and she would be indifferent if $\mu_0 = \mu_1$. While this can capture risk aversion (Y_t might be a utility), it does presumes a utilitarian social welfare function. Of course, the decision maker does not know s; the challenge is to appropriately use the signal ω. Formally, the decision maker may specify a *statistical treatment rule* $\delta : \Omega \mapsto [0,1]$

[5] Compactification of $\mathcal{Y}$ is a simple way of ensuring that expectations, and the differences of expectations that enter the definition of minimax regret below, are well defined.

that maps possible sample realizations ω onto treatment assignments $\delta(\omega) \in [0,1]$, where the value of δ is interpreted as probability of assigning treatment 1. Nonrandomized decision rules take values only in $\{0,1\}$, but randomization is allowed and will be used. The set of all decision rules will be denoted by $\mathcal{D}$. In the case of free treatment assignment, the decision maker can additionally specify an assignment scheme in $\Delta\{0,1\}^N$ before applying δ to the resultant sample; this extension will be suppressed in notation for simplicity. Finally, δ is called a *no-data rule* if it is constant on Ω (in words, if no sample information is used).

The expected outcome generated by δ given s is

$$u(\delta, s) = \mu_0 (1 - \mathbb{E}\delta(\omega)) + \mu_1 \mathbb{E}\delta(\omega),$$

where expectations are taken with respect to the sampling distribution of ω. Seen as a function of s, $u(\delta, s)$ is (the negative of) the *risk function* of treatment rule δ. Absent an objective distribution on s, attempts to optimize $u(\delta, s)$ induce a decision problem under ambiguity (Manski (2000)). The two most prominent resolutions of this problem are the Bayesian approach, i.e. to rank decision rules according to $\int u(\delta, s) d\pi$, where π is a prior on $\mathcal{S}$ (see Chamberlain (2011) for an elucidation and Dehejia (2005) for an application to treatment choice) and the maximin utility approach, i.e. to rank decision rules according to $\min_{s \in \mathcal{S}} u(\delta, s)$. Manski (2004) initiated reconsideration of minimax regret in this context. To understand this criterion, first define the regret incurred by decision rule δ in state s,

$$R(\delta, s) = \max_{d \in \mathcal{D}} u(d, s) - u(\delta, s),$$

the difference between the expected outcome induced by δ and the outcome that could have been achieved, had s been known. Minimax regret is a maximin ranking with respect to regret loss, thus it recommends to pick

$$\delta^* \in \arg\min_{\delta \in \mathcal{D}} \max_{s \in \mathcal{S}} R(\delta, s)$$

if such a δ^* exists, as is the case in most examples below. Minimax regret was originally introduced by Savages's (1951) reading of Wald (1950); see Stoye (2009c) for further references on history, axiomatizations, and applications. In order to avoid redundancy, I will jump directly to the core result that is this paper's starting point.

2.2 Existing Results

Exact solutions to the above treatment choice problems are available for both maximin utility and minimax regret. Indeed, one motivation for investigating minimax regret is a pathology of maximin utility: Every decision rule achieves maximin utility because $\min_{s \in \mathcal{S}} u(\delta, s) = 0$ for all δ, generated if $\mu_0 = \mu_1 = 0$, a degeneracy problem that was diagnosed already by Savage (1954) and, for the present problem or close variations of it, by Hirano and Porter (2009), Manski (2004), Schlag (2006), and Stoye (2009a).[6] Minimax regret avoids this trap.

For the binomial experiment, define

$$\delta^B(\omega^B) \equiv \begin{cases} 0, & I_N < 0 \\ 1/2, & I_N = 0 \\ 1, & I_N > 0 \end{cases},$$

where

$$\begin{aligned} I_N &\equiv N_1(\overline{y}_1 - 1/2) - N_0(\overline{y}_0 - 1/2) \\ &\propto [\#\,(\text{observed successes of treatment 1}) + \#\,(\text{observed failures of treatment 0})] \\ &\quad - [\#\,(\text{observed successes of treatment 0}) + \#\,(\text{observed failures of treatment 1})] \end{aligned}$$

with N_t the number of sample subjects assigned to treatment t, $\overline{y}_t$ a sample mean that conditions on $T = t$, and the convention that $N_t(\overline{y}_t - 1/2) = 0$ if $N_t = 0$. Thus, treatment is assigned according to the sign of the score I_N, with even randomization if $I_N = 0$. To get an intuition for this, note that $\mathbb{E}I_N/N = \mu_1 - \mu_0$, that is, I_N/N is an unbiased estimator of the average treatment effect (ATE) and its sign therefore a reasonable estimator of the better treatment's identity. Indeed, if sample design is by either matched pairs or constrained randomization, then δ^B simplifies to

$$\delta^B(\omega^B) \equiv \begin{cases} 0, & \overline{y}_1 < \overline{y}_0 \\ 1/2, & \overline{y}_1 = \overline{y}_0 \\ 1, & \overline{y}_1 > \overline{y}_0 \end{cases}.$$

Finally, for the normal experiment, $\delta^G \equiv 1\{\omega^G > 0\}$. Recall that ω^G can be intuited as estimating $\mu_1 - \mu_0$, thus δ^G is really very similar to δ^B.

The following result is an amalgam of Canner (1970), Hirano and Porter (2009), Schlag (2006), Stoye (2009a), and Tetenov (2009a).[7]

Proposition 1 ***(i)*** *In the binomial experiment, let sample design be any of matched pairs, constrained randomization, or independent randomization. Then δ^B achieves minimax regret.*

[6] In the abstract, the same types of examples can be constructed for minimax regret (Parmigiani (1992)). They do not seem to occur naturally in models of treatment choice; although, see the section on covariates below.

[7] Part (i) is found in Canner (1970, for matched pairs), Stoye (2009a, for matched pairs and independent randomization), and is also a corollary of results in Schlag (2006, for matched pairs and constrained randomization). Part (ii) minimally expands on Schlag (2006), who recommends constrained randomization. Part (iii) follows from results in both Hirano and Porter (2009) and Tetenov (2009a). Part (i) of the corollary was independently established in Stoye (2009a) and Schlag (2006), and both parts of the corollary were generalized by Tetenov (2009a).

(ii) *In the binomial experiment, let sample design be a choice variable. Then minimax regret is achieved by any of matched pairs (if N is even), constrained randomization, and independent randomization in conjunction with δ^B.*

(iii) *In the Gaussian experiment, δ^G achieves minimax regret.*

Corollary 2 ***(i)*** *The binomial decision problem has value*

$$\begin{aligned} R^B(N) &= \max_{a\in[1/2,1]} \left\{ (2a-1) \sum_{n<N/2} \binom{N'}{n} a^n (1-a)^{N'-n} \right\} \\ N' &= \max_{M\in\mathbb{N}} \{M \leq N : M \text{ is odd}\}, \end{aligned}$$

where $R^B(0) = 1/2$.

(ii) *The Gaussian decision problem has value*

$$R^G(\sigma) = \max_{\Delta\in[0,1]} \left\{ \Delta\Phi(-\Delta; 0, \sigma^2) \right\},$$

where $\Phi\left(\cdot;\mu,\sigma^2\right)$ is the normal c.d.f. with indicated parameters.

Proposition 1 establishes that in the benchmark problem, minimax regret is achieved by rather reasonable looking decision rules.[8] An important technique to establish this kind of result is game theoretic analysis (Wald (1945)). Consider a fictitious zero-sum game in which the decision maker picks a decision rule $\delta \in \mathcal{D}$ (possibly at random; note, though, that the result of any randomization over decision rules can itself be expressed as element of $\mathcal{D}$) and Nature picks a state of the world s (possibly at random, meaning that her strategy space is the set $\Delta\mathcal{S}$ of priors π over $\mathcal{S}$). After both players moved, s is drawn according to π, ω is drawn from s according to the relevant sampling scheme, and δ is operated on ω. The decision maker then pays $\max_{d\in\mathcal{D}} u(d,s) - u(\delta,s)$ to Nature. Any Nash equilibrium (δ^*, π^*) of this game characterizes a minimax regret treatment rule δ^*, and Nature's equilibrium strategy π^* can be interpreted as the least favorable prior that is implicitly selected by the minimax regret criterion.

The least favorable priors underlying proposition 1 are as follows. $R(\delta, s)$ can be shown to depend on s only through (μ_0, μ_1), thus identify states with couplets (μ_0, μ_1). Then the least favorable prior randomizes evenly over two symmetric states $(a^*, 1-a^*)$ and $(1-a^*, a^*)$, where computation of a^* leads to the value functions provided in corollary 2.

[8] A remark on uniqueness: With sample designs fixed, these rules are unique up to subtleties about tie-breaking. However, there are many sampling schemes other than the ones stated that, together with δ^B, achieve minimax regret if sample design is a choice parameter.

3 Introducing Covariates

3.1 A Problem with Minimax Regret

Most real-world economic decision and prediction problems involve covariates, and the according modification of the benchmark problem is of obvious interest. Thus, assume that some covariate $X \in \mathcal{X}$ is observable for both sample and treatment subjects. Then the decision maker can attempt to (implicitly) estimate the mean regression $(\mu_{0x}, \mu_{1x}) = \mathbb{E}(Y_0, Y_1 | X = x)$ and can provide accordingly conditioned treatment recommendations.

While this section's main result, proposition 3 below, applies under very weak assumptions on X, I will focus on two cases to simplify the subsequent discussion. First, assume that $\mathcal{X}$ has $K < \infty$ elements, a case that will be called "discrete X" henceforth, and that the distribution $P(X)$ is known. (Analysis of the Gaussian experiment will be limited to this case.) This also covers the case of finitely many discrete covariates $X_1, \ldots X_J$; these can be thought of as one discrete covariate supported on $\mathcal{X}_1 \times \ldots \times \mathcal{X}_J$. Second, assume that X is distributed uniformly on $\mathcal{X} = [0,1]$, understood to be endowed with the Borel sigma-algebra. This case is reasonably general because it could reflect an inverse quantile transformation. Importantly, knowledge of the distribution of X is presumed however. In either case, a state of the world s is a (measurable) function that takes covariate values x into conditional distributions of potential outcomes $s_x = P(Y_0, Y_1 | X = x)$. For simplicity, henceforth let (Y_{0x}, Y_{1x}) denote a random variable distributed according to s_x. If X is discrete, states reduce to vectors $(P(Y_{0x}, Y_{1x}))_{x \in \mathcal{X}}$ of length K.

Complete specification of the problem also requires state spaces and signals. Initially consider a state space that embodies no substantive restriction on the effect of X. For the binomial experiment, this state space collects all (measurable) mappings from $\mathcal{X}$ into $\Delta\{0,1\}^2$, specializing to $\mathcal{S}^B = \times_{k=1}^K \Delta \{0,1\}^2$ if X is discrete. A sample ω^B then collects realizations (t_n, x_n, y_n), where the distributions of both T and X depend on the sample design and y_n is an independent realization of $Y_{t_n x_n}$. For the Gaussian experiment, the state space is $\mathcal{S}^G = \times_{k=1}^K \Delta [0,1]^2$, and a signal $\omega^G \in \mathbb{R}^K$ collects covariate-specific signals ω_x^G distributed according to $N(\mu_{1x} - \mu_{0x}, \sigma_x^2)$, where $\sigma_x^2 = \sigma^2 / \Pr(X = x)$ for some known σ; this is the pattern of variances that would arise if a large sample were random with respect to X.

A statistical treatment rule maps samples ω into (measurable) treatment assignment functions $\delta : \mathcal{X} \mapsto [0,1]$, where $\delta_x(\omega)$ denotes the probability with which treatment 1 is assigned to subjects with covariate x if signal ω was observed. For discrete X, this specializes to $\delta(\omega) \in [0,1]^K$. A treatment rule's risk function is $u(\delta, s) = \int (\mu_{0x}(1 - \mathbb{E}\delta_x(\omega)) + \mu_{1x}\mathbb{E}\delta_x(\omega))\, dP(X)$, and regret is $R(\delta, s) = \max_{d \in \mathcal{D}} u(d, s) - u(\delta, s)$ as before.

An obvious question is whether in this extended problem, minimax regret is achieved by pooling observations across covariates, by conditioning on covariates, or by something in between. The intuitive trade-off is between the resolution of a decision rule and its sensitivity to sampling variation. Using large deviations bounds, Manski (2004) discovered that a lower bound on regret incurred by pooling exceeds an upper bound incurred by conditioning on covariates for rather small sample sizes. The tentative conclusion was that prevailing practice errs in the direction of too much pooling.

Stoye (2009a) re-analyzed the issue in terms of finite sample regret and found that Manski's result merely approximates a much stronger, and pathological, one. Minimax regret recommends to conduct inference separately for all values of *all* available covariates, even if this leads to empty sample cells. This conclusion extends to cases of many sample cells and small samples, where the resulting decision rule is essentially a no-data rule. Indeed, if X is continuous, then a no-data rule achieves minimax regret. This result reverses the thrust of previous findings, raising more questions about minimax regret than about prevailing practice. It motivates this section's analysis.

Here is an intuition for why the problem obtains. Both Manski (2004) and Stoye (2009a) use the maximally permissive state spaces just defined. These accommodate priors under which s_x and $s_{x'}$ are independent random variables, and Nature will choose just such priors in the fictitious game. Observations of treatment outcomes for covariate x are then uninformative about potential outcomes for covariate x', thus treatment rules which are Bayes against said priors separate inference across covariates. But the additive separability by covariate of $R(\delta, s)$ can be used to show that priors of this sort, in turn, best respond to these decision rules. An alternative, non-game theoretic intuition is as follows: Ceteris paribus, minimax regret selects for scenarios that make it hard to learn. Cross-covariate learning is hardest if cross-covariate signals are vacuous, and uniform such learning is impossible for continuous X because the mean regressions (μ_{0x}, μ_{1x}) can be arbitrarily "wiggly" in x.

However, these problems obtain because some perhaps exotic states are being contemplated. The least favorable prior driving the above result places equal weight on states that render (μ_{0x}, μ_{1x}) and $(\mu_{0x'}, \mu_{1x'})$ equal and states that render them symmetrically opposed. Depending on the substantive meaning of X, the latter states can be highly implausible. Consider a medical trial in which the race of sample subjects or, even worse, their date of birth is recorded. Clearly, there exists prior information to the effect that these covariates will matter only to a limited degree. The problem may, therefore, lie with (technically) a too permissive state space that (substantively) reflects underspecification of prior information. One might hope that it can be alleviated by introducing plausible prior constraints. The next subsection introduces such a constraint and shows that it can have dramatic effects.

3.2 Limiting the Effect of Covariates

This section's core result allows for both discrete and continuous X. The innovation is to shrink $\mathcal{S}$ by restricting the effect that X can have on $P(Y_{tx})$. This will be done through bounding $||P(Y_{tx}), P(Y_{tx'})||$, where $||\cdot||$ can stand for one of several metrics. I begin by analyzing the effect of the following assumption:

Assumption 1

$$|\mu_{tx} - \mu_{tx'}| \leq \kappa$$

for all $x, x' \in \mathcal{X}$, $t \in \{0, 1\}$.

Some such restriction will be exceedingly plausible in many cases. Consider again medical trials. Researchers would probably be willing to bound the effect of race on outcomes and to quite severely bound the effect of birthdays. At the same time, it will rarely be honest to bound a covariate's effect at precisely zero. Even birthdays might have a very small effect on reaction to medication (because they proxy for season of birth), so that restricting their effect to be exactly zero is likely an approximation. This is why this section's main result may be of some interest. In words, it states that for every N, there exists κ that should be thought of as "small but positive" s.t. covariates whose cumulative effect can be bounded by κ should be completely ignored. A precise statement is somewhat more long-winded and can be differentiated depending on how much of the sample design is exogenous.

