

Statistical Treatment Choice Based on Asymmetric Minimax Regret Criteria*

Aleksey Tetenov[†]

Department of Economics

Northwestern University

November 5, 2007

(JOB MARKET PAPER)

Abstract

This paper studies the problem of treatment choice between a status quo treatment with a known outcome distribution and an innovation whose outcomes are observed only in a finite sample randomized experiment. I evaluate statistical decision rules, which are functions that map sample outcomes into the planner's treatment choice for the population, based on regret, which is the expected welfare loss due to assigning inferior treatments. I extend previous work that applied the minimax regret criterion to treatment choice problems by considering decision criteria that asymmetrically treat Type I regret (due to mistakenly choosing an inferior new treatment) and Type II regret (due to mistakenly rejecting a superior innovation). I derive exact finite sample solutions to these problems for experiments with normal, Bernoulli, or bounded distributions of individual outcomes. In conclusion, I discuss approaches to the problem for other classes of distributions. Along the way, the paper compares asymmetric minimax regret criteria with statistical decision rules based on classical hypothesis tests.

*I thank Chuck Manski, Elie Tamer, Ben Jones, and seminar participants at Northwestern for very helpful comments.

[†]Department of Economics, Northwestern University, 2001 Sheridan Rd, Evanston, IL 60208, USA. *E-mail address:* a-tetenov@northwestern.edu

1 Introduction

Consider a planner who has to choose which one of two mutually exclusive treatments should be assigned to members of a population. One treatment is the status quo, whose effects are well known. The other is a promising innovation, whose exact effects have yet to be determined. The treatments in question may be, for example, two alternative drugs or therapies for a medical condition, or two different unemployment assistance programs. Suppose that a randomized clinical trial or some other experiment will be conducted and its results will be used to choose which treatment population members will receive.

The planner faces two problems. First, she has to know what experiment (in particular, what sample size) should be chosen to get a sufficiently accurate estimate of the treatment effect. Second, she has to select how treatment choices will be determined based on the statistical evidence obtained from the experiment. Often, treatment choice is based on the results of a statistical hypothesis test, which is constructed to keep the probability of mistakenly assigning an inferior innovation (a Type I error) below a specified level (usually .05 or .01). Then, the sample size is selected to obtain a high probability (usually .8 or .9) that the innovation will be chosen if its positive effect exceeds some value of interest.

Following Wald's (1950) formulation of statistical decision theory, I analyze the performance of alternative statistical methods based on their expected welfare over different realizations of the sampling process, rather than just their probabilities of error. In particular, I continue a recent line of work advocating and investigating treatment choice procedures that minimize maximum regret by Manski (2004, 2005, 2007a, 2007b), Hirano and Porter (2006), Stoye (2006, 2007a, 2007b) and Schlag (2006, 2007). Regret is the difference between the maximum welfare that could be achieved given full knowledge of the effects of both treatments (by assigning the treatment that is actually better) and the expected welfare of treatment choices based on experimental outcomes. The latter is smaller, because experimental outcomes generally do not allow the decision maker to choose the best treatment 100 percent of the time.

This paper's main departure from previous literature on the subject is asymmetric consideration of Type I regret (due to mistakenly using an inferior new treatment) and Type II regret (due to missing out on using a superior innovation). The persistent use in treatment choice problems of

the hypothesis testing approach, which allows Type II errors to occur with higher probability than Type I errors, suggests that many decision makers want to place the burden of proof on the new treatment. Most do so by selecting a low hypothesis test level, such as $\alpha = .05$. It is not clear what principles, besides convention, are there to guide the selection of hypothesis test level for the circumstances of a particular decision problem. Values of maximum Type I and maximum Type II regret of a statistical procedure could provide the decision maker with more relevant characteristics of its performance than the traditional hypothesis testing measures (test level and power), since regret takes into account both the probability of making an error and its economic magnitude.

How to balance Type I and Type II regret in a particular problem is up to the decision maker. In this paper I consider three criteria. First, the traditional minimax regret criterion gives equal consideration to Type I and Type II regret. It seeks to minimize the larger of the two, thus minimax regret solutions have equal maximum Type I and Type II regret.

The second criterion is minimax regret with an asymmetric linear reference-dependent welfare function. This criterion gives larger weight to maximum Type I regret, thus the maximum Type II regret of asymmetric minimax regret solutions is larger than their maximum Type I regret by a given factor. When the treatment effect estimate is normally distributed, hypothesis test based solutions with a given level α correspond to asymmetric minimax regret solutions for some asymmetry factor $K(\alpha)$ for any sample size and variance. In a sense, the minimax regret criterion with an asymmetric welfare function provides a decision-theoretic rationalization of hypothesis tests that is based on expected welfare.

The third criterion is limited Type I regret. Many decision makers face the problem of making treatment choices based on existing statistical evidence, without any control over its sample size and precision. Symmetric and asymmetric minimax regret, as well as hypothesis testing, often lead to solutions whose maximum Type I regret is proportional to the standard error of the estimate of average treatment effect. This may not be appealing to decision makers primarily interested in "safety" of new treatments, which I interpret as low Type I regret. The limited Type I regret criterion seeks to minimize maximum regret subject to an explicit constraint that maximum Type I regret should not exceed a given acceptable level. This approach guarantees limited expected welfare losses due to Type I errors regardless of the decision process that underlies sample size selection.

Instead of looking at maximum regret values, a Bayesian decision maker would assert a subjective probability distribution over the set of feasible treatment outcome distributions, use sample realizations to derive an updated posterior probability distribution, and maximize expected welfare with regard to that posterior distribution (which is equivalent to minimizing expected regret). Unfortunately, in many situations decision makers do not have any information that would form a reasonable basis for asserting a prior distribution. In group decision making, members of the group may disagree in their prior beliefs. These problems lead to frequent use of conventional prior distributions in applied Bayesian analysis. Bayesian treatment choice based on a conventional prior distribution, rather than on a subjective distribution reflecting the decision maker's prior information, does not have a clear economic justification. Decision making based on maximum regret is a conservative approach to dealing with the lack of reasonable prior beliefs, since maximum regret is the sharp upper bound on expected regret for decision makers with any prior distributions.

