Minimax Regret Treatment Choice with Finite Samples

Jörg Stoye*

New York University

October 3, 2006

Abstract

I use the minimax regret criterion to analyze choice between two treatments conditional on observation of a finite sample. The analysis is based on exact finite-sample regret and does not use asymptotic approximations nor finite-sample bounds. Core results are the following: (i) Minimax regret treatment rules are well approximated by empirical success rules in many cases, but differ from them significantly for small sample sizes and certain sample designs. (ii) Without imposing additional prior information, they prescribe inference that is completely separate across covariates. This result can be avoided by imposing sufficient prior information. (iii) The relative performance of empirical success rules can be evaluated and is significantly lacking in very small samples. (iv) Manski's (2004) analysis of optimal sample stratification can be reproduced, with somewhat different implications, in terms of finite sample results as opposed to large deviations bounds. I conclude by offering some methodological thoughts on minimax regret.

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

*I thank Gary Chamberlain, Chuck Manski, Karl Schlag, and seminar participants at Cornell, Harvard/MIT, Mannheim, Northwestern, NYU, and Pittsburgh for helpful comments. Financial support through the Robert Eisner Memorial Fellowship, Department of Economics, Northwestern University, as well as a Dissertation Year Fellowship, The Graduate School, Northwestern University, is gratefully acknowledged. Address: Jörg Stoye, Department of Economics, New York University, 269 Mercer Street 3FL, New York, NY 10003, j.stoye@nyu.edu.

1 Introduction

In this paper, the minimax regret criterion is used to analyze choice between two treatments based on a sample of subjects that have been subjected to one treatment each. This setting is similar to the one recently analyzed by Manski (2004). The main difference to Manski's approach is technical: I consider some extensions of the setting and, more importantly, base the analysis entirely on exact finite sample regret. As it turns out, moving from from finite-sample bounds to exact results leads to significant adjustments of substantive results.

Minimax regret as a criterion for statistical decisions has recently attracted renewed interest (Brock 2004, Hirano and Porter 2005, Manski 2004, 2005, 2006a, 2006b, Schlag 2006, Stoye 2006a). Unfortunately, derivation of finite sample minimax regret decision rules tends to be extremely hard. As a result, most of the existing literature either focuses on identification and altogether abstracts from sampling uncertainty (Brock 2004, Manski 2006b, Stoye 2006), states the finite-sample problem without attempting to solve it (Manski 2006a, section 4), derives bounds on finite-sample regret (Manski 2004), or employs large-sample approximations (Hirano and Porter 2005, Manski 2006a). To my knowledge, the only exact results for finite samples so far are found in a related paper by Schlag (2006), in Manski's (2006a, section 5) analysis of a case that he calls "curiously simple," and in his numerical analysis of the setup considered in proposition 1(iii) below (2005, chapter 3).

One important agenda of this paper is, therefore, to show that much can be learned from exact finite-sample analysis. On a substantive level, perhaps the most interesting finding is that some conclusions here refine findings in Manski (2004) in ways that might be considered surprising, or even controversial. The results also allow one to assess, for the decision situations considered here, the small sample performance of asymptotic approximations that have been proposed by Hirano and Porter (2005).

The paper is structured as follows. After setting up the notation and offering some motivation for minimax regret, I analyze the treatment choice problems without covariates, differentiating the anmalysis depending on whether one or both treatments are unknown, and in the latter case, how treatments were assigned to sample subjects. In some cases, the minimax regret rules are similar to empirical success rules, i.e. simple comparisons of sample means, although significant differences are uncovered as well. Minimax regret decision rules are generally quite different from ones informed by classical statistics.

The analysis is then extended to the situation where treatment outcomes may depend on a covariate x. This is a central concern in Manski (2004), and the core result here is maybe the most surprising one: Minimax regret completely separates inference between covariates for any sample size. In section 4 of the paper, I derive this result, comment on it, and suggest a partial solution to the problems it raises. In section 5, theoretical results are used to evaluate the finite sample performance of simple empirical success rules and to replicate parts of Manski's (2004) analysis but in terms of exact regret. Section 6 concludes with reflections on some interesting features of the results. All proofs are collected in the appendix.

2 Preliminaries and Motivation

A binary treatment, $t \in \{0,1\}$ induces distributions of potential outcomes $Y_t \in \mathbb{R}$. (Y_t may have a utility interpretation.) Assume that a priori bounds on Y_0 and Y_1 are finite and coincide, then it is w.l.o.g. to set $Y_0, Y_1 \in [0,1]$. A state of nature s specifies the distributions of potential outcomes and can, therefore, be identified with a probability measure $P(Y_0, Y_1) \in \Delta[0,1]^2$. The set of possible states of nature $\mathcal{S}$ is a (not necessarily proper) subset of $\Delta[0,1]^2$ that depends on maintained assumptions. It is worth noting that Y_0 and Y_1 are never restricted to be independent. I will use the following notational conventions: If Y_i is a random variable, then μ_i denotes its expectation and $\overline{y}_i$ a sample mean.

For any sample point, one treatment t is assigned and the according outcome is recorded. Thus, the decision-maker observes realizations of (t, Y_t). Let $\mathcal{S}_N$, with typical element s_N, denote the sample space induced by a sample of size N. (This implies a slight abuse of notation since I allow for N to be random.) The sample distribution over $\mathcal{S}_N$ then depends on s, on the sampling rule, and on the within-sample treatment assignment rule; both of the latter will differ across the scenarios considered below. The decision maker's strategy set is given by the set of statistical treatment rules $\mathcal{D}$ with typical element $\delta : \mathcal{S}_N \to [0,1]$, where δ maps possible sample outcomes into probabilities of assigning treatment 1. In particular, the decision maker is allowed to randomize.[1]

Any combination of state and decision rule induces an expected outcome, namely

$$u(\delta, s) \equiv E\left(\delta(s_N)\right)\mu_1 + \left[1 - E\left(\delta(s_N)\right)\right]\mu_0, \tag{1}$$

where all expectations are taken conditional on s. The efficacy of δ will be measured in terms not of u, but of expected regret relative to u, that is,

$$R(\delta, s) \equiv \max_{\delta' \in \mathcal{D}} \left\{u(\delta', s)\right\} - u(\delta, s),$$

[1] The decision rule may appear underspecified because I leave open whether randomization is independent across members of the treatment population, perfectly correlated (i.e., the entire population is assigned to treatment 1 with probability $\delta(s_N)$ and to treatment 0 otherwise), or anything in between. It is immediate from the next pargraph's definitions that only the expected fraction of the population assigned to treatment 1 matters for regret, hence all of these possibilities are equivalent.

the expected loss relative to the ex-post utility frontier. In fact, $R(\delta, s)$ is interpreted as risk function, and the decision maker aims to achieve uniform performance with respect to this risk, thus she wants to implement a minimax regret treatment rule

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

Although this paper is focused on the implications of minimax regret, I will briefly comment on its a priori motivation here.[2] The function $u : \mathcal{S} \mapsto \mathbb{R}$ defined in (1) is the usual risk function of decision rule δ, and the problem of ranking such risk functions is well known in statistical decision theory. Notice in particular that given a state s, the outcomes experienced by treatment recipients are uncertain, but this uncertainty is described by objective probabilities, and the definition of (1) implies a decision to resolve it by taking expectations. However, no objective probabilities of states s are given, i.e. the true state is ambiguous, and in general, many treatment rules will be admissible. To arrive at well defined policy recommendations, one must, therefore, commit to an attitude toward the ambiguity about s.

Two well-known such attitudes are embodied in the Bayesian and maximin utility approach. The former aggregates $u(\delta, s)$ by means of a subjective prior $\Pi \in \Delta\mathcal{S}$, the latter ranks decision rules according to their worst-case performance $\min_{s \in \mathcal{S}} u(\delta, s)$. For the purpose of this paper, and without meaning to take a general stance pro or contra Bayesianism, I will assume that imposition of priors is not desired. This leaves maximin utility as the best-known decision rule.

Maximin utility is frequently perceived as extremely pessimistic. In particular, it may be counter-intuitive to focus one's attention on the state that is objectively worst, even if one's choice of action matters hardly or not at all in that state. This consideration is aggravated by the fact that in the scenarios considered here, such a focus leads to trivial results. Applied to the case with Y_0 and Y_1 unknown, maximin utility induces indifference between all treatment rules because they all have the same worst-case performance (induced by the state of the world where Y_0 and Y_1 are identically equal to zero). If the distribution of Y_0 is known, then the maximin utility strategy is to always pick Y_0 if $\mu_0 > 0$, which is more determinate but hardly reasonable. Notice especially that in both cases, maximin utility fails to exclude "no-data rules" which do not take sample informatuon into account.

In contrast, minimax regret focuses on states in which one's decision matters. Since these will be states in which the distributions of Y_0 and Y_1 differ, one would expect no-data rules to be avoided, an expectation that will be borne out, albeit with caveats. Consequently, all examples but one in this paper show the minimax regret rule to be consistent in the sense of asymptotically choosing the correct

[2] In-depth discussions of the following issues, as well as axiomatic characterizations of different decision rules, are found in Stoye (2005); see also Hayashi (2005), Stoye (2006b), and the classic paper by Milnor (1954). For statisticians' discussions, compare Savage (1954) and Berger (1985).

treatment almost surely. This feature is obviously not shared by no-data rules, and hence not by the maximin utility approach.

To be sure, these observations do not imply that the (re-)discovery of minimax regret solves all worries of statistical decision theory. For example, the axiomatic analysis in Stoye (2005) reveals a trade-off: Minimax regret, unlike maximin utility, is consistent with the independence axiom even though it avoids priors. Thus, it is insensitive to risk that can be considered extraneous to a decision problem. Yet minimax regret, unlike maximin utility, violates independence of irrelevant alternatives and thus renders the ordering of two decision rules potentially sensitive to the "menu" of decision rules. Furthermore, one of this paper's contributions is the construction of realistic examples in which minimax regret may be counterintuitive, as well as one case in which it permits no-data rules. My claim, therefore, is not that minimax regret is the only good decision rule, but merely that it is worthy of further investigation – especially since it has, so far, attracted much less attention than its competitors.

The remainder of this paper is mainly concerned with finding at least one δ^*, and sometimes the exact set of minimax regret treatment rules, for different decision scenarios. The proofs will exploit the fact that δ^* can be represented as the decision maker's Nash equilibrium strategy in a fictitious zero-sum game against Nature. The game is explained in detail in the appendix; Nature's equilibrium strategy, also known as "worst-case prior," is sometimes of interest as well and will be labelled σ^*.

