



Minimax regret treatment choice with finite samples

Jörg Stoye*

Department of Economics, New York University, 19 W. 4th Street, New York, NY 10012, United States

ARTICLE INFO

Article history:

Received 18 October 2006
 Received in revised form
 30 January 2009
 Accepted 26 February 2009
 Available online 20 March 2009

JEL classification:
 C44

Keywords:

Finite sample theory
 Statistical decision theory
 Minimax regret
 Treatment response
 Treatment choice

ABSTRACT

This paper applies the minimax regret criterion to choice between two treatments conditional on observation of a finite sample. The analysis is based on exact small sample regret and does not use asymptotic approximations or finite-sample bounds. Core results are: (i) Minimax regret treatment rules are well approximated by empirical success rules in many cases, but differ from them significantly – both in terms of how the rules look and in terms of maximal regret incurred – for small sample sizes and certain sample designs. (ii) Absent prior cross-covariate restrictions on treatment outcomes, they prescribe inference that is completely separate across covariates, leading to no-data rules as the support of a covariate grows. I conclude by offering an assessment of these results.

© 2009 Elsevier B.V. All rights reserved.

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 exposed to one treatment each. This problem was recently analyzed by Manski (2004). The difference to Manski's approach is technical: I consider several extensions and, more importantly, base the analysis entirely on exact small sample regret as opposed to large deviation bounds. This adjustment qualitatively affects substantive results. It also illustrates the potential for closed-form small sample analysis in problems of this type.

Minimax regret as a criterion for treatment choice has recently attracted renewed interest (Brock, 2006; Eozenou et al., 2006; Hirano and Porter, 2008; Manski, 2004, 2005, 2006, 2007a,b, 2008; Schlag, 2006; Stoye, 2007a, 2009). Unfortunately, derivation of finite sample minimax regret decision rules appears extremely hard. As a result, most of the existing literature either focuses on identification and altogether abstracts from sampling uncertainty (Brock, 2006; Manski, 2006, 2008; Stoye, 2007a), states the finite sample problem without attempting to solve it (Manski, 2007a, Section 4), derives bounds on finite sample regret (Manski,

2004), or estimates minimax regret treatment rules (Manski (2007a,b); Stoye (2009)). To my knowledge, the only exact results for finite samples so far are found in related work by Canner (1970) and Schlag (2006), in Manski's (2007a Section 5) analysis of a case that he calls "curiously simple", and in his brute force numerical analysis of the setup considered in Proposition 1(iii) (2005, chapter 3).¹

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 refine those of Manski (2004) in ways that might be considered surprising, or even controversial. The results also allow one to improve numerical analyses presented in Manski (2004) and to gauge the similarity of small-sample decision problems to limit experiments as in Hirano and Porter (2008).

The paper is structured as follows. After setting up the notation and explaining minimax regret, I analyze treatment choice without covariates, differentiating the analysis 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 those informed by classical statistics.

* Corresponding address: Department of Economics, New York University, 269 Mercer Street, 10003 New York, NY, United States. Tel.: +1 212 9928682.

E-mail address: j.stoye@nyu.edu.

URL: <http://homepages.nyu.edu/~js3909>.

¹ Results that are subsequent to earlier versions of this paper are acknowledged in the conclusion.

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). The core result here may be the most surprising one, and refines Manski's (2004) finding in a way that overturns its interpretation. Specifically, in the setting considered by Manski (2004) and here, minimax regret completely separates inference across covariates for any sample size, leading to no-data rules as the support of a covariate grows large. This result will be established in Section 4 of the paper. Section 5 concludes with reflections on some interesting features of the results. All proofs are collected in an Appendix. A web Appendix on the author's homepage contains additional numerical results, including exact counterparts of bounding analyses in Manski (2004).

2. Setting the stage

2.1. The decision problem

The decision problem is as in Manski (2004), and notation is largely his, with slight modifications to align it with the literature on (statistical) decision theory. A decision maker has to 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 $y^j(t) : \{0, 1\} \rightarrow [0, 1]$ that maps treatments onto outcomes. Substantively, I therefore assume 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 population is a probability space (J, Σ, 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 $P(y^j(t))$. (Covariates will be introduced later.)

It will be instrumental to focus on the distribution $P(Y_0, Y_1)$ as unknown quantity. Specifically, $P(Y_0, Y_1)$ will be identified with a *state of the world* s , and the set \mathcal{S} will collect all states of the world that are considered feasible. I will analyze both a situation of complete ignorance and the problem of testing an innovation, in which the behavior of treatment 0 is well understood. Formally, complete ignorance means that $\mathcal{S} = \Delta[0, 1]^2$, the set of distributions over $[0, 1]^2$; testing an innovation means that $\mathcal{S} = \{Q(Y_0, Y_1) \in \Delta[0, 1]^2 : Q(Y_0) = P(Y_0)\}$, where $P(Y_0)$ is known. Further restrictions on potential outcome distributions could be imposed by restricting \mathcal{S} ; such analysis is undertaken in ongoing work.

If s were known, the decision maker would face a decision problem under risk. Assume that she would resolve this problem by maximizing expected outcome, 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 $\mu_0 = \mu_1$, where $\mu_t \equiv \mathbb{E}Y_t$. This does not presume risk neutrality because Y_t might be a utility; it does, however, presume a utilitarian social welfare function.

The decision maker observes treatment outcomes experienced by a random sample of N members of the treatment population. This statistical experiment generates a sample space $\Omega \equiv (\{0, 1\} \times [0, 1])^N$ with typical element $\omega = (t_n, y_n)_{n=1}^N$. The sampling distribution of T depends on the *sample design*, and different such designs will be considered. Conditional on a realization t_n, y_n is an independent realization of Y_{t_n} and, therefore, informative about s .

Treatment choice may condition on the outcome of the statistical experiment. Thus, the decision maker can specify a *statistical treatment rule* $\delta : \Omega \mapsto [0, 1]$ 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. In words, $\delta(\omega)$ specifies the probability with which treatment

1 will be assigned to members of the treatment population if the sample is ω . 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} .

The expected outcome generated by δ given s is

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

i.e. an average of μ_0 and μ_1 , weighted according to the probability that treatment 1 will be assigned. Seen as a function of s , $u(\delta, s)$ is (the negative of) the *risk function* of treatment rule δ . If s were known, the decision problem would be easy – the decision maker would, by assumption, use the no-data rule that assigns the better treatment independently of ω . But with s unknown, one now encounters a decision problem under ambiguity: Different treatment rules will be best for different states s , and there is no obvious probability distribution according to which different states should be weighted.²

Many decision criteria have been suggested for this situation. The two most prominent ones are the Bayesian approach, i.e. to place a subjective distribution on \mathcal{S} and then rank decision rules by the according expectation of $u(\delta, s)$, and maximin utility, i.e. to rank decision rules according to $\min_{s \in \mathcal{S}} u(\delta, s)$. In contrast to either, I follow Manski (2004) and other aforementioned references and evaluate treatment rules by their minimax regret. To understand this criterion, first define the regret incurred by decision rule δ in state s ,

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

the difference between the expected outcome induced by δ and the outcome that could have been achieved if s had been known. A minimax regret decision maker will minimize this quantity over all possible states, i.e. she will pick

$$\delta^* \in \arg \min_{\delta \in \mathcal{D}} \max_{s \in \mathcal{S}} R(\delta, s). \tag{1}$$

Minimax regret was originally introduced by Savage's (1951) reading of Wald (1950). Its recent reconsideration in the treatment choice literature is due to Manski (2004); see also Berger (1985, chapter 5) for a statistician's discussion. An in-depth historical as well as axiomatic discussion of minimax regret is found in Stoye (2007b); see also Hayashi (2008) and Stoye (2007c). Readers who are interested in extensive motivations of minimax regret are referred to this literature. Three brief remarks are as follows:

- Minimax regret has in common with maximin utility that it avoids the explicit use of priors. Whether this is an advantage or a weakness is a judgment call that will be avoided here. It is worth noting, though, that minimax regret implicitly selects a prior, hence one could think of it as a prior selection device motivated by a specific notion of uniform quality of decisions.
- Minimax regret differs markedly from maximin utility by fulfilling the von Neumann–Morgenstern independence axiom. At the same time, it is menu dependent: Adding decision rules to \mathcal{D} can affect the relative ranking of previously available ones, intuitively because it can alter the benchmark against which regret is evaluated.

² This problem was connected to the literature on ambiguity by Manski (2000). Except for a difference in labels, the risk function (interpreted as function of s) is the expected utility functional u of from Stoye's (2007b) axiomatization of maximin utility and minimax regret as well as Gilboa and Schmeidler's (1989) axiomatization of multiple prior maximin utility.

- Related papers (Hirano and Porter, 2008; Manski, 2004; Schlag, 2003) emphasize that the maximin utility criterion leads to trivial results in treatment choice problems.³ They advertise minimax regret as a prior-free criterion that avoids this feature. The triviality result obtains here as well. Under complete ignorance, every decision rule achieves maximin utility because $\min_{s \in \mathcal{S}} u(\delta, s) = 0$ for all δ , generated by the distribution $P(Y_0, Y_1)$ that is entirely concentrated at $(0, 0)$. For testing an innovation, the no-data rule $\delta \equiv 0$ achieves maximin utility, generated by the distribution $P(Y_1)$ that is concentrated at (0) . However, the conclusion regarding minimax regret is only partially confirmed. In fact, the present paper contains the first (to my knowledge) realistic example in which minimax regret is shown to admit a no-data rule.

2.2. The game theoretic approach

Exactly solving (1) appears extremely hard, mainly because evaluation of $R(\delta, s)$ requires integration over finite sample distributions. This difficulty is indirectly illustrated by much of the related literature: Manski (2004) works with finite sample bounds on $R(\delta^{ES}, s)$, where δ^{ES} is the empirical success rule, to be defined below; he does not assert that δ^{ES} achieves either small sample or asymptotic minimax regret. Manski (2007a,b) and Stoye (2009) estimate minimax regret treatment rules by sample analogs; a justification for this in terms of asymptotic efficiency is given by Hirano and Porter (2008). However, many exact results can be generated by framing the decision problem as a statistical game, a technique that will be explained in this section.