The paper proceeds in the following order. Section 2 expounds the decision-theoretic formulation of the problem and the criteria used to address it. In section 3, I consider a simple but instructive case where the experiment generates a normally distributed random variable with known or bounded variance. I analyze conventional treatment choice rules based on hypothesis testing and sample size choice based on power analysis in light of their maximum regret and compare them with minimax regret, asymmetric minimax regret and limited Type I regret solutions. Section 4 analyzes treatment choice when treatment outcomes are either binary or bounded random variables. Exact minimax regret results were obtained for these problems by Stoye (2006) and Schlag (2007). I extend their results to derive asymmetric minimax regret and limited Type I regret solutions using a different technique. I also demonstrate that the minimax-regret solution proposed by these authors for bounded outcomes does not minimize maximum regret if the decision maker can place an informative upper bound on the variance of the outcome distribution, which is the case in many applications. In the concluding section 5, I discuss the use of approximations, bounds, and numerical methods for problems that do not yet have convenient analytical solutions and illustrate their performance in a hypothetical clinical trial problem with rare dangerous side effects.

2 Statistical treatment rules, welfare and regret

The basic setting is the same as in Manski (2004, 2005) and in Manski and Tetenov (2007). The planner's problem is to assign members of a large population to one of two available treatments $t \in T, T = \{0, 1\}$. Let $t = 0$ denote the status quo treatment and $t = 1$ the innovation. Each member j of the population, denoted J , has a response function $y_j(t)$ describing that individual's potential outcome under each treatment t . The population is a probability space (J, Ω, P) and the probability distribution $P[y(\cdot)]$ of the random function $y(\cdot)$ describes treatment response across the population. The population is "large," in the sense that J is uncountable and $P(j) = 0, j \in J$.

The planner does not know the probability distribution P , but knows that it belongs to a set of feasible treatment response distributions $\{P_\gamma, \gamma \in \Gamma\}$. γ will be called the *state of the world*. I assume that average treatment outcomes $E_\gamma[y(t)]$ are finite for all t and γ .

All population members are observationally identical to the planner, thus the planner's treatment assignment decision can be fully described by an action $a \in A, A = [0, 1]$, where a denotes the proportion of the target population assigned by the planner to the innovative treatment $t = 1$. Proportion $1 - a$, then, is assigned to the status quo treatment $t = 0$. I assume that fractional treatment assignment ($0 < a < 1$) is carried out randomly.

I consider planners whose welfare from taking action a in state of the world γ is the average treatment outcome across the population:

$$\begin{aligned} U(a, \gamma) &\equiv (1 - a) \cdot E_\gamma[y(0)] + a \cdot E_\gamma[y(1)] \\ &= E_\gamma[y(0)] + \theta_\gamma \cdot a. \end{aligned}$$

The second line expresses the welfare function in terms of the average treatment effect

$$\theta_\gamma \equiv E_\gamma[y(1)] - E_\gamma[y(0)],$$

which is the primary population statistic of interest to the planner.

The planner conducts an experiment and observes its outcome – a random vector $X \in \mathcal{X}$. The probability distribution of X depends on the unknown state of the world γ and will be denoted by Q_γ . A (random) function δ mapping feasible experimental outcomes from \mathcal{X} into actions from

A will be called a *statistical treatment rule* (or simply a *decision rule*). The action chosen by a planner with statistical treatment rule δ when X is observed will be denoted by $\delta(X)$. The set of all such functions (feasible statistical treatment rules) will be labeled \mathcal{D} .

I follow Wald's (1950) approach and evaluate alternative statistical treatment rules based on the expected welfare they yield across repeated samples in each state of the world γ . If the planner's welfare function is $U(a, \gamma)$, then the expected welfare from using statistical treatment rule δ in state of the world γ equals

$$(1) \quad \begin{aligned} W(\delta, \gamma) &\equiv \int_{X \in \mathcal{X}} U(\delta(X), \gamma) dQ_\gamma \\ &= E_\gamma[y(0)] + \theta_\gamma E_\gamma[\delta(X)], \end{aligned}$$

where $E_\gamma[\delta(X)]$ denotes $\int_{X \in \mathcal{X}} \delta(X) dQ_\gamma$.

Statistical treatment rule δ_2 dominates δ_1 if $W(\delta_2, \gamma) \geq W(\delta_1, \gamma)$ for all $\gamma \in \Gamma$ with strict inequality at least for one value of γ . Statistical treatment rule δ_1 is said to be *admissible* if there does not exist any $\delta_2 \in \mathcal{D}$ that dominates δ_1 , otherwise δ_1 is called *inadmissible*.

The analysis of this paper is based on a normalization of the expected welfare called *regret*. The regret of statistical treatment rule δ is the difference between the highest expected welfare achievable by any feasible statistical treatment rule in state of the world γ and the expected welfare of statistical treatment rule δ :

$$R(\delta, \gamma) \equiv \sup_{\delta' \in \mathcal{D}} W(\delta', \gamma) - W(\delta, \gamma).$$

The highest welfare in state of the world γ is achieved by statistical treatment rule $\delta_\gamma^*(X) = 1$ $|\theta_\gamma > 0|$ that selects the optimal (in state γ) treatment regardless of experimental outcomes. The regret function, then, equals

$$(2) \quad R(\delta, \gamma) = W(\delta_\gamma^*, \gamma) - W(\delta, \gamma) = \begin{cases} \theta_\gamma \cdot (1 - E_\gamma[\delta(X)]) & \text{if } \theta_\gamma > 0 \\ -\theta_\gamma \cdot E_\gamma[\delta(X)] & \text{if } \theta_\gamma \leq 0. \end{cases}$$

The regret of a statistical treatment rule, thus, is the product of the probability of making an error (assigning an individual to the wrong treatment) and the magnitude of the welfare loss suffered

from that error.