Proposition 3 ***(i)*** *Consider the binomial experiment. Let X be discrete or continuous. Let $\overline{\delta}^B$ be the decision rule that mimics δ^B and ignores the existence of X, thus $\overline{\delta}^B_x(\omega^B) = \delta^B(\omega^B)$ for all (x, ω^B). If the sample design is any of matched pairs, constrained randomization, or independent randomization, the sample is a simple random sample with respect to X, and assumption 1 holds with some $\kappa \leq 2R^B(N)$, then $\overline{\delta}^B$ achieves minimax regret.*

(ii) *In part (i), if sample design is a choice variable, then any of matched pairs, constrained randomization, and independent randomization in conjunction with $\overline{\delta}^B$ achieves minimax regret. If sample stratification with respect to X is a choice variable, simple random sampling with respect to X achieves minimax regret. If both are choice variables, then simple random sampling with respect to X in conjunction with any of matched pairs, constrained randomization, and independent randomization and use of $\overline{\delta}^B$ achieves minimax regret.*

(iii) *Consider the Gaussian experiment. Let $\overline{\delta}^G$ be the decision rule that mimics δ^G and ignores the existence of X, thus $\overline{\delta}^G_x(\omega) = \delta^G\left(\sum_{x \in \mathcal{X}} \omega^G_x \Pr(X = x)\right)$ for all (x, ω). If assumption 1 holds with some $\kappa \leq 2R^G(N)$, then $\overline{\delta}^G$ achieves minimax regret.*

(iv) *In all of (i)-(iii), if assumption 1 holds with some $\kappa < 1$, then any no-data rule incurs strictly more than minimax regret.*

While the proof of this proposition involves some tedious algebra, there is an instructive intuition to most of it. Consider the problem of maximizing $R(\overline{\delta}^B, s)$, the regret incurred by a decision maker who employs $\overline{\delta}^B$. Initially assume that there exist values x, x' with $\mu_{1x} > \mu_{0x}$ yet $\mu_{1x'} < \mu_{0x'}$, thus the decision maker should ideally take the covariate into account. Define $(x)^+ = \max\{x, 0\}$ and write

$$\begin{aligned}
R(\overline{\delta}^B, s) &= (1 - \mathbb{E}\overline{\delta}^B(\omega^B)) \cdot \int (\mu_{1x} - \mu_{0x})^+ dP(X) + \mathbb{E}\overline{\delta}^B(\omega^B) \cdot \int (\mu_{0x} - \mu_{1x})^+ dP(X) \\
&\leq \sup_{x \in \mathcal{X}} (\mu_{1x} - \mu_{0x}) + \sup_{x \in \mathcal{X}} (\mu_{0x} - \mu_{1x}) \\
&= \sup_{x, x' \in \mathcal{X}} (\mu_{1x} - \mu_{0x} + \mu_{0x'} - \mu_{1x'}) \qquad (1) \\
&= \sup_{x, x' \in \mathcal{X}} (\underbrace{\mu_{1x} - \mu_{1x'}}_{\leq \kappa} + \underbrace{\mu_{0x'} - \mu_{0x}}_{\leq \kappa}) \\
&\leq 2\kappa,
\end{aligned}$$

where the first inequality is true because both $\sup_{x \in \mathcal{X}}(\mu_{1x} - \mu_{0x}) > 0$ and $\sup_{x \in \mathcal{X}}(\mu_{0x} - \mu_{1x}) > 0$ by assumption and $\mathbb{E}\overline{\delta}^B(\omega^B) \in [0, 1]$. Hence, if $\max_{s \in \mathcal{S}^B} R(\overline{\delta}^B, s)$ is attained on a state where the decision maker ought to make use of X, then its value cannot exceed 2κ.

Yet in fact, $\max_{s \in \mathcal{S}^B} R(\overline{\delta}^B, s) \geq R^B(N)$ because by choosing states in which $P(Y_{0x}, Y_{1x})$ is constant across x, Nature can mimic the least favorable prior from proposition 1. It follows that if $\kappa \leq R^B(N)/2$, then Nature's best responses to $\overline{\delta}^B$ are supported on states s.t. $sg(\mu_{0x} - \mu_{1x})$ is constant in x. In addition, it can be shown that among such states, one can as well restrict attention to states in which μ_{0x} and μ_{1x} are constant. Hence, one best response among these states is the aforementioned prior that renders μ_{0x} and μ_{1x} constant and entirely mimics the least favorable prior from proposition 1. As $\overline{\delta}^B$, in turn, best responds to this prior, an equilibrium of the fictitious game has been found.

This argument pretty much establishes the result for $\kappa \leq R^B(N)/2$. The improvement to $\kappa \leq 2R^B(N)$ obtains because the bound in (1) is far from tight. An improvement by a factor of 2 is rather easily obtainable. The first inequality in (1) separately bounds $\mathbb{E}\overline{\delta}^B(\omega^B)$ and $\left(1 - \mathbb{E}\overline{\delta}^B(\omega^B)\right)$ by 1, yet in fact their sum equals 1, thus the first line really displays a weighted average of $\int (\mu_{1x} - \mu_{0x})^+ dP(X)$ and $\int (\mu_{0x} - \mu_{1x})^+ dP(X)$. What's more, inspection of the definition of δ^B reveals that

$$\mathbb{E}\overline{\delta}^B(\omega^B) \geq 0 \text{ iff } \int \mu_{1x} dP(X) \geq \int \mu_{0x} dP(X) \text{ iff } \int (\mu_{1x} - \mu_{0x})^+ dP(X) \geq \int (\mu_{0x} - \mu_{1x})^+ dP(X),$$

thus the weighted average must favor the smaller of the integrals. One could easily formalize this observation to improve the bound to κ. A more intricate exercise in functional analysis reveals that if the maximum of $R(\overline{\delta}^B, s)$ is achieved by a state s.t. both $\int (\mu_{1x} - \mu_{0x})^+ dP(X) > 0$ and $\int (\mu_{0x} - \mu_{1x})^+ dP(X) > 0$, then – and only then – can its value be bounded at $\kappa/2$. Furthermore, this bound is in general attainable and, therefore, best possible. The argument for the Gaussian experiment differs only in the algebraic detail of the functional analysis part.

3.3 Extending the Result to other Distance Metrics

Proposition 3 can be extended to other notions of distance between probability distributions. Fix any distributions P and Q on the Borel sets $\mathcal{B}(\mathbb{R})$ on the real line. Define the Total Variation distance $||\cdot||_{TV}$ by $||P,Q||_{TV} = \max_{E\in\mathcal{B}(\mathbb{R})} |P(E) - Q(E)|$, the log odds ratio distance by $||P,Q||_{LOR} = \max_{E\in\mathcal{B}(\mathbb{R})} |\log(P(E)/Q(E))|$, and the Kullback-Leibler divergence (a.k.a. relative entropy) by $D_{KL}(P||Q) = \int \log(dP/dQ)dP$ (assuming existence of these quantities).[9] Then the following obtains.

Proposition 4 *The conclusions of proposition 3 obtain if any of the following bounds apply:*

$$\begin{aligned} ||P(Y_{tx}), P(Y_{tx'})||_{TV} &\leq 2R^B(N), \\ ||P(Y_{tx}), P(Y_{tx'})||_{LOR} &\leq 2\log\frac{1+2R^B(N)}{1-2R^B(N)}, \\ D_{KL}(P(Y_{tx})||P(Y_{tx'})) &\leq 2R^B(N)\log\frac{1+2R^B(N)}{1-2R^B(N)}, \end{aligned}$$

respectively the same bounds in terms of R^G.

This result is established by showing that bounds in these different metrics induce bounds on $|\mu_{tx} - \mu_{tx'}|$ and then invoking proposition 3 (modulo some changes in variables). This might seem inefficient because bounds on Kullback-Leibler divergences and log odds ratios are strictly stronger than the implied bounds on differences in means, thus there might be slack in the resulting threshold values. But this is not so: The least favorable priors which render proposition 3 tight also maximize $|\mu_{tx} - \mu_{tx'}|$ subject to the respective bound on D_{KL} or $||\cdot||_{LOR}$. One could, therefore, think of bounding $|\mu_{tx} - \mu_{tx'}|$ as an efficient approach in the sense of getting the desired effect through the weakest among a set of restrictions. Also, as $|\mu_{tx} - \mu_{tx'}| = 1$ implies infinite log odds ratio distance as well as Kullback-Leibler divergence, one can eliminate no-data rules by imposing any finite bound on either quantity.

3.4 Additional Results for Special Cases

Proposition 3 is tight in the sense that if $R^B(N) < 2\kappa$, then it cannot be claimed in general that minimax regret is achieved by pooling information – Nature's best response to pooling might be a prior under which the decision maker should take X into account, breaking the fictitious game's equilibrium. However, while a complete, closed-form analysis of the modified decision problem for discrete X seems elusive, additional results are available for special cases. First, a continuous covariate should be ignored for *any* κ.

[9] If P and Q are Bernoulli with parameters μ_P and μ_Q, then these simplify to $||P,Q||_{TV} = |\mu_P - \mu_Q|$ (and this then also coincides with the Kolmogorov-Smirnov and other distances), $||P,Q||_{LOR} = \left|\log\left((\mu_P(1-\mu_Q))/(\mu_Q(1-\mu_P))\right)\right|$, and $D_{KL}(P||Q) = \mu_P\log(\mu_P/\mu_Q) + (1-\mu_P)\log((1-\mu_P)/(1-\mu_Q))$. The proof of proposition 6 uses these simplifications, but the result does not materially depend on binary outcomes.

Proposition 5 *Consider the binomial experiment. Let X be continuous. Then $\overline{\delta}^B$ achieves minimax regret for all κ. The minimax regret value of the decision problem is* $\max\{R^B(N), \kappa/2\}$.

The core observation here is that if X is continuous, Nature can always enforce a regret that is arbitrarily close to $\kappa/2$. The idea is as follows: Pick some event $E \in \mathcal{B}([0,1])$ s.t. $\int 1\{x \in E\}dx = 1/2$ and let (μ_{0x}, μ_{1x}) equal $((1-\kappa)/2, (1+\kappa)/2)$ on E and $((1+\kappa)/2, (1-\kappa)/2)$ otherwise. It is easy to see that $\overline{\delta}^B$ and any other decision rule that fail to take X into account will attain expected regret of $\kappa/2$. Of course, if the decision maker knew E, her best response would be otherwise. However, Nature can randomize over states of this type that vary by selection of E, and $\mathcal{B}([0,1])$ is a sufficiently rich domain for uniform learning about E to be elusive. On the other hand, an intermediate step in establishing proposition 3 was to show that $\max_{s \in \mathcal{S}^B} R(\overline{\delta}^B, s) \leq \max\{R^B(N), \kappa/2\}$. If this maximum equals $R^B(N)$, proposition 3 established equilibrium. If it equals $\kappa/2$, then $\overline{\delta}^B$ achieves the lower bound on minimax regret just identified, thus the result.

This result is not true if the distribution of X has mass points, most notably if it is discrete, in which case consideration of covariates will frequently be minimax regret optimal. Quite to the contrary, if X is finite, then the older result about sample splitting has its own "robustness region": For any sample size N and distribution of covariates $P(X)$, minimax regret continues to advise against cross-covariate inference if κ is large enough, where "large enough" is defined through a threshold number that exceeds 1 in a few cases involving very noisy signals, but is generally (much) below 1 and decays to zero at a "parametric inference" rate as N increases or σ decreases.

Proposition 6 *Let X be discrete and fix $P(X)$.*

(i) *Consider the binomial experiment. There exists a function $\overline{\kappa}^B : \mathbb{N} \mapsto (0,1]$ with $\overline{\kappa}^B(n) = O(n^{-1/2})$ s.t. the following holds: Fix any sample stratification $(N_x)_{x \in \mathcal{X}}$, where N_x is the number of sample subjects with covariate value x. If assumption 1 can be imposed but only with $\kappa \geq \overline{\kappa}^B(\min_{x \in \mathcal{X}} N_x)$, then the minimax regret rule applies δ^B separately for each value of X.*

(ii) *Consider the binomial experiment. There exists a function $\widetilde{\kappa}^B : \mathbb{N} \mapsto (0,1]$ with $\widetilde{\kappa}^B(n) = O(n^{-1/2})$ s.t. the following holds: Assume random sampling with respect to X, then if assumption 1 can be imposed but only with $\kappa \geq \widetilde{\kappa}^B(N)$, the minimax regret rule applies δ^B separately for each value of X.*

(iii) *Consider the Gaussian experiment. There exists a number $\overline{\kappa}^G \in (0,1]$ s.t. if assumption 1 can be imposed but only with $\kappa \geq \sigma\overline{\kappa}^G$, the minimax regret rule applies δ^G separately for each value of X.*

The basic intuition for this is that for any given experiment and set of parameter values, as κ grows large, the prior that supported Stoye's (2009a) "no cross-covariate inference" finding eventually becomes consistent with assumption 1. Analogous statements for other distance metrics are omitted for brevity.

These results can also be thought of in terms of holding κ fixed and varying cell sizes by changing sample size (in the binomial experiment) or variance (in the Gaussian experiment). For any given κ, proposition 3 states that samples should be completely pooled if N is small enough [σ^2 is large enough]. For continuous X, proposition 5 completes the analysis. For discrete X, proposition 6 establishes what happens as N grows [σ^2 becomes small]. It stands to reason that this will make cross-covariate inference "fade out" in some way, but the actual result is much stronger: Complete separation of inference by covariate value is reinstated beyond some finite sample size. Indeed, there is an intriguing symmetry between propositions 3 and 6: Minimax regret treatment rules do not only converge to the extremes of "complete pooling of samples" and "complete splitting of samples" as $\kappa \to 0$ respectively $\kappa \to 1$ (holding N fixed) or as $N \to 0$ respectively $N \to \infty$ (holding κ fixed), but precisely attain those limits for κ or N large [small] enough and then do not change any more as κ or N become even larger [smaller].

To illustrate these findings, consider an example from Stoye (2009a). The experiment is binomial, there is a binary covariate $X \in \{m, f\}$ with $\Pr(X = m) = 1/2$, the decision maker can freely choose the sample design, and both treatments are unknown. Then if the effect of X is unrestricted, the decision maker will want to separately sample males and females and to separately apply δ^B within each group, no matter how small the overall sample size.

Table 1 illustrates how this conclusion is affected by limiting the effect of X. Columns refer to different overall sample sizes N. The first row, labelled, $\underline{\kappa}^B(N)$, displays the threshold at which proposition 3 begins to apply. In words, if one is willing to assume $\left|\mu_{tm} - \mu_{tf}\right| \leq \underline{\kappa}^B(N)$, then X should be ignored. Assume now for simplicity that the sample is stratified by covariate, with exactly half of subjects male. Then the second row of table 1 displays upper thresholds $\overline{\kappa}^B(N)$ beyond which proposition 6(i) applies, that is, minimax regret decision makers will split samples. Note that $\underline{\kappa}^B(N)$ does not depend on the specifics of the example (i.e., X being binary or the sample design), whereas $\overline{\kappa}^B(N)$ does. Rows three and four display the analogous quantities for the normal experiment with $\sigma^2 = 1$ and rescaled by $N^{-1/2}$, a localization that allows one to approximate the globally least favorable prior. The point is to show, by means of comparison with the first two rows, whether this approximation performs well.

The "robustness region" from proposition 3 is reasonably large. For $N = 10$, if one is willing to bound $\left|\mu_{tm} - \mu_{tf}\right|$ by 10% – which would frequently appear plausible –, the finite sample minimax regret recommendation is to completely ignore gender. A bound on $\left|\mu_{tm} - \mu_{tf}\right|$ of 1% – still reasonable for some treatments and certainly if X were date of birth, say – would justify ignoring the covariate at a rather substantive sample size of $N = 1000$. At the same time, as N increases, separate inference

$\mathbf{N}$	2	4	6	10	20	50	100	200	500	1000
$\underline{\boldsymbol{\kappa}}^B(\mathbf{N})$	0.25	0.176	0.141	0.109	0.076	0.048	0.034	0.024	0.015	0.011
$\overline{\boldsymbol{\kappa}}^B(\mathbf{N})$	0.5	0.5	0.366	0.302	0.235	0.147	0.106	0.075	0.048	0.034
$\mathbf{N}^{-1/2}\underline{\boldsymbol{\kappa}}^G(\mathbf{N})$	0.340	0.170	0.138	0.107	0.076	0.048	0.034	0.024	0.015	0.011
$\mathbf{N}^{-1/2}\overline{\boldsymbol{\kappa}}^G(\mathbf{N})$	0.752	0.532	0.434	0.336	0.238	0.150	0.106	0.075	0.048	0.034

Table 1: Numerical illustration of propositions 3 and 6: Threshold values for kappa below which a covariate should be ignored and above which samples should be split.

becomes relevant for moderate κ, thus it is not to be dismissed altogether. For $N = 500$, if one believes an effect of 5% to be plausible, one should (from the minimax regret point of view) split samples. While the intermediate cases are, unfortunately, extremely hard to solve, table 1 illustrates that the particular minimax regret treatment choice problem is solved for a rather wide range of parameter values. It is also seen that the normal approximation performs very well, with agreement to three and more significant digits beginning at $N = 100$.