Consider the following simultaneous move zero-sum game. The decision maker picks a statistical decision rule $\delta \in \mathcal{D}$; Nature picks a state of the world s . Both agents may randomize, and Nature's mixed strategies will be designated by $\pi \in \Delta \mathcal{S}$. Since \mathcal{D} is closed under probabilistic mixture, no new notation is needed to accommodate randomization by the decision maker. Nature's payoff (and the decision maker's loss) is given by $R(\delta, s)$, and both players maximize expected payoff. Assume that the game has a Nash equilibrium (δ^*, π^*) , then the following facts are well known⁴: (i) δ^* is a minimax regret treatment rule. The distribution π^* , which can be of independent interest, is usually called *least favorable prior*. (ii) Any minimax regret decision rule δ' is a best response to π^* , and any least favorable prior π' is a best response to δ^* . In particular, if (δ', π') is a Nash equilibrium as well, then so are (δ', π^*) and (δ^*, π') .

The upshot of (i) is that optimality of δ^* can be established by "guessing and verifying" (δ^*, π^*) . Given a guess of (δ^*, π^*) , one merely needs to verify Nature's best-response condition

$$s^* \in \arg \max_{s \in \mathcal{S}} R(\delta^*, s), \quad \forall s^* \text{ in the support of } \pi^* \quad (2)$$

and the decision maker's best-response condition

$$\delta^* \in \arg \min_{\delta \in \mathcal{D}} \int R(\delta, s) d\pi^*. \quad (3)$$

Condition (3) can be further simplified as follows:

$$\begin{aligned} \arg \min_{\delta \in \mathcal{D}} \int R(\delta, s) d\pi^* &= \arg \min_{\delta \in \mathcal{D}} \int \left(\max_{d \in \mathcal{D}} u(d, s) - u(\delta, s) \right) d\pi^* \\ &= \arg \max_{\delta \in \mathcal{D}} \int u(\delta, s) d\pi^*, \end{aligned}$$

using that $\max_{d \in \mathcal{D}} u(d, s)$ does not depend on δ . Hence, a minimax regret decision maker will behave like a Bayesian with utility function $u(\delta, s)$ and prior π^* . Specifically, let $\mathbb{E}(\cdot|\omega)$ denote posterior expectations induced by prior π^* and data ω , then any minimax regret treatment rule must assign treatment 1 whenever $\mathbb{E}(Y_1|\omega) > \mathbb{E}(Y_0|\omega)$ and treatment 0 if the opposite inequality holds. This condition is usually easy to verify. The proof difficulty, if any, lies in establishing (2); the trick is that this can frequently be done without fully evaluating R or π^* .

The upshot of (ii) is that beyond establishing minimaxity of δ^* , one can frequently show that δ^* is nearly unique, meaning that any minimax regret treatment rule must be a relatively minor modification of it. All in all, the game theoretic approach provides a powerful proof technique. Its fundamental limitation lies in the fact that no advice is given on how to find δ^* .⁵

3. Treatment choice without covariates

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

(i) *Matched pairs*: Let both treatments be unknown. Sample size N is even, and $N/2$ sample points are assigned to either treatment.

(ii) *Random assignment*: Assume again that both treatments are unknown and let within-sample treatment assignment be by independent tosses of a fair coin.

(iii) *Testing an innovation*: Assume that treatment 0 is well understood, i.e. the distribution of Y_0 is known. Obviously, all sample points will be assigned to treatment 1.

3.1. Binary outcomes

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. Hence, the possible states of the world simplify to $\mathcal{S}' = \Delta\{0, 1\}^2$ respectively $\mathcal{S}' = \{P(Y_0, Y_1) \in \Delta\{0, 1\}^2 : \mathbb{E}(Y_0) = \mu_0\}$ with μ_0 known. This restriction allows to isolate some core issues and to generate (essentially) "if and only if"-statements. Specifically, minimax regret treatment rules can be characterized as follows.⁶

Proposition 1. (i) *In the case of matched pairs, minimax regret is achieved by*

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

where \bar{y}_t is a sample average that conditions on $T = t$, and with the understanding that if $N = 0$, then $\delta_1^* = 1/2$. Furthermore, any minimax regret treatment rule must agree with δ_1^* except when $\bar{y}_0 = \bar{y}_1$, and δ_1^* is the unique minimax regret treatment rule that is measurable with respect to $\text{sgn}(\bar{y}_1 - \bar{y}_0)$.

(ii) *In the case of random assignment, let N_t denote the number of sample subjects assigned to treatment t . Then minimax regret is achieved by*

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

⁵ Note that the proof technique requires different probabilistic thinking than the original setting. The original setting is frequentist, hence s is presumed fixed. But within the fictitious game, s is a random variable with distribution π^* (in equilibrium). Consequently, the decision maker's "within-game inference problem" is Bayesian, namely to form beliefs about s given prior π^* and signal ω .

⁶ Part (i) of this result extends, and abbreviates the proof of, a previous finding by Canner (1970, Section 4).

³ This observation was already anticipated, although not in a treatment choice context, by Savage (1954).

⁴ See, for example, Berger (1985, Section 5). The original source for this insight is Wald (1945).

where

$$I_N \equiv N_1(\bar{y}_1 - 1/2) - N_0(\bar{y}_0 - 1/2) \\ \propto [\# \text{ (observed successes of treatment 1)} \\ + \# \text{ (observed failures of treatment 0)} \\ - [\# \text{ (observed successes of treatment 0)} \\ + \# \text{ (observed failures of treatment 1)}]$$

with the understanding that $N_t(\bar{y}_t - 1/2) = 0$ if $N_t = 0$. Furthermore, any minimax regret treatment rule must agree with δ_2^* except when $I_N = 0$, and δ_2^* is the unique minimax regret treatment rule that is measurable with respect to $\text{sgn}(I_N)$.

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

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

where $n^* \in \{1, \dots, N\}$ and $\lambda^* \in [0, 1]$ are characterized as follows⁷:

$$\max_{a \in [0, \mu_0]} (\mu_0 - a) \left[\sum_{n > n^*} \binom{N}{n} a^n (1 - a)^{N-n} \right. \\ \left. + \lambda^* \binom{N}{n^*} a^{n^*} (1 - a)^{N-n^*} \right] \\ = \max_{a \in [\mu_0, 1]} (a - \mu_0) \left[\sum_{n < n^*} \binom{N}{n} a^n (1 - a)^{N-n} \right. \\ \left. + (1 - \lambda^*) \binom{N}{n^*} a^{n^*} (1 - a)^{N-n^*} \right]. \quad (4)$$

Furthermore, any minimax regret treatment rule must agree with δ_3^* except when $N\bar{y}_1 = n^*$, and δ_3^* is the unique minimax regret treatment rule that is measurable with respect to $N\bar{y}_1$. If $N = 0$, then $(n^*, \lambda^*) = (0, 1 - \mu_0)$.

Proposition 1 not only identifies minimax regret rules, but establishes their near uniqueness in the following sense: In every case, any minimax regret rule must agree with δ^* whenever the value of δ^* is 0 or 1; this is because any other minimax regret rule must be Bayes against the same prior. Furthermore, δ^* is always the simplest possible minimax regret rule in the sense that it uses a data-independent tie-breaking rule, whereas any other minimax regret rule would have to rely on additional (but ancillary) sample information.⁸

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

⁷ Expression (4) may look as if λ^* is defined only in terms of n^* . Indeed, (4) can be solved for λ^* given any conjectured n^* . However, the proof shows that generically, $\lambda^* \in [0, 1]$ for exactly one choice of n^* . The exception is that the expression may be solved by $(n^*, 0)$ as well as $(n^* + 1, 1)$, which describe the same decision rule.

⁸ Here is an example: In the setting of part (i), ties could also be broken by sequentially dropping the last matched pair from the sample until δ_1^* is determinate when applied to the remaining sample, with even randomization if ω is reduced to the empty set. The resulting decision function is formally deterministic for many ω where δ_1^* is randomized, but is distributed as δ_1^* in every state s , thus has the same risk function.

Proposition 1 is readily extended to the case where N is a random variable with known distribution. Inspection of the proofs reveals that statements (i) and (ii) would go through unchanged. In part (iii), (n^*, λ^*) would become a tedious function of the realization of N . An explicit presentation is omitted to economize on notation, especially in part (iii).

the “simple empirical success rule” $\delta^{ES} \equiv \mathbb{I}\{\bar{y}_1 > \bar{y}_0\}$.⁹ This is clearly a simplification—in the spirit of the paper, one would want to use a minimax regret treatment rule if it were known. Proposition 1(i) shows that for binary outcomes and matched pairs, δ^{ES} is reasonably close, the modification being that tie-breaking must be symmetric. For randomized treatment assignment, the minimax regret decision rule will asymptotically agree with δ^{ES} but differ from it markedly for small samples and also for specific realizations of rather large ones; see the conclusion for an example.

The characterization of δ_3^* is implicit, but numerical evaluation is easy. Table 1 illustrates the result for a selection of sample sizes N and values of μ_0 .¹⁰ Specifically, the table displays $\alpha \equiv n^* + 1 - \lambda^*$, a smooth index of the treatment rule's conservatism, with higher values indicating more conservative rules: The number to the left of the decimal point is the critical number of observed successes that leads to randomized treatment assignment, and the number to its right gives the probability with which this randomization will pick treatment 0.

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 (e.g., Berger (1985, Chapter 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, with randomization on the threshold to maximize the test's power. The table can be read in exact analogy to Table 1. For example, if $\mu_0 = 0.25$ and $N = 10$, then the minimax regret decision rule prescribes to adopt treatment 1 with probability 0.18 if 2 successes were observed and with probability 1 if more than 2 successes were observed. The hypothesis test will reject H_0 (hence, recommend adoption of treatment 1) with probability 0.52 if 5 successes were observed and with probability 1 if even more successes were observed.

Although it must eventually resemble an empirical success rule, the hypothesis testing rule is much more conservative than minimax regret. The reason is that it emphasizes avoidance of type I errors over avoidance of type II errors. In contrast, the minimax regret rule equalizes regret between two worst-case scenarios that correspond to the two different error types, thus type I and type II errors are treated roughly symmetrically. To be sure, while the classical decision rule will incur high worst-case regret, the minimax regret rule will fail on classical terms. Continuing the above example, simple computations establish that for $\mu_0 = .25$ and $N = 10$, the size of the implicit hypothesis test is 53%. Furthermore, it can be shown that for any $\mu_0 \in (0, 1)$, this size converges to 50% as N grows large.¹¹ Which consideration matters depends on one's objective function; the point here is to demonstrate that the choice makes a big difference.