2.1 Treatment choice based on hypothesis testing

The most common framework used for treatment choice between a status quo treatment and an innovation is hypothesis testing. Typically, the researcher poses two mutually exclusive statistical hypotheses – a null hypothesis $H_0 : \theta_\gamma \leq 0$, that the innovation is no better than the status quo treatment, and an alternative hypothesis $H_1 : \theta_\gamma > 0$, that the innovation is superior. If the null hypothesis is rejected, then treatment $t = 1$ is assigned to the population. If it is not rejected, the status quo treatment $t = 0$ is assigned.

Rejecting the null hypothesis when it is, in fact, true (assigning an inferior innovation $t = 1$ to the population) is called a *Type I error*. Not rejecting the null hypothesis when it is, in fact, false (assigning the status quo treatment instead of the superior innovation) is called a *Type II error*. Hypothesis testing procedures are designed to have a certain *significance level*, which is the probability of making a Type I error (the maximum probability over states of the world γ that fall under the null hypothesis). The significance level (also called α -level) is usually set at conventional values $\alpha = 0.05$ or $\alpha = 0.01$.

The probability of not making a Type II error (assigning an innovation when it is superior to the status quo treatment) is called the *power of the test*. The power of the test is usually calculated for some specific value $\bar{\theta}_\gamma > 0$. The sample size of an experiment is selected so that a hypothesis test with a chosen significance level would have the desired power (typically .8 or .9) at $\bar{\theta}_\gamma$.

2.2 Treatment choice based on maximum regret

Savage (1951) introduced the criterion of minimizing maximum difference between potential and realized welfare (now called regret) in a review of Wald (1950) as a clarification of Wald's *minimax principle*. Under the *minimax regret* criterion, statistical treatment rule δ' is preferred to δ if

$$\max_{\gamma \in \Gamma} R(\delta', \gamma) < \max_{\gamma \in \Gamma} R(\delta, \gamma).$$

A planner who accepts the minimax-regret criterion should select a statistical treatment rule

that satisfies

$$(3) \quad \delta_M \in \arg \min_{\delta \in \mathcal{D}} \max_{\gamma \in \Gamma} R(\delta, \gamma)$$

and select a sample size such that the maximum regret $\max_{\gamma \in \Gamma} R(\delta_M, \gamma)$ is acceptable.

2.3 Asymmetric reference-dependent welfare

As a way to express the planner's desire to place the burden of proof on the innovation, I will also consider asymmetric reference-dependent welfare functions. For an asymmetry coefficient $K > 0$, let the welfare function $U_{A(K)}$ be linear in the average treatment outcomes with the same slope as U above the reference point $E_\gamma[y(0)]$ and a K times steeper slope below the reference point.

Formally, define $U_{A(K)}$ as:

$$\begin{aligned} U_{A(K)}(a, \gamma) &\equiv E_\gamma[y(0)] + \begin{cases} (U(a, \gamma) - E_\gamma[y(0)]) & \text{if } U(a, \gamma) > E_\gamma[y(0)], \\ K \cdot (U(a, \gamma) - E_\gamma[y(0)]) & \text{if } U(a, \gamma) \leq E_\gamma[y(0)], \end{cases} \\ &= E_\gamma[y(0)] + \begin{cases} \theta_\gamma \cdot a & \text{if } \theta_\gamma > 0, \\ K\theta_\gamma \cdot a & \text{if } \theta_\gamma \leq 0. \end{cases} \end{aligned}$$

The expected welfare for this kinked linear welfare function equals

$$(4) \quad \begin{aligned} W_{A(K)}(\delta, \gamma) &\equiv \int_{X \in \mathcal{X}} U_{A(K)}(\delta(X), \gamma) dQ_\gamma \\ &= E_\gamma[y(0)] + \begin{cases} \theta_\gamma E_\gamma[\delta(X)] & \text{if } \theta_\gamma > 0, \\ K \cdot \theta_\gamma E_\gamma[\delta(X)] & \text{if } \theta_\gamma \leq 0. \end{cases} \end{aligned}$$

Ordinal relationships between expected welfare of two statistical decision rules do not depend on the asymmetry factor $K > 0$. For any $\delta_1, \delta_2 \in \mathcal{D}$ and $\gamma \in \Gamma$:

$$W(\delta_2, \gamma) \begin{matrix} \geq \\ \leq \end{matrix} W(\delta_1, \gamma) \iff W_{A(K)}(\delta_2, \gamma) \begin{matrix} \geq \\ \leq \end{matrix} W_{A(K)}(\delta_1, \gamma).$$

Thus, the set of admissible statistical treatment rules is the same for all asymmetrical linear welfare functions (4) and for the standard linear welfare (1).

The regret function for expected welfare (4) equals

$$\begin{aligned}
R_{A(K)}(\delta, \gamma) &\equiv \sup_{\delta' \in \mathcal{D}} W_{A(K)}(\delta', \gamma) - W_{A(K)}(\delta, \gamma) \\
&= \begin{cases} \theta_\gamma \cdot (1 - E_\gamma[\delta(X)]) & \text{if } \theta_\gamma > 0, \\ -K\theta_\gamma \cdot E_\gamma[\delta(X)] & \text{if } \theta_\gamma \leq 0, \end{cases} \\
&= \begin{cases} R(\delta, \gamma) & \text{if } \theta_\gamma > 0, \\ KR(\delta, \gamma) & \text{if } \theta_\gamma \leq 0. \end{cases}
\end{aligned}$$