3 Treatment Choice Without Covariates

This section analyzes treatment choice when there are no covariates. I consider three sample designs that can be seen as stylized – and occasionally exact – descriptions of real-world data gathering procedures.[3]

(i) Stratified assignment: Let both treatments be unknown, hence $P(Y_0, Y_1)$ is unrestricted over $\Delta[0,1]^2$. The sample size N is known to be even; other than that, its distribution is known but arbitrary. (In particular, it might be degenerate so that N is known a priori.) Let within-sample treatment assignment be by even stratification, with $N/2$ sample points assigned to either treatment. The intuition for this is that sample subjects arrive in pairs and every pair is assigned to different treatments.

[3] The proof of proposition 1 is easily extended to show that all three sampling schemes minimize maximal regret over possible sampling schemes, given the respective assumptions about N. Further results on endogenous sample design are found in Schlag (2006).

(ii) Random assignment: Assume again that both treatments are unknown, but let the distribution of N be completely unrestricted (again, N may also be known). Let within-sample treatment assignment be by independent tosses of a fair coin.

(iii) Testing an innovation: Assume that one treatment, $t = 0$ say, is well understood, i.e. the marginal distribution of Y_0 is known. This should be thought of as testing an innovation against a status quo treatment. Obviously all sample points will be assigned to treatment 1 in this case. For simplicity, let N be known.

I also begin by assuming that outcomes are binary, i.e. $Y_0, Y_1 \in \{0, 1\}$, where a realization of $y_t = 1$ will be called a success. This restriction will later be dropped, but allows one to isolate some core issues and to generate "if and only if"-statements. Specifically, the set of minimax regret treatment rules is characterized as follows.[4]

Proposition 1 ***(i)*** *In the case of stratified sampling, minimax regret is achieved by*

$$\delta_1^*(s_N) \equiv \begin{cases} 0, & \overline{y}_1 < \overline{y}_0 \\ 1/2, & \overline{y}_1 = \overline{y}_0 \\ 1, & \overline{y}_1 > \overline{y}_0 \end{cases}.$$

Furthermore, any minimax regret treatment rukle must agree with δ_1^* except when $\overline{y}_0 = \overline{y}_1$, and δ^* is the unique minimax regret treatment rule that is measurable with respect to $(\overline{y}_0, \overline{y}_1)$.

(ii) *In the case of random assignment, minimax regret is achieved by*

$$\delta_2^*(s_N) \equiv \begin{cases} 0, & I_N < 0 \\ 1/2, & I_N = 0 \\ 1, & I_N > 0 \end{cases},$$

where

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

Furthermore, any minimax regret treatment rule must agree with δ_2^* excepot when $I_N = 0$, and δ_2^* is the unique minimax regret treatment rule that is measurable with respect to I_N.

[4] Part (i) of this result generalizes, and abbreviates the proof of, a previous finding by Canner (1970, section 4). The analysis of case (iii) has been subsequently extended by Manski and Tetenov (2006).

(iii) *In the case of testing an innovation, minimax regret is achieved by*

$$\delta_3^*(s_N) \equiv \begin{cases} 0, & N\overline{y}_1 < n^* \\ \lambda^*, & N\overline{y}_1 = n^* \\ 1, & N\overline{y}_1 > n^* \end{cases},$$

where $n^* \in \{1, \ldots, N\}$ *and* $\lambda^* \in [0, 1)$ *are uniquely characterized as follows:*

$$\max_{a\in[0,\mu_0]} \{(\mu_0 - a)\left[\Pr\left(N\overline{y}_1 > n^* | \mu_1 = a\right) + \lambda^* \Pr(N\overline{y}_1 = n^* | \mu_1 = a)\right]\} =$$
$$\max_{a\in[\mu_0,1]} \{(a - \mu_0)[\Pr(N\overline{y}_1 < n^* | \mu_1 = a) + (1 - \lambda^*) \Pr(N\overline{y}_1 = n^* | \mu_1 = a)]\}.$$

Furthermore, any minimax regret treatment rule must agree with δ_3^* *except when* $N\overline{y}_1 = n^*$*, and* δ_3^* *is the unique minimax regret treatment rule that is measurable with respect to* $N\overline{y}_1$*, the count of successes of treatment* 1*.*

Proposition 1 not only identifies minimax regret rules, but is essentially an "iff"-statement: In all three cases, minimax regret rules are pinned down up to behavior conditional on some threshold, and they are the simplest possible rules in the sense that any other rule must use additional (but, as it turns out, irrelevant) sample information on that threshold.[5]

Other than this, part (i) of the proposition is not surprising, given the ex ante symmetry of both decision scenario and risk function. The case might be somewhat different with part (ii). As the second, more verbal characterization of I_N suggests, an intuition for it goes as follows: Every observation of either a success of treatment 0 or a failure of treatment 1 constitutes some kind of signal in favor of treatment 0. Every other observation is a signal in the opposite direction. The decision maker should just keep a score of these signals, that is, she should decide according to which type of signal occurred more often.

Parts (i) and (ii) of proposition 1 provide limited support for an aspect of Manski's (2004) analysis. To estimate the regret incurred by different sample designs, he restricts attention to "empirical success" decision rules of the form $\delta^{ES}(\overline{y}_1, \overline{y}_0) = 1_{\overline{y}_1 > \overline{y}_0}$. This is clearly a simplification – in the spirit of the paper, one would want to use the minimax regret treatment rule if it were known. Proposition 1(i) shows that for binary outcomes and even stratification of samples, δ^{ES} is reasonably close, the modification being that the tie-breaking rule must be ex ante symmetric. Differences are larger for the other scenarios; in case (ii), the decision rule will asymptotically agree with an empirical success rule

[5] To give a flavor, alternative minimax regret rules can in all cases be constructed as follows: Divide $\mathcal{S}_N$ into pairs of sample realizations that are equiprobable conditional on any $s \in \mathcal{S}$ – e.g., elements of a pair might be permutations of each other. Fix a function $\gamma : \mathcal{S}_N \mapsto \{0, 1\}$ that maps exactly one element of every pair onto 1 and set $\delta^*(s_N) = \gamma$ conditional on the threshold event.

	N = 1	**N = 2**	**N = 3**	**N = 4**	**N = 5**	**N = 10**	**N = 20**	**N = 50**	**N = 100**	**N = 500**
μ_0= **.05**	0 (0.33)	0 (0.48)	0 (0.59)	0 (0.44)	0 (0.74)	1 (0.22)	1 (0.18)	2 (0.68)	5 (0.18)	25 (0.18)
μ_0= **.25**	0 (0.64)	0 (0.82)	1 (0.07)	1 (0.33)	1 (0.58)	2 (0.82)	5 (0.32)	12 (0.82)	25 (0.32)	125 (0.32)
μ_0= **.50**	1 (0)	1 (0.5)	2 (0)	2 (0.5)	3 (0)	5 (0.5)	10 (0.5)	25 (0.5)	50 (0.5)	250 (0.5)
μ_0= **.75**	1 (0.36)	2 (0.18)	2 (0.91)	3 (0.67)	4 (0.42)	8 (0.18)	15 (0.68)	38 (0.18)	75 (0.68)	375 (0.68)
μ_0= **.95**	1 (0.67)	2 (0.52)	3 (0.41)	4 (0.56)	5 (0.26)	9 (0.78)	19 (0.82)	48 (0.32)	95 (0.82)	475 (0.82)

Table 1: Testing an innovation: The minimax regret decision rule.

	N = 1	**N = 2**	**N = 3**	**N = 4**	**N = 5**	**N = 10**	**N = 20**	**N = 50**	**N = 100**	**N = 500**
μ_0= **.05**	1(0)	1(0.5)	1(0.68)	1(0.79)	1(0.87)	2(0.48)	3(0.43)	5(0.81)	9(0.38)	33(0.78)
μ_0= **.25**	1(0.8)	2(0.2)	2(0.76)	3(0.02)	3(0.61)	5(0.48)	8(0.85)	18(0.19)	32(0.78)	141(0.58)
μ_0= **.50**	1(0.9)	2(0.8)	3(0.6)	4(0.2)	4(0.88)	8(0.11)	14(0.21)	31(0.35)	58(0.74)	268(0.89)
μ_0= **.75**	1(0.93)	2(0.91)	3(0.88)	4(0.84)	5(0.79)	10(0.11)	18(0.62)	42(0.90)	82(0.51)	391(0.29)
μ_0= **.95**	1(0.95)	2(0.94)	3(0.94)	4(0.94)	5(0.94)	10(0.92)	20(0.86)	50(0.35)	98(0.84)	483(0.27)

Table 2: Testing an innvoation: The classical decision rule (5 percent significance, one-tailed test).

but differ from it markedly for small samples. An analysis of the relative performance of the empirical success rule, i.e. the degree of its inefficiency in terms of regret, will be conducted in section 5.

The characterization of δ^* in part (iii) is implicit, but numerical evaluation is easy. Table 1 illustrates the result for a selection of sample sizes N and values of μ_0.[6] Specifically, every cell has the format "$n^*(1-\lambda^*)$," thus the left-hand number is the critical number of observed successes that leads to randomized treatment assignment, and the number in parentheses gives the probability with which this randomization will pick treatment 0. The conversion from λ^* to $(1-\lambda^*)$ has been performed because the sum of the two numbers, $(n^*+1-\lambda^*)$, can be seen as a smooth index of the treatment rule's conservatism, with higher values indicating more conservative rules.

The minimax regret rule approximates an empirical success rule for rather small samples. This renders it akin to Bayesian decision rules derived from noninformative priors (see Berger 1985, section 3), but puts it in stark contrast to decision criteria informed by classical statistics. To illustrate this, table 2 displays the decision rule employed by a statistician who chooses treatment 1 if the data reject $H_0 : \mu_1 \leq \mu_0$ at 5% significance. The table can be read in exact analogy to table 1 – for example, if $\mu_0 = 0.25$ and $N = 10$, then decision rule prescribes to adopt treatment 1 with probability 0.52 if 5 successes were observed and with probability 1 if even more successes were observed.[7]

Although it must eventually converge to an empirical success rule, the classical statistician's decision

[6] This table extends table 3.1 in Manski (2005).

[7] The randomization conditional on a critical number of successes maximizes the test's power given its size.

rule is much more conservative than minimax regret. The reason is that it exclusively guards against the possibility of an inferior Y_1 "looking good" in a favorable sample. In contrast, the minimax regret rule equalizes regret between two worst-case scenarios, one of which is similar but the other one concerns the possibility of a superior Y_1 "looking bad."