4 Introducing Imperfect Experiments

The benchmark example presumes a perfectly valid experiment: Treatment assignment is externally randomized, so causal inference from treatment to effect is not hindered by selective noncompliance and the like. In short, the experiment is internally valid. Furthermore, the sampling universe coincides with the population to which treatments will be futurely assigned, thus causal inference can be extrapolated to the treatment population. In short, the experiment is externally valid. I will now analyze models that allow for some failure of either internal or external validity. While it would be conceptually feasible to do so by writing explicit, fully identified models of selection and then conducting minimax regret analysis, this road will not be taken. One interesting feature of maximin criteria is that by offering a decision theoretic resolution to ambiguity, they do not force the user to resolve all ambiguity via identifying assumptions. This feature will here be exploited by models that do not imply statistical identifiability of the average treatment effect or even of its sign and hence, of the better treatment's identity. The analysis thereby connects decision theoretic analysis of treatment choice with another literature pioneered by Manski, namely partial identification.[10]

Contrary to the preceding section, I will provide minimax regret treatment rules for all possible values of underlying parameters. Two notable findings are as follows: First, as was the case with covariates, minimax regret treatment rules completely ignore small but positive departures from the

[10] See Manski (2003) for a survey. Minimax regret treatment choice has also been linked to partial identification analysis in Brock (2006), Stoye (2007, 2009b), Manski (2007a, 2007b, 2009), and Tetenov (2009b).

benchmark model. Second, for departures that take the decision problem outside of this "robustness region," minimax regret is achieved by discarding sample data until one is back in the robustness region. The latter feature may initially appear surprising, but has a clear intuition that will be elaborated.

4.1 Limited Internal or External Validity: Some Partially Identified Models

Limitations to internal validity occur when causal inference toward the sampling universe is not warranted. Typical causes of this problem are selection into treatment and/or selective noncompliance with assigned treatment. For decision theoretic analysis of treatment choice, what is important is the maximal wedge such issues drive between observable treatment outcomes and treatment outcomes that would be observed in perfect experiments. To illustrate, consider the following examples.[11]

Example 1 ***Selective Noncompliance as a Partial Identification Problem***

Let treatment assignment be at random, but assume that a probability mass ε *of subjects are noncompliers, that is, they may not receive the assigned treatment. Assume that assigned treatment is observed but received treatment is not. Let* $D \in \{0, 1\}$ *denote assigned treatment. If* $\Delta = \mathbb{E}(Y|D = 1) - \mathbb{E}(Y|D = 0)$ *is the (readily observable) intention-to-treat effect (ITE) and* $\widetilde{\Delta} = \mathbb{E}(Y_1 - Y_0)$ *is the average treatment effect (ATE) of interest, then tight bounds on the latter in terms of the former are*

$$\widetilde{\Delta} \in [\Delta - 2\varepsilon, \Delta + 2\varepsilon] \cap [-1, 1].$$

The above bounds are achieved by assuming that all noncompliers are "anti-compliers" whose received treatment never agrees with their assigned treatment, and (unless $\widetilde{\Delta} \in [-1, 1]$ binds) that the average treatment effect for noncompliers is either -1 or $+1$.[12] These scenarios will frequently be implausible. Consider, therefore, a model in which noncompliers are either "always-takers" or "never-takers," that is they, ignore their treatment assignment but do not (or cannot) act against it. This assumption informs the literature on Latent Average Treatment Effects (Angrist, Imbens, and Rubin (1996)), where it is called monotonicity. It is essentially equivalent to imposing a threshold crossing model for selection into treatment, with assigned treatment functioning as an instrument (Vytlacil (2002)). In some cases, one could go even further and impose one-sided noncompliance, that is, all noncompliers are never-takers. This assumption makes sense when access to the treatment can be controlled. Imposing these restrictions substantially refines the bounds.

[11] The appendix contains proof of claims made in the examples.

[12] This scenario has observable implications, namely either $\mu_0 \in [0, 1 - \varepsilon]$ and $\mu_1 \in [\varepsilon, 1]$ or vice versa. This does not affect minimax regret analysis. A similar remark applies to example 2.

Example 2 ***Some Refinements of Example 1***

In the setting of example 1, assume monotonicity, then worst-case bounds on $\widetilde{\Delta}$ *improve to*

$$\widetilde{\Delta} \in [\Delta - \varepsilon, \Delta + \varepsilon] \cap [-1, 1].$$

Strengthening monotonicity to one-sided noncompliance does not further tighten these bounds.

Very roughly speaking, the length of the identified set is halved because for each of the two non-complier subpopulations, one of the two relevant potential outcomes is observed. For these particular bounds (but not for bounds on other quantities of potential interest), it is immaterial whether non-compliers are always-takers or never-takers, so excluding always-takers has no further effect.

Limitations to internal validity are not the only problem with experiments. In his recent monograph, Manski (2007) mentions external validity, i.e. the question whether conclusions that are valid within an experimental population can be extrapolated beyond the lab. In Manski's words, "an experiment is said to have external validity if the distribution of outcomes realized by a treatment group is the same as the distribution of outcomes that would be realized in an actual program" (p. 26). However, "participation in experiments ordinarily cannot be mandated in democracies. Hence experiments in practice usually draw subjects at random from a pool of persons who volunteer to participate. So one learns about treatment response within the population of volunteers rather than within the population of interest" (p. 138).

To explicitly model such a concern, let $Z \in \{0, 1\}$ indicate membership in the sampling universe, so the treatment population is characterized by a joint distribution of (Y_0, Y_1, Z). Assume the experiment is internally valid, thus the sample is a simple size N random sample from the sampling universe and treatment assignment is at random as well as enforced. Then the experiment identifies $\mathbb{E}(Y_1 - Y_0 | Z = 1)$, the true ATE in the sampling unverse. The problem is that if the sampling universe is a selective subset of the population, this quantity may not equal $\mathbb{E}(Y_1 - Y_0)$. Obviously, interesting results are available only if the former is at least somewhat informative about the latter,[13] and the relation between them will be accordingly constrained. Two partially identified models are as follows.

Example 3 ***Selection into Sampling as a Contamination Problem***

Assume $\Pr(Z = 1) = 1 - \varepsilon$*, where* ε *is known. Then tight bounds on* $\widetilde{\Delta}$ *are*

$$\widetilde{\Delta} \in [(1 - \varepsilon)\Delta - \varepsilon, (1 - \varepsilon)\Delta + \varepsilon],$$

where $\Delta = \mathbb{E}(Y_1 - Y_0 | Z = 1)$ *is the ATE in the sampling universe.*

[13] Perhaps this is not as obvious as one might have thought. Woorall (2007) relates the case of benoxaprofen, where a dramatic failure of external validity occurred because the medication was mainly used on elderly patients, but these were not at all sampled into the randomized clinical trials.

Example 4 ***Selection into Sampling as a Hidden Covariate Problem***

Let the decision to join the sampling population (e.g., to volunteer) be modelled via a logit model, i.e.

$$\Pr(Z = 1|X = x) = \frac{\exp(\alpha + \beta x)}{1 + \exp(\alpha + \beta x)},$$

where the unobservable covariate X is arbitrarily (Borel-measurably) distributed on $[0, 1]$. The effect of X on treatment response is unconstrained, but $\beta \in [0, c]$, where c is known. Then tight (up to weak vs. strict inequalities) bounds on $\widetilde{\Delta}$ are

$$\frac{\mathbb{E}(Y_1|Z = 1)}{\exp(c) + (1 - \exp(c))\mathbb{E}(Y_1|Z = 1)} - \frac{\exp(c)\mathbb{E}(Y_0|Z = 1)}{1 + (\exp(c) - 1)\mathbb{E}(Y_0|Z = 1)}$$
$$\leq \widetilde{\Delta} \leq \frac{\exp(c)\mathbb{E}(Y_1|Z = 1)}{1 + (\exp(c) - 1)\mathbb{E}(Y_1|Z = 1)} - \frac{\mathbb{E}(Y_0|Z = 1)}{\exp(c) + (1 - \exp(c))\mathbb{E}(Y_0|Z = 1)}.$$

In all examples, ε or c are assumed known. If they can only be bounded, then it is w.l.o.g. to set them equal to their upper bounds because identification decays monotonically as the parameter increases. It is conceivable that while not strictly known, ε or c are identifiable from data other than the ones under consideration. In other contexts, they are best thought of as user-specified "plausibility" or "sensitivity" parameters.

4.2 Results and an Interpretation

Examples 1-3 are similar in that bounds can be written as

$$\widetilde{\Delta} \in [a\Delta - b, a\Delta + b] \cap [-1, 1] \tag{2}$$

for some $a \in (0, 1]$ and $b > 0$. The easiest way to model the corresponding decision problem is to think of $(\mathcal{S}^B, \Omega^B)$ respectively $(\mathcal{S}^G, \Omega^G)$ as describing population distribution and sample space of observable quantities. However, in addition to choosing s from $\mathcal{S}^B$ or $\mathcal{S}^G$ and thereby defining Δ, Nature can now also choose the true treatment effect $\widetilde{\Delta}$ subject to (2). Choice of $\widetilde{\Delta}$ is easy to concentrate out and will therefore be omitted in notation below, but it makes closed-form expressions for regret much more complicated. This section's first main result provides minimax regret treatment rules for any such model.

Proposition 7 *Assume partial identification of $\widetilde{\Delta}$ is as in (2). Then:*

(i) *Consider the binomial experiment with independent randomization. Let δ^B_M denote treatment rule δ^B applied on the first M data points only. Let $N' = \max\{M \leq N : M \text{ odd}\}$ and $\gamma_N = 2^{-N'} \sum_{n > N'/2} \binom{N'}{n}(2n - N')$, noting that $\gamma_N \geq 1/2$, with equality if $N = 1$. Minimax regret is* *achieved by*

$$\widetilde{\delta}^B(\omega^B) \equiv \begin{cases} \delta^B(\omega^B) & \text{if } b \le \frac{a}{2\gamma_N}, \\ \alpha^* \delta^B_{N^*}(\omega^B) + (1-\alpha^*)\delta^B_{N^*-2}(\omega^B) & \text{if } \frac{a}{2\gamma_N} < b < a, \\ \frac{a}{b} \cdot \delta^B_1(\omega^B) + \left(1 - \frac{a}{b}\right) \cdot \frac{1}{2} & \text{if } a \le b < 1, \\ \frac{1}{2} & \text{if } b \ge 1. \end{cases} \tag{3}$$

Here

$$\begin{aligned} N^* &= \min\{M : \gamma_M > a/2b\}, \\ \alpha^* &= \frac{a/2b - \gamma_{N^*-2}}{\gamma_{N^*} - \gamma_{N^*-2}}. \end{aligned}$$

(ii) *Consider the binomial experiment with free within-sample treatment assignment. Then minimax regret is achieved by independent randomization in conjunction with δ^* as just defined.*

(iii) *Consider the Gaussian experiment. Minimax regret is achieved by*

$$\widetilde{\delta}^G(\omega) \equiv \begin{cases} \delta^G(\omega^G) & \text{if } b \le (\pi/2)^{1/2}\sigma a, \\ \Phi(\omega^G; 0, (2/\pi)(b/a)^2 - \sigma^2) & \text{if } (\pi/2)^{1/2}\sigma a < b < 1, \\ \frac{1}{2} & \text{if } b \ge 1. \end{cases} \tag{4}$$

Corollary 8 ***(i)*** *The minimax regret value of the binomial decision problem equals*

$$\widetilde{R}^B(N; a, b) = \begin{cases} 2^{-N'} \max_{\Delta \in [0,(1-b)/a]} \left\{ (a\Delta + b) \sum_{n < N'/2} \binom{N'}{n} (1+\Delta)^n (1-\Delta)^{N'-n} \right\} & \text{if } b \le \frac{a}{2\gamma_N}, \\ \frac{1}{2}\min\{b, 1\} & \text{otherwise.} \end{cases}$$

(ii) *The minimax regret value of the Gaussian decision problem equals*

$$\widetilde{R}^G(\sigma; a, b) = \begin{cases} \max_{\Delta \in [0,(1-b)/a]} \left\{ (a\Delta + b)\Phi(-\Delta; 0, \sigma^2) \right\} & \text{if } b \le (\pi/2)^{1/2}\sigma a, \\ \frac{1}{2}\min\{b, 1\} & \text{otherwise.} \end{cases}$$

To gain an intuition for this result, inspect (3) and imagine increasing b while holding a and N fixed, thus moving downward in the display. This corresponds to reducing the (worst-case) information contained in the data. For b small enough, δ^B applies unchanged, i.e. small enough distortions should be completely ignored. This robustness result links this section's analysis to the preceding section's, although it comes with the caveat that it applies for independent randomization only; the other sample designs do not in general achieve minimax regret here.

As b increases, the worst-case wedge between Δ and $\widetilde{\Delta}$ increases, thus the signal generated by the sampling universe becomes potentially less and less informative regarding the treatment population. Once b exceeds a certain threshold, this has the following consequence: If the decision maker too predictably follows the signal generated by the sampling universe, she becomes vulnerable to exploitation by Nature through states of the world in which Δ and $\widetilde{\Delta}$ have different signs, thus the signal has a misleading tendency. The intuition is clear in an extreme case: Imagine that $a \approx 0$ and b is large, yet sample size is such that the sign of Δ is correctly represented in the sample with probability near 1. Then $\max_{s \in \mathcal{S}^B} R(\delta^B, s) \approx b$, this value being attained by states where $\Delta = 1$ yet $\widetilde{\Delta} = -b$. Clearly the decision maker wants to avoid this trap; after all, even the no-data rule $\delta^{no-data} \equiv 1/2$ incurs expected regret of only 1/2. One salient way out is to ensure that the signal is not followed too closely. More specifically, for every (a, b), there exists a maximal sample size s.t. δ^B still achieves minimax regret. The mixture rule in the second line of (3) discards data until this threshold is reached, with randomization occurring because the relevant threshold "sample size" is generically a mixture of two adjacent ones. The expression in the third line of (3) reflects the case where even a sample size of $N = 1$ induces a "too informative" signal. Close inspection reveals that the third case prescribes mixing the benchmark rule, applied to the first sample point only, with $\delta^{no-data}$, thus it can be thought of as using an effective sample size of less than 1. An extreme is reached in the last line: $\widetilde{\delta}^B$ here coincides with $\delta^{no-data}$, i.e. the effective sample size is zero. All in all, moving downward in (3) reflects a reduction in effective sample size used.

In terms of the underlying equilibrium and also the decision problem's value function, the most important change occurs as $b = a/2\gamma_N$. Indeed, the reader may wonder how minimax regret can ever prescribe to discard data, given that the decision rules have to be Bayes against some prior. The answer is that as soon as δ^B fails to achieve minimax regret, the least favorable prior implies $\Delta = 0$, thus sample data are noise and the decision maker's equilibrium strategy in the fictitious game is unconstrained by the mechanics of Bayesian updating.

The intuition for the Gaussian case is just the same: Moving down in (4) corresponds to a decay of identification, countered by an increase of variance of the signal used. The middle line of the display really summarizes a two step procedure, whereby the decision maker first adds a realization of $N(0, \left((2/\pi)(b/a)^2 - \sigma^2\right))$ to the signal, thus artificially increasing its variance to the threshold value of $(2/\pi)(b/a)^2$, and then applies δ^G.

It is again instructive to alternatively think about the result in terms of holding (a, b) fixed and changing N. For typical (a, b), there exists a threshold sample size N such that imperfections of the experiments are completely ignored if sample size is below the threshold. At the same time, as soon as sample size exceeds the threshold, the particular minimax regret rule identified here effectively ceases to change because it ignores all data points beyond the threshold anyway. Thus, one again finds that minimax regret treatment rules "lock in" at certain limits as N becomes small or large enough.

The theme that using too much information eventually makes the decision maker vulnerable is connected to an independent finding by Tetenov (2009b), who establishes that in certain situations of partial identification, the returns – in terms of minimax regret – to improving identification completely dominate returns to increasing sampling precision once the latter exceeds a certain threshold. The result is driven by a similar mechanism, a possible intuition being that if very much sample information is available, then fully exploiting it makes the planner too predictable in scenarios where said information may be misleading. The same theme also emerges in the detail that proposition 7 generally recommends independent randomization and not matched pairs or constrained randomization. While these do achieve minimax regret for $\varepsilon = 0$ and also for sufficiently small ε, numerical evaluation reveals that there are parameter values where they fail to do so, specifically if b is just below $a/2\gamma_N$. The reason is that under some off-equilibrium priors, sample design by matched pairs is more efficient than independent randomization in estimating Δ. This added precision is arguably attractive in the benchmark problem, where it implies that constrained randomization weakly dominates independent randomization; in the extended scenario, it can make the decision maker vulnerable to manipulation.

All of these intuitions carry over to the hidden covariate model. The formal statement is as follows.