The only difference between this regret function and the regret function for standard linear welfare (2) is the factor K for $\theta_\gamma \leq 0$. Maximum regret under the asymmetrical welfare function can be expressed through the regret function for linear welfare as

$$\max_{\gamma \in \Gamma} R_{A(K)}(\delta, \gamma) = \max(K \cdot \bar{R}_{Type\ I}(\delta), \bar{R}_{Type\ II}(\delta)),$$

where

$$\bar{R}_{Type\ I}(\delta) \equiv \max_{\gamma: \theta_\gamma \leq 0} R(\delta, \gamma)$$

is the *maximum Type I regret* (maximum regret across states of the world in which the innovation is inferior) under the linear welfare function and

$$\bar{R}_{Type\ II}(\delta) \equiv \max_{\gamma: \theta_\gamma > 0} R(\delta, \gamma)$$

is the *maximum Type II regret* (maximum regret across states of the world in which the innovation is superior). The names Type I and Type II regret are given in analogy to Type I and Type II errors in hypothesis testing. Type I regret is the welfare loss due to Type I errors, while Type II regret is the welfare loss due to Type II errors under the null hypothesis $H_0 : \theta_\gamma \leq 0$.

Since the asymmetry factor K does not affect admissibility, I will only consider asymmetrical welfare functions indirectly, by solving the weighted minimax regret problem

$$(5) \quad \min_{\delta \in \mathcal{D}} \max(K \cdot \bar{R}_{Type\ I}(\delta), \bar{R}_{Type\ II}(\delta))$$

for the linear expected welfare (1). In problem (5) the planner gives K times greater weight to regret from Type I errors.

2.4 Treatment rules with limited Type I regret

Using minimax regret treatment rules may pose a particular problem for decision makers who do not have a choice over the precision of statistical evidence on which they have to base their decisions. Consider an extreme example. Suppose that a medical regulatory agency (the Food and Drug Administration in the United States or the European Agency for the Evaluation of Medicinal Products) has to choose whether to approve an innovative treatment for a common disease Z . The status quo medical treatment for disease Z has a proven record of curing the disease with probability .5. Proponents of the innovative treatment provide the regulator with results of a clinical trial in which five randomly selected patients with disease Z received the innovative treatment and in all five cases the disease has been cured. The minimax regret and hypothesis testing (at .05 level) statistical treatment rules both prescribe that all population members should be assigned to the innovation based on this experimental result. Both rules, however, imply much higher expected welfare losses if the innovation is inferior than clinical trials of usual size; the decision maker may not find this acceptable.

For decision makers who are primarily concerned with welfare loss due to mistakenly assigning an inferior innovation and cannot control the precision of experimental evidence, I propose the *limited Type I regret* criterion:

$$(6) \quad \delta_{L(\bar{r})} \in \arg \min_{\delta \in \mathcal{D}} \max_{\gamma \in \Gamma} R(\delta, \gamma),$$

$$s.t. \bar{R}_{Type I}(\delta) \leq \bar{r}.$$

The criterion selects a statistical treatment rule with minimal maximum regret, subject to an explicit constraint that Type I regret (regret from mistakenly assigning the innovation) should not exceed a given value \bar{r} . This criterion is similar to the classical hypothesis testing criterion in that both aim to limit the damage from Type I errors. Limited Type I regret, however, expresses the desired level of "safety" in terms of the maximum possible welfare loss from Type I errors, rather than just the maximum probability of making them.

3 Simple normal experiment

I will first consider a very simple experiment whose outcome $X \in \mathbb{R}$ is a scalar normally distributed random variable with unknown mean $\theta_\gamma \in \mathbb{R}$ and known variance σ^2 :

$$X \sim \mathcal{N}(\theta_\gamma, \sigma^2).$$

While X is a scalar, it need not originate from an experiment with sample size one. For example, X could be a sample average $X = \frac{1}{N} \sum_{i=1}^n x_i$ of N independent random observations. If observations (x_1, \dots, x_N) all have a normal distribution $\mathcal{N}(\theta_\gamma, \sigma_0^2)$, then X is a sufficient statistic for (x_1, \dots, x_N) with variance $\sigma^2 = \frac{\sigma_0^2}{N}$. Comparing single normal draw experiments with different values of σ , then, is equivalent to comparing experiments with different sample sizes.

More importantly, the probability distribution of many commonly used statistical estimators of average treatment effect converges to a normal distribution as sample size grows $\sqrt{N}(\hat{\theta} - \theta_\gamma) \xrightarrow{D} \mathcal{N}(0, \sigma_0^2)$. Then the asymptotic distribution of $\hat{\theta}$ is said to be $\mathcal{N}(\theta_\gamma, \frac{\sigma_0^2}{N})$. Heuristically, studying experiments with a single normally distributed outcome for different values of σ will suggest what effect different types of decision rules and sample sizes have on regret in more general settings.

It follows from the results of Karlin and Rubin (1956, Theorem 1) that if the distribution of X exhibits the monotone likelihood ratio property (which is true for normal and binomial distributions) and the welfare function is (1), then the class of *monotone decision rules*

$$\delta_{T,\lambda}(X) \equiv \begin{cases} 1 & X > T \\ \lambda & X = T \\ 0 & X < T \end{cases}, \lambda \in [0, 1], T \in \mathbb{R},$$

is essentially complete (for any decision rule δ' there exists $\delta_{T,\lambda}$ such that $W(\delta', \gamma) \leq W(\delta_{T,\lambda}, \gamma)$ in all states of the world). Since the probability of observing $X = T$ equals zero for the normal distribution, it follows that a smaller class of *threshold decision rules*

$$\delta_T(X) \equiv 1 |X > T|, T \in \mathbb{R}$$

is also essentially complete. Thus, considering other rules is not necessary in this problem.