I conclude this section by allowing for Y_0 and Y_1 to be distributed arbitrarily on the unit interval. This case seems much more complex than the above, but minimax regret treatment rules for it follow immediately from proposition 1 via an important observation due to Schlag (2003, 2006). To formalize it, call a state s a Bernoulli state if it implies Bernoulli distributions of both Y_0 and Y_1. Observe that Bernoulli states are fully characterized by triples $E(Y_0, Y_1, Y_0 Y_1)$. For any state s, call its Bernoulli equivalent the Bernoulli state s' such that s and s' induce the same value of $E(Y_0, Y_1, Y_0 Y_1)$. Finally, for any state space $\mathcal{S}$, let $\mathcal{S}'$ denote its Bernoulli equivalent, generated by replacing every state $s \in \mathcal{S}$ with its Bernoulli equivalent s'. ($\mathcal{S}'$ will in general have lower cardinality than $\mathcal{S}$.) Then the following is true.

Proposition 2 ***(Schlag 2003, 2006)*** *Consider any state space $\mathcal{S}$ that contains its Bernoulli equivalent, i.e. $\mathcal{S}' \subseteq \mathcal{S}$. Let δ' achieve minimax regret over $\mathcal{S}'$. Define the decision rule δ^* for $\mathcal{S}$ as follows: (i) Replace every observation $y_t \in [0, 1]$ with one independent realization of a Bernoulli variable with parameter y_t. (ii) Operate δ'.*

Then δ^* achieves minimax regret.

The proposition's hypothesis is fulfilled by all state spaces considered here. Thus, the findings of proposition 1 can be applied to the general case by preceding them with an information coarsening in which outcome observations other than $\{0, 1\}$ are replaced by independent tosses of biased coins. The intuition for why this works is as follows: In the fictitious game, the coarsening removes any incentive for Nature to use non-Bernoulli states, thus one can as well presume that she does – but then the coarsening does not matter. One price of the generalization is that "only if" will generally be lost. Also, randomization must accord with y_t even if the latter is a utility, hence δ^* is implementable only if the decision maker knows her utility function.

A notable feature of this result is that for general treatment outcomes, minimax regret treatment rules may differ much from empirical success rules. For one thing, the information coarsening causes δ^* to be randomized even if $\overline{y}_0 \neq \overline{y}_1$. As an example, if one realization of either treatment has been observed and the values are $y_0 = .5$ respectively $y_1 = .6$, then an empirical success rule would assign all future subjects to treatment 1, but any minimax regret rule will do so only with probability 55%.[8]

[8] I write "any" because for the special cases of $N = 2$ with stratified sampling, $N = 1$ with random assignment, or $N = 1$ with testing an innovation, "only if"-statements along the line of proposition 1 can be shown. See the web appendix for details.

Furthermore, $\overline{y}_0$ and $\overline{y}_1$ need not be sufficient statistics for the sample under minimax regret. If a sample of 4 is evenly stratified by treatment, then observations of $(1,1)$ on treatment 1 and $(1,0)$ on treatment 0 lead to $\delta^* = 1$; observations of $(1,1)$ respectively $(1/2, 1/2)$ will, in contrast, lead to treatment 1 being adopted with probability 0.875. To be sure, these differences vanish as samples become large. For all cases of proposition 1, one can verify that for every state s and hence also under every prior (although not uniformly over $\mathcal{S}$ or $\Delta\mathcal{S}$), δ^* is asymptotically equivalent to an empirical success rule.

4 Treatment Choice with Covariates

Assume now that there exists a finite-valued covariate $x \in \mathcal{X} = \{1, \ldots, X\}$ with known distribution, and regret is evaluated conditional on it. One might then wonder whether treatment rules should take x at least partially into account or rather pool information across covariate values. This question is at the core of Manski's (2004) analysis. The answer is intuitively non-obvious for finite samples because one encounters a trade-off between the decision rule's resolution and the size of relevant sample cells. Using bounds on regret, Manski (2004) establishes that for surprisingly small sample sizes, covariate-wise empirical success rules dominate any decision rules that pool information. This conclusion becomes much more stark under exact finite-sample analysis: In the natural extension of the previous propositions' setup, the minimax regret treatment rule will separate inference across covariates for *any* sample size.

I establish this by showing a more general result which requires some more notation. Outcomes are now contingent on covariates as well as treatments, thus the outcome variable is $Y_{tx} \in [0,1]$, and a state of the world is a distribution $s = P\left((Y_{0x}, Y_{1x})_{x \in \mathcal{X}}\right)$. A statistical treatment rule maps samples s_N into vectors of treatment assignment probabilities $\delta(s_N) \in [0,1]^X$, whose components δ_x are identified with probabilities of assigning treatment 1 to subjects with covariate x.

The natural extension of the previous section's state space to this situation, and also the extension analyzed by Manski (2004), is to identify the set of states of the world $\mathcal{S}$ with the set of distributions over $[0,1]^{2X}$. The present result, however, only needs the following structure: Let $\mathcal{S}_x$ denote the set of possible distributions $P(Y_{0x}, Y_{1x})$, then $\mathcal{S} = \times_x \mathcal{S}_x$, that is, the overall state space is the Cartesian product of covariate-wise state spaces.

Fix any decision problem that consists of such a state space $\mathcal{S}$ and a sampling rule, implying a sample space $\mathcal{S}_N$ with typical element $s_N = \{(t_n, x_n, y_n)\}_{n=1}^N$. Then for any pre-assigned covariate x, one can think of the problem of a decision maker who only knows $\mathcal{S}_x$, only sees the restricted sample $s_{Nx} \equiv \{(t_n, x_n, y_n) \in s_N : x_n = x\}$, and only needs to assign treatment to future subjects with covariate x. Let δ_x^* be a minimax regret treatment rule for this restricted problem. In many cases, $\{\delta_x^*\}_{x \in \mathcal{X}}$

will be known from previous propositions. For example, if the covariate is $x \in \{male, female\}$ with $\Pr(x = male) = 1/2$, the sample is a size-N simple random sample from the population, and sample treatments are assigned by coin toss, then δ^*_{female} and δ^*_{male} are characterized by proposition 1(ii). The following proposition, however, is applicable beyond these examples because it holds independently of whether the δ^*_x are known.

Proposition 3 *Let there be a finite-valued covariate x. Then a decision rule δ^* achieves minimax regret if it can be written as $\delta^*(s_N) = (\delta^*_x(s_{Nx}))_{x\in\mathcal{X}}$, where the δ^*_x are minimax regret solutions to covariate-wise decision problems as just defined.*

In words, proposition 3 establishes that the decision maker may completely separate inference across covariates. The intuition for this is as follows: The fictitious game's utility function is additively separable across treatments. With this in mind, consider a prior constructed such that (i) its marginal for any covariate is a worst-case prior for the accordingly restricted decision problem and (ii) treatment outcomes are independent across covariates. Then δ^* is easily seen to be a best response; in particular, nothing is to be gained from cross-covariate inference. But the prior can also shown to be a best response to δ^*.

Notice that this idea requires independence of treatment outcome distributions across covariates, but not independence of sampling distributions. For example, it applies to simple random samples from the population, even though cell sizes will then correlate. Indeed, this is the example given just before the proposition.

Paradoxically, the recommendation based on exact finite-sample analysis does not seem desirable in a world of finite samples. It requires one to condition on covariates even if this leads to extremely small or empty sample cells. Whilst medical researchers might want to consider the effect of race on treatment outcomes, they will hardly want to altogether ignore experiences made with white subjects when considering treatment for black subjects unless samples are extremely large. And I have not yet mentioned the idea of conditioning on birthdays!

What's more, proposition 3 suggests that as the support of a covariate grows, minimax regret treatment rules will approach no-data rules, because the proportion of covariate values that have been observed in the sample vanishes. Indeed, it is possible to extend the result to a continuous covariate, and it is then true that a no-data rule achieves minimax regret. To formalize this, let $\mathcal{X} = (0, 1]$. Then a state s can be identified with a pair of functions $(\mu_{0x}, \mu_{1x}) : \mathcal{X} \to [0, 1]^2$, and a decision rule maps the sample space onto decision functions $\delta : (0, 1] \to [0, 1]$.

Proposition 4 *Let x be continuous. Then $\min_{\delta\in\mathcal{D}} \{\sup_{s\in\mathcal{S}} R(\delta, s)\} = 1/2$. This value is achieved by the decision rule that sets $\delta^*(s_N) \equiv 1/2$, independent of sample observations.*

The above no-data rule is weakly dominated – if a covariate value observed in the sample ever recurs exactly, the according sample information should be used. However this happens with zero probability, so δ^* achieves minimax regret.

Propositions 3 and 4 pose a difficulty to advocates of minimax regret because no-data rules are obviously undesirable. What's more, the existence of easy examples in which the maximin utility criterion generates them is frequently used against maximin utility, and sometimes in favor of minimax regret as an alternative (Savage 1954, Schlag 2003, Manski 2004, and section 2 of this paper).

But perhaps, the problem lies not with minimax regret per se, but with underspecification of prior information. For example, the worst-case prior supporting proposition 3 replicates previous worst-case priors conditional on any given covariate and furthermore renders (Y_{0x}, Y_{1x}) independent of $(Y_{0x'}, Y_{1x'})$. Hence, Y_{1x} is allowed to vary extremely widely across covariates; furthermore, signals conditional on x are presumed to be entirely uninformative for treatment outcomes conditional on x'. One or both of these implications are inappropriate in many applications. The decision situation should then be respecified so as to exclude them.

One possibility would be to restrict σ to a set $\Pi \subset \Delta\mathcal{S}$ of admissible priors. Obviously, this makes sense only if one indeed interprets σ as a prior, but users who do so might find this approach very helpful. For example, say a medical researcher wants to accommodate the possibility that reactions to treatments differ widely by race and/or gender, yet also believes that across the universe of treatments, this is the exception rather than the rule.

Alternatively, one can attempt to constrain the state space by means of plausible restrictions that destroy the product structure of $\mathcal{S}$. This approach should be attractive to non-Bayesians because it does not rely on committing to priors. One such idea would be to endow $\mathcal{S}$ with some distance metric and restrict attention to states in which (Y_{0x}, Y_{1x}) and $(Y_{0x'}, Y_{1x'})$ are not too far apart. To illustrate the promise of such restrictions, I will now partially analyze one example along these lines.

Assumption 1 ***Bounded Effect of Covariates (BEC)***

$$|\mu_{xt} - \mu_{x't}| \leq k, \forall x, x' \in \mathcal{X}, t \in \{0, 1\}.$$

This assumption means that treatment outcomes cannot vary too much across covariates. It can be scaled from "no effect" to vacuousness as k ranges from 0 to 1. For those covariates where propositions 3 and 4 appear absurd, one would certainly be able to specify k rather low. Nonetheless such a specification will be conceptually subjective even when it is substantively uncontroversial. This matter will be taken up in the concluding discussion.