Proposition 9 *Assume partial identification of $\widetilde{\Delta}$ is through the hidden covariate model. Then:*

(i) *Consider the binomial experiment with independent randomization. Minimax regret is achieved by*

$$\widetilde{\delta}^B(\omega^B) \equiv \begin{cases} \delta^B(\omega^B) & \text{if } c < \log\left(\gamma_N^{-1} + \sqrt{1+\gamma_N^{-2}}\right), \\ \alpha^* \delta^B_{N^*} + (1-\alpha^*)\delta^B_{N^*-2} & \text{if } \log\left(\gamma_N^{-1} + \sqrt{1+\gamma_N^{-2}}\right) \leq c < \log\left(2+\sqrt{5}\right), \\ \frac{4\exp(c)}{\exp(2c)-1}\delta^B_1(\omega^B) + \left(1 - \frac{4\exp(c)}{\exp(2c)-1}\right)\cdot\frac{1}{2} & \text{if } c \geq \log\left(2+\sqrt{5}\right). \end{cases}$$

(ii) *Consider free within-sample treatment assignment. Then minimax regret is achieved by independent randomization in conjunction with the above decision rule.*

(iii) *Consider the Gaussian experiment. Minimax regret is achieved by*

$$\widetilde{\delta}^G(\omega^G) \equiv \begin{cases} \delta^G(\omega^G) & \text{if } c \leq \log\left(\sqrt{2\pi}\sigma + \sqrt{1+2\pi\sigma^2}\right), \\ \Phi\left(\omega^G; 0, \left(\frac{\exp(2c)-1}{\sqrt{8\pi}\cdot\exp(c)}\right)^2 - \sigma^2\right) & \text{if } c > \log\left(\sqrt{2\pi}\sigma + \sqrt{1+2\pi\sigma^2}\right). \end{cases}$$

Corollary 10 ***(i)*** *The binomial experiment's minimax regret value is*

$$\widetilde{R}^B(N;c) =$$

$$\begin{cases} 2^{-N'} \max_{\Delta\in[0,1]} \left\{ \frac{(1+\Delta)\exp(c)-(1-\Delta)}{(1+\Delta)\exp(c)+(1-\Delta)} \sum_{n<N'/2} \binom{N'}{n}(1+\Delta)^n(1-\Delta)^{N'-n} \right\} & \text{if } c \leq \log\left(\gamma_N^{-1} + \sqrt{1+\gamma_N^{-2}}\right), \\ \frac{\exp(c)-1}{2\exp(c)+1} & \text{otherwise.} \end{cases}.$$

(ii) *The Gaussian experiment's minimax regret value is*

$$\widetilde{R}^G(\sigma;c) = \begin{cases} \max_{\Delta\in[0,1]} \frac{(1+\Delta)\exp(c)-(1-\Delta)}{(1+\Delta)\exp(c)+(1-\Delta)} \Phi\left(-\Delta;0,\sigma^2\right) & \text{if if } c \leq \log\left(\sqrt{2\pi}\sigma + \sqrt{1+2\pi\sigma^2}\right), \\ \frac{\exp(c)-1}{2\exp(c)+1} & \text{otherwise.} \end{cases}.$$

Numerical values for the threshold distortions $\varepsilon^B(N)$ (for examples 1 through 3) or $c^B(N)$ (for example 4) below which δ^B applies are displayed in table 2.[14] For example, consider the contamination

[14] The table displays $\exp c^B(N)$ because it equals the maximally possible distortion of the probability to volunteer: $\Pr(Z=1|X=1)/\Pr(Z=1|X=0)$ is tightly bounded from above by $\exp c$.

$\mathbf{N}$	**1**	**3**	**5**	**10**	**20**	**50**	**100**	**200**	**500**	**1000**
$\boldsymbol{\varepsilon}^{\mathbf{B}}(\mathbf{N})$ (ex. 1)	0.5	0.333	0.267	0.203	0.142	0.089	0.063	0.044	0.028	0.020
$\boldsymbol{\varepsilon}^{\mathbf{B}}(\mathbf{N})$ (ex. 2)	1	0.667	0.533	0.406	0.284	0.178	0.126	0.089	0.056	0.040
$\boldsymbol{\varepsilon}^{\mathbf{B}}(\mathbf{N})$ (ex. 3)	0.5	0.4	0.348	0.289	0.221	0.151	0.112	0.082	0.053	0.038
$\exp \mathbf{c}^{\mathbf{B}}(\mathbf{N})$ (ex. 4)	4.23	3.00	2.53	2.10	1.71	1.42	1.28	1.19	1.12	1.08

Table 2: Numerical illustration: Threshold values from corollaries 8 and 10. Minimax regret treatment rules ignore disortions below these thresholds.

scenario from example 3. Then if $N = 10$, minimax regret prescribes to ignore contaminations of up to 29%, a scenario in which the sample information could be rather misleading. The corresponding number for a sample size of $N = 100$ still exceeds 10%. The rescaled, corresponding numbers for the Gaussian experiment are omitted for brevity but again provide very good approximations.

Some caveats to this section's analysis are as follows. First, minimax regret decision rules are essentially unique (up to subtleties about tie-breaking) only as long as they coincide with δ^B or δ^G, i.e. as long as learning occurs. Once a noninformative equilibrium has been reached, there is a plethora of minimax regret treatment rules. For an obvious example, the decision maker could use any appropriately sized subset of sample points rather than the first ones. Earlier versions of this paper contained yet other rules which are available from the author.[15]

Second, closed-form analysis is feasible because the bounds on $\widetilde{\Delta}$ are highly tractable. In related scenarios where they are characterized as solutions to linear programming problems (Balke and Pearl (2000)) or complicated integrals (Hill and Kreider (2009)), closed form analysis would seem a formidable challenge. This is true quite independently of whether one analyzes the binomial or the Gaussian experiment; thus, the difficulty is not mainly due to a quest for exact finite sample results.

5 Conclusion

Manski (2004) placed minimax regret treatment choice on the agenda of a small but active literature in econometrics. This literature was quite successful in characterizing minimax regret experimental designs and treatment rules for stylized problems, but recently also contained some findings that

[15] A notable case occurs in proposition 7 whenever $\widetilde{\delta}^B = \frac{a}{b} \cdot \delta_1^B + \left(1 - \frac{a}{b}\right) \cdot \frac{1}{2}$. Here $\mathbb{E}\widetilde{\delta}^B(\omega)$ is linear in $(\mu_1 - \mu_0)$, and any unbiased estimator of $(\mu_1 - \mu_0)$ can substitute for δ_1^B. A salient example is to rather use $\frac{1+I_N/N}{2}$. In the missing data scenario with $\varepsilon > 1/2$, this is an unbiased estimator of a rule discovered by Manski (2007a) for the case of known Δ. Furthermore, Manski (2007a) observed that under conditions very similar to $\varepsilon > 1/2$ (in particular, the identity of the better treatment is known to be unlearnable), this specific rule is finite sample minimax regret optimal. Manski's finding is, therefore, fully consistent with proposition 7.

should worry proponents of minimax regret. In particular, no-data recommendations, which are a core argument against maximin utility treatment choice, affect minimax regret as soon as covariates are introduced.

This paper extended previous findings in the direction of more realism. It did so in two ways that are substantively different but are amenable to similar analyses, down to specific algebraic tricks reappearing across proofs. Results contained some technical lessons as well as some substantive findings that may be more charitable to minimax regret. Regarding the technical lessons, the scope of game theoretic analysis of minimax regret problems was substantially enhanced, and in some cases, its limits may have been encountered. Even assuming that a Nash equilibrium of the fictitious game has been guessed, the hardest step in verification is to establish Nature's best-response condition, i.e. that the prior is least favorable given the decision rule. This generally requires to evaluate (and maximize over) the sampling expectation $\mathbb{E}\delta(\omega)$. In the benchmark problem, $\mathbb{E}\delta(\omega)$ is a value of either a binomial or a normal c.d.f. and, therefore, exceedingly tractable. In all cases where a minimax regret solution was provided, the least favorable prior was either such that the optimal rules were unchanged or such that the sample data were noise. In the former case, $\mathbb{E}\delta(\omega)$ is as in the benchmark problem; in the latter case, one has some latitude to construct decision rules whose sampling distribution is again tractable. If δ is more complex, one would worry that closed form solutions will quickly become infeasible. Minimax regret treatment choice with discrete X is very possibly an example.

Regarding substantive findings, the introduction of covariates with arbitrarily large effects as well as of unrestricted disagreement between response functions in sample and treatment population causes problems for minimax regret, yet minimax regret is locally insensitive to these features. To repeat, what is interesting is less that the effect of these features on minimax regret treatment rules approaches zero as their potential magnitude approaches zero; such a result should be expected from any reasonable treatment rule. It is not obvious, however, that said effect becomes exactly zero for some positive (and in some cases, arguably quite large) magnitude of the effect.

Another intriguing feature is the "stop collecting data" result for situations of sufficiently large limitations to data. The result is robust in that similar issues appear in related but different contexts (Tetenov (2009b)), and it may at first raise eyebrows – econometricians are surely not used to discarding data. But on second thought, it may not be so clear that this result is unreasonable. Without fully identifying assumptions, there is a limit – precisely delineated by partial identification analysis – on how informative the signal generated by the data can be at *any* sample size. This fact holds true irrespective of which decision criterion one uses. The discovery is merely that for the purpose of minimax regret treatment choice, the limit is fully attained for some finite sample size.

Regarding their methodological significance, I finally emphasize that these results are driven by a specific feature of minimax regret not shared by many other decision criteria. Bayesianism and maximin utility attempt to maximize some increasing function of the risk function $u(\delta, s)$, thus they intrinsically care about "good" versus "bad" states but not about how much efficiency loss – meaning underperformance relative to what could have been achieved ex post – a decision rule causes in a given state. Tick loss functions, which induce risk functions along the lines of $u(\delta, s) = \Pr(\delta$ assigns the better treatment$)$, have the opposite intuition: They are only sensitive to whether the ex post correct decision was made and not to the magnitude of loss caused. Minimax regret seems unique in combining aspects of both, reflecting a consideration for ex post optimality but also for the stakes at play. The local robustness features discovered here can be linked to this feature, that is, proof intuitions rely on it. This is some cause for optimism that similar features could be uncovered with respect to other modifications of the problem. It is hoped that along with axiomatic discussions and empirical applications, findings like these advance our understanding of the trade-offs involved in choosing among treatment rules and experimental designs for real-world decision problems.

A Proofs

Proposition 1 and Corollary 2

Preliminaries. This proposition synthesizes existing results, so a detailed proof will be omitted. However, it will be helpful to explain the least favorable priors and to explicitly recapitulate some arguments that will be heavily drawn upon. Recall that results are established by analyzing the following game: The decision maker (DM) picks a statistical decision rule $\delta \in \mathcal{D}$, and Nature picks a state of the world s. Both players may randomize. As $\mathcal{D}$ is closed under randomization, mixed strategies for DM can be identified with elements of $\mathcal{D}$. A mixed strategy for Nature is a distribution π over states. After both players moved, s is drawn from $\mathcal{S}$ according to π, ω is drawn from a distribution determined by s and the sample design, and δ is operated on ω. Nature's payoff (and DM's loss) is

$$R(\delta, s) = \max_{d \in \mathcal{D}} u(d, s) - u(\delta, s) = \max_{d \in \mathcal{D}} \{\mu_1 \mathbb{E} d(\omega) + \mu_0 (1 - \mathbb{E} d(\omega))\} - (\mu_1 \mathbb{E}\delta(\omega) + \mu_0(1 - \mathbb{E}\delta(\omega))).$$

Any Nash equilibrium (δ^*, π^*) of this game characterizes a minimax regret treatment rule δ^* and least favorable prior π^*. Note also the following simplification: As $\max_{d \in \mathcal{D}} \{\mu_1 \mathbb{E} d(\omega) + \mu_0(1 - \mathbb{E} d(\omega))\}$ does not depend on δ, a treatment rule δ minimizes the above iff it maximizes $\mu_1 \mathbb{E}\delta(\omega) + \mu_0(1 - \mathbb{E}\delta(\omega))$. Thus, δ^* must be Bayes given prior π^*, a condition that is typically easy to verify. The tricky part in most proofs is to verify Nature's best-response condition, namely that all states s^* in the support of π^* maximize $R(\delta^*, s)$.

It turns out that regret depends on states of the world s only through (μ_0, μ_1). Anticipating this observation, I will sometimes identify states with vectors $(\mu_0, \mu_1) \in [0, 1]^2$ right away.

(i) Let the sample design be independent randomization. Then a Nash equilibrium of the fictitious game is given by (δ^B, π^*), where π^* randomizes evenly over the two states

$$\{(a^*, 1-a^*), (1-a^*, a^*)\}$$

for some number $a^* > 1/2$. To verify the DM's best response condition, recall a best response must be Bayes against π^*, hence it must assign treatment 1 whenever $\mathbb{E}(Y_1|\omega^B) > \mathbb{E}(Y_0|\omega^B)$, where $\mathbb{E}(\cdot|\omega^B)$ denotes posterior expectation given prior π^* and data ω^B. $\mathbb{E}(Y_1|\omega^B) > \mathbb{E}(Y_0|\omega^B)$ obtains iff ω^B is more likely under $(\mu_0, \mu_1) = (1-a^*, a^*)$ than under $(\mu_0, \mu_1) = (a^*, 1-a^*)$. Letting N_t denote the number of sample subjects assigned to treatment 1 and n_t the number of recorded successes of treatment t, one can write out binomial likelihoods to find

$$\begin{aligned}
\mathbb{E}(Y_1|\omega^B) &> \mathbb{E}(Y_0|\omega^B) \\
\iff \binom{N_0}{n_0}(1-a^*)^{n_0}(a^*)^{N_0-n_0}\binom{N_1}{n_1}(a^*)^{n_1}(1-a^*)^{N_1-n_1} &> \binom{N_0}{n_0}(a^*)^{n_0}(1-a^*)^{N_0-n_0}\binom{N_1}{n_1}(1-a^*)^{n_1}(a^*)^{N_1-n_1} \\
\iff (a^*)^{I_N} &> (1-a^*)^{I_N} \\
\iff I_N &> 0.
\end{aligned}$$

To verify Nature's best-response condition, assume initially that N is odd and write

$$\delta^B = 1\{I_N > 0\} = 1\left\{\frac{N+I_N}{2} > \frac{N}{2}\right\},$$

yet under independent randomization, $(N + I_N/N)/2$ is distributed as a binomial variable with parameters $(N, (1+\mu_1-\mu_0)/2)$. Thus, Nature's best-response problem against δ^B can be written as

$$\max_{s \in \mathcal{S}} R(\delta^B, s) = \max_{s \in \mathcal{S}} \max \left\{ \begin{array}{c} (\mu_1-\mu_0)F\left(\frac{N-1}{2}; N, \frac{1+\mu_1-\mu_0}{2}\right), \\ (\mu_0-\mu_1)\left(1-F\left(\frac{N-1}{2}; N, \frac{1+\mu_1-\mu_0}{2}\right)\right) \end{array} \right\}, \tag{5}$$

where $F(\cdot; N, p)$ denotes the binomial c.d.f. with indicated parameters. As

$$1-F\left(\frac{N-1}{2}; N, \frac{1+\mu_1-\mu_0}{2}\right) = F\left(\frac{N-1}{2}; N, \frac{1+\mu_0-\mu_1}{2}\right),$$

this objective depends on s only through $|\mu_0-\mu_1|$. It follows that one best response to δ^B is π^* as described above, where $a^* = (1+\Delta^*)/2$ and Δ^* solves

$$\max_{\Delta \in [0,1]} \phi(\Delta), \phi(\Delta) = \Delta F\left(\frac{N-1}{2}; N, \frac{1+\Delta}{2}\right). \tag{6}$$

This establishes equilibrium. The value of the game is $R^B(N) = \phi(\Delta^*)$, thus solving (6) leads to the corollary.

Now let N be even. I_N is necessarily even in this case, so if $I_N \neq 0$, $\delta^B(\omega)$ is invariant to dropping the last observation from the sample. Conditionally on $I_N = 0$, I_{N-1} (i.e. I_N applied after dropping the last observation from the sample) is equally likely to equal either of $\{-1, 1\}$; thus, the 50/50 tie-breaking can be implemented by dropping the last observation from the sample. In sum, if DM plays δ^B, Nature's best-response problem is as if δ^B were applied to a sample of size $(N-1)$, and π^* (with Δ^* computed for sample size $N-1$) is a best response. On the other hand, the argument that δ^B best responds to π^* is unchanged. Thus an equilibrium has been found, and $R^B(N) = R^B(N-1)$ in this case.

Consider now constrained randomization, then it is easy to see that δ^B continues to be a best response to π^*. Brute force algebra in Schlag (2006) establishes that $(a^*, 1-a^*)$ maximizes $\mathbb{E}\delta^B(\omega^B)$ among states of the form $(a, a-\Delta^*)$, thus one can restrict attention to states of the form $(a, 1-a)$; but then it is clear that π^* remains a best response to δ^B. Matched pairs coincide with constrained randomization if N is even.