Given that X is normally distributed, the regret of a threshold decision rule δ_T in state of the world γ equals

$$R(\delta_T, \gamma) = \begin{cases} \theta_\gamma \cdot P_\gamma(X \leq T) = \theta_\gamma \cdot \Phi\left(\frac{T - \theta_\gamma}{\sigma}\right) & \text{if } \theta_\gamma > 0, \\ -\theta_\gamma \cdot P_\gamma(X > T) = -\theta_\gamma \cdot \Phi\left(\frac{\theta_\gamma - T}{\sigma}\right) & \text{if } \theta_\gamma \leq 0, \end{cases}$$

which is the probability of making an incorrect decision multiplied by $|\theta_\gamma|$, the magnitude of the loss incurred from the mistake. Φ denotes the standard normal cumulative distribution function.

Maximum Type I and Type II regret equal

$$(7) \quad \begin{aligned} \bar{R}_{Type\ I}(\delta_T) &= \max_{\gamma: \theta_\gamma \leq 0} \left\{ -\theta_\gamma \cdot \Phi\left(\frac{\theta_\gamma - T}{\sigma}\right) \right\} = \sigma \cdot \max_{h \leq 0} \left\{ -h\Phi\left(h - \frac{T}{\sigma}\right) \right\}, \\ \bar{R}_{Type\ II}(\delta_T) &= \max_{\gamma: \theta_\gamma > 0} \left\{ \theta_\gamma \cdot \Phi\left(\frac{T - \theta_\gamma}{\sigma}\right) \right\} = \sigma \cdot \max_{h > 0} \left\{ h\Phi\left(\frac{T}{\sigma} - h\right) \right\}. \end{aligned}$$

The right-hand equalities are derived by substituting $h = \frac{\theta_\gamma}{\sigma}$. These functions have finite positive values for every $T \in \mathbb{R}$. Since $R(\delta_T, \theta_\gamma) = R(\delta_{-T}, -\theta_\gamma)$, it follows that $\bar{R}_{Type\ II}(\delta_T) = \bar{R}_{Type\ I}(\delta_{-T})$. Lemma 1 shows that the decision maker faces a trade off between maximum Type I and maximum Type II regret. Higher threshold values imply lower Type I regret, but necessarily higher Type II regret.

Lemma 1 a) $\bar{R}_{Type\ I}(\delta_T)$ is a continuous, strictly decreasing function of T ,

$$\lim_{T \rightarrow -\infty} \bar{R}_{Type\ I}(\delta_T) = \infty \text{ and } \lim_{T \rightarrow \infty} \bar{R}_{Type\ I}(\delta_T) = 0,$$

b) $\bar{R}_{Type\ II}(\delta_T)$ is a continuous, strictly increasing function of T ,

$$\lim_{T \rightarrow -\infty} \bar{R}_{Type\ II}(\delta_T) = 0 \text{ and } \lim_{T \rightarrow \infty} \bar{R}_{Type\ II}(\delta_T) = \infty.$$

Figure 1 displays the maximum Type I and maximum Type II regret as functions of the decision rule threshold T . The scale of both axes is normalized by σ . The maximum regret $\max_{\gamma \in \Gamma} R(\delta_T, \theta_\gamma) = \max(\bar{R}_{Type\ I}(\delta_T), \bar{R}_{Type\ II}(\delta_T))$ is minimized when $\bar{R}_{Type\ I}(\delta_T) = \bar{R}_{Type\ II}(\delta_T)$, which happens only at $T = 0$. The minimax regret treatment rule in this problem is δ_0 . This is sometimes called the *plug-in rule* (a plug-in rule takes the estimated value of the average treatment effect and assigns

treatments as if it were the true value).

Similarly, the minimax regret statistical treatment rule under asymmetric welfare function $W_{A(K)}$ is uniquely characterized by the equation

$$K \cdot \bar{R}_{Type I}(\delta_T) = \bar{R}_{Type II}(\delta_T).$$

By substituting right-hand expressions from (7), this characterization can be rewritten as

$$K \cdot \max_{h \leq 0} \left\{ -h\Phi \left(h - \frac{T}{\sigma} \right) \right\} = \max_{h > 0} \left\{ h\Phi \left(\frac{T}{\sigma} - h \right) \right\}.$$

Since only one value of $\frac{T}{\sigma}$ solves the equation for a given K , the threshold of the minimax regret statistical treatment rule is proportional to σ .

A conventional one-sided hypothesis test with significance level α rejects the null hypothesis ($\theta \leq 0$) and assigns the innovative treatment if $X > \sigma\Phi^{-1}(1 - \alpha)$. This critical value guarantees that the probability of a Type I error does not exceed α for any $\theta_\gamma \leq 0$. Since $\frac{X - \theta_\gamma}{\sigma}$ has a standard normal distribution,

$$\begin{aligned} P(X > \sigma\Phi^{-1}(1 - \alpha)) &= 1 - P\left(\frac{X - \theta_\gamma}{\sigma} \leq \Phi^{-1}(1 - \alpha) - \frac{\theta_\gamma}{\sigma}\right) = \\ &= 1 - \Phi\left(\Phi^{-1}(1 - \alpha) - \frac{\theta_\gamma}{\sigma}\right) \leq \\ &\leq 1 - \Phi(\Phi^{-1}(1 - \alpha)) = \alpha. \end{aligned}$$

The statistical treatment rule based on results of a hypothesis test with level α is a threshold rule $\delta_{H(\alpha)}$ with threshold $H(\alpha) \equiv \sigma\Phi^{-1}(1 - \alpha)$. For a given test level α , the threshold T is proportional to the standard error σ . Thus a hypothesis test based treatment rule can be rationalized as a solution to an asymmetrical minimax regret problem with asymmetry factor

$$K(\alpha) = \frac{\max_{h > 0} \{h\Phi(H(\alpha)/\sigma - h)\}}{\max_{h \leq 0} \{-h\Phi(h - H(\alpha)/\sigma)\}} = \frac{\max_{h > 0} \{h\Phi(\Phi^{-1}(1 - \alpha) - h)\}}{\max_{h \leq 0} \{-h\Phi(h - \Phi^{-1}(1 - \alpha))\}}.$$