3.2. General outcome distributions

Now return to the case where (Y_0, Y_1) is distributed arbitrarily on $[0, 1]^2$. Minimax regret treatment rules for this case can be generated from Proposition 1 as follows. Call a state s a Bernoulli state if it implies Bernoulli distributions of both Y_0 and Y_1 (i.e., binary outcomes). Observe that Bernoulli states are fully characterized by triplets $\mathbb{E}(Y_0, Y_1, Y_0Y_1)$. For any state s , call its Bernoulli equivalent the Bernoulli state s' such that s and s' induce

⁹ Here and henceforth, $\mathbb{I}\{E\}$ denotes the indicator function for event E .

¹⁰ This table extends Table 3.1 in Manski (2005).

¹¹ By size of the implicit hypothesis test, I mean the probability of rejecting H_0 – hence adopting treatment 1 – even though H_0 is true, evaluated at $\mu_1 = \mu_0$, which is the parameter value within H_0 that maximizes this probability.

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

μ_0	$N = 1$	$N = 2$	$N = 3$	$N = 4$	$N = 5$	$N = 10$	$N = 20$	$N = 50$	$N = 100$	$N = 500$
.05	0.33	0.48	0.59	0.44	0.74	1.22	1.18	2.68	5.18	25.18
.25	0.64	0.82	1.07	1.33	1.58	2.82	5.32	12.82	25.32	125.32
.50	1.0	1.5	2.0	2.5	3.0	5.5	10.5	25.5	50.5	250.5
.75	1.36	2.18	2.91	3.67	4.42	8.18	15.68	38.18	75.68	375.68
.95	1.67	2.52	3.41	4.56	5.26	9.78	19.82	48.32	95.82	475.82

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

μ_0	$N = 1$	$N = 2$	$N = 3$	$N = 4$	$N = 5$	$N = 10$	$N = 20$	$N = 50$	$N = 100$	$N = 500$
.05	1.0	1.5	1.68	1.79	1.87	2.48	3.43	5.81	9.38	33.78
.25	1.8	2.2	2.76	3.02	3.61	5.48	8.85	18.19	32.78	141.58
.50	1.9	2.8	3.6	4.2	4.88	8.11	14.21	31.35	58.74	268.89
.75	1.93	2.91	3.88	4.84	5.79	10.11	18.62	42.90	82.51	391.29
.95	1.95	2.94	3.94	4.94	5.94	10.92	20.86	50.35	98.84	483.27

the same value of $\mathbb{E}(Y_0, Y_1, Y_0Y_1)$. 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' . Finally, for any decision rule $\delta \in \mathcal{D}$, define $\tilde{\delta} \in \tilde{\mathcal{D}}$ as follows: (i) Replace every observation $y_n \in [0, 1]$ with one independent realization \tilde{y}_n of a Bernoulli variable Y_n with parameter y_n . (ii) Operate δ on $\tilde{\omega} \equiv (t_n, \tilde{y}_n)_{n=1}^N$. In words, $\tilde{\delta}$ is generated from δ by preceding the latter with an information coarsening in which outcome observations other than $\{0, 1\}$ are replaced by independent tosses of biased coins. Perhaps surprisingly, this coarsening does not affect minimax analysis.

Proposition 2.

- (i)

$$\min_{\delta \in \mathcal{D}} \max_{s' \in \mathcal{S}'} R(\delta, s') = \min_{\tilde{\delta} \in \tilde{\mathcal{D}}} \max_{s \in \mathcal{S}} R(\tilde{\delta}, s)$$

$$\delta \in \arg \min_{\delta \in \mathcal{D}} \max_{s' \in \mathcal{S}'} R(\delta, s') \implies \tilde{\delta} \in \arg \min_{\tilde{\delta} \in \tilde{\mathcal{D}}} \max_{s \in \mathcal{S}} R(\tilde{\delta}, s).$$
- (ii) If $\mathcal{S}' \subseteq \mathcal{S}$, then furthermore

$$\min_{\delta \in \mathcal{D}} \max_{s' \in \mathcal{S}'} R(\delta, s') = \min_{\delta \in \mathcal{D}} \max_{s \in \mathcal{S}} R(\delta, s)$$

$$\delta \in \arg \min_{\delta \in \mathcal{D}} \max_{s' \in \mathcal{S}'} R(\delta, s') \implies \tilde{\delta} \in \arg \min_{\tilde{\delta} \in \tilde{\mathcal{D}}} \max_{s \in \mathcal{S}} R(\tilde{\delta}, s).$$

The last statement contains the most important insight: If $\mathcal{S}' \subseteq \mathcal{S}$ and δ^* achieves minimax regret over \mathcal{S}' , then $\tilde{\delta}^*$ achieves minimax regret over \mathcal{S} . This idea is not original to this paper—it was previously and independently discovered by [Cucconi \(1968\)](#), [Gupta and Hande \(1992\)](#), and [Schlag \(2003\)](#). The intuition for why it 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.

If \mathcal{S} equals any one specification from Section 2.1, and \mathcal{S}' is the according specification from Section 3.1, then $\mathcal{S}' \subseteq \mathcal{S}$. Hence, $\tilde{\delta}_1^*$, $\tilde{\delta}_2^*$, respectively $\tilde{\delta}_3^*$ – randomized versions of the rules identified in [Proposition 1](#) – achieve minimax regret for the respective sample designs but with outcomes unrestricted over $[0, 1]$. Some clarifications are in order: First, a price of the generalization is that the “near uniqueness” statements are lost. Second, the disentangling of parts (i) and (ii), which is original to this paper, highlights that the generalization depends on Bernoulli states being possible. Third, the deliberate discarding of information in the binary randomization step raises the possibility of the resulting rules being inadmissible.

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 $\tilde{\delta}_1^*$ to be randomized even if $\bar{y}_0 \neq \bar{y}_1$. As an example, $\delta^{ES}((0, .5), (1, .6)) = 1$ but $\tilde{\delta}_1^*((0, .5), (1, .6)) = 0.55$. Furthermore, \bar{y}_0 and \bar{y}_1 cease to be sufficient statistics for matched pairs. For example, $\tilde{\delta}_1^*((0, 0), (0, 1), (1, 1), (1, 1)) = 1$ but $\tilde{\delta}_1^*((0, 1/2), (0, 1/2), (1, 1), (1, 1)) = 0.875$.

To be sure, these differences vanish as samples become large. For all cases of [Proposition 1](#), one can verify that $\tilde{\delta}_i^*$ is asymptotically equivalent to δ^{ES} for every state s (although not uniformly over \mathcal{S}). Also, it is not intuitively obvious that $\tilde{\delta}_1^*$ is more attractive than δ^{ES} . However one feels about $\tilde{\delta}_1^*$, though, [Proposition 2](#) is of interest because it yields finite sample, nonparametric minimax regret efficiency bounds that apply to all feasible statistical treatment rules. I now turn to these bounds.

3.3. Value of the decision problem

The minimax regret value of a decision problem (that is, the maximal regret incurred by δ^*) equals the fictitious game’s value. Thus, understanding the equilibrium of the fictitious game allows one to compute this value. Let $R_1^*(N)$, $R_2^*(N)$, and $R_3^*(N)$ denote the minimax regret values of the above decision problems; e.g., $R_1^*(N) = \min_{\delta \in \mathcal{D}} \max_{s \in \mathcal{S}} R(\delta, s)$ given matched pairs.¹² Then minimal expansion of the previous proofs yields the following result.

Corollary 1.

$$R_1^*(N) = R_2^*(N) = \max_{a \in [1/2, 1]} \left\{ (2a - 1) \sum_{n < N/2} \binom{N'}{n} a^n (1 - a)^{N' - n} \right\}$$

$$N' = \max_{M \in \mathbb{N}} \{M \leq N : M \text{ is odd}\},$$

where $R_1^*(0) = R_2^*(0) = 1/2$.

$$R_3^*(N) = \max_{a \in [\mu_0, 1]} \left\{ (a - \mu_0) \left[\sum_{n < n^*} \binom{N}{n} a^n (1 - a)^{N - n} + (1 - \lambda^*) \binom{N}{n^*} a^{n^*} (1 - a)^{N - n^*} \right] \right\}$$

with (n^*, λ^*) as in [Proposition 1](#) (iii), and with $R_3^*(0) = \mu_0(1 - \mu_0)$. These expressions apply to both binary and general outcome distributions.

¹² Dependence of R_3^* on μ_0 is suppressed to achieve unified notation.

This corollary establishes finite sample minimax regret efficiency bounds that apply to all treatment rules, it allows one to assess the small sample performance of the empirical success rule δ^{ES} , and it implies the exact gain (in terms of minimax regret) from increasing the sample size. All of these applications are carried out in the web Appendix. An interesting finding is that for general outcomes and very small samples, δ^{ES} incurs at least double the best possible maximal regret.

4. Treatment choice with covariates

This section extends the analysis to treatment choice with covariates. To begin, assume that there exists a finite-valued covariate X supported on $\mathcal{X} = \{x_1, \dots, x_K\}$, where all values of X occur with non-zero probability. The covariate is observable in both the sample data and the treatment population; the decision maker can, therefore, condition treatment choice on covariates. One might 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 small 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 outperform any rule that pools information. This conclusion becomes much more stark under exact analysis: In the natural extension of the previous propositions’ setup, minimax regret is achieved by separating inference across covariates for any sample size.

Formalizing this insight requires some additional notation. Potential outcomes are now random variables Y_{tx} that depend on treatment as well as covariate. A state of the world s is a distribution $P((Y_{0x}, Y_{1x})_{x \in \mathcal{X}})$ with marginals $s_x \equiv P(Y_{0x}, Y_{1x})$; define also the conditional state space $\mathcal{S}_x \equiv \{s_x : s \in \mathcal{S}\}$. A sample ω 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 $Y_{t_n x_n}$. To keep the argument simple, exclude sequential sample designs, thus (T, X) cannot depend on lagged realizations of $Y_{t_n x_n}$. A statistical treatment rule maps samples ω into vectors of treatment assignment probabilities $\delta(\omega) \in [0, 1]^K$, whose components $\delta_x(\omega)$ are identified with probabilities of assigning treatment 1 to subjects with covariate value x . A treatment rule’s risk function is $u(\delta, s) \equiv \sum_{x \in \mathcal{X}} \Pr(X = x) (\mu_{0x}(1 - \mathbb{E}\delta_x(\omega)) + \mu_{1x}\mathbb{E}\delta_x(\omega))$, where $(\mu_{0x}, \mu_{1x}) \equiv \mathbb{E}(Y_{0x}, Y_{1x})$. Regret is $R(\delta, s) \equiv \max_{d \in \mathcal{D}} u(d, s) - u(\delta, s)$ as before. Finally, it will turn out that judicious specification of the state space \mathcal{S} is of utmost importance. The most natural choice, entertained by Manski (2004) and here, is to set $\mathcal{S} = \Delta[0, 1]^{2K}$ (with restrictions on $P((Y_{0x}, Y_{1x})_{x \in \mathcal{X}})$ in the case of testing an innovation). But this section’s results only require the following, weaker richness condition: For any vector $(s'_x)_{x \in \mathcal{X}} \in \times_{x \in \mathcal{X}} \mathcal{S}_x$, there exists a state $s \in \mathcal{S}$ s.t. $(s_x)_{x \in \mathcal{X}} = (s'_x)_{x \in \mathcal{X}}$. In words, the state space is not constrained by cross-covariate restrictions.