A full analysis of the implications of BEC is currently beyond my reach, but the following lemmata are instructive.

Lemma 1 *Fix any one sampling scheme from proposition 1 but with covariate $x \in \mathcal{X}$ and with $(N_x)_{x \in \mathcal{X}}$ known. Then for any $k > 0$, there exist N_k s.t. proposition 3 applies if $\min_{x \in \mathcal{X}} N_x \geqq N_k$. ($N_k$ can differ across sampling schemes.)*

Lemma 2 *For any case of proposition 1 and sample size N, there exists $\underline{k}_N > 0$ s.t. if assumption BEC holds with $k < \underline{k}_N$, then there exists a minimax regret rule δ^* with $\delta^*_\times(s_N) \equiv \delta^*_{x'}(s_N)$, i.e. treatment assignment does not vary across covariates.*

Lemma 1 can be seen as "robustification" of proposition 3. It shows that the proposition does not require the full Cartesian product structure of $\mathcal{S}$ but only an environment that is sufficiently permissive with respect to $|\mu_{tx} - \mu_{tx'}|$. At the same time, the lemma implies an interesting observation regarding the asymptotic behavior of minimax regret. Obviously, minimax regret will asymptotically condition on every covariate in the sense of converging to some decision rule that does so. Lemma 1 implies a much stronger finding: As long as k is strictly positive, full separation of inference across covariates will be achieved for some sample size that is large enough (uniformly over x) but finite.

Lemma 2 similarly robustifies a rationale for ignoring covariates. To see this, note that if assumption BEC is imposed with $k = 0$, then the covariate should trivially be ignored. This observation may suffice to resolve the problem with respect to subjects' birthdays, but not with respect to covariates like race and gender, whose influence may sometimes believed to be small but rarely exactly zero. Lemma 2 establishes that complete ignorance of covariates will achieve minimax regret for k positive but small enough. Thus, the policy of ignoring a covariate is robust to minor deviations from the assumption that this covariate does not matter. Furthermore, inspection of the proof reveals that lemma 2 does not require x to be discrete; hence, this result also applies to the problem pointed out in proposition 4.

5 Applications

I will now use finite-sample results to compare the minimax regret value of different decision situations and sample designs. These computations have numerous applications. Consider, therefore, a decision maker who has to assign treatments contingent on observations from a sample of size N and a certain pre-specified design. This situation can be varied in several ways: one or both treatments may be unknown, and there may or may not be a finite-valued covariate. All in all, I will analyze the following scenarios.

Case 1: Both treatments are unknown, and there is no covariate.

Case 2: Both treatments are unknown, there exists a finite-valued covariate with known distribution, and the sample is stratified according to covariate with $N_x \geq 0$.

Case 3: One treatment is known, and there is no covariate.

Case 4: One treatment is known, there exists a finite-valued covariate with known distribution, and the sample is stratified according to covariate with $N_x \geq 0$.

All of these scenarios are covered by previous propositions, so that their value in terms of finite sample minimax regret can be computed. Results are displayed below. To keep algebraic expressions manageable, define

$$B(n, N, p) \equiv \begin{pmatrix} N \\ n \end{pmatrix} p^n (1-p)^{N-n},$$

the probability of n successes in N draws under a binomial distribution with success rate p.

Proposition 5 ***Exact Minimax Regret Value of Different Sample Designs***

(i) *Assume that case 1 applies, then minimax regret is*

$$\begin{aligned} R_1(N) &= \max_{a \in [1/2,1]} \left\{ (2a-1) \left[\sum_{0<n\leq N/2'} B(n, N', a) \right] \right\} \\ N' &= \max_{M \in \mathbb{N}} \{ M \leq N : M \textit{ is even} \}. \end{aligned}$$

(ii) *Assume that case 2 applies, then minimax regret is*

$$R_2(N_1, \ldots, N_X) = \sum_{x=1,\ldots,X} \Pr(x) R_1(N_x).$$

(iii) *Assume that case 3 applies, then minimax regret is*

$$R_3(N) = \max_{a \in [0,\mu_0]} \left\{ (\mu_0 - a) \left[\sum_{n^*<n\leq N'} B(n, N, a) + \lambda^* B(n^*, N, a) \right] \right\}$$

with (n^*, λ^*) *as in proposition 3.*

(iv) *Assume that case 4 applies, then minimax regret is*

$$R_4(N_1, \ldots, N_X) = \sum_{x \in \mathcal{X}} \Pr(x) R_3(N_x).$$

In all cases, $R_i(0) = 1/2$.

It is conceptually trivial to determine minimax regret values along the lines of (i) and (iii) for random sample sizes, but the algebraic expressions will be very cumbersome since the objective functions must involve an additional integration step. Similarly, (ii) and (iv) could be extended to random stratifications but with significant notational and computational effort.

N_0	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
R^*	.1250	.0870	.0706	.0609	.0543	.0495	.0458	.0428	.0403	.0382
$\underline{R}$	.25	.1481	.1066	.0883	.0765	.0681	.0617	.0568	.0528	.0495
N_0	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**
R^*	.0364	.0348	.0335	.0322	.0311	.0301	.0292	.0284	.0276	.0269
$\underline{R}$	.0467	.0442	.0422	.0403	.0387	.0372	.0359	.0347	.0336	.0326
N_0	**21**	**22**	**23**	**24**	**25**	**50**	**100**	**200**	**500**	**1000**
R^*	.0263	.0257	.0251	.0246	.0241	.0170	.0120	.0085	.0054	.0038
$\underline{R}$	.0317	.0308	.0300	.0293	.0286	.0193	.0132	.0091	.0056	.0039

Table 3: Selected values of R* and of a lower bound for regret incurred by the empirical success rule (two unknown treatments).

5.1 Performance of Empirical Success Rules

The empirical success rule, $\delta^{ES} \equiv 1_{\overline{y}_1 \geq \overline{y}_0}$, is of special interest for at least three reasons. First, it is probably the most obvious decision rule to employ. Second, it is the one used by Manski (2004). Third, Hirano and Porter (2005) show how the machinery of Le Cam (1986) can be brought to minimax regret problems. If their approach is applied to the scenarios considered here, then δ^{ES} emerges as approximation to δ^*. Yet at the same time, δ^{ES} was seen to significantly differ from the minimax regret rule for small samples.

How much do these differences matter in terms of regret incurred? By searching over a restricted state space, one can numerically bound from below $\max_{s \in \mathcal{S}} R(\delta^{ES}, s)$, the maximum regret incurred by δ^{ES}. I do this first for the setting of proposition 1(i). If anything, this choice biases results in favor of δ^{ES} because compared to proposition 1(ii), it renders δ^{ES} relatively similar to δ^*.

In table 3, the lower bound, labelled $\underline{R}$, is contrasted with the exact maximal regret R^* – as evaluated from proposition 5 – for different sample cell sizes $N_0 = N/2$. The table is justified in more detail in appendix B.

The approximation to δ^* comes at significant cost when samples are very small, incurring at least double the true minimax regret for $N_0 = 1$. But the table also reveals convergence: For $N_0 = 20$, i.e. a sample size of 40, δ^{ES} incurs an inefficiency of 21% of true minimax regret, and this percentage becomes very small for large samples.

A similar exercise can be performed for the case that the expected outcome of one treatment is known. The results are tables 4 through 6, which use the three values of μ_0 that were also considered by Manski (2004). In these tables, the lower bound $\underline{R}$ sometimes displays erratic behavior, e.g. it increases in N. This may occasionally reflect movement in the true regret, but probably also attests to

N	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
R^*	.0408	.0358	.0315	.0279	.0248	.0221	.0197	.0176	.0158	.0143
$\underline{R}$	.2256	.1270	.0859	.0634	.0492	.0396	.0326	.0273	.0232	.0199
N	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**
R^*	.0130	.0120	.0111	.0103	.0096	.0090	.0087	.0087	.0086	.0086
$\underline{R}$	.0193	.0195	.0197	.0198	.0199	.0200	.0201	.0202	.0203	.0203
N	**21**	**22**	**23**	**24**	**25**	**50**	**100**	**200**	**500**	**1000**
R^*	.0085	.0084	.0083	.0081	.0080	.0054	.0037	.0026	.0017	.0012
$\underline{R}$	.0187	.0172	.0159	.0147	.0136	.0084	.0059	.0037	.0021	.0014

Table 4: Selected values of R* and of a lower bound for regret incurred by the empirical success rule (one unknown treatment, ţ0=0.05).

N	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
R^*	.0900	.0516	.0389	.0380	.0345	.0299	.0268	.0265	.0252	.0232
$\underline{R}$	.1406	.0625	.0705	.0790	.0517	.0431	.0455	.0473	.0353	.0336
N	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**
R^*	.0217	.0215	.0208	.0196	.0187	.0186	.0181	.0173	.0167	.0166
$\underline{R}$	.0346	.0354	.0279	.0280	.0285	.0289	.0238	.0242	.0245	.0248
N	**21**	**22**	**23**	**24**	**25**	**50**	**100**	**200**	**500**	**1000**
R^*	.0162	.0157	.0152	.0151	.0149	.0104	.0074	.0052	.0033	.0023
$\underline{R}$	.0213	.0215	.0218	.0220	.0193	.0132	.0090	.0060	.0036	.0025

Table 5: Selected values of R* and of a lower bound for regret incurred by the empirical success rule (one unknown treatment, ţ0=0.25).

the difficulty of finding the regret-maximizing state of nature; certain values of N and μ_0 allow for more efficient guesses than others. Nonetheless, it becomes apparent that the relative underperformance, in terms of regret, of δ^{ES} is quite large for small and moderate sample sizes.