(ii) Let sample design be a choice variable. Then δ^B combined with *any* sample design is a best response to π^*, but π^* is a best response to δ^B in combination with the above (and other) sample designs, thus a Nash equilibrium has been found.

(iii) The equilibrium is much like the one in (i), in particular $\overline{\pi}^*$ changes only by the way in which a^* is computed. Nature's best-response problem is

$$\begin{aligned}\max_{s \in \mathcal{S}} R(\delta^G, s) &= \max_{s \in \mathcal{S}} \max \left\{ (\mu_1 - \mu_0) \left(1 - \mathbb{E}\delta^G(\omega^G)\right), (\mu_0 - \mu_1)\mathbb{E}\delta^G(\omega^G) \right\} \\ &= \max_{s \in \mathcal{S}} \max \left\{ \Delta\Phi(-\Delta; 0, \sigma^2), -\Delta\Phi(\Delta; 0, \sigma^2) \right\}\end{aligned}$$

where the simplification defines $\Delta = \mu_1 - \mu_0$ and uses that $\delta^G = 1\{\omega^G > 0\}$, thus $1 - \mathbb{E}\delta^G(\omega^G) = \Pr(\omega^G < 0) = \Phi(0; \Delta, \sigma^2) = \Phi(-\Delta; 0, \sigma^2)$ (and similarly for $\mathbb{E}\delta^G(\omega^G)$). This objective depends on s only through Δ, and it is symmetric around $\Delta = 0$. Hence, a best response is a prior just like π^* from part (i), except that Δ^* solves

$$\max_{\Delta \in [0,1]} \Delta\Phi(-\Delta; 0, \sigma^2). \tag{7}$$

The argument that δ^G best responds do π^* is similar to but easier than before. The value of the game can be computed by solving (7), which leads to the corollary.

Proposition 3

(i)

Preliminaries. Suppose that $R^B(N) \geq 2\kappa$. Then an equilibrium of the game is given by $(\overline{\delta}^B, \overline{\pi}^*)$, where the prior $\overline{\pi}^*$ is characterized by $s_x = s_{x'}$ for all x, x', thus treatment response is constant in x, and s_x randomizes evenly over the same $\{(a^*, 1-a^*), (1-a^*, a^*)\}$ as in proposition 1(i). In words, the equilibrium mimics the equilibrium from proposition 1(i) (and has the same value $R^B(N)$), with both players ignoring the existence of X. The claim will be established by verifying the equilibrium. Assume first independent randomization with N odd. Recall that states are (measurable) mappings from $\mathcal{X}$ to $\Delta\{0,1\}^2$ that assign to every covariate value x a distribution of potential outcomes $P(Y_{0x}, Y_{1x})$ with expectations $(\mu_{0x}, \mu_{1x}) = \mathbb{E}(Y_{0x}, Y_{1x})$. Define the marginal expectations $\mu_t = \int \mu_{tx} dP(X)$.

Step 1: Simplifying Nature's best-response problem. Nature's response to $\overline{\delta}^B$ must be supported on states that solve

$$
\begin{aligned}
& \max_{s\in\mathcal{S}^B} R(\overline{\delta}^B, s) \\
= & \max_{s\in\mathcal{S}^B} \int \left(\max\{\mu_{0x},\mu_{1x}\} - \mu_{0x}\left(1-\mathbb{E}\overline{\delta}^B\left(\omega^B\right)\right) - \mu_{1x}\mathbb{E}\overline{\delta}^B\left(\omega^B\right)\right) dP(X) \\
= & \max_{s\in\mathcal{S}^B} \int \left(\mathbb{E}\overline{\delta}^B\left(\omega^B\right)\cdot(\mu_{0x}-\mu_{1x})\cdot 1\{\mu_{0x}>\mu_{1x}\} + \left(1-\mathbb{E}\overline{\delta}^B\left(\omega^B\right)\right)\cdot(\mu_{1x}-\mu_{0x})\cdot 1\{\mu_{1x}\geq\mu_{0x}\}\right) dP(X) \\
= & \max_{s\in\mathcal{S}^B} \left\{\mathbb{E}\overline{\delta}^B\left(\omega^B\right)\int(\mu_{0x}-\mu_{1x})\cdot 1\{\mu_{0x}>\mu_{1x}\}dP(X) + \left(1-\mathbb{E}\overline{\delta}^B\left(\omega^B\right)\right)\int(\mu_{1x}-\mu_{0x})\cdot 1\{\mu_{1x}\geq\mu_{0x}\}dP(X)\right\},
\end{aligned}
$$

where the first step substitutes in for the definition of $R(\overline{\delta}^B, s)$ and the next steps rearrange terms, using that $\overline{\delta}^B\left(\omega^B\right)$ does not depend on x.

Define

$$
\begin{aligned}
a &= \frac{\int(\mu_{0x}-\mu_{1x})1\{\mu_{0x}>\mu_{1x}\}dP(X)}{\int 1\{\mu_{0x}>\mu_{1x}\}dP(X)}, \\
b &= \frac{\int(\mu_{1x}-\mu_{0x})1\{\mu_{0x}\leq\mu_{1x}\}dP(X)}{\int 1\{\mu_{0x}\leq\mu_{1x}\}dP(X)}, \\
p &= \int 1\{\mu_{0x}>\mu_{1x}\}dP(X),
\end{aligned}
$$

with the convention that $a=0$ if $p=0$ and $b=0$ if $p=1$. Recall that with N odd, one has $\mathbb{E}\overline{\delta}^B\left(\omega^B\right) = 1 - F\left(\frac{N-1}{2}; N, \frac{1+\mu_1-\mu_0}{2}\right)$, which depends on s only through $\Delta = \mu_1 - \mu_0 = (1-p)b - pa$. Thus, for the remainder of this proof only, define $f((1-p)b-pa) = 1 - F\left(\frac{N-1}{2}; N, \frac{1+(1-p)b-pa}{2}\right)$, then Nature's best-response problem can be written as

$$
\max_{s\in\mathcal{S}^B} R(\overline{\delta}^B, s) = \max_{s\in\mathcal{S}}\left\{paf((1-p)b-pa) + (1-p)b\left[1-f((1-p)b-pa)\right]\right\} \tag{8}
$$

This objective depends on s only through (a,b,p).

Step 2: Maximizing (8) over aligned states. Call a state "aligned" if $(\mu_{0x}-\mu_{1x})(\mu_{0x'}-\mu_{1x'}) \geq 0$ for all $x, x' \in \mathcal{X}$, thus the optimal treatment is the same across covariate values. Suppose the maximum in (8) is attained by an aligned state that induces values (a^*, b^*, p^*). The state being aligned implies that either $a^* = 0$ or $b^* = 0$, thus the same value of the objective is attained by setting

$$(\mu_{0x}, \mu_{1x}) = ((1 + a^*p^* - b^*(1-p^*))/2, (1 - a^*p^* + b^*(1-p^*))/2)$$

for all x. Hence, for maximizing (8) over aligned states, one can restrict attention to states s.t. (μ_{0x}, μ_{1x}) is constant. But then (8) simplifies to (5), hence $\overline{\pi}^*$ is a best response. On the other hand, $\overline{\delta}^B$ is Bayes against $\overline{\pi}^*$ by the algebra from proposition 1(i), thus the equilibrium is verified.

Step 3: Bounding the value of (8) over misaligned states. The constraint that $|\mu_{1x} - \mu_{0x}| \leq \kappa$ for all x implies that

$$\mu_{1x} - \mu_{0x} + \mu_{0x'} - \mu_{1x'} = \underbrace{\mu_{1x} - \mu_{1x'}}_{\leq\kappa} + \underbrace{\mu_{0x'} - \mu_{0x}}_{\leq\kappa} \leq 2\kappa \tag{9}$$

for all $x, x' \in \mathcal{X}$. Consider now "misaligned" states s in which $(\mu_{0x} - \mu_{1x})(\mu_{0x'} - \mu_{1x'}) < 0$ for some $x, x' \in \mathcal{X}$. Then

$$2\kappa \geq \sup_{x\in\mathcal{X}} \{\mu_{0x} - \mu_{1x}\} + \sup_{x\in\mathcal{X}} \{\mu_{1x} - \mu_{0x}\} \geq a + b,$$

where the first inequality uses (9) and the second one uses that in misaligned states, both $\sup_{x\in\mathcal{X}} \{\mu_{0x} - \mu_{1x}\} > 0$ and $\sup_{x\in\mathcal{X}} \{\mu_{1x} - \mu_{0x}\} > 0$, thus $a \leq \sup_{x\in\mathcal{X}} \{\mu_{0x} - \mu_{1x}\}$ and $b \leq \sup_{x\in\mathcal{X}} \{\mu_{1x} - \mu_{0x}\}$. Hence, the value of (8) is bounded from above by the value of

$$\begin{aligned} & \max_{a,b,p\in[0,1]^3: a+b\leq 2\kappa} \phi(a,b,p), \\ \phi(a,b,p) &= pa \cdot f((1-p)b - pa) + (1-p)b \cdot [1 - f((1-p)b - pa)] \\ &= (pa - (1-p)b) \cdot f((1-p)b - pa) + (1-p)b. \end{aligned} \tag{10}$$

This generally provides an upper bound on $\max_{s\in\mathcal{S}^B} R(\overline{\delta}^B, s)$ because measurability of s generates constraints on p that are not incorporated in (10). However, the value of (10) is attainable if there exist (a^*, b^*, p^*) that maximize (10) and a measurable event $E \subseteq \mathcal{X}$ s.t. $\Pr(X \in E) = p^*$, in which case one can set

$$(\mu_{0x}, \mu_{1x}) = \begin{cases} \left(\frac{1+a^*}{2}, \frac{1-a^*}{2}\right) & \text{if } x \in E, \\ \left(\frac{1-b^*}{2}, \frac{1+b^*}{2}\right) & \text{if } x \notin E. \end{cases} \tag{11}$$

It follows that if the value of (10) is attained by some (a^*, b^*, p^*) with $p^* \in \{0, 1\}$, then it is attained by an aligned state, and step 2 applies.

The remainder of this proof shows that if the value of (10) is attained by some (a^*, b^*, p^*) with $p^* \in (0, 1)$, then this value equals $\kappa/2$. It follows that if $R^B(N) \geq \kappa/2$, then $\overline{\pi}^*$ best responds to $\overline{\delta}^B$, and $(\overline{\delta}^B, \overline{\pi}^*)$ is an equilibrium. Thus, assume that a solution (a^*, b^*, p^*) to (10) has $p^* \in (0, 1)$. Partial derivatives of ϕ with respect to a and b are

$$\begin{aligned}
\phi_a &= p \cdot f((1-p)b - pa) - p(pa - (1-p)b) \cdot f'((1-p)b - pa) \\
&= p\left[f((1-p)b - pa) - (pa - (1-p)b) \cdot f'((1-p)b - pa)\right], \\
\phi_b &= -(1-p) \cdot f((1-p)b - pa) + (1-p)(pa - (1-p)b) \cdot f'((1-p)b - pa) + 1 - p \\
&= (1-p)\left[1 - f((1-p)b - pa) + (pa - (1-p)b) \cdot f'((1-p)b - pa)\right].
\end{aligned}$$

Recalling that f is nonnegative and strictly increasing, it follows that $p^*a - (1-p^*)b \leq 0 \Rightarrow \phi_a > 0$ and $p^*a - (1-p)b^* \geq 0 \Rightarrow \phi_b > 0$. One of those always obtains, hence the constraint $a + b \leq 2\kappa$ binds. Substituting in for $a = 2\kappa - b$ leads to the simplified problem

$$\begin{aligned}
&\max_{b \in [0,2\kappa], p \in [0,1]} \widetilde{\phi}(b, p), \\
\widetilde{\phi}(b, p) &= (2\kappa p - b) \cdot f(b - 2\kappa p) + (1-p)b.
\end{aligned}$$

Suppose $b^* = 2\kappa$, thus $a^* = 0$. Then the value of the problem can be attained by setting $(a, b, p) = (0, 2\kappa p^*, 1)$, so step 2 applies once again. A symmetric argument applies if $b^* = 0$. Suppose, therefore, that $b^* \in (0, 2\kappa)$. Together with $p^* \in (0,1)$, this requires first-order conditions to hold:

$$\begin{aligned}
\widetilde{\phi}_b &= -f(b - 2\kappa p) + (2\kappa p - b) \cdot f'(b - 2\kappa p) + 1 - p \overset{!}{=} 0, \\
\widetilde{\phi}_p &= 2\kappa \cdot f(b - 2\kappa p) - 2\kappa(2\kappa p - b) \cdot f'(b - 2\kappa p) - b \overset{!}{=} 0.
\end{aligned}$$

Substituting in and observing cancellations leads to $b^* = 2\kappa(1-p^*)$. Substituting this into the objective leads to a further concentrated objective function $\overline{\phi}(p) = (4\kappa p - 2\kappa) \cdot f(2\kappa - 4\kappa p) + 2\kappa(1-p)^2$. Noting the linear relation $\Delta = 2\kappa - 4\kappa p$, one can express this as optimization in Δ, leading to the ultimately simplified problem

$$\max_{\Delta \in [-2\kappa, 2\kappa]} \overline{\phi}(\Delta), \overline{\phi}(\Delta) = -\Delta f(\Delta) + \frac{(2\kappa + \Delta)^2}{8\kappa}. \tag{12}$$

The objective function has derivatives

$$\begin{aligned}
\overline{\phi}_\Delta &= -\Delta f'(\Delta) - f(\Delta) + \frac{1}{2} + \frac{\Delta}{4\kappa}, \\
\overline{\phi}_{\Delta\Delta} &= -2f'(\Delta) - \Delta f''(\Delta) + \frac{1}{4\kappa}.
\end{aligned}$$

Recalling that $f(0) = 1/2$, inspection of $\overline{\phi}_\Delta$ reveals a critical point at $\Delta = 0$ with value $\overline{\phi}(0) = \kappa/2$. Further analysis of (12) requires computation of analytic derivatives. If $N = 1$, then $f(\Delta) = (1+\Delta)/2$, $f'(\Delta) = 1/2$, and $f''(\Delta) = 0$, leading to $\overline{\phi}_{\Delta\Delta} = 1/4\kappa - 1$. Hence, $\overline{\phi}$ is either convex throughout or concave throughout, and the critical point at $\Delta = 0$ is the only interior critical point of $\overline{\phi}$, establishing the claim. Suppose now that $N \geq 3$. Recalling that $F(n; N, \pi) = (N-n)\binom{N}{n} \int_0^{1-\pi} t^{N-n-1}(1-t)^n dt$, write

$$
\begin{aligned}
f(\Delta) &= 1-\frac{N+1}{2}\binom{N}{(N-1)/2}\int_0^{(1-\Delta)/2} t^{(N-1)/2}(1-t)^{(N-1)/2}dt \\
f'(\Delta) &= (N+1)\binom{N}{(N-1)/2}2^{-N-1}(1-\Delta^2)^{(N-1)/2} \\
f''(\Delta) &= -(N+1)\binom{N}{(N-1)/2}2^{-N-1}(N-1)\Delta(1-\Delta^2)^{(N-3)/2},
\end{aligned}
$$

leading to

$$
\overline{\phi}_{\Delta\Delta}=\frac{1}{4\kappa}+(N+1)\binom{N}{(N-1)/2}2^{-N-1}(1-\Delta^2)^{(N-3)/2}\left((N+1)\Delta^2-2\right). \tag{13}
$$

This function is symmetric around $\Delta = 0$ and is positive whenever $|\Delta| \geq \sqrt{2/(N+1)}$. If $|\Delta| < \sqrt{2/(N+1)}$, then inspecting (13) and recalling $N \geq 3$ reveals that $\overline{\phi}_{\Delta\Delta}$ increases in $|\Delta|$. It follows that $\overline{\phi}$ is either convex throughout, or concave on an interval centered at $\Delta = 0$ and convex otherwise. Thus if there is an interior optimum, then it coincides with the critical point identified at $\Delta = 0$ and, therefore, has value $\kappa/2$.

The extension to N even is just as before. For constrained randomization, it again follows from algebra in Schlag (2006) that Nature is best off playing only states which are symmetric in the sense that $\mu_0 + \mu_1 = 1$. After restricting attention to those states, the proof applies unchanged and thereby also extends to matched pairs.

(ii) Consider the extended game in which the decision maker may also choose the sample design. If Nature plays $\overline{\pi}^*$, then $\overline{\delta}^B$ in conjunction with any of the sample designs is a best response. Yet if DM chooses any of those strategies, then analysis of Nature's best-response problem is just as in part (i), hence the same result obtains. The same argument applies to the extended game in which the decision maker may also choose a sample design with respect to X, and these arguments can also be combined.