To formalize the notion of “no cross-covariate inference,” I need to define a notion of conditional (on a covariate) decision problems. Thus, for any $x \in \mathcal{X}$ and sample realization ω , let $\omega_x \equiv \{(t_n, x_n, y_n) : x_n = x\}$ collect those observations where $x_n = x$. (ω_x may be the empty set.) Let $\Omega_x \equiv \{\omega_x : \omega \in \Omega\}$ collect possible realizations of ω_x . Then $\mathcal{D}_x : \Omega_x \rightarrow [0, 1]$ collects decision rules d_x which map conditional sample realizations ω_x onto probabilities of assigning treatment 1. Substantively, $d_x(\omega_x) \in [0, 1]$ should be thought of as assigning treatment for subjects with covariate value x , all other assignments remaining unspecified. Thus, the decision rules in \mathcal{D}_x are available to a decision maker who only sees sample realizations for a given covariate value x and only needs to assign treatment conditional on that same x . Call this decision problem the conditional problem, and call d_x a conditional decision

rule. Then conditional decision problems can be analyzed just like unconditional ones: They induce conditional expected outcomes $u_x(d_x, s_x) \equiv \mu_{0x}(1 - \mathbb{E}d_x(\omega_x)) + \mu_{1x}\mathbb{E}d_x(\omega_x)$ and conditional expected regret $R_x(d_x, s_x) \equiv \max_{\delta_x \in \mathcal{D}_x} u_x(\delta_x, s_x) - u_x(d_x, s_x)$; the notation here reflects that u_x and R_x depend on s only through s_x . In particular, for a given x , there might exist a conditional minimax regret treatment rule $d_x^* \in \arg \min_{d_x \in \mathcal{D}_x} \max_{s_x \in \mathcal{S}_x} R_x(d_x, s_x)$, and it might be supported by a conditional least favorable prior $\pi_x^* \in \Delta \mathcal{S}_x$.

Any collection $(d_x)_{x \in \mathcal{X}}$ of conditional decision rules can be used to define a statistical decision rule δ that uses no cross-covariate information—just let $\delta_x(\omega) = d_x(\omega_x)$ for every x . This decision rule might appear inefficient in small samples, even if its components make efficient use of conditional information. In other words, even $\delta^*(\omega) = (d_x^*(\omega_x))_{x \in \mathcal{X}}$, with each d_x^* a conditional minimax regret treatment rule as just defined, might not appear compelling in the overall decision problem. This section’s core result is to refute this intuition. In fact, δ^* achieves finite sample minimax regret.

Proposition 3. Assume that d_x^* , supported by least favorable prior $\pi_x^* \in \Delta \mathcal{S}_x$, exists for every $x \in \mathcal{X}$. Then minimax regret is achieved by δ^* as just defined.

I will first give an illustration for the proposition’s use, then an intuition for why it is true, and then discuss it. As to its use, Proposition 3 does not require that $(d_x^*)_{x \in \mathcal{X}}$ is known, but when it is, an explicit minimax regret treatment rule emerges. Furthermore, when “near uniqueness” results are available, the proposition can be strengthened. As a concrete example, here is a decision problem analyzed in Manski’s (2004) discussion of optimal stratification.

Example 1. Let $\mathcal{X} = \{m, f\}$ and let the sample be stratified into 4 cells as follows: $N_m/2 \geq 0$ men receive treatment 0, $N_m/2$ men receive treatment 1, $N_f/2 \geq 0$ women receive treatment 0, and $N_f/2$ women receive treatment 1. Then minimax regret is achieved by δ^* , characterized by first transforming the data as in Proposition 2 and then operating δ^* defined by:

$$\delta_m^*(\omega) \equiv \begin{cases} 0, & \bar{y}_{1m} < \bar{y}_{0m} \\ 1/2, & \bar{y}_{1m} = \bar{y}_{0m} \\ 1, & \bar{y}_{1m} > \bar{y}_{0m} \end{cases}$$

$$\delta_f^*(\omega) \equiv \begin{cases} 0, & \bar{y}_{1f} < \bar{y}_{0f} \\ 1/2, & \bar{y}_{1f} = \bar{y}_{0f} \\ 1, & \bar{y}_{1f} > \bar{y}_{0f} \end{cases}$$

where \bar{y}_{tx} is the sample average that conditions on $(X = x, T = t)$, with the understanding that $\bar{y}_{tx} = 0$ if $N_x = 0$.

If outcomes are binary, then one can directly apply δ^* and furthermore has: (i) Any minimax regret treatment rule must agree with δ^* up to tie-breaking. (ii) A modification of δ^* whereby $[\bar{y}_{1m} = \bar{y}_{0m}] \Rightarrow \delta_m^*(\omega) = \delta_f^*(\omega)$ and $[\bar{y}_{1f} = \bar{y}_{0f}] \Rightarrow \delta_f^*(\omega) = \delta_m^*(\omega)$ fails to achieve minimax regret unless $N_m = N_f = 0$.

These statements hold even if $N_m = 0$ or $N_f = 0$.

Together with Corollary 1, this example can be used to conduct Manski’s (2004) analysis of optimal stratification in terms of exact regret as opposed to large deviations bounds; numerical results are provided in the web Appendix. However, the example also highlights the counterintuitive nature of Proposition 3 and, with binary outcomes, allows for a stronger claim: Any minimax regret treatment rule must agree with δ^* up to tie-breaking, and even the tie-breaking cannot be the commonsensical one, namely according to the comparison of treatment outcomes conditional on the other covariate value.¹³

¹³ Using Proposition 1(ii), the example is easily extended – with the same message – to randomized treatment assignment. By the extension of Proposition 1(ii) to random N , one can furthermore extend it to the case where the sample is a simple random sample from the population (as opposed to being stratified by gender).

To see why the proposition holds, observe that $R(\delta^*, s) = \sum_{x \in \mathcal{X}} \Pr(X = x) R_x(d_x^*, s_x)$. Nature's best response condition (2) is, therefore, fulfilled by any prior π whose conditional distributions π_x coincide with the conditional least favorable priors π_x^* . Among all such priors, let π^* be the one that furthermore renders s_x independent of $s_{x'}$ for any $x \neq x'$, so that ω_x is uninformative about $s_{x'}$. δ^* is Bayes against π^* , thus an equilibrium of the fictitious game has been found. The richness condition on \mathcal{S} is needed to ensure existence of π^* .

Paradoxically, the recommendation based on small sample analysis does not seem desirable in a world of small samples. It requires one to condition on covariates even if this leads to extremely small or empty sample cells. While 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 surely, they would not want to consider an arbitrary covariate that happens to be in the data set.

What's more, Proposition 3 implies that as the support of a covariate grows, minimax regret treatment rules approach no-data rules, because the proportion of covariate values that have been observed in the sample vanishes. Indeed, one can extend the result to a continuous covariate, and a no-data rule then achieves minimax regret. To formalize this, let $\mathcal{X} = [0, 1]$ and assume without further loss of generality that the distribution of X is uniform. Then a decision rule maps the sample space onto decision functions $\delta : [0, 1] \rightarrow [0, 1]$. In this setup, $\max_{s \in \mathcal{S}} R(\delta, s)$ may not be attained, but the following result holds.¹⁴

Proposition 4. Consider the setup for continuous X as just defined. Fix an arbitrary sample size $N < \infty$ and an arbitrary sampling scheme for T and X . Then:

(i) If (μ_{0x}, μ_{1x}) is unknown, then $\min_{\delta \in \mathcal{D}} \sup_{s \in \mathcal{S}} R(\delta, s) = 1/2$. This value is achieved by $\delta_x^*(\omega) = 1/2, \forall \omega$.

(ii) If μ_{0x} is a known, Lebesgue measurable function on $[0, 1]$, then $\min_{\delta \in \mathcal{D}} \sup_{s \in \mathcal{S}} R(\delta, s) = \int \mu_{0x}(1 - \mu_{0x}) dx$. This value is achieved by $\delta_x^*(\omega) = 1 - \mu_{0x}, \forall \omega$.

To be sure, Proposition 4 does not say that no-data rules outperform sensible data-dependent ones with respect to pointwise regret. In fact, they are weakly dominated and hence inadmissible. Two examples of dominating rules are to use the no-data rule except if a covariate value observed in the sample ever recurs exactly, in which case the according sample information should be used; or (as pointed out by a referee) to completely ignore the covariate and use δ_1^* , δ_2^* , or δ_3^* if they apply to the resulting problem. However, no treatment rule outperforms the no-data rule uniformly over \mathcal{S} and, therefore, not in terms of the worst-case analysis that informs minimax regret. Intuitively, the problem is that (μ_{0x}, μ_{1x}) respectively μ_{1x} can be infinitely "wiggly" functions of x , so that the amount of information revealed by a sample of any given size N cannot be bounded away from zero. This intuition relates Proposition 4 to classic impossibility results regarding uniform performance of statistical decision procedures (e.g., Bahadur and Savage (1956)).

5. Concluding remarks

This paper contributes to a rather young literature in which a rather old criterion is applied to models of real-world decisions. I characterized finite sample minimax regret rules for the scenario analyzed by Manski (2004) and variations thereof. Important results include the comparison between exact minimax regret rules and empirical success rules or rules informed by classical statistics. Perhaps most interestingly, the analysis of covariates leads to a reassessment of Manski's (2004) findings on the subject. On a more general level, it was seen that exact analysis is more feasible than might have been anticipated and can generate some significant and rather general insights.

Numerous extensions of the results might be of interest, and some already exist. For example, I mentioned the generalization to N being a random variable with known distribution. The case where N is ambiguous (i.e., its distribution is unknown) is not interesting because regret is trivially maximized by minimizing N . Proofs are also easily extended to show that if the sample design itself is a choice variable, then all three designs considered in Proposition 1 are minimax regret optimal under their respective information assumptions. Further results on endogenous sample design are found in Schlag (2006). Subsequently to research reported here, Manski and Tetenov (2007) generalized Proposition 1(iii) to regret functionals based on nonlinear transformations of $u(\delta, s)$; Tetenov (2007) generalized it to asymmetric regret functionals that differentiate between type I and type II errors; and Stoye (2007d) solved certain missing data problems. Other questions remain open, prominently among them the extension to multivalued treatments.