5.2 Comparing Sample Stratifications

Proposition 5 can also be used to compare the value, in terms of minimax regret incurred by the optimal decision rule, of different sample stratifications. This replicates some of Manski's (2004) analysis but eliminates two elements of approximation. Firstly, computations are based on exact regret and not an upper bound on it; secondly, they presume that conditional on sample designs, exact minimax regret rules are chosen, whereas Manski restricts attention δ^{ES}. For the case of a binary covariate with

N	1	2	3	4	5	6	7	8	9	10
R^*	.0625	.0625	.0435	.0435	.0353	.0353	.0304	.0304	.0272	.0272
$\underline{R}$	.1250	.1099	.0531	.0697	.0462	.0533	.0401	.0441	.0356	.0382
N	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**
R^*	.0247	.0247	.0229	.0229	.0214	.0214	.0201	.0201	.0191	.0191
$\underline{R}$	.0321	.0340	.0294	.0309	.0273	.0284	.0255	.0264	.0240	.0247
N	**21**	**22**	**23**	**24**	**25**	**50**	**100**	**200**	**500**	**1000**
R^*	.0182	.0182	.0174	.0174	.0167	.0120	.0085	.0060	.0038	.0027
$\underline{R}$	.0227	.0233	.0216	.0221	.0206	.0143	.0096	.0066	.0040	.0028

Table 6: Selected values of R* and of a lower bound for regret incurred by the empirical success rule (one unknown treatment, ţ0=0.5).

stratification according to N_{tx}, the resulting optimal stratifications are shown in table 7, which can be compared to Manski's (2004, table II) display of what he calls "quasi-optimal stratifications." The according comparison values are reproduced in parentheses.[9]

The table reveals that previous bounds had considerable slack. (Its two sources – suboptimal performance of δ^{ES} and slack of large deviations bounds – are not separated here.) Accordingly, optimal stratifications frequently differ from the quasi-optimal ones, although not by very much. Furthermore, proposition 5 allows for very fast computation of the solutions, so that the table can be extended to much larger sample sizes.

6 Concluding Remarks

This paper contributes to a rather young literature in which a rather old criterion, namely minimax regret, is applied to models of real-world decisions.[10] The aim was to characterize finite sample minimax regret rules for the scenario analyzed by Manski (2004) and variations thereof. Important results include the comparison between δ^* and δ^{ES}, not least because it illuminates the finite sample precision of asymptotic approximations as in Hirano and Porter (2005), and the "separate inference" result on covariates, which reinforces a finding in Manski (2004) to the point that contrary to its original interpretation, the result may cast doubt on the minimax regret principle rather than on conventional practice. On a more general level, I believe to have shown that exact analysis can generate some significant and rather general insights.

A natural question is whether the minimax regret criterion manages to convince the user. There

[9] As in Manski (2004) and in proposition 5 above, I restrict attention to deterministic stratifications.

[10] The formal introduction of minimax regret is generally attributed to Savage (1951).

	$\Pr(\mathbf{X}=\mathbf{0})=\mathbf{.05}$			$\Pr(\mathbf{X}=\mathbf{0})=\mathbf{.25}$		
N	$\mathbf{N}_{00},\mathbf{N}_{10}$	$\mathbf{N}_{01},\mathbf{N}_{11}$	**R**	$\mathbf{N}_{00},\mathbf{N}_{10}$	$\mathbf{N}_{01},\mathbf{N}_{11}$	**R**
4	0(0)	2(2)	.108(.338)	1(1)	1(1)	.125(.423)
8	1(0)	3(4)	.073(.250)	1(1)	3(3)	.084(.293)
12	1(1)	5(5)	.058(.203)	2(2)	4(4)	.067(.234)
16	1(1)	7(7)	.050(.173)	3(3)	5(5)	.058(.205)
20	1(2)	9(8)	.045(.154)	3(3)	7(7)	.052(.182)
24	2(2)	10(10)	.041(.143)	4(4)	8(8)	.047(.162)
28	2(2)	12(12)	.037(.133)	5(5)	9(9)	.044(.153)
32	2(2)	14(14)	.035(.124)	5(6)	11(10)	.041(.144)
36	2(2)	16(16)	.033(.115)	6(6)	12(12)	.038(.137)
40	3(2)	17(18)	.031(.108)	7(7)	13(13)	.037(.129)
44	3(2)	19(20)	.030(.101)	7(8)	15(14)	.035(.122)
48	3(2)	21(22)	.029(.094)	8(8)	16(16)	.033(.116)
52	3(3)	23(23)	.027(.088)	8(8)	18(18)	.032(.110)
60	4	26	.025	10	20	.030
80	5	35	.022	13	27	.026
100	6	44	.020	16	34	.023
200	12	88	.014	32	68	.016
500	31	219	.009	81	169	.010
1000	62	438	.006	163	337	.007

Table 7: Optimal (deterministic) sample stratifications; values from Manski (2004a) in brackets.

are two ways to investigate this question: One is axiomatic analysis, the other one is to look whether actual minimax regret rules make sense. I will conclude by reflecting on the latter in the light of this paper's findings. In particular, I will reconsider propositions 1(ii) and 3, partly because based on reception of early drafts of this paper, they may contradict readers' intuitions, and partly because I think that the discussion highlights significant points.

With proposition 3, the case is rather easy. A minimax regret decider acts as if she had probabilistic beliefs according to the worst-case prior σ^*. This need not mean that she actually believes σ^*, but it will lead to satisfactory results only if σ^* is not entirely implausible. Hence, minimax regret treatment rules – and maximin-type decision rules more generally – require proper specification of all available prior information. The most obvious way to do this is to restrict $\mathcal{S}$; users who use minimax regret as a prior selection device could alternatively make the set of admissible priors a subset of $\Delta\mathcal{S}$. Indeed, I showed how one intuitive restriction on $\mathcal{S}$ will alleviate the problem by robustifying the rationale for ignoring implausible covariates. It must be borne in mind, however, that this resolution is achieved by introducing additional assumptions that may be uncontroversial but are subjective. One argument for maximin-type rules as opposed to Bayesianism, namely that the former avoid the imposition of subjective information, accordingly loses some of its force. Indeed, the imposition of BEC with subjective parameter κ blurs the distinction between the "classical" approach considered in this paper and the robust Bayesians' Γ-minimax regret approach, which imposes a set of priors and then uses minimax regret to generate point-valued decision rules.

Proposition 1(ii) poses a more subtle challenge to intuitions and may even appear perfectly acceptable to many readers. But consider the following example (due to Chuck Manski): A sample of size 1100 has been drawn, 1000 of which were allocated to treatment 0. Among these observations, 550 successes and 450 failures were observed. Among the 100 subjects assigned to treatment 1, 99 successes and 1 failure were observed. Then the minimax regret treatment rule prescribes to assign all future subjects to treatment 0. Intuitively, this conclusion is less than obvious, and indeed, many other decision criteria will prescribe treatment 1. What is going on here?

Some understanding can be gained by reconsidering the maximin utility criterion. Here as well as in other papers, maximin utility generates trivial decision rules. The reason is that it optimizes against a worst-case prior σ^* under which treatments are uniformly catastrophic *even if sample evidence overwhelmingly shows that this prior cannot be right.* I suspect that this unresponsiveness to likelihoods is the true problem of maximin utility, and that the triviality results are just symptoms. Furthermore, the problem with proposition 1(ii) might be just the same. Given the worst-case prior that supports δ^* ("either $(\mu_0, \mu_1) = (a, 1-a)$ or $(\mu_0, \mu_1) = (1-a, a)$"), the prescription to choose treatment 0 is doubtlessly correct. But the example presumes a sample outcome which is an extreme tail event under either support point of this prior, thus it renders this prior overwhelmingly unlikely. If proposition

1(ii) is seen as a problem, then I suggest that its cause lies in minimax regret's selective ignorance of likelihoods, or in other words, in the fact that the worst-case prior can be dogmatic on some dimensions.

These arguments are not intended to "prove" minimax regret "wrong." On the contrary, I believe that it deserves much further investigation, and also that other criteria, from classical statistics to Bayesianism, have significant drawbacks too. But obviously, these should not keep proponents of minimax regret from being candid about potential drawbacks of this rule.

A Proofs

Preliminaries Most proofs proceed by analyzing the following zero-sum game: (i) The decision maker chooses a statistical treatment rule $\delta : s_N \to [0,1]$, Nature chooses a mixed strategy $\sigma \in \Delta(\mathcal{S})$ over states $s \in \mathcal{S}$. (ii) A neutral meta-player draws s according to σ, then s_N according to s. (iii) The decision maker's payoff is $E(R(\delta, s))$. By standard results, the game has Nash equilibria in all cases analyzed below. It is useful because of the following, well-known implication of the Maximin Theorem (e.g., Berger 1985, section 5):[11]

Lemma 3 ***(i)*** *A treatment rule δ^* achieves minimax regret iff there exists $\sigma^* \in \Delta(\mathcal{S})$ s.t. (δ^*, σ^*) is a Nash equilibrium of the above game.*

(ii) *If the pair (δ^*, σ^*) has this property, then δ^{**} [σ^{**}] achieves minimax regret iff (δ^{**}, σ^*) [(δ^*, σ^{**})] has the property as well.*

Whilst I will not give a complete proof, here is an intuition for why this works. A minimax regret treatment rule must solve

$$\delta^* \in \arg\min_{\delta e \mathcal{D}} \left\{ \max_{s \in \mathcal{S}} R(\delta, s) \right\} = \arg\min_{\delta e \mathcal{D}} \left\{ \max_{\sigma \in \Delta \mathcal{S}} \int R(\delta, s) d\sigma \right\},$$

where the equality holds because the expectation is a linear operator, hence both maximization problems have the same value. But then one can write

$$\delta^* \in \arg\min_{\delta e \mathcal{D}} \left\{ \int R(\delta, s) d\sigma^* \right\},$$

where

$$\sigma^* \in \arg\max_{\sigma \in \Delta \mathcal{S}} \left\{ \int R(\delta^*, s) d\sigma \right\}.$$

Yet these two conditions just say that (δ^*, σ^*) must be a fixed point of the fictitious game's best-response correspondence.

[11] The statement of the result presumes that $\mathcal{S}$ is compact, which is given in all applications in this paper; a more general formulation would require the use of limit operators.

In general, a state s is given by a distribution $P\left((Y_{0x}, Y_{1x})_{x=1}^{X}\right) \in \Delta[0,1]^{2X}$, where Y_{tx} is the outcome induced by treatment t and covariate x. Since joint realizations of (Y_{0x}, Y_{1x}) are never observed, the distribution of s_N depends on the distribution of (Y_{0x}, Y_{1x}) only through the latter's marginals. Furthermore, regret directly depends on this distribution only through (μ_{0x}, μ_{1x}). Hence, it is w.l.o.g. to impose that Y_{0x} and Y_{1x} are independent. When covariates are not considered, x is dropped from notation. When outcomes are binary, the distribution of Y_{tx} is Bernoulli and fully characterized by $\mu_{tx} \equiv E(Y_{tx})$. Thus, in the setup of proposition 1(i)-(ii), a state is really a couplet $s = (\mu_0, \mu_1)$, and in the setup of proposition 1(iii), a scalar $s = \mu_1$.