(iii) Here, the claim is that if $R^G(\sigma) \geq \kappa/2$, then there is an equilibrium entirely analogous to the one in part (i), except that a^* is computed as in proposition 1(iii). Observe that $\overline{\delta}^G$ can be written as $\overline{\delta}^G(\omega^G) = 1\left\{\overline{\omega}^G > 0\right\}$, where $\overline{\omega}^G = \sum_{x\in\mathcal{X}} \omega_x^G \Pr(X = x)$ is distributed according to $N\left(\mu_1 - \mu_0, \sigma^2\right)$. Hence, $\mathbb{E}\overline{\delta}^G\left(\omega^G\right)$ depends on s only through $\Delta = \mu_1 - \mu_0$, and step 1 of the proof goes through as before, except that one now needs to define $f((1-p)b - pa) = 1 - \Phi(0; (1-p)b - pa, \sigma^2)$. Step 2 goes through precisely as before; observe in particular that $\overline{\omega}^G$ weights the signals ω_x^G according to their precision, thus it is Bayes against $\overline{\pi}^*$. Step 3 is unchanged up to the computation of derivatives of f following expression (12). To complete the argument from there, recall

$$
f(\Delta) = 1 - \Phi(0; \Delta, \sigma^2) = \Phi(0; -\Delta, \sigma^2) = \Phi(\Delta; 0, \sigma^2) \tag{14}
$$

to write

$$
\begin{aligned}
f'(\Delta) &= \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{\Delta^2}{2\sigma^2}\right), \\
f''(\Delta) &= -\frac{1}{\sqrt{2\pi\sigma^2}} \frac{\Delta}{\sigma^2} \exp\left(-\frac{\Delta^2}{2\sigma^2}\right),
\end{aligned}
$$

thus

$$
\overline{\phi}_{\Delta\Delta} = \frac{1}{4\kappa} + \frac{1}{\sqrt{2\pi}} \sigma^{-3} \exp\left(-\frac{\Delta^2}{2\sigma^2}\right) \cdot \left(2\Delta^2 - \sigma^2\right),
$$

which is amenable to the same analysis that was applied to (13).

(iv) It was shown above that $\max_{s \in \mathcal{S}^B} R(\overline{\delta}^B, s) \leq \max\{R^B(N), \kappa/2\}$ respectively $\max_{s \in \mathcal{S}^G} R(\overline{\delta}^G, s) \leq \max\{R^G(N), \kappa/2\}$, both of which are strictly less than $1/2$ if $\kappa < 1$. Yet the best no-data rule incurs maximal expected regret of

$$
\min_{d \in [0,1]} \max_{s \in \mathcal{S}} \int \left(\max\{\mu_{0x}, \mu_{1x}\} - \mu_{0x}(1-d) - \mu_{1x} d\right) dP(X) = \min_{d \in [0,1]} \max\{d, 1-d\} = 1/2
$$

for both $\mathcal{S} = \mathcal{S}^B$ and $\mathcal{S} = \mathcal{S}^G$.

Proposition 4 In all cases, the core observation is that choosing $(Y_{tx}, Y_{tx'})$ to be Bernoulli with expectations $(a, 1-a)$ maximizes $|\mu_{tx} - \mu_{tx'}|$ subject to constraints in terms of other distance metrics on $(P(Y_{tx}), P(Y_{tx'}))$. Thus, one can derive bounds on $|\mu_{tx} - \mu_{tx'}|$ from bounds in other metrics and then invoke proposition 3. The resulting bound cannot in general be improved upon because if the bound on expected regret over misaligned states of $\kappa/2$ is attained, then by inspection of (11), it is attained by a state that solves this maximization problem, hence the constraint in terms of another metric binds.

Total Variation Distance: For Bernoulli distributions, $||P(Y_{tx}), P(Y_{tx'})||_{TV} = |\mu_{tx} - \mu_{tx'}|$, thus the claim. For general distributions on $[0,1]$, consider maximization of $(\mu_{tx} - \mu_{tx'})$ for simplicity. Suppose that $||P(Y_{tx}), P(Y_{tx'})||_{TV} \leq \gamma$. Let $E \in \mathcal{B}([0,1])$ be the largest event on which $P(Y_{tx})$ dominates $P(Y_{tx'})$, then $|\Pr(Y_{tx} \in E) - \Pr(Y_{tx'} \in E)| \leq \gamma$. It follows that if $(P^*(Y_{tx}), P^*(Y_{tx'}))$ are Bernoulli distributions with parameters $(\Pr(Y_{tx} \in E), \Pr(Y_{tx'} \in E))$, then $||P(Y_{tx}), P(Y_{tx'})||_{TV} \leq \gamma$ as well, yet $(\mu_{tx} - \mu_{tx'})$ has been increased. Thus, the solution of the problem is attained by Bernoulli distributions.

Log Odds Ratio Distance: For Bernoulli distributions,

$$
||P(Y_{tx}), P(Y_{tx'})||_{LOR} \leq \gamma \Longleftrightarrow \log\left((\mu_{tx}(1-\mu_{tx'}))/(\mu_{tx'}(1-\mu_{tx}))\right) \leq \gamma.
$$

This, in turn, implies that

$$2\log\frac{1+|\mu_{tx}-\mu_{tx'}|}{1-|\mu_{tx}-\mu_{tx'}|}\leq\gamma$$

and therefore the result through proposition 3. While the above implication is one-sided, it will be seen that the result is tight. Thus, consider the problem

$$\max_{(\mu,\rho)\in[0,1]^2}|\mu-\rho|\quad\text{s.t.}\ \log\frac{\mu(1-\rho)}{\rho(1-\mu)}\leq\gamma.$$

Noting the problem's symmetry, replace $|\mu-\rho|$ with $(\mu-\rho)$ in the objective for tractability. Letting λ denote the Lagrange multiplier on the constraint, relevant partial derivatives are

$$\begin{aligned}\nabla_\mu(x) &= 1-\lambda\left(\frac{1}{x}+\frac{1}{1-x}\right)\\ \nabla_\rho(x) &= -1+\lambda\left(\frac{1}{x}+\frac{1}{1-x}\right).\end{aligned}$$

Clearly $\nabla_\mu(x)=\nabla_\mu(1-x)=-\nabla_\rho(x)=-\nabla_\rho(1-x)$. As $\mu=\rho$ corresponds to minimizing the objective, a solution (μ^*,ρ^*) must have the feature that $\mu^*=1-\rho^*$. Using this constraint to eliminate ρ and reparameterizing the problem by setting $\mu=(1+\Delta)/2\Leftrightarrow\Delta=|\mu-\rho|$, one has the new problem

$$\max_{\Delta\in[0,1]}\Delta\quad\text{s.t.}\ \gamma\geq\log\frac{\left(\frac{1+\Delta}{2}\right)^2}{\left(\frac{1-\Delta}{2}\right)^2}=2\log\frac{1+\Delta}{1-\Delta}.$$

This is solved by $\gamma=2\log((1+\Delta)/(1-\Delta))$, establishing the claim. To see that the bound cannot be improved upon, note that the misaligned states that potentially achieve the regret bound of $\kappa/2$ in proposition 3 have the feature that if $\mu_{0x}\neq\mu_{0x'}$, then $\mu_{0x}=1-\mu_{0x'}$, thus they represent solutions to the above problem for some γ. The extension to general distributions is similar to before.

Kullback-Leibler Divergence: For Bernoulli distributions,

$$D_{KL}(P(Y_{tx})||P(Y_{tx'}))\leq\gamma\iff\mu_{tx}\log(\mu_{tx}/\mu_{tx'})+(1-\mu_{tx})\log((1-\mu_{tx})/(1-\mu_{tx'})).$$

The argument is then similar to the one for log odds ratio distance, except that it is based on analyzing the maximization problem

$$\max_{(\mu,\rho)\in[0,1]^2}|\mu-\rho|\quad\text{s.t.}\ \mu\log\frac{\mu}{\rho}+(1-\mu)\log\frac{1-\mu}{1-\rho}\leq\gamma.$$

This problem has the same symmetry properties as the one for log odds ratio distance, so a similar reparameterization leads to

$$\max_{\Delta\in[0,1]}\Delta\quad\text{s.t.}\ \gamma\geq\frac{1+\Delta}{2}\log\frac{1+\Delta}{1-\Delta}+\frac{1-\Delta}{2}\log\frac{1-\Delta}{1+\Delta}=\Delta\log\frac{1+\Delta}{1-\Delta},$$

and the argument is concluded just as before.

Proposition 5 The core insight, shown in the next paragraph, is that if X is continuous, then the minimax regret value of the decision problem cannot be less than $\kappa/2$. Now assume the decision maker plays $\overline{\delta}^B$. Proposition 3 established that $\max R(\overline{\delta}^B, s)$ is bounded above by $\max\{R^B(N), \kappa/2\}$ and that in those cases where the value equals $R^B(N)$, $\overline{\delta}^B$ achieves minimax regret. For the other cases, it now follows that $\overline{\delta}^B$ achieves the lower bound of $\kappa/2$ on the decision problem's minimax regret value and must, therefore, achieve minimax regret (though more properly called infsup regret in this case).

To bound the value of the decision problem by $\kappa/2$ from below, construct a sequence of priors $\pi_i, i = 1, 2, \ldots$ as follows. Define the partition $W_i \equiv \{[0, 1/i], (1/i, 2/i], \ldots ((i-1)/i, 1]\}$ of the unit interval. Let $(w_i^j)_{j=1}^{2^i}$ collect the subsets of W_i in arbitrary order. Define the collection of states $\left(s_i^j\right)_{j=1}^{2^i}$ by identifying s_i^j with the degenerate distribution concentrated at

$$(\mu_{0x}, \mu_{1x}) = \\ \left((1+\kappa/2)\,1\left\{x \in w_i^j\right\} + (1-\kappa/2)\,1\left\{x \in w_i^j\right\}, (1-\kappa/2)\,1\left\{x \in w_i^j\right\} + (1+\kappa/2)\,1\left\{x \in w_i^j\right\}\right)$$

for every x. Let π_i be the uniform distribution over states $\left(s_i^j\right)_{j=1}^{2^i}$, i.e. π_i assigns probability 2^{-i} to every s_i^j. Note the following features of π_i: (i) The prior expectation of (Y_{0x}, Y_{1x}) equals $(1/2, 1/2)$. (ii) With slight abuse of notation, let $w_i(x)$ be the element of W_i that contains x. Then s_x and $s_{x'}$ are independent whenever $w_i(x) \neq w_i(x')$. Now minimal adaptation of algebra from Stoye (2009a, proposition 4) shows that $\lim_{i\to\infty} \min_{\delta \in D} \int R(\delta, s) d\pi_i = \kappa/2$ (in particular, the values of Bayes responses to priors π_i converge to $\kappa/2$), thus $\inf_{\delta\in\mathcal{D}} \sup_{s\in\mathcal{S}} R(\delta, s) \geq \kappa/2$.

Proposition 6

(i) For this proof, write Δ_N^* for the quantity Δ^* from proposition 1(i) to emphasize its dependence on N. Stoye (2009a, proposition 3) shows that the least favorable prior which supports separated inference across covariates can be described as follows: For every $x \in \mathcal{X}$, (μ_{0x}, μ_{1x}) equals either $\left((1+\Delta_{N_x}^*)/2, (1-\Delta_{N_x}^*)/2\right)$ or $\left((1-\Delta_{N_x}^*)/2, (1+\Delta_{N_x}^*)/2\right)$, where the two possibilities are equally likely and are realized independently of each other for any $x \neq x'$. To establish the claim, it therefore suffices to show that $\Delta_N^* = O(N^{-1/2})$. Thus, use (6) and a change of variables $\widetilde{\Delta} = \sqrt{N}\Delta_N^*$ to write

$$\begin{aligned} \Delta_N^* &= \arg\max_{\Delta\in[0,1]} \Delta F\left(\frac{N'-1}{2}; N', \frac{1+\Delta}{2}\right) \\ \iff \sqrt{N}\Delta_N^* = & \quad \arg\max_{\widetilde{\Delta}\in[0,\sqrt{N}]} \widetilde{\Delta}N^{-1/2} F\left(\frac{N'-1}{2}; N', \frac{1+\widetilde{\Delta}N^{-1/2}}{2}\right) \\ &= \arg\max_{\widetilde{\Delta}\in[0,\sqrt{N}]} \widetilde{\Delta} F\left(\frac{N'-1}{2}; N', \frac{1+\widetilde{\Delta}N^{-1/2}}{2}\right). \end{aligned}$$

Now, the Berry-Esseen theorem yields

$$F\left(\frac{N'-1}{2}; N, \frac{1+\widetilde{\Delta}_N N^{-1/2}}{2}\right) \to \Phi(0, -\widetilde{\Delta}_N, 1)$$

for any sequence $\widetilde{\Delta}_N$ s.t. $\widetilde{\Delta}_N N^{-1/2} \to 0$ (thus the approximated sequence of binomial distributions has nonvanishing variance). At the same time, arguments extremely similar to the evaluations after (12) imply that the maximand is quasiconcave on $\left[0, \sqrt{N}\right]$. It follows from standard arguments for consistency of extremum estimators that

$$\sqrt{N}\Delta_N^* \to \arg \max_{\Delta \in [0,\infty)} \Delta\Phi(0, -\Delta, 1),$$

thus $\Delta_N^* = O(N^{-1/2})$.

(ii) This follows much like part (i), building on the observation that after writing $\Delta_\sigma^* = \arg\max \Delta\Phi(-\Delta; 0, \sigma)$ and noting $\Phi(-\Delta; 0, \sigma) = \Phi(-2\Delta; 0, 4\sigma^2)$, one finds $\Delta_\sigma^* = \sigma\Delta_1^*$.

(iii) Proposition 1 is easily extended to the case where N is a random variable with known distribution Q. In this case, δ^B combined with either independent or constrained randomization (or even matched pairs if N only takes even values) continues to achieve minimax regret. The Nash equilibrium supporting δ^B is just as before, except that Δ^* solves

$$\Delta^* = \arg\max \Delta \cdot \int F\left(\frac{N'-1}{2}; N', \frac{1+\Delta}{2}\right) dQ.$$

With this extension of proposition 1 in place, the proof is much like part (i). Specifically, let $i = 1, 2, \ldots$ index any sequence of distributions Q_i of N s.t. there exists a nonstochastic sequence N_i with $N_i \to \infty$ and $\int N/N_i dQ_i(N) \to 1$. Then $\Delta_{N_i}^* = O(N_i^{-1/2})$ by essentially the same argument as before.

Example 1 Let $D \in \{0,1\}$ indicate assigned treatment, let $T \in \{0,1\}$ denote received treatment, and let $C \in \{00, 01, 11, 10\}$ indicate a subject's compliance type, where $C = 00$ indicates never-takers, $C = 01$ indicates compliers, and so on. Then one can write

$$\begin{aligned}\mathbb{E}(Y|D=1) = {} & \mathbb{E}(Y_1|D=1, C=01)\Pr(C=01|D=1) + \mathbb{E}(Y_0|D=1, C=00)\Pr(C=00|D=1) \\ & + \mathbb{E}(Y_1|D=1, C=11)\Pr(C=11|D=1) + \mathbb{E}(Y_0|D=1, C=10)\Pr(C=10|D=1).\end{aligned}$$

Treatment was assigned at random, thus $\mathbb{E}(Y_1|D=1, C=c) = \mathbb{E}(Y_1|C=c)$ and $\Pr(C=c|D=1) = \Pr(C=c)$, in particular $\Pr(C=01|D=1) = 1-\varepsilon$. Also using that $Y_0, Y_1 \in [0,1]$, write

$$\begin{aligned}& \mathbb{E}(Y|D=1) \in [(1-\varepsilon)\mathbb{E}(Y_1|C=01), (1-\varepsilon)\mathbb{E}(Y_1|C=01)+\varepsilon] \\ \iff & (1-\varepsilon)\mathbb{E}(Y_1|C=01) \in [\mathbb{E}(Y|D=1)-\varepsilon, \mathbb{E}(Y|D=1)].\end{aligned}$$

Similar reasoning establishes that

$$(1-\varepsilon)\mathbb{E}(Y_0|C=01) \in \left[\mathbb{E}(Y|D=0)-\varepsilon, \mathbb{E}(Y|D=0)\right].$$

Use $\Delta = \mathbb{E}(Y|D=1) - \mathbb{E}(Y|D=0)$ to conclude

$$(1-\varepsilon)\left(\mathbb{E}(Y_1|C=01) - \mathbb{E}(Y_0|C=01)\right) \in \left[\Delta-\varepsilon, \Delta+\varepsilon\right].$$

Now write

$$\begin{aligned}
\widetilde{\Delta} &= \mathbb{E}Y_1 - \mathbb{E}Y_0 \\
&= \left(\mathbb{E}(Y_1|C=01) - \mathbb{E}(Y_0|C=01)\right)\Pr(C=01) + \left(\mathbb{E}(Y_1|C\neq 01) - \mathbb{E}(Y_0|C\neq 01)\right)\Pr(C\neq 01) \\
&= \underbrace{(1-\varepsilon)\left(\mathbb{E}(Y_1|C=01) - \mathbb{E}(Y_0|C=01)\right)}_{\in[\Delta-\varepsilon,\Delta+\varepsilon]} + \varepsilon\underbrace{\left(\mathbb{E}(Y_1|C\neq 01) - \mathbb{E}(Y_0|C\neq 01)\right)}_{\in[-1,1]} \\
&\in [\Delta-2\varepsilon, \Delta+2\varepsilon].
\end{aligned}$$

Together with the trivial $\widetilde{\Delta} \in [-1,1]$, this establishes validity of the bounds. They are achieved by the scenario indicated in the text, i.e. all noncompliers are of type $Z=10$ and have individual treatment effects of either $+1$ (for the upper bound) or -1 (for the lower bound).