A natural question is whether the minimax regret criterion is an attractive alternative to existing approaches. There are two ways to investigate this: One is axiomatic analysis, the other one is to look whether actual minimax regret rules make sense. The former has been done elsewhere (Hayashi, 2008; Milnor, 1954; Stoye, 2007b,c). The latter was one purpose of this paper, and the results raise some concerns. Consider specifically the findings on covariates. Manski (2004) discovered that for surprisingly small sample sizes, lower bounds on expected regret incurred by pooling the sample exceed upper bounds incurred by conditioning on covariates. This finding was used to criticize prevailing practice, tentatively suggesting that there is too much pooling of observations across covariate values. However, it is now seen that the result was merely an approximation of a much stronger, and pathological, one. The exact finding cannot any more inform a critique of prevailing practice, but rather raises questions about minimax regret. What's more, the existence of easy examples in which the maximin utility criterion generates no-data rules is frequently used to argue against maximin utility, and sometimes in favor of minimax regret as an alternative. Proposition 4 undermines this line of reasoning as well.

I will conclude by offering some thoughts on where the problem lies. Specifically, a minimax regret decider will act as if she had probabilistic beliefs described by the least favorable prior π^* . This need not mean that she actually believes π^* ; minimax regret may serve as pragmatic prior selection device for users who are reluctant to specify informative priors. But it means that to understand the workings of minimax regret treatment rules, it can be instructive to investigate π^* .

This is especially salient with respect to Proposition 3. In those applications where completely ignoring cross-covariate information sounds absurd, it is noted that the corresponding prior – specifically, $P(Y_{1x})$ and $P(Y_{1x'})$ may be extremely different – appears excessively conservative with respect to what can be learned from data. Perhaps the problem can be alleviated by properly specifying available prior information. The most obvious

¹⁴ I avoid an explicit definition of \mathcal{S} here to minimize notational burden. The proof only requires the following, minimal richness: For every finite, Lebesgue measurable step function $f : [0, 1] \rightarrow [0, 1]^2$, there exists a state $s \in \mathcal{S}$ s.t. $(\mu_{0x}, \mu_{1x}) = f(x)$ for all x .

way to do this is to restrict δ ; for example, appropriate restrictions on $(P(Y_{0k}))_{k \in \mathcal{X}}$ might exclude states in the support of π^* . Users who explicitly use minimax regret as a prior selection device could also make the set of admissible priors a subset of $\Delta\delta$ and directly exclude π^* . Ongoing research indicates that some natural such restrictions lead to reasonable minimax regret treatment rules. However, even when they are substantively uncontroversial, and although they may avoid probabilistic language, these restrictions will typically be of a subjective nature. Introducing them may lead to convergence between the frequentist methods analyzed here and approaches that are “robust Bayesian” in the sense of admitting multiple priors (e.g., Berger (1985)).

Proposition 1(ii) poses a more subtle challenge to intuitions and may appear reasonable to many readers. But consider the following example (due to Charles Manski): The randomized treatment design was applied to a sample of size 1100, and 1000 subjects 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 *any* minimax regret treatment rule prescribes to assign *all* future subjects to treatment 0. This conclusion is less than obvious, and 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 contexts, maximin utility generates trivial decision rules. The reason is that it optimizes against a least favorable prior π^* 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 π_2^* (“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 realization which renders this prior empirically implausible. If Proposition 1(ii) is seen as a problem, then its cause may lie in minimax regret’s selective ignorance of likelihoods, or in other words, in the fact that the least favorable prior is typically 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 that other criteria have significant downsides as well. But obviously, these should not keep proponents of minimax regret from being candid about potential drawbacks of this criterion.

Acknowledgments

I thank Gary Chamberlain, Bo Honoré, Chuck Manski, Karl Schlag, seminar participants at Cornell, Harvard/MIT, Mannheim, Munich, Northwestern, Nuffield College, NYU, Pittsburgh, Princeton, Queen’s, and UCL, conference participants at EUI and Minnesota, as well as a co-editor and three referees for helpful comments. Of course, any and all errors are mine. 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.

Appendix. Proofs

A.1. Proposition 1

Preliminaries.

Observe the following simplifications: $R(\delta, s)$ can be expressed as

$$R(\delta, s) = \max \{ \mu_0, \mu_1 \} - (\mu_0 \mathbb{E}(1 - \delta(\omega)) + \mu_1 \mathbb{E}\delta(\omega)) \tag{5}$$

$$= (\mu_1 - \mu_0)^+ \mathbb{E}(1 - \delta(\omega)) + (\mu_0 - \mu_1)^+ \mathbb{E}\delta(\omega), \tag{6}$$

where $Y^+ \equiv \max\{Y, 0\}$ is the positive restriction of Y . Because ω collects independent realizations of Y_t , this expression depends on $P(Y_0, Y_1)$ only through the marginal distributions $P(Y_0)$ and $P(Y_1)$. Since Y_0 and Y_1 are binary, $P(Y_0)$ and $P(Y_1)$ are characterized by (μ_0, μ_1) . In this proof, I therefore identify states with couplets $(\mu_0, \mu_1) \in [0, 1]^2$. Also, define n_t as the number of successes recorded for treatment t , thus $n_t = N/2 \cdot \bar{y}_t$ in part (i), $n_t = N_t \bar{y}_t$ in part (ii), and $n_1 = N \bar{y}_1$ in part (iii). Finally, I follow game theoretic conventions by using asterisks to denote the equilibrium values of choice parameters.

(i) I will show that the fictitious game has a Nash equilibrium (δ_1^*, π_1^*) , where δ_1^* is as defined in the proposition and where π_1^* is the uniform distribution over $\{(a, b), (b, a)\}$ for some constants (a, b) with $a > b$. (Computation of the constants leads to Corollary 1 but is not needed here.) Notice that in the definition of δ_1^*, \bar{y}_t can be replaced with n_t as just defined.

Nature’s best-response condition (2) is fulfilled if π_1^* is supported on

$$\begin{aligned} & \arg \max_{s \in [0, 1]^2} R(\delta_1^*, s) \\ &= \arg \max_{(\mu_0, \mu_1) \in [0, 1]^2} \left\{ (\mu_1 - \mu_0)^+ \left(\Pr(n_0 > n_1) + \frac{1}{2} \Pr(n_0 = n_1) \right) \right. \\ & \quad \left. + (\mu_0 - \mu_1)^+ \left(\Pr(n_1 > n_0) + \frac{1}{2} \Pr(n_1 = n_0) \right) \right\}. \end{aligned} \tag{7}$$

Here, the probabilities refer to the sampling distribution of (n_0, n_1) , e.g.

$$\begin{aligned} \Pr(n_0 > n_1) &= \sum_{n_1=0}^{N/2-1} \sum_{n_0=n_1+1}^{N/2} \binom{N/2}{n_0} \\ & \quad \times \mu_0^{n_0} (1 - \mu_0)^{N/2-n_0} \binom{N/2}{n_1} \mu_1^{n_1} (1 - \mu_1)^{N/2-n_1}. \end{aligned}$$

The objective is continuous and the feasible set is compact, hence the arg max is nonempty. Furthermore, the objective is symmetrical in μ_0 and μ_1 , hence the arg max contains (a, b) iff it contains (b, a) . This establishes existence of π_1^* .

It remains to show that δ_1^* is Bayes against any such π_1^* . This requires that $\delta_1^*(\omega) = 1$ if

$$\begin{aligned} & \mathbb{E}(Y_1 | \omega) > \mathbb{E}(Y_0 | \omega) \\ & \iff \frac{\frac{1}{2} a \Pr(\omega | s = (b, a)) + \frac{1}{2} b \Pr(\omega | s = (a, b))}{\frac{1}{2} \Pr(\omega | s = (b, a)) + \frac{1}{2} \Pr(\omega | s = (a, b))} \\ & \quad > \frac{\frac{1}{2} b \Pr(\omega | s = (b, a)) + \frac{1}{2} a \Pr(\omega | s = (a, b))}{\frac{1}{2} \Pr(\omega | s = (b, a)) + \frac{1}{2} \Pr(\omega | s = (a, b))} \\ & \iff \Pr(\omega | s = (b, a)) > \Pr(\omega | s = (a, b)) \\ & \iff \binom{N/2}{n_0} b^{n_0} (1 - b)^{N/2-n_0} \binom{N/2}{n_1} a^{n_1} (1 - a)^{N/2-n_1} \\ & \quad > \binom{N/2}{n_0} a^{n_0} (1 - a)^{N/2-n_0} \binom{N/2}{n_1} b^{n_1} (1 - b)^{N/2-n_1} \\ & \iff \left(\frac{a}{1 - a} \right)^{n_1 - n_0} > \left(\frac{b}{1 - b} \right)^{n_1 - n_0} \\ & \iff n_1 - n_0 > 0. \end{aligned}$$

Here, the expectations are posterior expectations induced by prior π_1^* and data ω . The first step expands these by means of standard

Bayesian formulas, and the remaining steps are algebra. Similarly, $\delta_1^*(\omega)$ must be 0 whenever $n_1 - n_0 < 0$ and can be arbitrary if $n_1 - n_0 = 0$, including the case that $N = 0$.

This establishes the Nash equilibrium. As any minimax regret decision rule must be Bayes against π_1^* , δ_1^* is unique whenever it is a strict best response, that is, except when $n_0 = n_1$. Restrict attention to decision rules that depend only on (n_0, n_1) , then it follows that δ_1^* is unique up to the tie-breaking probability. Assume this probability favors treatment 0 [1], then Nature will want to deviate to the pure strategy concentrated on (b, a) $[(a, b)]$. Thus tie-breaking must be even.

(ii) I will show that the fictitious game has a Nash equilibrium (δ_2^*, π_2^*) , where δ_2^* is as defined in the proposition and where π_2^* is the uniform distribution over $\{(a, 1 - a), (1 - a, a)\}$ for some constant $a > 1/2$. In other words, π_2^* is like π_1^* with the additional feature that $a + b = 1$.