In any given state s, one treatment is better, and that treatment's expected outcome defines the utility frontier. Absent covariates, regret is therefore given by

$$\begin{aligned} R(\delta, s) &= \max\{\mu_0, \mu_1\} - [\mu_0 E(1-\delta(\theta)) + \mu_1 E\delta(\theta)] \\ &= (\mu_1 - \mu_0)^+ E(1-\delta(\theta)) + (\mu_0 - \mu_1)^+ E(\delta(\theta)), \end{aligned}$$

where $Y^+ \equiv \max\{Y, 0\}$ is the positive restriction of Y. (The second expression has a direct intuition: For a given s, only one summand is nonzero, and it equals the utility loss caused by choosing the wrong treatment, weighted by the probability of doing so.)

I will frequently use the fact that any s in the support of σ^* must individually maximize $R(\delta^*, s)$.

Proposition 1

(i) Let δ^* be as stated in the proposition, then σ^* is a best response iff it is supported on

$$\begin{aligned} \arg\max_{(\mu_0,\mu_1)\in[0,1]^2} \Bigg\{ &\int (\mu_1 - \mu_0)^+ \left[\left(\Pr(\overline{y}_0 > \overline{y}_1 | N) + \frac{1}{2}\Pr(\overline{y}_0 = \overline{y}_1 | N)\right)\right] dP(N) \\ &+ \int (\mu_0 - \mu_1)^+ \left[\Pr(\overline{y}_1 > \overline{y}_0 | N) + \frac{1}{2}\Pr(\overline{y}_1 = \overline{y}_0 | N)\right] dP(N) \Bigg\} \end{aligned}$$

where $P(N)$ is the distribution of N.

This $\arg\max$ is nonempty since the objective is continuous and the feasible set is compact, and it is symmetrical in μ_0 and μ_1. Hence, it contains (a, b) iff it contains (b, a), and σ^* can be chosen to randomize evenly between $s_0 \equiv (a, b)$ and $s_1 \equiv (b, a)$ for some $a > b$. The best response to this σ^* will equal 1 if the posterior probability of s_1 exceeds the one of s_0 and 0 otherwise, where the tie-breaking rule is not restricted. Since s_0 and s_1 have equal prior probability, the posterior probability of s_1 exceeds that of s_0 iff the sample is more likely given s_1. The following argument shows that this is the case iff $\overline{y}_1 > \overline{y}_0$:

$$\begin{aligned}
\Pr((N\overline{y}_0, N\overline{y}_1) = (n_0, n_1)|s_0) &> \Pr((N\overline{y}_0, N\overline{y}_1) = (n_0, n_1)|s_1) \\
\iff B(n_0, N, a)B(n_1, N, b) &> B(n_0, N, b)B(n_1, N, a) \\
\iff \begin{pmatrix} N \\ n_0 \end{pmatrix} a^{n_0}(1-a)^{N-n_0} \begin{pmatrix} N \\ n_1 \end{pmatrix} b^{n_1}(1-b)^{N-n_1} &> \begin{pmatrix} N \\ n_0 \end{pmatrix} b^{n_0}(1-b)^{N-n_0} \begin{pmatrix} N \\ n_1 \end{pmatrix} a^{n_1}(1-a)^{N-n_1} \\
\iff \left(\frac{a}{1-a}\right)^{n_0-n_1} &> \left(\frac{b}{1-b}\right)^{n_0-n_1} \\
\iff n_0 - n_1 &> 0.
\end{aligned}$$

Thus (δ^*, σ^*) is a Nash equilibrium. As any other minimax regret decision rule δ^{**} must be best against σ^*, δ^* is unique whenever it is a strict best response, that is, except when $\overline{y}_0 = \overline{y}_1$. Restrict attention to decision rules that depend only on $(\overline{y}_0, \overline{y}_1)$, then it follows that δ^* is unique up to the tie-breaking probability. Assume this probability favors treatment t, then Nature will want to deviate to the pure strategy concentrated on s_{1-t}. Thus tie-breaking must be even.

(ii) Assume σ^* randomizes evenly between $s_0 \equiv (a, 1-a)$ and $s_1 \equiv (1-a, a)$ for some $a \in [1/2, 1]$, thus treatment t is the correct choice in state s_t. Conditional on a sample of size N having been observed, the best response to σ^* equals 0 if the likelihood ratio between s_0 and s_1 exceeds 1. Every realization of $y_0 = 1$ or $y_1 = 0$ increases this likelihood ratio by a factor $a/(1-a)$, whereas every other event reduces it by the same factor, thus the likelihood ratio induced by the sample is

$$L_N = \left(\frac{a}{1-a}\right)^{\substack{\#(\text{observations of } y_0=1) \\ +\#(\text{observations of } y_1=0)}} \left(\frac{1-a}{a}\right)^{\substack{\#(\text{observations of } y_1=1) \\ +\#(\text{observations of } y_0=0)}} = \left(\frac{a}{1-a}\right)^{I_N}$$

and δ^* is indeed a best response to σ^*. It remains to verify that some σ^* of the hypothesized form is a best response to δ^*, i.e. it is supported on

$$\begin{aligned}
\arg \max_{(\mu_0,\mu_1)\in[0,1]^2} \Bigg\{ &(\mu_1 - \mu_0)^+ \left[\Pr(I_N > 0) + \frac{1}{2}\Pr(I_N = 0)\right] \\
&+ (\mu_0 - \mu_1)^+ \left[\Pr(I_N < 0) + \frac{1}{2}\Pr(I_N = 0)\right]\Bigg\}.
\end{aligned}$$

As before, this arg max is nonempty and symmetric in the sense that it contains (a, b) iff it contains (b, a). It therefore suffices to show that it contains an element $(a, 1-a)$. This is done by establishing that the objective function depends on (μ_0, μ_1) only via $(\mu_1 - \mu_0)$.

It clearly suffices to show that I_N is independent of μ_0 given $(\mu_1 - \mu_0)$, for any N. The proof will be by induction over N, thus assume the result for $N = n$ and consider $\Pr(I_{n+1} = x)$. This event can occur in four ways: Either the first n sample points induced $I_n = x - 1$ and the last observation was

$y_0 = 1$ or $y_1 = 0$, or $I_n = x+1$ and one of the other events happened. Thus

$$\begin{aligned}\Pr(I_{n+1} = x) &= \Pr(I_n = x-1) \cdot \left(\frac{1}{2}\mu_0 + \frac{1}{2}(1-\mu_1)\right) + \Pr(I_n = x+1) \cdot \left(\frac{1}{2}(1-\mu_0) + \frac{1}{2}\mu_1\right) \\ &= \Pr(I_n = x-1) \cdot \left(\frac{1}{2} + \frac{1}{2}(\mu_0 - \mu_1)\right) + \Pr(I_n = x+1) \cdot \left(\frac{1}{2} + \frac{1}{2}(\mu_1 - \mu_0)\right).\end{aligned}$$

The proof is concluded by observing that I_0 is deterministically equal to zero and therefore independent of μ_0 given $(\mu_1 - \mu_0)$. Uniqueness can be shown as in (i).

(iii) Assume that δ^* is as stated, then σ^* must be supported on

$$\begin{aligned}\arg\max_{a \in [0,1]} \{&(\mu_0 - a)^+ \left[\Pr(N\overline{y}_1 > n^* | \mu_1 = a) + \lambda^* \Pr(N\overline{y}_1 = n^* | \mu_1 = a)\right] \\ &+ (a - \mu_0)^+ \left[\Pr(N\overline{y}_1 < n^* | \mu_1 = a) + (1-\lambda^*) \Pr(N\overline{y}_1 = n^* | \mu_1 = a)\right]\}.\end{aligned}$$

Assume by contradiction that $\mu_0 \geq a$ a.s. given σ^*, then the decision maker's best response would be to always choose treatment 0, but then σ^* would cause zero regret and therefore fail to maximize regret (strictly positive regret can obviously be enforced). A similar argument excludes the possibility that $\mu_0 \leq a$ a.s. given σ^*. Thus σ^* must randomize over at least one a s.t. $\mu_0 \geq a$ and one a' s.t. $\mu_0 \leq a'$. This, however, requires that both of these be in the above arg max and therefore that

$$\begin{aligned}\max_{a \in [0,1]} \{&(\mu_0 - a) \left[\Pr(N\overline{y}_1 > n^* | \mu_1 = a) + \lambda^* \Pr(N\overline{y}_1 = n^* | \mu_1 = a)\right] = \\ &\max_{a \in [0,1]} \{(a - \mu_0) \left[\Pr(N\overline{y}_1 < n^* | \mu_1 = a) + (1-\lambda^*) \Pr(N\overline{y}_1 = n^* | \mu_1 = a)\right]\}.\end{aligned}$$

This is the condition stated. Define $\alpha \equiv n+1-\lambda \in [0, N+1]$. This is a smooth indicator of a treatment rule's conservatism, with $\alpha = 0$ respectively $\alpha = N+1$ indicating no-data rules that always respectively never assign treatment 1. It is easy to see that $\Pr(N\overline{y}_1 > n | \mu_1 = a) + \lambda \Pr(N\overline{y}_1 = n | \mu_1 = a)$ strictly decreases in α for any a; thus the l.h.s. of the above equation strictly decreases in α and similarly, the r.h.s. increases in it. Furthermore, both l.h.s. and r.h.s. are continuous in α, the l.h.s. is zero for $\alpha = 0$ and one for $\alpha = N+1$, and the r.h.s. is zero for $\alpha = N+1$ but one for $\alpha = 0$. It follows that equality obtains at exactly one intermediate value of α and, by implication, of (n, λ).

To prove that this value is (n^*, λ^*), it remains to show that σ^* can be found s.t. δ^* is a best response. Conjecture that σ^* randomizes over two states a and b with $a < \mu_0 < b$. Denote by σ_a^* the probability of state a. For the decision maker to randomize at $N\overline{y}_1 = n^*$, this observation must induce indifference between the treatments conditional on Bayesian updating of σ^*, hence

$$\begin{aligned}&(b - \mu_0) \Pr(\mu_1 = b | N\overline{y}_1 = n^*) = (\mu_0 - a) \Pr(\mu_1 = a | N\overline{y}_1 = n^*) \\ \Longrightarrow\ &(b - \mu_0) \Pr(N\overline{y}_1 = n^* | \mu_1 = b) \cdot (1 - \sigma_a^*) = (\mu_0 - a) \Pr(N\overline{y}_1 = n^* | \mu_1 = a) \cdot \sigma_a^* \\ \Longrightarrow\ &\sigma_a^* = \frac{(b - \mu_0) \Pr(N\overline{y}_1 = n^* | \mu_1 = b)}{(b - \mu_0) \Pr(N\overline{y}_1 = n^* | \mu_1 = b) + (\mu_0 - a) \Pr(N\overline{y}_1 = n^* | \mu_1 = a)} \in (0, 1),\end{aligned}$$

so an appropriate σ_a^* can always be found. Since the Binomial distribution has the Monotone Likelihood Ratio property, it is immediate that for this σ_a^*, the decision maker will strictly prefer treatment 1 [0] iff $N\overline{y}_1 > [<] n^*$. Thus, uniqueness can be shown as before.

