Journal of Econometrics 151 (2009) 70–81

Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

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.

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.

0304-4076/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.jeconom.2009.02.013

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 $y^j(t) : \{0, 1\} \to [0, 1]$ 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 $\mathcal{S} = \left\{Q(Y_0, Y_1) \in \Delta[0, 1]^2 : Q(Y_0) = P(Y_0)\right\}$, 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 aforecited 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 $u \circ f$ 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.