Nature's best response condition (2) requires that π_2^* is supported on

$$\begin{aligned} & \arg \max_{s \in [0,1]^2} R(\delta_2^*, s) \\ &= \arg \max_{(\mu_0, \mu_1) \in [0,1]^2} \left\{ (\mu_1 - \mu_0)^+ \left[\Pr(I_N < 0) + \frac{1}{2} \Pr(I_N = 0) \right] \right. \\ & \quad \left. + (\mu_0 - \mu_1)^+ \left[\Pr(I_N > 0) + \frac{1}{2} \Pr(I_N = 0) \right] \right\}, \end{aligned}$$

where the probabilities refer to the sampling distribution of $I_N = n_1 - n_0 - (N_1 - N_0)/2$. For example,

$$\begin{aligned} \Pr(I_N > 0) &= \sum_{N_1=0}^N \binom{N}{N_1} 2^{-N} \\ & \times \sum_{(n_0, n_1): n_1 - n_0 - (N_1 - N_0)/2 > 0} \binom{N_0}{n_0} \mu_0^{n_0} (1 - \mu_0)^{N_0 - n_0} \\ & \times \binom{N_1}{n_1} \mu_1^{n_1} (1 - \mu_1)^{N_1 - n_1}. \end{aligned}$$

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 suffices to show that I_N is independent of μ_0 given $(\mu_1 - \mu_0)$. The proof will be by induction over N , thus assume the result for 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 $(t_{N+1}, y_{N+1}) \in \{(0, 0), (1, 1)\}$, or $I_N = x + 1$ and $(t_{N+1}, y_{N+1}) \in \{(0, 1), (1, 0)\}$. Thus

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

The argument is concluded by observing that I_0 is deterministically equal to zero.

To verify that δ_2^* is Bayes against any such π_2^* , write

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

$$\begin{aligned} & > \binom{N_0}{n_0} a^{n_0} (1 - a)^{N_0 - n_0} \binom{N_1}{n_1} (1 - a)^{n_1} a^{N_1 - n_1} \\ \iff & a^{n_1 + N_0 - n_0} (1 - a)^{N_1 - n_1 + n_0} > a^{N_1 - n_1 + n_0} (1 - a)^{n_1 + N_0 - n_0} \\ \iff & a^{2n_1 - N_1 - (2n_0 - N_0)} > (1 - a)^{2n_1 - N_1 - (2n_0 - N_0)} \\ \iff & a^{I_N} > (1 - a)^{I_N} \\ \iff & I_N > 0, \end{aligned}$$

where the algebra omits some initial steps that are completely analogous to the similar argument in (i). Arguments for near uniqueness of δ_2^* are as in (i).

(iii) For consistency of notation, I continue to identify states with couplets $s = (\mu_0, \mu_1)$. Recall, however, that μ_0 is now given and Nature can only choose μ_1 . I will show that there exists a Nash equilibrium (δ_3^*, π_3^*) , where δ_3^* is as in the proposition and π_3^* is supported on $\{(\mu_0, b), (\mu_0, a)\}$ for some constants $a > \mu_0 > b$.

Nature's best-response condition (2) requires π_3^* to be supported on

$$\begin{aligned} & \arg \max_{s \in \{\mu_0\} \times [0,1]} R(\delta_3^*, s) \\ &= \arg \max_{\mu_1 \in [0,1]} \left\{ (\mu_0 - \mu_1)^+ (\Pr(n_1 > n^*) + \lambda^* \Pr(n_1 = n^*)) \right. \\ & \quad \left. + (\mu_1 - \mu_0)^+ (\Pr(n_1 < n^*) + (1 - \lambda^*) \Pr(n_1 = n^*)) \right\}, \quad (8) \end{aligned}$$

where the probabilities refer to the sampling distribution of n_1 , e.g.

$$\Pr(n_1 > n^*) = \sum_{n > n^*} \binom{N}{n} \mu_1^n (1 - \mu_1)^{N - n}.$$

For π_3^* to be a best response, it is necessary that the arg max in (8) contains some state (μ_0, b) as well as some state (μ_0, a) . Inspection of (8) reveals that this requires

$$\begin{aligned} & \max_{\mu_1 \in [0, \mu_0]} \left\{ (\mu_0 - \mu_1) (\Pr(n_1 > n^*) + \lambda^* \Pr(n_1 = n^*)) \right\} \\ &= \max_{\mu_1 \in [\mu_0, 1]} \left\{ (\mu_1 - \mu_0) (\Pr(n_1 < n^*) + (1 - \lambda^*) \Pr(n_1 = n^*)) \right\}, \quad (9) \end{aligned}$$

which equals condition (4) upon writing out the probabilities.

I now show that (n^*, λ^*) can be chosen so that (9) holds. More specifically, condition (9) uniquely determines (n^*, λ^*) as a function of μ_0 , with the caveat that $(n^*, 0)$ and $(n^* + 1, 1)$ represent the same decision rule. To see this, define $\alpha \equiv n^* + 1 - \lambda^* \in [0, N + 1]$, the variable whose values are displayed in Table 1. The mapping from distinct decision rules (n^*, λ^*) to α is one-to-one. Indeed, α is an indicator of a treatment rule's conservatism, with $\alpha = 0$ $[N + 1]$ indicating the no-data rule that always [never] assigns treatment 1. To take care of some trivial cases, note that if $\mu_0 = 0$ [1], then condition (9) holds with $\alpha = 0$ $[N + 1]$.

Now let $\mu_0 \in (0, 1)$ and consider the l.h.s. of (9). For any fixed μ_1 , $(\mu_0 - \mu_1) (\Pr(n_1 > n^*) + \lambda^* \Pr(n_1 = n^*))$ decreases in α , strictly so except when $\mu_1 = 0$ and $\alpha \geq 1$, in which case it is constant at 0. Since the l.h.s. of (9) cannot be solved by any μ_1 that sets it equal to 0, it strictly decreases in α . Furthermore, it is continuous in α by inspection, and straightforward computations show that it equals μ_0 if $\alpha = 0$ and 0 if $\alpha = N + 1$. By similar arguments, the r.h.s. of (9) strictly and continuously increases from 0 to $(1 - \mu_0)$ as α increases from 0 to $N + 1$. It follows that (9) holds for exactly one $\alpha \in [0, N + 1]$.

It remains to show that δ_3^* is a best response to π_3^* . As in the preceding two cases, the monotone likelihood ratio property of the binomial distribution implies that there exists some threshold \tilde{n} s.t. a best response to π_3^* assigns all subjects to treatment 1 $[0]$ if $n_1 > [<] \tilde{n}$. Thus, equilibrium is obtained if the prior probability $p^* \equiv \Pr(s = (\mu_0, a))$ is chosen such that $\tilde{n} = n^*$. This requires that the decision maker is indifferent between treatments conditional

on $n_1 = n^*$, hence that the posterior expectation of Y_1 conditional on $n_1 = n^*$ fulfils

$$\begin{aligned} \mu_0 &= \mathbb{E}(Y_1 | n_1 = n^*) \\ &= \frac{ap^* \Pr(n_1 = n^* | s = (\mu_0, a)) + b(1 - p^*) \Pr(n_1 = n^* | s = (\mu_0, b))}{p^* \Pr(n_1 = n^* | s = (\mu_0, a)) + (1 - p^*) \Pr(n_1 = n^* | s = (\mu_0, b))}. \end{aligned}$$

Clearly $\mathbb{E}(Y_1 | n_1 = n^*)$ continuously increases from b to a as p^* increases from 0 to 1, hence p^* can always be found.

For $N = 0$, it is easily established that $\delta_3^* = 1 - \mu_0$, supported by prior π_3^* that puts weight μ_0 on $(\mu_0, 1)$ and weight $(1 - \mu_0)$ on $(\mu_0, 0)$. The “near uniqueness” statements follow as before.

A.2. Proposition 2

(i) The crucial observation is that $R(\tilde{\delta}, s) = R(\delta, s')$, which immediately implies both claims. To see it, consider first the unconditional (on \tilde{Y}_n) distribution of Y_n . This distribution is Bernoulli because $\tilde{Y}_n \in \{0, 1\}$. Furthermore, a simple application of the law of iterated expectations shows that its parameter is μ_{t_n} . It follows that the distribution of \tilde{Y}_n , and hence of $\tilde{\omega}$, depends on s only through (μ_0, μ_1) .

Now recall that $R(\tilde{\delta}, s) = \max\{\mu_0, \mu_1\} - (\mu_0 \mathbb{E}(1 - \delta(\tilde{\omega})) + \mu_1 \mathbb{E}\delta(\tilde{\omega}))$. In view of the preceding paragraph’s insight, this expression depends on s only through (μ_0, μ_1) . Since the transformation from s to s' preserves (μ_0, μ_1) , one finds that $R(\tilde{\delta}, s) = R(\tilde{\delta}, s')$. But $\tilde{\omega} = \omega$ a.s. if the state is s' , hence $R(\tilde{\delta}, s') = R(\delta, s')$.

(ii) Write

$$\min_{\delta \in \mathcal{D}} \max_{s' \in \mathcal{S}'} R(\delta, s') \leq \min_{\delta \in \mathcal{D}} \max_{s \in \mathcal{S}} R(\delta, s) \leq \min_{\tilde{\delta}: \tilde{\delta} \in \tilde{\mathcal{D}}} \max_{s \in \mathcal{S}} R(\tilde{\delta}, s),$$

where the l.h. inequality obtains because $\mathcal{S}' \subseteq \mathcal{S}$ and the r.h. inequality obtains because $\{\tilde{\delta} : \tilde{\delta} \in \tilde{\mathcal{D}}\} \subseteq \mathcal{D}$. Now, (i) implies that both inequalities bind, hence the first claim obtains. The second claim follows because $R(\tilde{\delta}, s) = R(\delta, s')$.