Example 2 The monotonicity assumption implies that $C \in \{00, 01, 11\}$. With other variables as in example 1, elementary probability calculus yields

$$\mathbb{E}(Y|D=1) = \mathbb{E}(Y_1|D=1, C\in\{01,11\})\Pr(C\in\{01,11\}|D=1) + \mathbb{E}(Y_0|D=1,C=00)\Pr(C=00|D=1),$$
$$\mathbb{E}(Y|D=0) = \mathbb{E}(Y_0|D=0, C\in\{00,01\})\Pr(C\in\{00,01\}|D=0) + \mathbb{E}(Y_1|D=0,C=11)\Pr(C=11|D=0).$$

Writing $\varepsilon_0 = \Pr(C=00)$ and $\varepsilon_1 = \Pr(C=11)$ and using independence of D and C, conclude

$$\begin{aligned}
&\Delta = \mathbb{E}(Y|D=1) - \mathbb{E}(Y|D=0) \\
&= (1-\varepsilon_0)\,\mathbb{E}(Y_1|C\in\{01,11\}) + \varepsilon_0\mathbb{E}(Y_0|C=00) - (1-\varepsilon_1)\,\mathbb{E}(Y_0|C\in\{00,01\}) - \varepsilon_1\mathbb{E}(Y_1|C=11),
\end{aligned}$$

thus (using that $Y_0, Y_1 \in [0,1]$)

$$(1-\varepsilon_0)\,\mathbb{E}(Y_1|C\in\{01,11\}) - (1-\varepsilon_1)\,\mathbb{E}(Y_0|C\in\{00,01\}) \in \left[\Delta-\varepsilon_0, \Delta+\varepsilon_1\right].$$

Yet one can also write

$$\begin{aligned}
\widetilde{\Delta} &= \mathbb{E}Y_1 - \mathbb{E}Y_0 \\
&= (1-\varepsilon_0)\,\mathbb{E}(Y_1|C\in\{01,11\}) + \varepsilon_0\mathbb{E}\left(Y_1|C=00\right) - (1-\varepsilon_1)\,\mathbb{E}(Y_0|C\in\{00,01\}) - \varepsilon_1\mathbb{E}\left(Y_0|C=11\right) \\
&= \underbrace{(1-\varepsilon_0)\,\mathbb{E}(Y_1|C\in\{01,11\}) - (1-\varepsilon_1)\,\mathbb{E}(Y_0|C\in\{00,01\})}_{\in[\Delta-\varepsilon_0,\Delta+\varepsilon_1]} + \underbrace{\varepsilon_0\mathbb{E}\left(Y_1|C=00\right) - \varepsilon_1\mathbb{E}\left(Y_0|C=11\right)}_{\in[-\varepsilon_1,\varepsilon_0]} \\
&\in [\Delta-\varepsilon_0-\varepsilon_1, \Delta+\varepsilon_0+\varepsilon_1] = [\Delta-\varepsilon, \Delta+\varepsilon].
\end{aligned}$$

Together with $\widetilde{\Delta} \in [-1,1]$, this establishes validity of the bounds; that they are best possible is again easy to verify. Excluding always-takers means to additionally constrain $(\varepsilon_0, \varepsilon_1) = (\varepsilon, 0)$, which does not affect this result.

Example 3 Write

$$\widetilde{\Delta} = \mathbb{E}(Y_1 - Y_0) = (1-\varepsilon)\mathbb{E}(Y_1 - Y_0|Z=1) + \varepsilon\mathbb{E}(Y_1 - Y_0|Z=0) \in [(1-\varepsilon)\Delta - \varepsilon, (1-\varepsilon)\Delta + \varepsilon].$$

Example 4 Consider the problem of maximizing the true value of $\mathbb{E}Y_0$ conditional on $\mathbb{E}(Y_0|Z=1)$ being fixed at its observed value μ_0. Formally, solve

$$\max \int \mu_{0x} dP(X) \quad \text{s.t.} \quad \frac{\int \mu_{0x} \exp(\alpha + \beta x)/(1+\exp(\alpha+\beta x))dP(X)}{\int \exp(\alpha+\beta x)/(1+\exp(\alpha+\beta x))dP(X)} = \mu_0,$$

where $\mu_{0x} = \mathbb{E}(Y_0|X=x)$ as in the section on covariates and where the function μ_{0x} as well as $(\alpha, \beta, P(X))$ are choice parameters. Clearly the solution will have $X \in \{0,1\}$ a.s. and $\mu_{0x} = 1\{X=0\}$. This means that $\mathbb{E}(Y_0) = \Pr(\mu_{0x} = 1) = \Pr(X=0)$ and $\mathbb{E}(Y_0|Z=1) = \Pr(\mu_{0x} = 1|Z=1) = \Pr(X = 0|Z=1)$, where probabilities are taken w.r.t. the distribution of (Z, X). Now,

$$\begin{aligned}\frac{\Pr(X=1|Z=1)}{\Pr(X=0|Z=1)} &= \frac{\Pr(Z=1|X=1)\Pr(X=1)}{\Pr(Z=1|X=0)\Pr(X=0)} \\ &\leq \frac{\exp(\alpha+\beta)(1+\exp(\alpha))}{(1+\exp(\alpha+\beta))\exp(\alpha)}\frac{\Pr(X=1)}{\Pr(X=0)} < \exp(\beta)\frac{\Pr(X=1)}{\Pr(X=0)} \leq \exp(c)\frac{\Pr(X=1)}{\Pr(X=0)},\end{aligned}$$

where this bound is not attainable but tight in the sense that Nature can make both weak inequalities bind and can make the strict one arbitrarily close to binding by driving $\alpha \to -\infty$. Now some mild algebra reveals that the problem's value can be bounded above by

$$\mathbb{E}(Y_0) < \frac{\exp(c)\mathbb{E}(Y_0|Z=1)}{1 + (\exp(c)-1)\mathbb{E}(Y_0|Z=1)}$$

with the same remark about tightness. This bound is used to bound $\widetilde{\Delta}$ from below in the example. The upper bound on $\mathbb{E}Y_0$ and the bounds on $\mathbb{E}Y_1$ can be derived similarly. Finally, the upper [lower] bound on $\mathbb{E}Y_1$ and the lower [upper] bound on $\mathbb{E}Y_0$ can be attained (subject to the remark about tightness) by choosing the same $(\alpha, \beta, P(X))$ and, therefore, in the same state of the world.

Proposition 7

(i) Let N be odd; the extension to N even is as before. The result does not extend to matched pairs or constrained randomization.

Step 1: Describing the equilibrium. Every state s implies expected values of latent outcomes in the sampling universe (μ_0, μ_1) as well as true expected potential outcomes $(\widetilde{\mu}_0, \widetilde{\mu}_1)$, where $\widetilde{\mu}_1 - \widetilde{\mu}_0 \in [a(\mu_1 - \mu_0) - b, a(\mu_1 - \mu_0) + b]$ in accordance with (2). It will be seen that regret depends on the state of the world only through $(\mu_0, \mu_1, \widetilde{\mu}_0, \widetilde{\mu}_1)$, so states will be identified with these parameters. The least favorable prior π^* randomizes evenly over

$$\left(\frac{1+\Delta^*}{2}, \frac{1-\Delta^*}{2}, \min\left\{\frac{1+a\Delta^*+b}{2}, 1\right\}, \max\left\{\frac{1-a\Delta^*-b}{2}, 0\right\}\right)$$
$$, \left(\frac{1-\Delta^*}{2}, \frac{1+\Delta^*}{2}, \max\left\{\frac{1-a\Delta^*-b}{2}, 0\right\}, \min\left\{\frac{1+a\Delta^*+b}{2}, 1\right\}\right),$$

where the implied value of $\Delta^* = \mu_1^* - \mu_0^*$ will be characterized later. The implied ATE $\widetilde{\Delta}^*$ achieves its upper bound in one state and its lower bound in the other.

Depending on parameter values, exactly one of two types of equilibrium obtains. In the "informative equilibrium," one has $\Delta^* > 0$, and δ^B is a best response by minimal adaptation of the according argument in proposition 1(i). In the "noninformative equilibrium," one has $\Delta^* = 0$, thus the sample data are noise, and any treatment rule – including $\widetilde{\delta}^B$ – is Bayes against π^*.

Step 2: Equilibrium when $b \geq 1$. With $b \geq 1$, Nature sets $\Delta^* = 0$ and $\widetilde{\Delta}^* = 1$. Thus the sample data are noise, and any decision rule incurs expected regret of 1/2. At the same time, the no-data rule $\delta^{no-data} \equiv 1/2$ incurs maximal expected regret of 1/2 (a proof of this is just as in proposition 3(iv)). Hence, an equilibrium has been discovered. For the remainder of the proof, assume $b < 1$.

Step 3: Simplifying Nature's best-response problem. Nature's response must be supported on states s that maximize

$$R(\widetilde{\delta}^B, s) = \max\left\{(\widetilde{\mu}_1 - \widetilde{\mu}_0)(1 - \mathbb{E}\widetilde{\delta}^B(\omega^B)), (\widetilde{\mu}_0 - \widetilde{\mu}_1)\mathbb{E}\widetilde{\delta}^B(\omega^B)\right\}.$$

Note that $\widetilde{\delta}^B$ depends on ω^B only through (I_{N-2}, I_N). Recalling that $(N + I_N/N)/2$ is distributed as a binomial variable with parameters $(N, (1 + \mu_1 - \mu_0)/2))$, it follows that $\mathbb{E}\widetilde{\delta}^B(\omega^B)$ depends on (μ_0, μ_1) only through Δ. Furthermore, a state s solves Nature's best-response problem if its implied values of $(\Delta, \widetilde{\Delta})$ solve

$$\max_{\substack{\Delta \in [-1,1] \\ \widetilde{\Delta} \in [a\Delta - b, a\Delta + b] \cap [-1,1]}} \max\left\{\widetilde{\Delta}(1 - f(\Delta)), -\widetilde{\Delta} f(\Delta)\right\},$$

where $f(\Delta) = \mathbb{E}\widetilde{\delta}^B(\omega^B)$ as before. As f increases in Δ and $\widetilde{\Delta}$ enters the objective linearly, this problem can be simplified to

$$\max_{\Delta \in [-1,1]} \max\left\{\min\{a\Delta + b, 1\}(1 - f(\Delta)), -\max\{a\Delta - b, -1\} f(\Delta)\right\}.$$

For a further simplification, note that $f(\Delta) = 1 - f(-\Delta)$ by the construction of $\widetilde{\delta}^B$, thus

$$\min\{a(-\Delta) + b, 1\}(1 - f(-\Delta)) = \min\{-a\Delta + b, 1\} f(\Delta) = -\max\{a\Delta - b, -1\} f(\Delta).$$

Hence, the problem is solved by some Δ^* iff it is solved by $-\Delta^*$ as well. With this in mind, there always exists a best response of form π^*, where the number Δ^* must solve

$$\max_{\Delta \in [-1,1]} \{\min\{a\Delta + b, 1\}(1 - f(\Delta))\}$$

and $\widetilde{\Delta}^* = \min\{a\Delta^* + b, 1\}$. Finally, this objective is negative whenever $\Delta < -b/a$ and decreases in Δ whenever $a\Delta + b > 1 \Leftrightarrow \Delta > (1-b)/a$, thus the above expression can be further simplified to

$$\max_{\Delta \in [-b/a,(1-b)/a] \cap [-1,1]} \phi(\Delta), \phi(\Delta) = (a\Delta + b)\,(1 - f(\Delta)), \tag{15}$$

the problem that will be considered henceforth. Note that as $a > 0$ and $b < 1$ are assumed, the feasible set for Δ (which will be suppressed below for brevity) always contains 0.

Step 4: Equilibrium when $a \leq b$. In this case, $\widetilde{\delta}^B(\omega^B) = \frac{a}{b} \cdot \delta_1^B(\omega^B) + \left(1 - \frac{a}{b}\right) \cdot \frac{1}{2}$, thus $f(\Delta) = \frac{1}{2} + \frac{a}{2b}\Delta$, and (15) turns into

$$\max_{\Delta} (a\Delta + b)\left(\frac{1}{2} - \frac{a}{2b}\Delta\right) = \max_{\Delta} \frac{1}{2}\left(b - \frac{a^2}{b}\Delta^2\right),$$

which is maximized by $\Delta^* = 0$, confirming the noninformative equilibrium.

Step 5: Uniqueness of Δ^*. Let now $a > b$. This step shows that (15) is solved by a unique Δ^*; furthermore, this Δ^* has the same sign as $\phi'(0)$ and in particular equals zero iff $\phi'(0) = 0$. For $N = 1$, direct evaluation of derivatives reveals that ϕ is concave. For $N \geq 3$, write

$$\phi''(\Delta) = -2af'(\Delta) - (a\Delta + b)f''(\Delta) \tag{16}$$

and (recalling again that $F(n; N, p) = (N - n)\binom{N}{n}\int_0^{1-p} t^{N-n-1}(1-t)^n dt$)

$$\begin{aligned}
f(\Delta) &= 1 - \alpha \underbrace{\frac{N+1}{2}\binom{N}{\frac{N-1}{2}}}_{\varphi} \int_0^{\frac{1-\Delta}{2}} t^{\frac{N-1}{2}}(1-t)^{\frac{N-1}{2}} dt - (1-\alpha)\underbrace{\frac{N-1}{2}\binom{N-2}{\frac{N-3}{2}}}_{\psi} \int_0^{\frac{1-\Delta}{2}} t^{\frac{N-3}{2}}(1-t)^{\frac{N-3}{2}} dt \\
f'(\Delta) &= \alpha\varphi 2^{-N}(1-\Delta^2)^{\frac{N-1}{2}} + (1-\alpha)\psi 2^{2-N}(1-\Delta^2)^{\frac{N-3}{2}} \\
f''(\Delta) &= -\alpha\varphi 2^{-N}(N-1)\Delta(1-\Delta^2)^{\frac{N-3}{2}} - (1-\alpha)\psi 2^{2-N}(N-3)\Delta(1-\Delta^2)^{\frac{N-5}{2}}.
\end{aligned}$$

Some algebra yields

$$\phi''(\Delta) = 2^{-N}(1-\Delta^2)^{\frac{N-5}{2}} \underbrace{\begin{bmatrix} -2a\alpha\varphi(1-\Delta^2)^2 - 8a(1-\alpha)\psi(1-\Delta^2) + \\ \alpha\varphi(N-1)\Delta(1-\Delta^2)\,(a\Delta+b) + 4(1-\alpha)\psi(N-3)\Delta\,(a\Delta+b) \end{bmatrix}}_{\gamma(\Delta)}.$$

On $(-b/a, 1)$, the sign of this expression equals the sign of γ. Using $a\Delta + b \geq 0$, γ is easily seen to be negative on $[-b/a, 0]$. Furthermore, γ is a polynomial of fourth degree in Δ with the following properties (using $N \geq 3$ and $a\Delta + b \geq 0$): The coefficient on Δ^4 is negative, thus γ is concave as $\Delta \to \pm\infty$. The coefficient on Δ^2 is positive, thus γ is convex at $\Delta = 0$. Thus, γ is first convex then concave on $[0, \infty)$. But is it also easily verified that $\gamma(0) < 0$ and $\gamma(1) \geq 0$, which is consistent with γ being first convex then concave on $[0, \infty)$ only if γ is first negative then nonnegative on $[0, 1]$. It follows that ϕ is first concave then convex on $[-b/a, 1]$. As also $\phi(1) = 0$, ϕ is quasiconcave on the relevant range of Δ. As ϕ is differentiable at 0, the claim follows.