Proposition 2 Consider the fictitious game in which Nature's action space is restricted to $\mathcal{S}'$. Let (δ', σ') be a Nash equilibrium. Conditional on σ', δ^* as defined in the proposition coincides with δ' and is, therefore, a best response to it as well. But on the other hand, $R(\delta^*, s)$ depends on s only through (μ_0, μ_1), implying that $R(\delta^*, s) = R(\delta^*, s')$. This is obvious for the direct effect of s; as to its indirect effect via $\delta^*(s_N)$, it is achieved by stage (i) in the construction of δ^*. Since $\mathcal{S}' \subseteq \mathcal{S}$, it follows that σ' is a best response to δ^* on $\Delta\mathcal{S}$. Hence, (δ^*, σ') is a Nash equilibrium of the original game.

Proposition 3 Fix a sampling scheme and presume that for every covariate x, $\{(\delta_x^*, \sigma_x^*)\}$ is a Nash equilibrium of the appropriately restricted game. Assume now that the decision maker uses δ^* as defined in the proposition. Nature then seeks to maximize

$$R(\delta^*, s) = \sum_{x \in \mathcal{X}} \Pr(X = x) \cdot \Big((\mu_{1x} - \mu_{0x})^+ E(1 - \delta_x^*(s_{Nx})) + (\mu_{0x} - \mu_{1x})^+ E(\delta_x^*(s_{Nx})) \Big),$$

Due to its additive separability, this objective function will be maximized by every element of $\times_{x \in \mathcal{X}} \text{supp}(\sigma_x^*)$. Let σ^* be the prior that has marginals $(\sigma_x^*)_{x \in \mathcal{X}}$ and furthermore renders (Y_{0x}, Y_{1x}) independent of $(Y_{0x'}, Y_{1x'})$ for every $x \neq x'$. Technically, if $\sigma_x^*(s_x)$ denotes the probability that σ_x^* assigns to s_x, then σ^* assigns probability $\sigma^*(s) \equiv \prod_{x \in \mathcal{X}} \sigma_x^*(s_x)$ to $s \equiv (s_x)_{x \in \mathcal{X}}$. This prior is supported on $\times_{x \in \mathcal{X}} \text{supp}(\sigma_x^*)$ and hence a best response to δ^*. Furthermore, since draws from σ^* generate no signals across covariates, δ^* is a best response to σ^*.

Proposition 4 Fix N and assume for simplicity that X is distributed uniformly over $(0, 1]$; this can always be achieved by reparameterization. The proof will be by construction of a sequence of priors $\{\sigma_i\}$ s.t. $\lim_{i \to \infty} \min_{\delta \in \mathcal{D}} \left\{ \int R(\delta, s) d\sigma \right\} = 1/2$. Fix $i \in N$ and define the partition $P_i \equiv \{(0, 1/i], \ldots ((i-1)/i, 1]\}$. Consider the set of mappings $\left\{ f_i^j \right\}_{j=1}^{2^i}$ from $(0, 1]$ into $\{0, 1\}$ that are measurable on P_i. (There are 2^i different such mappings since they can be identified with the set of i-tuples of zeros and ones. Let j index them in an arbitrary order.) Identify the state s_i^j with the degenerate distribution $P(Y_{0x}, Y_{1x})$ concentrated at

$$\begin{aligned} Y_{0x} &= 1_{f_i^j(x) = 1} \\ Y_{1x} &= 1 - Y_{0x}. \end{aligned}$$

Let σ_i be the uniform distribution over states $\left\{ s_i^j \right\}_{j=1}^{2^i}$.

Consider now the Bayes act for σ_i contingent on observation of a sample $\{(t_n, x_n, y_n)\}_{n=1}^N$. To write it down, define P_i^x to be the element of P_i that contains x. Then

$$\delta^* \equiv \begin{cases} 0, & \left[P_i^x \in \cup_{n=1}^N P_i^{x_n}\right] \wedge [x_n \in P_i^x \Rightarrow y_n = 1 - t_n] \\ 1, & \left[P_i^x \in \cup_{n=1}^N P_i^{x_n}\right] \wedge [x_n \in P_i^x \Rightarrow y_n = t_n] \\ 1/2 & \text{otherwise.} \end{cases}$$

is a Bayes act. This is obvious from a verbal description: The sample has revealed $P(Y_{0x}, Y_{1x})$ for all x s.t. P_i^x occurred in the sample, i.e. whenever $P_i^x \in \cup_{n=1}^N P_i^{x_n}$. In this case, δ^* follows the signal in the obvious way. The sample is completely uninformative otherwise; δ^* reflects this by the (arbitrary) tie-breaking probability of 1/2. The expected regret incurred by δ^* can be computed by evaluating $E\left(R(\delta^*, s)\right)$ over any state s_i^j. The numerical result is

$$E\left(R\left(\delta^*, s_i^j\right)\right) = 0 \cdot \Pr\left(P_i^x \in \cup_{n=1}^N P_i^{x_n}\right) + \frac{1}{2} \cdot \Pr\left(P_i^x \notin \cup_{n=1}^N P_i^{x_n}\right) \geq \frac{1}{2} \cdot \frac{i-N}{i} \to \frac{1}{2}$$

as $i \to \infty$. Here, the inequality uses the fact that a sample of size N can be distributed over at most N elements of P_i. This establishes the claim. Notice finally that although I use step functions s_i^j, the result holds if (μ_{0x}, μ_{1x}) is restricted to be continuous in x since the s_i^j could easily be smoothed.

Lemma 1 I show the proof for stratified samples only. From propositions 1(i) and 3, the minimax regret treatment rule without assumption BEC is supported by a worst case prior σ^* that randomizes evenly over $\times_{x \in \mathcal{X}} \{(a_x, 1 - a_x), (1 - a_x, a_x)\}$, where the notation suppresses that $a_x \in [1/2, 1]$ is a function of N_x. The regret incurred conditional on x can be written as

$$\begin{aligned} R(\delta_x^*, \sigma_x^*) &= (a_x - (1 - a_x))E(1 - \delta_x^*(s_{Nx})) \\ &= (2a_x - 1)\left(\sum_{n<N_x/2} B(n, N_x, a_x) + \frac{1}{2}B(n, N_x/2, a_x)\right) \\ &\overset{d}{\to} (2a_x - 1)\,\Phi\left(-N_x^{1/2}\frac{a_x}{1 - a_x}(a_x - 1/2)\right), \end{aligned}$$

where the last step approximates the binomial distribution by a Normal one. Because $a_x > 1/2$, $a_x/(1 - a_x)$ does not vanish. Suppose by contradiction that $\lim_{N_x \to \infty} a_x > 1/2$, then $R(\delta_x^*, \sigma_x^*)$ vanishes at exponential rate (according to the tails of the Normal distribution). In contrast, assume that $a_x - 1/2 = O\left(N_x^{-1/2}\right)$, then $\Phi(\cdot)$ converges to a positive number and $R(\delta_x^*, \sigma_x^*) = O\left(N_x^{-1/2}\right)$, i.e. the convergence is slower. It follows that $\lim_{N_x \to \infty} a_x = 1/2$ and hence $\lim_{N_x \to \infty} (a_x - (1 - a_x)) = 0$. Clearly the result holds uniformly over x, so if $\min_{x \in \mathcal{X}} N_x$ is large enough, then the worst-case prior is consistent with assumption BEC, and the fictitious equilibrium is unchanged.

Lemma 2 Consider the fictitious game in which (i) Nature is restricted to play states which fulfil $(\mu_{x1} - \mu_{x0})(\mu_{x'1} - \mu_{x'0}) \geq 0, \forall x, x' \in \mathcal{X}$, i.e. optimal treatment does not vary across covariates, (ii) the decision maker is restricted to play decision rules with $\delta^*_x(s_N) = \delta^*_{x'}(s_N), \forall x, x' \in \mathcal{X}$. This game has a Nash equilibrium which supports a minimax regret decision rule as described in the lemma, and which has some value $R^* > 0$. It is easy to see that at this equilibrium, (ii) does not constrain the decision maker. It remains to show that for k small enough, (i) does not constrain Nature.

Consider, therefore, a deviation by Nature to a state where the correct treatment differs across covariates. Clearly, the regret such achieved is bounded above by $\max_{x\in\mathcal{X}}\{|\mu_{x1} - \mu_{x0}|\}$. For any $x \in \mathcal{X}$ with $\mu_{x1} \geq \mu_{x0}$, pick $x' \in \mathcal{X}$ with $\mu_{x'1} \leq \mu_{x'0}$ (x' exists due to the presumed nature of the deviation) and write

$$\mu_{x1} - \mu_{x0} = \underbrace{\mu_{x1} - \mu_{x'1}}_{\leq k} + \underbrace{\mu_{x'1} - \mu_{x'0}}_{\leq 0} + \underbrace{\mu_{x'0} - \mu_{x0}}_{\leq k} \leq 2k.$$

Clearly the same bound applies to $\mu_{x0} - \mu_{x1}$ and hence to expected regret, thus Nature will not want to deviate from the original equilibrium if $k \leq R^*/2$.

Proposition 5

(i) I first establish the claim for the setting of proposition 1(ii). The minimax regret value of the game is the value function of Nature's best-response problem at the equilibrium, hence

$$R_1(N) = \max_{\mu_0,\mu_1\in[0,1]} \left\{ (\mu_1 - \mu_0)^+ \left[\Pr(I_N > 0) + \frac{1}{2}\Pr(I_N = 0) \right] \right\},$$

where I can omit the term $(\mu_0 - \mu_1)^+ \left[\Pr(I_N < 0) + \frac{1}{2}\Pr(I_N = 0)\right]$ because of the objective function's symmetry. Recalling also that in equilibrium, $\mu_1 = 1 - \mu_0$, one immediately finds that for odd N,

$$R_1(N) = \max_{a\in[1/2,1]} \{(2a-1)\Pr(I_N > 0)\} = \max_{a\in[1/2,1]} \left\{ (2a-1) \sum_{n=0}^{(N-1)/2} B(n, N, a) \right\}.$$

If N is even, then I_N must be even as well. Consider ignoring the last sample point. If $I_N \neq 0$, this will not affect the treatment rule. If $I_N = 0$, then since the treatment assignment rule is symmetric, this amounts to an even ex ante randomization (although not, in general, an ex post randomization). Hence, the treatment rule that ignores the last sample point achieves minimax regret, hence $R(N) = R(N-1)$ for even N.