$K(\alpha)$ is the ratio of maximum Type II to maximum Type I regret of the hypothesis test based decision rule, which depends only on the test level α . In this normal model, the correspondence

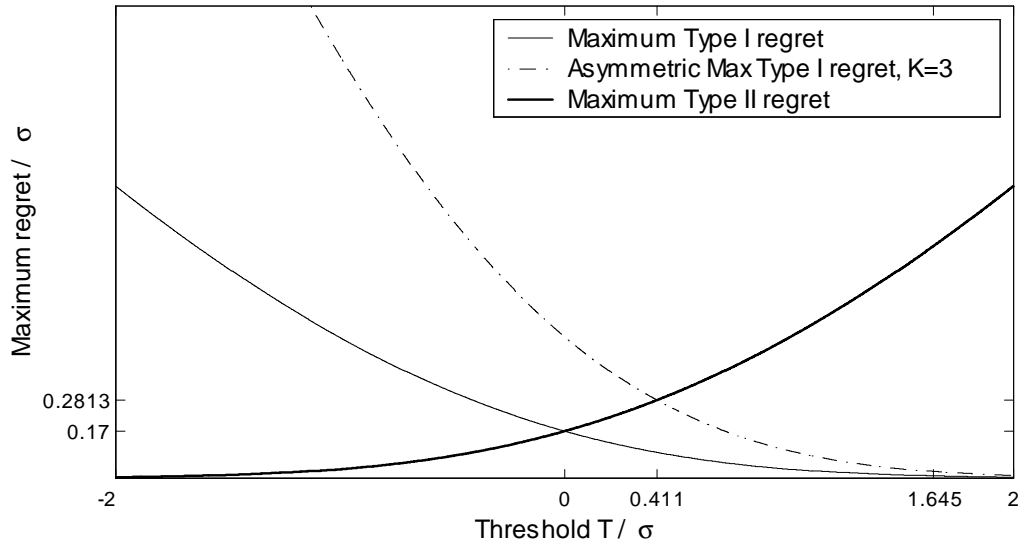


Figure 1: Maximum Type I and Type II regret as functions of the decision rule threshold.

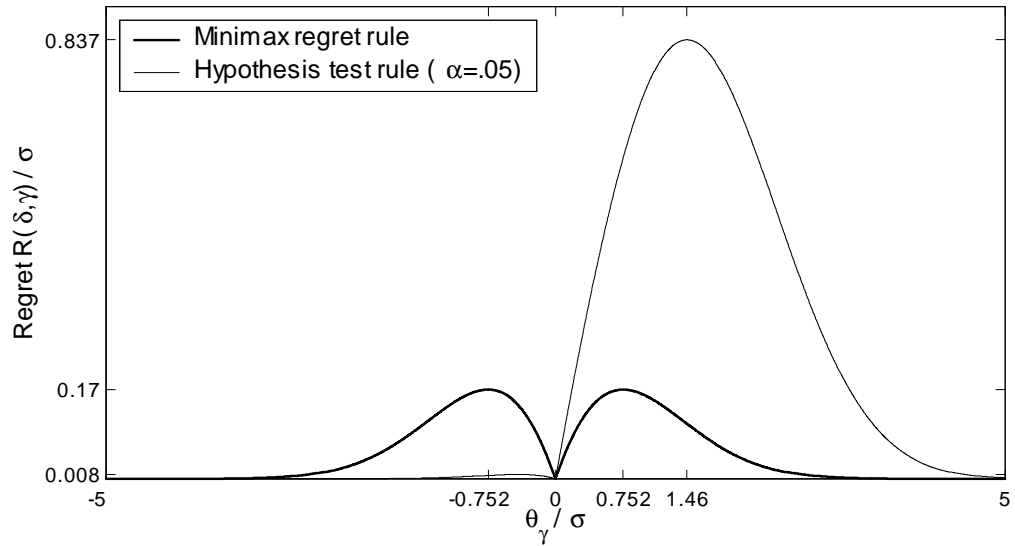


Figure 2: Regret functions of minimax regret and hypothesis test based decision rules.

Test significance level	Threshold	Max Type I regret	Max Type II regret	$K(\alpha)$
$\alpha = .5$ (minimax regret)	$T = 0$	$.17\sigma$	$.17\sigma$	1
$\alpha = .25$	$T = .6745\sigma$	$.0608\sigma$	$.3724\sigma$	6.125
$\alpha = .1$	$T = 1.282\sigma$	$.01877\sigma$	$.6409\sigma$	34.15
$\alpha = .05$	$T = 1.645\sigma$	$.008178\sigma$	$.8371\sigma$	102.4
$\alpha = .025$	$T = 1.96\sigma$	$.003665\sigma$	1.026σ	279.9
$\alpha = .01$	$T = 2.326\sigma$	$.001304\sigma$	1.264σ	969.6

Table 1: Maximum Type I and Type II regret of statistical treatment rules induced by hypothesis tests based on a normally distributed estimate with variance σ^2 .

between a hypothesis test based rule with level α and an asymmetric minimax regret rule with level $K(\alpha)$ does not depend on the standard error of σ , and thus on sample size. This feature is specific to the normal example. For example, if X is a binomial variable, then hypothesis test based rules with the same level correspond to different asymmetric minimax regret treatment rules for different sample sizes.

Table 1 provides maximum Type I and II regret values and the asymmetry factors corresponding to commonly used hypothesis test levels. Decision rules based on the one-sided $\alpha = .05$ level hypothesis test minimize maximum regret for decision makers who place 102 times greater weight on Type I regret than on Type II regret. Decision rules based on $\alpha = .01$ level tests are minimax regret for decision makers who place nearly 970 times greater weight on Type I regret. The trade off between Type I and Type II regret is markedly different from the trade off between raw Type I and Type II error rates (an $\alpha = .05$ level test has a 95% maximum probability of Type II error, which is 19 times higher than the maximum probability of the test's Type I error).