Step 6: Developing the informative equilibrium. Initially restrict attention to decision rules δ^* with $\alpha^* = 1$. Evaluation of the combinatorial terms labeled (φ, ψ) in step 5 reveals that among these rules, $f'(0)$ strictly increases in N^*, hence

$$\phi'(0) = \frac{a}{2} - bf'(0)$$

strictly decreases in N^*. Since α^* interpolates between neighboring N^*, $\phi'(0)$ can be thought of as a continuously decreasing function of $((N^* - 1)/2 + \alpha^*)$. Hence, if $\phi'(0) > 0$ for $(N^*, \alpha^*) = (1, 1)$, then $\phi'(0) > 0$ for $((N^* - 1)/2 + \alpha^*)$ below some threshold, $\phi'(0) = 0$ for $((N^* - 1)/2 + \alpha^*)$ on the threshold, and $\phi'(0) < 0$ for $((N^* - 1)/2 + \alpha^*)$ above the threshold. Let N^* correspond to the threshold. If $N < N^*$, then existence of an informative equilibrium follows: The planner plays δ^B, Nature's best response is characterized by $\Delta^* \geq 0$, hence equilibrium obtains. If $N \geq N^*$, existence of a noninformative equilibrium follows: The planner plays $\alpha^* \delta^B_{N^*} + (1 - \alpha^*)\delta^B_{N^*-2}$, and Nature's best response is characterized by $\Delta^* = 0$. It remains to verify whether $\phi'(0) > 0$ for $(N^*, \alpha^*) = (1, 1)$. Direct evaluation (namely, $f'(0) = 1/2$ if $N = 1$) shows that this is the case iff $a > b$.

(ii) Follows as in proposition 1.

(iii) The equilibrium supporting this minimax regret treatment rule is very similar to the one from part (i). In particular, π^* is the same except that Δ^* solves

$$\Delta^* = \begin{cases} \arg\max_\Delta (a\Delta + b)\, \Phi(-\Delta; 0, \sigma^2) & \text{in the informative equilibrium,} \\ \arg\max_\Delta (a\Delta + b)\, \Phi\left(-\Delta; 0, \frac{b^2}{a^2}\frac{2}{\pi}\right) & \text{in the noninformative equilibrium.} \end{cases} \tag{17}$$

Steps 1 through 4 of the proof go through essentially unchanged, with the observations that $\widetilde{\delta}^G$ is Bayes given π^* and that $\widetilde{\delta}^G$ depends on s only through $\mu_1 - \mu_0$ becoming easier. Step 5 is unchanged up to (16). To continue from there, use (14) and the expressions immediately after it in the text to write

$$\phi''(\Delta) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{\Delta^2}{2\sigma^2}\right) \cdot \left(\frac{a}{\sigma^2}\Delta^2 + \frac{b}{\sigma^2}\Delta - 2a\right).$$

Using standard formulae for quadratic equations, this expression is negative on

$$\left(-\frac{b}{2a}-\sqrt{\frac{b^2}{4a^2}+2\sigma^2},-\frac{b}{2a}+\sqrt{\frac{b^2}{4a^2}+2\sigma^2}\right),$$

an interval that contains $[-b/a,0]$, and nonnegative otherwise. Step 5 can then be completed as before.

Regarding step 6, $f'(0)=(2\pi)^{-1/2}\sigma^{-1}$ continuously decreases in σ and approaches 0 as $\sigma\to\infty$. Thus, if $\phi'(0)=\frac{a}{2}-bf'(0)<0$, then the first-order condition $\phi'(0)=0$ can be achieved by increasing σ^2. Specifically, algebra shows that

$$\phi'(0)\geq 0 \Longleftrightarrow \sigma^2\geq \frac{b^2}{a^2}\frac{2}{\pi}.$$

If σ^2 is below this threshold, the decision maker can increase it artificially by replacing ω with $\omega+\xi$, where $\xi\sim N\left(0,\frac{b^2}{a^2}\frac{2}{\pi}-\sigma^2\right)$ and then applying δ^G. Simple algebra reveals that this is equivalent to the rule given in the proposition.

Corollary 8 Follows by evaluating (15) respectively (17).

Proposition 9 This is very similar to proposition 7, and I only point out necessary adaptations to part (i).

Step 1: Describing the equilibrium. As before, except that the algebraic expression for the bounds on $\widetilde{\Delta}$ in terms of Δ changed.

Step 2: No analog to this step is needed.

Step 3: Simplifying Nature's best-response problem. $R(\widetilde{\delta}^B,s)$ depends on s only through

$$(\mu_0,\mu_1,\widetilde{\mu}_0,\widetilde{\mu}_1)=(\mathbb{E}(Y_0|Z=1),\mathbb{E}(Y_1|Z=1),\mathbb{E}(Y_0),\mathbb{E}(Y_1)),$$

and Nature's best-response problem boils down to

$$\begin{aligned}
& \max_{(\mu_0,\mu_1,\widetilde{\mu}_0,\widetilde{\mu}_1)\in[0,1]^4} \max\left\{(\widetilde{\mu}_1-\widetilde{\mu}_0)(1-f(\Delta)),(\widetilde{\mu}_0-\widetilde{\mu}_1)f(\Delta)\right\} \\
\text{s.t.}\quad & \frac{\mu_0}{\exp(c)+(1-\exp(c))\mu_0}\leq\widetilde{\mu}_0\leq\frac{\exp(c)\mu_0}{1+(\exp(c)-1)\mu_0} \\
& \frac{\mu_1}{\exp(c)+(1-\exp(c))\mu_1}\leq\widetilde{\mu}_1\leq\frac{\exp(c)\mu_1}{1+(\exp(c)-1)\mu_1}.
\end{aligned}$$

Inspection of the way in which $(\widetilde{\mu}_0,\widetilde{\mu}_1)$ enter the objective reveals immediate simplification to

$$\max_{(\mu_0,\mu_1)\in[0,1]^2}\max\left\{\begin{array}{c}\left(\frac{\exp(c)\mu_1}{1+(\exp(c)-1)\mu_1}-\frac{\mu_0}{\exp(c)+(1-\exp(c))\mu_0}\right)(1-f(\Delta)),\\ \left(\frac{\exp(c)\mu_0}{1+(\exp(c)-1)\mu_0}-\frac{\mu_1}{\exp(c)+(1-\exp(c))\mu_1}\right)f(\Delta)\end{array}\right\}.$$

This objective does *not* depend on (μ_0, μ_1) only through Δ. However, evaluation of first and second derivatives reveals that the problem

$$\max_{\substack{(\mu_0,\mu_1)\in[0,1]^2: \\ \mu_1=\mu_0+\Delta}} \left\{ \frac{\exp(c)\mu_1}{1+(\exp(c)-1)\mu_1} - \frac{\mu_0}{\exp(c)+(1-\exp(c))\mu_0} \right\}$$

is solved by $\mu_0 = (1-\Delta)/2 = 1-\mu_1$, and similarly for $\left(\frac{\exp(c)\mu_0}{1+(\exp(c)-1)\mu_0} - \frac{\mu_1}{\exp(c)+(1-\exp(c))\mu_1}\right)$. It follows that a best response for Nature must have $\mu_0 = 1-\mu_1$, in which case the maximization problem can be further simplified to

$$\max_{\Delta\in[-1,1]} \max \left\{ \underbrace{\frac{(1+\Delta)\exp(c)-(1-\Delta)}{(1+\Delta)\exp(c)+(1-\Delta)}}_{\rho(\Delta)} (1-f(\Delta)), \frac{1+\Delta-(1-\Delta)\exp(c)}{1+\Delta+(1-\Delta)\exp(c)} f(\Delta) \right\}.$$

This objective exhibits the same symmetry as before, thus a least favorable prior can be constructed as before. Also taking into account that $\rho(\Delta) < 0$ if $\Delta < (1-\exp c)/(1+\exp c)$, whereas $\rho(\Delta) < 1$ for any $\Delta < 1$, Δ^* must solve

$$\max_{\Delta\in[(1-\exp c)/(1+\exp c),1]} \phi(\Delta), \phi(\Delta) = \rho(\Delta)(1-f(\Delta)), \tag{18}$$

the problem that will be analyzed henceforth.

Step 4: Equilibrium when $c \geq \ln\left(2+\sqrt{5}\right)$. In this case, $f(\Delta) = \frac{1}{2} - \frac{2\exp(c)}{\exp(2c)-1}\Delta$, thus (18) simplifies to

$$\max_{\Delta} \rho(\Delta)\left(\frac{1}{2} - \frac{2\exp(c)}{\exp(2c)-1}\Delta\right).$$

By evaluating derivatives, this can be verified to be solved by $\Delta^* = 0$ as required. (Readers who wish to verify algebra should take note that $\rho'(\Delta) = 4\exp(c)/((1+\Delta)\exp(c)+(1-\Delta))^2$.)

Step 5: Uniqueness of Δ^*. The argument for quasiconcavity of ϕ becomes more complex, the remainder is as before. For $N = 1$, one can verify concavity of ϕ by evaluating derivatives. For $N \geq 3$, observe first that ρ is increasing and concave, whereas algebra from the preceding proof revealed that $1-f(\Delta)$ is decreasing and concave, over $\Delta \in [-1,0]$. These facts jointly imply that ϕ is concave on $[(1-\exp c)/(1+\exp c), 0]$. It is also easily verified that $\phi(0) = \frac{\exp c - 1}{2(\exp c+1)} > 0$ and that $\phi(1) = 0$. It follows that if $\phi'(0) \geq 0$, then $\phi(\Delta)$ has a critical point Δ_0 on $[0,1]$. Let $k = \rho(\Delta_0)$ and $l = \rho'(\Delta_0)$, then $\rho_0(\Delta) = \rho(\Delta_0) + \rho'(\Delta_0)(\Delta-\Delta_0)$ is the tangent to ρ at Δ_0. As ρ is concave with positive intercept, ρ_0 is positive affine, thus $\phi(\Delta) \leq \rho_0(\Delta)(1-f(\Delta))$. But with ρ_0 being positive affine, algebra from proposition 7, step 5, can be entirely mimicked to show that $\rho_0(\Delta)(1-f(\Delta))$ is quasiconcave on $[0,1]$, is positive and increasing at 0, and equals 0 at 1. It follows that $\rho_0(1-f)$ has at most one maximum on $[0,1]$, and any such maximum obtains iff a first-order condition holds. But $\frac{d}{d\Delta}\rho_0(\Delta_0)(1-f(\Delta_0)) =$

$\rho_0(\Delta_0)(-f'(\Delta_0)) + \rho_0'(\Delta_0)(1-f(\Delta_0)) = \rho(\Delta_0)(-f'(\Delta_0)) + \rho'(\Delta_0)(1-f(\Delta_0)) = 0$, so the maximum occurs at Δ_0. Since also $\rho_0(\Delta_0)(1-f(\Delta_0)) = \rho(\Delta_0)(1-f(\Delta_0))$, it follows that $\phi(\Delta)$ is uniquely maximized at Δ_0.

Step 6: Developing the informative equilibrium. This is essentially as before. The crucial condition for $\phi'(0) \geq 0$ at $(N^*, \alpha^*) = (1,1)$ now is

$$\rho'(0) - \rho(0) \geq 0,$$

which can be verified to obtain iff $c \leq \ln\left(2+\sqrt{5}\right)$.

Corollary 10 Follows by evaluating (18). or an analogous expression for the Gaussian case.

References

[1] Angrist, J.D., G.W. Imbens, and D.B. Rubin, 1996. Identification of Causal Effects Using Instrumental Variables. *Journal of the American Statistical Association*, 91: 444-455.

[2] Balke, A. and J. Pearl, 1997. Bounds on Treatment Effects from Studies with Imperfect Compliance. *Journal of the American Statistical Association*, 92, 1172–6.

[3] Bergemann, D. and K.H. Schlag, 2007. Robust Monopoly Pricing. Cowles Foundation Discussion Paper 1527R.

[4] —, 2008. Pricing Without Priors. *Journal of the European Economic Association (Papers and Proceedings)*, 6, 560-569.

[5] Brock, W.A., 2006. Profiling Problems with Partially Identified Structure. *Economic Journal*, 116, F427-F440.

[6] Brock, W.A., and C.F. Manski, 2011. Competitive Lending with Partial Knowledge of Loan Repayment: Some Positive and Normagive Analysis. *Journal of Money, Credit, and Banking*, 43, 441-459.

[7] Canner, P.L., 1970. Selecting One of Two Treatments when the Responses are Dichotomous. *Journal of the American Statistical Association*, 65, 293-306.

[8] Chamberlain, G., 2000. Econometrics and Decision Theory. *Journal of Econometrics*, 95, 255-283.

[9] —, 2011. Bayesian Aspects of Treatment Choice. In J.Geweke, G. Koop, and H. van Dijk (eds.), *The Oxford Handbook of Bayesian Econometrics*. Oxford University Press.

[10] 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.

[11] Dehejia, R., 2005. Program Evaluation as a Decision Problem. *Journal of Econometrics*, 125, 141-173.

[12] Eozenou, P., J. Rivas, and K.H. Schlag, 2006. Minimax Regret in Practice – Four Examples on Treatment Choice. Discussion paper, European University Institute.

[13] Gupta, S.S., and S.N. Hande, 1992. On Some Nonparametric Selection Procedures. In A.K.Md.E.Saleh (ed.), *Nonparametric Statistics and Related Topics*. Elsevier.

[14] Hill, S., and B. Kreider, 2009. Partially Identifying Treatment Effects with an Application to Covering the Uninsured. *Journal of Human Resources*, 44, 409-449.

[15] Hirano, K. and J.R. Porter, 2009. Asymptotics for Statistical Treatment Rules. *Econometrica*, 77, 1683-1701.

[16] Manski, C.F., 1989. Anatomy of the Selection Problem. *Journal of Human Resources*, 24, 343-360.

[17] —, 2000. Identification Problems and Decisions under Ambiguity: Empirical Analysis of Treatment Response and Normative Analysis of Treatment Choice. *Journal of Econometrics*, 95, 415-442.

[18] —, 2003. *Partial Identification of Probability Distributions*. Springer Verlag.

[19] —, 2004. Statistical Treatment Rules for Heterogeneous Populations. *Econometrica*, 72: 1221-1246.

[20] —, 2006. Search Profiling with Partial Knowledge of Deterrence. *Economic Journal*, 116, F385-F401.

[21] —, 2007a. Minimax-Regret Treatment Choice with Missing Outcome Data. *Journal of Econometrics*, 139, 105-115.

[22] —, 2007b. *Identification for Prediction and Decision*. Harvard University Press.

[23] —, 2009. Diversified Treatment under Ambiguity. *International Economic Review*, 50, 1013-1041.

[24] —, 2010. Identification of Treatment Effects with Social Interactions. Discussion paper, Northwestern University.

[25] Manski, C.F. 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.

[26] Parmigiani, G., 1992. Minimax, Information, and Ultrapessimism. *Theory and Decision*, 33, 241-252.

[27] Savage, L.J., 1951. The Theory of Statistical Decisions. *Journal of the American Statistical Association*, 46, 55-67.

[28] —, 1954. *The Foundations of Statistics*. Wiley.

[29] Schlag, K.H., 2003. How to Minimize Maximum Regret under Repeated Decision-Making. Discussion paper, European University Institute.

[30] —, 2006. Eleven. Discussion paper, European University Institute.

[31] Stoye, J., 2007. Minimax Regret Treatment Choice with Incomplete Data and Many Treatments. *Econometric Theory*, 23, 190-199.

[32] —, 2009a. Minimax Regret Treatment Choice with Finite Samples. *Journal of Econometrics* 151: 70-81.

[33] —, 2009b. Minimax Regret Treatment Choice with Missing Data: An Application to Young Offenders. *Journal of Statistical Theory and Practice*, 3, 239-254.

[34] —, 2009c. Minimax Regret. In S.N. Durlauf and L.E. Blume (eds.), *The New Palgrave Dictionary of Economics*.

[35] Tetenov, A., 2009a. Statistical Treatment Choice Based on Asymmetric Minimax Regret Criteria. Discussion paper, Collegio Carlo Alberto.

[36] —, 2009b. Measuring Precision of Statistical Inference on Partially Identified Parameters. Discussion paper, Collegio Carlo Alberto.

[37] Vytlacil, E., 2002. Independence, Monotonicity, and Latent Index Models: An Equivalence Result. *Econometrica*, 70, 331-341.

[38] Wald, A., 1945. Statistical Decision Functions Which Minimize the Maximum Risk. *Annals of Mathematics*, 46, 265-280.

[39] —, 1950. *Statistical Decision Functions*. Wiley.

[40] Worrall, J., 2007. Evidence in Medicine and Evidence-Based Medicine. *Philosophical Compass* 2: 981-1022.