A.3. Corollary 1

I compute the value functions under the assumption that outcomes are binary; results extend to general outcomes by Proposition 2. First compute $R_2^*(N)$. The minimax regret value of the fictitious game is the value function of Nature’s best-response problem at the equilibrium, hence

$$\begin{aligned} R_2^*(N) &= \max_{(\mu_0, \mu_1) \in [0, 1]^2} R(\delta_2^*, s) \\ &= \max_{(\mu_0, \mu_1) \in [0, 1]^2} \left\{ (\mu_1 - \mu_0)^+ \left[\Pr(I_N < 0) + \frac{1}{2} \Pr(I_N = 0) \right] \right. \\ &\quad \left. + (\mu_0 - \mu_1)^+ \left[\Pr(I_N > 0) + \frac{1}{2} \Pr(I_N = 0) \right] \right\} \\ &= \max_{(\mu_0, \mu_1) \in [0, 1]^2} \left\{ (\mu_1 - \mu_0)^+ \left[\Pr(I_N < 0) + \frac{1}{2} \Pr(I_N = 0) \right] \right\}, \end{aligned}$$

where the last step uses the objective function’s symmetry. Now, observe that $I_N \propto 2n - N$, where $n \equiv n_1 + N_0 - n_0$ counts successes of treatment 1 as well as failures of treatment 0. Hence, $I_N < 0$ iff $n < N/2$. By the proof of Proposition 1(ii), one can set $\mu_0 = 1 - \mu_1$, implying that the distribution of n is binomial with parameters (N, μ_1) . This immediately implies that for odd N ,

$$\begin{aligned} R_2^*(N) &= \max_{\mu_1 \in [1/2, 1]} \{(2\mu_1 - 1) \Pr(I_N < 0)\} \\ &= \max_{a \in [1/2, 1]} \left\{ (2a - 1) \sum_{n < N/2} \binom{N}{n} a^n (1 - a)^{N-n} \right\}. \end{aligned}$$

If N is even, then $2n - N$ must be even as well. Consider ignoring the last sample point. If $2n - N \neq 0$, this will not affect $\delta_2^*(\omega)$. If $2n - N = 0$, then due to the sample design’s symmetry, it amounts to an even randomization. Hence, δ_2^* based on the first $(N - 1)$ sample points is a best response to π_2^* when N is even. This implies that $R_2^*(N) = R_2^*(N - 1)$ when N is even.

Next, $\pi_1^* = \pi_2^*$, hence $R_1^*(N) = R_2^*(N)$. To see this, recall that π_1^* randomizes evenly over states $\{(a, b), (b, a)\}$ with $a > b$. I will show that among such priors, one can restrict attention to those that furthermore have $a + b = 1$. Thus, fix some $\Delta \in (0, 1]$ and consider the class Γ_Δ of priors π_b that randomize evenly over $\{(b + \Delta, b), (b, b + \Delta)\}$. Within Γ_Δ , minimax regret is achieved by setting $b = (1 - \Delta)/2$, so that $a + b = 1$ obtains. To see this, use (6) to write $\max_{b \in [0, 1 - \Delta]} \min_{\delta \in \mathcal{D}} \int R(\delta, s) d\pi_b = \Delta \max_{b \in [0, 1 - \Delta]} \min_{\delta \in \mathcal{D}} \int \mathbb{E} |\delta(\omega) - t(s)| d\pi_b$, where $t(s) \equiv \mathbb{I}\{\mu_1 > \mu_0\}$ denotes the treatment that is in fact better in state s . Now, $\min_{\delta \in \mathcal{D}} \int \mathbb{E} |\delta(\omega) - t(s)| d\pi_b$ is the value of the statistical experiment in which $t(s)$ is estimated under 0/1-loss with prior π_b . This value depends on b only through the distribution of ω . Assume w.l.o.g. that within-sample treatment assignment was in alternating sequence, then $(y_{2n} - y_{2n-1})^{N/2}$ is sufficient for ω (observation of a matched pair that experiences the same outcome does not lead to any updating of π_b). Compare the distribution of $(y_{2n} - y_{2n-1})$ if $b = (1 - \Delta)/2$ with its distribution under any other choice of b . Simple algebra shows that the former distribution is a mixture of the latter one with a point mass at 0. As the DM can always perform this mixture herself, $\min_{\delta \in \mathcal{D}} \int \mathbb{E} |\delta(\omega) - t(s)| d\pi_b \leq \min_{\delta \in \mathcal{D}} \int \mathbb{E} |\delta(\omega) - t(s)| d\pi_{(1-\Delta)/2}$.

Regarding $R_3^*(N)$, it follows along the lines of the argument for $R_2^*(N)$ that

$$\begin{aligned} R_3^*(N) &= \max_{\mu_1 \in [0, 1]} \left\{ (\mu_0 - \mu_1)^+ \left[\sum_{n > n^*} \binom{N}{n} \mu_1^n (1 - \mu_1)^{N-n} \right. \right. \\ &\quad \left. \left. + \lambda^* \binom{N}{n^*} \mu_1^{n^*} (1 - \mu_1)^{N-n^*} \right] \right. \\ &\quad \left. + (\mu_1 - \mu_0)^+ \left[\sum_{n < n^*} \binom{N}{n} \mu_1^n (1 - \mu_1)^{N-n} \right. \right. \\ &\quad \left. \left. + (1 - \lambda^*) \binom{N}{n^*} \mu_1^{n^*} (1 - \mu_1)^{N-n^*} \right] \right\} \\ &= \max_{a \in [0, 1]} \left\{ (a - \mu_0) \left[\sum_{n < n^*} \binom{N}{n} a^n (1 - a)^{N-n} \right. \right. \\ &\quad \left. \left. + (1 - \lambda^*) \binom{N}{n^*} a^{n^*} (1 - a)^{N-n^*} \right] \right\}, \end{aligned}$$

where the last step uses (9).

A.4. Proposition 3

As in Proposition 1, $R(\delta, s)$ depends on s only through marginal distributions $P(Y_{x'})$, so I identify states with collections of marginal distributions $s = (s_x)_{x \in \mathcal{X}}$. Let π^* be the prior that assigns probability $\pi^*(s) \equiv \prod_{x \in \mathcal{X}} \pi_x^*(s_x)$ to s . This prior induces conditional priors $(\pi_x^*)_{x \in \mathcal{X}}$ and renders s_x independent of $s_{x'}$ for $x \neq x'$. Note that π^* exists whenever \mathcal{S} meets the richness condition stated in the text. The proof will verify that (δ^*, π^*) is a Nash equilibrium of the fictitious game. To begin, for any $x' \neq x$, independence of s_x and $s_{x'}$ implies that ω_x is uninformative about $s_{x'}$. That δ^* is Bayes against π^* is then elementary.

Nature's best-response condition (2) requires that any s in the support of π^* maximize

$$\begin{aligned} R(\delta^*, s) &= \max_{\delta \in \mathcal{D}} u(\delta, s) - u(\delta^*, s) \\ &= \max_{\delta \in \mathcal{D}} \sum_{x \in \mathcal{X}} \Pr(X = x) (\mu_{0x} (1 - \mathbb{E}\delta_x(\omega)) + \mu_{1x}\mathbb{E}\delta_x(\omega)) \\ &\quad - \sum_{x \in \mathcal{X}} \Pr(X = x) (\mu_{0x} (1 - \mathbb{E}\delta_x^*(\omega)) + \mu_{1x}\mathbb{E}\delta_x^*(\omega)) \\ &= \sum_{x \in \mathcal{X}} \Pr(X = x) \max\{\mu_{0x}, \mu_{1x}\} \\ &\quad - \sum_{x \in \mathcal{X}} \Pr(X = x) (\mu_{0x} (1 - \mathbb{E}d_x^*(\omega_x)) + \mu_{1x}\mathbb{E}d_x^*(\omega_x)) \\ &= \sum_{x \in \mathcal{X}} \Pr(X = x) (\max\{\mu_{0x}, \mu_{1x}\} \\ &\quad - (\mu_{0x} (1 - \mathbb{E}d_x^*(\omega_x)) + \mu_{1x}\mathbb{E}d_x^*(\omega_x))) \\ &= \sum_{x \in \mathcal{X}} \Pr(X = x) R_x(d_x^*, s_x). \end{aligned}$$

Here, the first two steps substitute definitions from the text, the third step explicitly solves the maximization problem and also uses that $\delta_x^*(\omega) = d_x^*(\omega_x)$, the fourth step collects terms, and the last step uses (5) applied to $R_x(d_x^*, s_x)$. By assumption, any s_x in the support of π_x^* maximizes $R_x(d_x^*, s_x)$, thus (2) is verified.

A.5. Example 1

The conditional decision problems are covered by Proposition 2 in conjunction with Proposition 1(i); hence, that δ^* achieves minimax regret follows from those results in conjunction with Proposition 3. To see the additional claims, let outcomes be binary. Using Proposition 1(i) and Corollary 1, the least favorable prior π^* from Proposition 3 randomizes evenly over $\{(a_m, b_m, a_f, b_f), (a_m, b_m, b_f, a_f), (b_m, a_m, a_f, b_f), (b_m, a_m, b_f, a_f)\}$ for some $(a_m, b_m, a_f, b_f) \in [0, 1]^4$, where states are identified with vectors $\mathbb{E}(Y_{0m}, Y_{1m}, Y_{0f}, Y_{1f})$ as in Proposition 1 and where $a_m > b_m, a_f > b_f$. Claim (i) follows because any minimax regret treatment rule must be a best response to π^* . To see (ii), let δ be the modified decision rule, then Nature will strictly prefer state (a_m, b_m, b_f, a_f) over (a_m, b_m, a_f, b_f) because in the former state, δ will tend to tie-break in the wrong direction. The exception is that in the no-data problem where $N_m = N_f = 0, (\delta_m, \delta_f) = (1/2, 1/2)$ achieves minimax regret.

A.6. Proposition 4

Preliminaries.

As before, only marginal distributions of Y_{ix} matter, thus states x can be identified with functions $s_x : [0, 1] \rightarrow \Delta[0, 1]^2$ where $s_x = P(Y_{0x}, Y_{1x})$. Intuitively, one might want to extend Proposition 3 to unknown (μ_{0x}, μ_{1x}) by verifying the following least favorable prior: Let s_x be supported on $\{(0, 1), (1, 0)\}$ for all x and let s_x be independent of $s_{x'}$ for all $x \neq x'$. But there is a problem: For any state s , let $w_s \equiv \{x : \mu_{0x} = 1\} \subseteq [0, 1]$, then the conjectured prior would induce the uniform distribution over all possible sets w_s , i.e. over the power set of $[0, 1]$ – but this distribution does not exist (e.g., chapter 3 in Billingsley (1995)). Part (i) below shows that the conjectured prior can be approximated. The intuition behind part (ii) is the same. Unlike previous proofs, the arguments do not establish Nash equilibria, which may not exist.