Figure 2 compares the regret functions of the minimax regret treatment rule δ_0 and the treatment rule $\delta_{H(.05)}$ induced by a hypothesis test with significance level $\alpha = .05$ over a range of feasible values of θ_γ . The scale of both axes is normalized by σ . The maximum regret of the hypothesis test rule is approximately $.837\sigma$, which is nearly five times higher than the maximum regret of the minimax regret treatment rule (approximately $.17\sigma$). The hypothesis test rule has lower regret over $\theta_\gamma \leq 0$, but it can only achieve it by greatly increasing the regret for $\theta_\gamma > 0$. The greatest expected welfare losses from using a hypothesis test rule occur when the innovative treatment is moderately effective.

3.1 Limited Type I regret

Compared to the minimax regret rule, hypothesis testing with significance level $\alpha = .05$ has a clear advantage in lower regret over $\theta_\gamma \leq 0$. This can make minimax regret unattractive for decision makers who are more concerned about negative consequences of accepting a potentially inferior new treatment than about its potential foregone benefits. I do not think, however, that hypothesis testing practices adequately address such concerns. It is common to see tests with the same significance level $\alpha = .05$ applied to treatment effect estimates with different variance and sample size. While such tests always limit the probability of Type I error to .05, the maximum Type I regret ($\approx .008\sigma$) is proportional to σ .

Many decision makers, no doubt, would like to sensibly adjust the test level to the circumstances of a particular problem. Considering maximum Type I regret of a threshold rule instead of its the maximum probability of Type I error simplifies this task. Table 1 provides maximum Type I and Type II regret values for threshold rules corresponding to hypothesis tests with different significance levels.

Instead of imposing a limit on the probability of Type I errors, the decision maker could directly impose a limit \bar{r} on maximum acceptable Type I regret and use the limited Type I regret criterion (6). The limited Type I statistical treatment rule coincides with the minimax regret rule if its maximum regret $.17\sigma$ does not exceed \bar{r} . Otherwise, it selects a treatment rule with the smallest threshold $T > 0$ that ensures that maximum Type I regret does not exceed \bar{r} . If the estimator has high variance σ^2 , reducing maximum Type I regret comes at a price of higher Type II regret. For example, if the decision maker finds that a threshold value $T = 1.645\sigma$ is required to bring maximum Type I regret to an acceptable level \bar{r} , she has to accept that such statistical treatment rule implies a maximum Type II regret that is over 100 times larger than \bar{r} . This underscores the importance of using estimators of treatment effect with low variance (high sample size), which allow the decision maker to attain acceptable maximum Type I regret with statistical treatment rules that have lower Type II regret.

3.2 Sample size selection

I will illustrate sample size selection based on maximum regret by comparing it with one of the conventional methods. The International Conference on Harmonisation formulated "Guideline E9: Statistical Principles for Clinical Trials" (1998), adopted by the US Food and Drug Administration and the European Agency for the Evaluation of Medicinal Products. The guideline provides researchers with the values of Type I and Type II errors typically used for hypothesis testing and sample size selection in clinical trials. For hypothesis testing, the limit on the probability of Type I errors is traditionally set at 5% or less. The trial sample size is typically selected to limit the probability of Type II errors to 10-20% for a minimal value of the treatment effect that is deemed to have "clinical relevance" or at the anticipated value of the effect of the innovative treatment.

Suppose that a researcher considers bearable the loss of public welfare due to a 10% probability that her innovative treatment could be rejected if its actual treatment effect equals $\bar{\theta} > 0$. Following the convention, she selects the sample size for which the variance of X equals $\bar{\sigma}^2$, where $\bar{\sigma}^2$ satisfies the condition that X will fall under the 5% hypothesis test threshold $H(.05) = \bar{\sigma}\Phi^{-1}(.95)$ with probability 10% if $\theta_\gamma = \bar{\theta}$:

$$P(X \leq H(.05) | \theta_\gamma = \bar{\theta}) = \Phi\left(\Phi^{-1}(.95) - \frac{\bar{\theta}}{\bar{\sigma}}\right) = .1,$$

$$\bar{\sigma} = \frac{\bar{\theta}}{\Phi^{-1}(.95) - \Phi^{-1}(.1)} = \frac{\bar{\theta}}{2.926}.$$

The value of regret that the researcher finds acceptable at $\theta_\gamma = \bar{\theta}$ thus equals $.1\bar{\theta}$. This procedure does not make apparent to the researcher that a much larger welfare loss will be suffered at a twice smaller value of $\theta_\gamma = 1.46\bar{\sigma} \approx .5\bar{\theta}$, where the regret function achieves its maximum of $.837\bar{\sigma} = .286\bar{\theta}$.

Consider now how the sample size would differ if it were selected by the researcher with an explicit objective that maximum regret should equal $.1\bar{\theta}$ in two scenarios. First, suppose that the researcher planning the experiment has to take for granted that the decision making will be carried out using a 5% hypothesis test rule. Since its maximum regret equals $.837\bar{\sigma}$, she would select

sample size such that

$$\begin{aligned} .837\sigma &= .1\bar{\theta} \\ \sigma &= \frac{.1\bar{\theta}}{.837} = \frac{.1 \cdot 2.926\bar{\sigma}}{.837} = .35\bar{\sigma}, \end{aligned}$$

which implies sample size that is over 8 times larger than the one selected by power calculations in the example above. In a second scenario, suppose that the researcher has control over treatment assignment and plans to use the minimax-regret decision rule δ_0 . Since the maximum regret of the minimax-regret decision rule equals $.17\sigma$, the sample size should be such that

$$\begin{aligned} .17\sigma &= .1\bar{\theta}_\gamma \\ \sigma &= 1.722\bar{\sigma}, \end{aligned}$$

which implies sample size that is almost 3 times smaller than the one selected by power calculations.