It remains to show that the same value applies to the setting of proposition 1(i). To do this, consider the fictitious game used to prove proposition 1. The worst-case prior is an even randomization between two states $s_0 \equiv (a + \Delta, a)$ and $s_1 \equiv (a, a + \Delta)$ for some $\alpha \leq 1/2$ and $\Delta \leq 1 - a$. If one restricts Nature to choose priors that can be parameterized in this form, the constraint on her strategy set does not bind at the equilibrium, hence the game's value is unaffected. I therefore analyze this restricted game.

Consider the signal generated by any pair of sample points that have been assigned different treatments. This signal can be written as (y_0, y_1) and can take values in $\{(0,0), (1,1), (0,1), (1,0)\}$. The first two of these induce the same likelihood ratio between s_0 and s_1 – namely, 1 – and therefore the same updating from any prior – namely, none. Hence, $\theta \equiv y_1 - y_0$ is a sufficient statistic for (y_0, y_1), and the game's equilibrium path is unaffected by replacing the signal $(y_{0n}, y_{1n})_{n=1}^{N/2}$ with $(\theta_n)_{n=1}^{N/2}$.

The distribution of θ_n is characterized as

$$(\Pr(\theta = -1), \Pr(\theta = 0), \Pr(\theta = 1)) = (a(1-a-\Delta), 1 - a(1-a-\Delta) - (a+\Delta)(1-a), (a+\Delta)(1-a)).$$

Evaluation of first- and second-order conditions shows that both $\Pr(\theta = -1)$ and $\Pr(\theta = 1)$ are maximized, and $\Pr(\theta = 0)$ is minimized, when $a = (1-\Delta)/2$. In particular, the minimized value of $\Pr(\theta = 0)$ equals $(1-\Delta^2)/2$. Now consider the following, further manipulation of the game: θ_n is replaced by θ'_n, which is distributed as θ_n except that

$$\left(\theta'_n | \theta_n = 0\right) \equiv \begin{cases} 0 & \text{with probability } \frac{(1-\Delta^2)}{2\Pr(\theta_n = 0)} \\ 1 & \text{with probability } \frac{1}{2}\left(1 - \frac{(1-\Delta^2)}{2\Pr(\theta_n=0)}\right) \\ -1 & \text{with probability } \frac{1}{2}\left(1 - \frac{(1-\Delta^2)}{2\Pr(\theta_n=0)}\right) \end{cases},$$

where all these probabilities lie on $[0,1]$ due to arguments immediately preceding the expression. (Intuitively, the neutral signal is sometimes replaced by a coin toss.) For any (a, Δ), the regret incurred by the decision maker's best response increases through this manipulation. To see this, fix any (a, Δ), then expected regret under the best response is the product of Δ and the probability that the decision maker will incorrectly infer which of $\{s_0, s_1\}$ was realized. But the latter probability increases as the signal $(\theta_n)_{n=1}^{N/2}$ is garbled. It follows that the manipulated game has a higher regret value than the original game.

I now solve the manipulated game. Straightforward algebra reveals that for any (a, Δ), the full distribution of θ'_n is as follows:

$$\theta'_n \equiv \begin{cases} 0 & \text{with probability } \frac{(1-\Delta^2)}{2} \\ 1 & \text{with probability } \left(\frac{1+\Delta}{2}\right)^2 \\ -1 & \text{with probability } \left(\frac{1-\Delta}{2}\right)^2 \end{cases}.$$

This distribution does not depend on a. Without loss of generality, I can therefore restrict nature to choose $a = \frac{1-\Delta}{2}$. But now, it is easily verified that a Nash equilibrium for this game is given by the objects (δ^*, σ^*) from proposition 2. It follows that the game's value is R_1. Recalling that the manipulated game must have a higher value than the original one, it follows that the minimax regret value for proposition 1 is bounded from above by R_1. But Nature can enforce R_1 in that game by playing δ^*; thus, its value is also bounded from below by R_1.

(ii)

Straightforward.

(iii)

This follows from proposition 3 just like part (i) follows from proposition 2 (without any adjustments for N being even).

(iv)

Straightforward.

B Explanation of Tables 3-6

In this section, I elaborate how the lower bounds on regret $\underline{R}$, displayed in tables 3 and 4, have been computed.

Consider the problem of maximizing regret against an empirical success rule,

$$\max_{P(Y_0,Y_1)\in\Delta[0,1]^2}\left\{(\mu_1-\mu_0)\left(\Pr\left(\overline{y}_0>\overline{y}_1\right)+\frac{1}{2}\Pr(\overline{y}_0=\overline{y}_1)\right)\right\}.$$

A full treatment of this problem appears very hard. But its value can be bounded from below by searching over a restriction of $\Delta[0,1]^2$. In particular, I consider the following cases:

- Both treatments induce Bernoulli distributions, i.e. the search space is restricted to $\Delta\{0,1\}^2$; similar to previous arguments, the search parameters are then just (μ_0,μ_1).

- Treatment 0 induces a degenerate distribution concentrated at some point μ_0, whereas treatment 1 has a Bernoulli distribution. This case has two subcases according as $\mu_1>[<]\mu_0$.

All of these cases incur an open set problem. This is easily seen when both outcomes are binary. In case that $\overline{y}_0=\overline{y}_1$, the decision rule will assign the correct treatment half the time. Regret can, therefore, be increased by letting the better of the two treatments be supported (with unchanged probabilities) on $\{0,1-\varepsilon\}$ rather than $\{0,1\}$. By letting $\varepsilon\to0$, one can approximate the effect of a decision rule whose tie-breaking goes in the wrong direction. A similar problem exists with case 2. To generate well-behaved problems, I therefore rig the tie-breaking rule against the decision maker in all cases. As a result, I strictly speaking establish that the bounds displayed in the tables are approachable as opposed to achievable.

With these remarks in mind, alternative lower bounds on the regret can be computed as follows. In these expressions, B stands for binomial probabilities as above, and F represents the binomial c.d.f., i.e. $F(n,N,\mu)\equiv\sum_{i=0}^{n}B(i,N,\mu)$.

- Case 1: $\max_{\mu_0,\mu_1\in[0,1]}\left\{(\mu_1-\mu_0)\sum_{n=0}^{N}B(n,N,\mu_0)\cdot F(n,N,\mu_1)\right\}$.

- Case 2, first subcase: $\max_{\mu_0,\mu_1\in[0,1]}\{(\mu_1-\mu_0)\,F(\mu_0N,N,\mu_1)\}$.

- Case 2, second subcase: $\max_{\mu_0,\mu_1\in[0,1]}\{(\mu_0-\mu_1)\,(1-F(\mu_0N-1,N,\mu_1)\}$.

Table 3 is generated by numerical evaluation of all of these. The two subcases of case 2 turn out to yield identical regrets. Cases 1 and 2 coincide for $N = 1$. Otherwise, case 2 binds, i.e. yield the highest regret, for $N = 2$, and case 1 binds thereafter.

Table 4 only uses case 1, but with the variation that the support of Y_1 is generalized to $\{0, x\}$. The idea here is the following: If one evaluates case 1 only, then $\underline{R}$ sometimes sharply increases in N; for example, this occurs if $\mu_0 = 0.25$ and N moves from 3 to 4. In the specific example, this happens because for $N = 4$, treatment 1 will be rejected even if one success is recorded. For any given Bernoulli distributed treatment, the probability of it being adopted therefore jumps downward, which means that for any such treatment with parameter exceeding μ_0 – i.e. one that should, in fact, be adopted –, regret increases.

This observation spawns an intuition: Perhaps for $N = 3$ and $\mu_0 = 0.25$, one might want to consider distributions Y_1 supported on $\{0, 3/4\}$ because in this case, one success in 4 trials will still lead to rejection. Indeed, this is how the according cell of table 4b has been found. More generally, I search over some salient guesses of the upper support point and also execute an algorithm in which this variable is handed down to the maximizer.

References

[1] Berger, J.O. (1985[1980]): *Statistical Decision Theory and Bayesian Analysis (2nd Edition).* Berlin, New York: Springer Verlag.

[2] Brock, W.A. (2004): "Profiling Problems with Partially Identified Structure," mimeo, University of Wisconsin.

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

[4] Hayashi, T. (2006): "Regret Aversion and Opportunity-dependence," mimeo, University of Texas-Austin.

[5] Hirano, K. and J.R. Porter (2005): "Asymptotics for Statistical Treatment Rules," mimeo, University of Arizona and University of Wisconsin-Madison.

[6] Le Cam, L. (1986): *Asymptotic Methods in Statistical Decision Theory.* Berlin, New York: Springer Verlag.

[7] Manski, C.F. (2004): "Statistical Treatment Rules for Heterogeneous Populations," *Econometrica* 72: 1221-1246.

[8] — (2005): *Social Choice with Partial Knowledge of Treatment Response.* Princeton, Oxford: Princeton University Press.

[9] — (2006a): "Minimax-Regret Treatment Choice with Missing Outcome Data," forthcoming in the *Journal of Econometrics*.

[10] — (2006b): "Search Profiling with Partial Knowledge of Deterrence," forthcoming in the *Economic Journal*.

[11] Manski, C.F. and A. Tetenov (2006): "Admissible Treatment Rules for a Risk-averse Planner with Experimental Data on an Innovation," forthcoming in the *Journal of Statistical Planning and Inference*.

[12] Milnor, J. (1954): "Games Against Nature," in R.M. Thrall, C.H. Coombs, R.L. Davis (Eds.), *Decision Processes*. New York: Wiley.

[13] Savage, L.J. (1951): "The Theory of Statistical Decisions," *Journal of the American Statistical Association* 46: 55-67.

[14] — (1954): *The Foundations of Statistics*. New York: Wiley.

[15] Schlag, K.H. (2003): "How to Minimize Maximum Regret under Repeated Decision-Making," mimeo, European University Institute.

[16] — (2006): "Eleven," mimeo, European University Institute.

[17] Stoye, J. (2005): "Statistical Decisions under Ambiguity," mimeo, New York Unversity.

[18] — (2006a): "Minimax Regret Treatment Choice with Incomplete Data and Many Treatments," forthcoming in *Econometric Theory*.

[19] — (2006b): "Axioms for Minimax Regret Choice Correspondences," mimeo, New York University.