(i) There exists a sequence of priors $\{\pi_i\}$ s.t. $\lim_{i \rightarrow \infty} \min_{\delta \in \mathcal{D}} \int R(\delta, s) d\pi_i = 1/2$, implying that $\lim_{i \rightarrow \infty} \min_{\delta \in \mathcal{D}} \max_{s \in \text{supp}(\pi_i)} R(\delta, s) \geq 1/2$ and hence the result. To construct π_i , define the partition $W_i \equiv \{[0, 1/i], (1/i, 2/i], \dots, ((i-1)/i, 1]\}$ 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 distributions $(s_i^j)_{j=1}^{2^i}$ by identifying s_i^j with the degenerate distribution concentrated at

$$(\mu_{0x}, \mu_{1x})_{x \in \mathcal{X}} = \left(\mathbb{I}\{x \in w_i^j\}, 1 - \mathbb{I}\{x \in w_i^j\} \right)_{x \in \mathcal{X}}.$$

Let π_i be the uniform distribution over states $(s_i^j)_{j=1}^{2^i}$, i.e. π_i assigns

probability 2^{-i} to every s_i^j . Notice 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 $w_i(x) \neq w_i(x')$.

Recall that $\min_{\delta \in \mathcal{D}} \int R(\delta, s) d\pi_i$ is achieved by any decision rule δ^* that is Bayes for $u(\delta, s)$ given prior π_i . By elementary calculations, the posterior induced by sample data $\omega = \{(t_n, x_n, y_n)\}_{n=1}^N$ is concentrated on the truth for any $x \in \bigcup_{n=1}^N w_i(x_n)$ and coincides with π_i otherwise. A Bayes rule δ^* then assigns the correct treatment conditional on any $x \in \bigcup_{n=1}^N w_i(x_n)$ and is unrestricted otherwise; let $\delta_x^*(\omega) = 1/2$ in the latter case. It follows that for any s in the support of π_i ,

$$\begin{aligned} R(\delta^*, s) &= \int \left(\max\{\mu_{0x}, \mu_{1x}\} - \Pr\left(x \in \bigcup_{n=1}^N w_i(x_n)\right) \right) \cdot \max\{\mu_{0x}, \mu_{1x}\} \\ &\quad - \Pr\left(x \notin \bigcup_{n=1}^N w_i(x_n)\right) \cdot \frac{\mu_{0x} + \mu_{1x}}{2} \Big) dx \\ &= \int \left(1 - \Pr\left(x \in \bigcup_{n=1}^N w_i(x_n)\right) - \frac{1}{2} \Pr\left(x \notin \bigcup_{n=1}^N w_i(x_n)\right) \right) dx \\ &= \frac{1}{2} - \frac{1}{2} \int \Pr\left(x \in \bigcup_{n=1}^N w_i(x_n)\right) dx \\ &\geq \frac{1}{2} - \frac{1}{2} \int \sum_{n=1}^N \Pr(x \in w_i(x_n)) dx \\ &= \frac{1}{2} - \frac{1}{2} \int \frac{N}{i} dx \rightarrow \frac{1}{2} \end{aligned}$$

as $i \rightarrow \infty$. Here, the first equation holds by definition, the second one follows by substituting in for (μ_{0x}, μ_{1x}) , the third one collects terms, the inequality follows from basic probability calculus, and the last equation uses the uniform distribution of X . Finally, the no-data rule incurs maximal regret of $1/2$, achieved by any Bernoulli state.

(ii) I will construct a similar sequence $\{\pi_i\}$ as in (i). For every i , partition $[0, 1]$ into level sets as follows:

$$W_i = \{\{x : 0 \leq \mu_{0x} \leq 1/i\}, \{x : 1/i < \mu_{0x} \leq 2/i\}, \dots, \{x : (i-1)/i < \mu_{0x} \leq 1\}\}.$$

This partition is legal because μ_{0x} is measurable. As before, define $w_i(x)$ to denote the element of W_i that contains x ; also, for each $w \in W_i$, define $\mu_{0w} = \inf\{\mu_{0x} : x \in w\}$. For future use, note in particular that $\mu_{0w_i(x)} \geq \mu_{0x} - 1/i$. Define

$(w_i^j)_{j=1}^{2^i}$ and $(s_i^j)_{j=1}^{2^i}$ as in (i). Then π_i is characterized by $\Pr(s = s_i^j) = \prod_{w \in W_i} (\mu_{0w} \cdot \mathbb{I}\{w \in w_i^j\} + (1 - \mu_{0w}) \cdot \mathbb{I}\{w \notin w_i^j\})$. Note the following features of π_i : (i) The prior expectation of Y_{1x} equals $\mu_{0w_i(x)} \leq \mu_{0x}$, (ii) s_x and $s_{x'}$ are independent whenever $w_i(x) \neq w_i(x')$. Thus, a Bayes rule δ^* assigns the better treatment whenever

$x \in \bigcup_{n=1}^N w_i(x_n)$ and sets $\delta_x^*(\omega) = 0$ otherwise. Now,

$$\int R(\delta^*, s) d\pi_i = \int \int \left(\max\{\mu_{0x}, \mu_{1x}\} - \Pr\left(x \in \bigcup_{n=1}^N w_i(x_n)\right) \times \max\{\mu_{0x}, \mu_{1x}\} - \Pr\left(x \notin \bigcup_{n=1}^N w_i(x_n)\right) \mu_{0x} \right) dx d\pi_i.$$

As before, $\int \Pr\left(x \in \bigcup_{n=1}^N w_i(x_n)\right) dx \rightarrow 0$, hence

$$\begin{aligned} \dots &= \int \int (\max\{\mu_{0x}, \mu_{1x}\} - \mu_{0x}) dx d\pi_i \\ &= \int \left(\int \max\{\mu_{0x}, \mu_{1x}\} d\pi_i - \mu_{0x} \right) dx \\ &= \int (\mu_{0w_i(x)} + (1 - \mu_{0w_i(x)})\mu_{0x} - \mu_{0x}) dx \\ &= \int \mu_{0w_i(x)}(1 - \mu_{0x}) dx \\ &\geq \int (\mu_{0x} - 1/i)(1 - \mu_{0x}) dx \\ &\rightarrow \int \mu_{0x}(1 - \mu_{0x}) dx \end{aligned}$$

as $i \rightarrow \infty$. Here, the second step evaluates the inner integral, using that by construction of π_i , $\max\{\mu_{0x}, \mu_{1x}\}$ equals 1 with probability $\mu_{0w_i(x)}$ and equals μ_{0x} otherwise. The other steps are algebra.

Now consider the no-data rule δ with $\delta_x(\omega) = 1 - \mu_{0x}$. Maximal regret incurred by this rule is

$$\begin{aligned} \max_{s \in \mathcal{S}} R(\delta, s) &= \max_{s \in \mathcal{S}} \left\{ \int ((\mu_{1x} - \mu_{0x})^+ \mu_{0x} + (\mu_{0x} - \mu_{1x})^+ (1 - \mu_{0x})) dx \right\} \\ &= \int \mu_{0x}(1 - \mu_{0x}) dx, \end{aligned}$$

where the first equality uses (6); the second equality follows by solving the maximization problem.

References

Bahadur, R.R., Savage, L.J., 1956. The nonexistence of certain statistical procedures in nonparametric problems. *Annals of Mathematical Statistics* 27, 1115–1122.
 Berger, J.O., 1985. *Statistical Decision Theory and Bayesian Analysis*, 2nd ed. Springer Verlag.

Billingsley, P., 1995. *Probability and Measure*, 3rd ed. Wiley.
 Brock, W.A., 2006. Profiling problems with partially identified structure. *Economic Journal* 116, F427–F440.
 Canner, P.L., 1970. Selecting one of two treatments when the responses are dichotomous. *Journal of the American Statistical Association* 65, 293–306.
 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.
 Eozenou, P., Rivas, J., Schlag, K.H., 2006. Minimax regret in practice—Four examples on treatment choice. Mimeo. European University Institute.
 Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18, 141–153.
 Gupta, S.S., Hande, S.N., 1992. On some nonparametric selection procedures. In: Saleh, A.K.Md.E. (Ed.), *Nonparametric Statistics and Related Topics*. Elsevier.
 Hayashi, T., 2008. Regret aversion and opportunity-dependence. *Journal of Economic Theory* 139, 242–268.
 Hirano, K., Porter, J.R., 2008. Asymptotics for statistical treatment rules. Mimeo. University of Arizona.
 Manski, C.F., 2000. Identification problems and decisions under ambiguity: Empirical analysis of treatment response and normative analysis of treatment choice. *Journal of Econometrics* 95, 415–442.
 Manski, C.F., 2004. Statistical treatment rules for heterogeneous populations. *Econometrica* 72, 1221–1246.
 Manski, C.F., 2005. *Social Choice with Partial Knowledge of Treatment Response*. Princeton University Press.
 Manski, C.F., 2006. Search profiling with partial knowledge of deterrence. *Economic Journal* 116, F385–F401.
 Manski, C.F., 2007a. Minimax-regret treatment choice with missing outcome data. *Journal of Econometrics* 139, 105–115.
 Manski, C.F., 2007b. Identification for Prediction and Decision. Harvard University Press.
 Manski, C.F., 2008. Diversified treatment under ambiguity. Mimeo. Northwestern University.
 Manski, C.F., Tetenov, A., 2007. Admissible treatment rules for a risk-averse planner with experimental data on an innovation. *Journal of Statistical Planning and Inference* 137, 1998–2010.
 Milnor, J., 1954. Games against nature. In: Thrall, R.M., Coombs, C.H., Davis, R.L. (Eds.), *Decision Processes*. Wiley, New York.
 Savage, L.J., 1951. The theory of statistical decisions. *Journal of the American Statistical Association* 46, 55–67.
 Savage, L.J., 1954. *The Foundations of Statistics*. Wiley.
 Schlag, K.H., 2003. How to minimize maximum regret under repeated decision-making. Mimeo. European University Institute.
 Schlag, K.H., 2006. Eleven. Mimeo. European University Institute.
 Stoye, J., 2007a. Minimax regret treatment choice with incomplete data and many treatments. *Econometric Theory* 23, 190–199.
 Stoye, J., 2007b. Statistical decisions under ambiguity. Mimeo. New York University.
 Stoye, J., 2007c. Axioms for minimax regret choice correspondences. Mimeo. New York University.
 Stoye, J., 2007d. Minimax regret treatment choice with finite samples and missing outcome data. In: G. de Cooman, J. Veinarová, M. Zaffalon (Eds.), *Proceedings of the Fifth International Symposium on Imprecise Probability: Theories and Applications*.
 Stoye, J., 2009. Partial identification and robust treatment choice: An application to young offenders. *Journal of Statistical Theory and Practice* 3, 239–254.
 Tetenov, A., 2007. Statistical treatment choice based on asymmetric minimax regret criteria. Mimeo. Northwestern University.
 Wald, A., 1945. Statistical decision functions which minimize the maximum risk. *Annals of Mathematics* 46, 265–280.
 Wald, A., 1950. *Statistical Decision Functions*. Wiley.