3.3 Normally distributed outcomes with unknown variance

So far in this section I have assumed that the planner knows the variance of the normally distributed average treatment effect estimate X . Suppose now, instead, that the data (x_1, \dots, x_N) consists of N independent normally distributed observations with unknown mean θ_γ and unknown variance σ_γ^2 . Let the set of feasible states of the world be

$$\Gamma \equiv \{\gamma : \theta_\gamma \in \mathbb{R}, \sigma_\gamma^2 \in [\underline{\sigma}^2, \bar{\sigma}^2]\},$$

where $\underline{\sigma}^2 > 0$ and $\bar{\sigma}^2 < \infty$ and let

$$\bar{\Gamma} \equiv \{\gamma : \theta_\gamma \in \mathbb{R}, \sigma_\gamma^2 = \bar{\sigma}^2\}$$

denote the subset of states of the world with the highest feasible outcome variance. Let $\bar{X} \equiv \frac{1}{N} \sum_{i=1}^N x_i$ be the sample mean and $S^2 \equiv \frac{1}{N} \sum_{i=1}^N (x_i - \bar{X})^2$ the sample variance. It is well known (cf. Berger, 1985) that the pair (\bar{X}, S^2) is a sufficient statistic for (x_1, \dots, x_N) , thus only decision rules that are functions of \bar{X} and S^2 need to be considered. It turns out, however, that decision

rules satisfying criteria based on maximum Type I and Type II regret could often be found within a smaller class of threshold decision rules that depend only on the sample mean \bar{X} .

Proposition 2 *Let $\delta_T \equiv 1 \mid \bar{X} > T \mid$ be a threshold statistical treatment rule such that $T^* \equiv \frac{\sqrt{N}}{\bar{\sigma}} |T|$ satisfies the condition*

$$(8) \quad \max_{h \in (0, T^*)} \left\{ h \cdot \Phi \left(\frac{\bar{\sigma}}{\underline{\sigma}} (T^* - h) \right) \right\} \leq \max_{h \geq T^*} \{ h \cdot \Phi (T^* - h) \},$$

then

a) *maximum Type I and Type II regret of δ_T over the set Γ is the same as over the set $\bar{\Gamma}$:*

$$\begin{aligned} \max_{\gamma \in \Gamma: \theta_\gamma \leq 0} R(\delta_T, \gamma) &= \max_{\gamma \in \bar{\Gamma}: \theta_\gamma \leq 0} R(\delta_T, \gamma), \\ \max_{\gamma \in \Gamma: \theta_\gamma > 0} R(\delta_T, \gamma) &= \max_{\gamma \in \bar{\Gamma}: \theta_\gamma > 0} R(\delta_T, \gamma), \end{aligned}$$

b) *there is no statistical treatment rule $\delta'(\bar{X}, S^2)$ that has both lower maximum Type I regret and lower maximum Type II regret than δ_T .*

Condition (8) ensures that the threshold decision attains maximum Type I and maximum Type II regret on the subset $\bar{\Gamma}$. If it is not satisfied, the maximum Type I or maximum Type II regret of δ_T could be higher on the set Γ than on $\bar{\Gamma}$, then there maybe exists a non-threshold decision rule that has both lower Type I and lower Type II regret than δ_T .

It follows from Proposition 2 that threshold decision rules that satisfy minimax regret, asymmetric minimax regret, and limited Type I regret criteria for outcomes with fixed variance (set of feasible states of the world $\bar{\Gamma}$) also satisfy the corresponding criteria for outcomes with bounded variance (set of feasible states Γ) if their threshold values satisfy condition (8). The range of thresholds for which condition (8) holds depends on the ratio $\frac{\bar{\sigma}}{\underline{\sigma}}$. For $\frac{\bar{\sigma}}{\underline{\sigma}} = 1$, it holds if $\frac{\sqrt{N}}{\bar{\sigma}} |T| \leq 1.25$. In the opposite extreme case when $\frac{\bar{\sigma}}{\underline{\sigma}} \rightarrow \infty$, it holds if $\frac{\sqrt{N}}{\bar{\sigma}} |T| \leq .22$.

4 Exact statistical treatment rules for binary and bounded outcomes

Exact solutions to the minimax regret and limited type I regret problems and exact maximum regret values are available when the data X consists of N independent random outcomes of treatment $t = 1$, provided that the outcomes are binary or have bounded values. I will first consider the case of binary outcomes and then its extension to outcomes with bounded values.

4.1 Binary outcomes

Let the treatment outcomes of the innovative treatment $t = 1$ be binary, w.l.o.g. let $y(1) \in \{0, 1\}$, and let the known average outcome of the status quo treatment $t = 0$ equal $p_0 \equiv E[y(0)] \in (0, 1)$. Let the set of feasible probability distributions of $y(1)$ be a set of Bernoulli distribution with means $p_\gamma \in [a, b]$, $0 \leq a < p_0 < b \leq 1$ (if p_0 is outside of the interval $[a, b]$, then the treatment choice problem is trivial). The experimental data consists of N independent random outcomes (x_1, \dots, x_N) , each having a Bernoulli distribution with mean p_γ . The sum of outcomes $X = \sum_{i=1}^n x_i$ has a binomial distribution with parameters N and p_γ . X is a sufficient statistic for (x_1, \dots, x_N) , so it is sufficient to consider statistical treatment rules that are functions of X .

It follows from the results of Karlin and Rubin (1956, Theorems 1 and 4) that monotone statistical treatment rules

$$\delta_{T,\lambda}(X) = \begin{cases} 1 & X > T \\ \lambda & X = T \\ 0 & X < T \end{cases}, T \in \{0, \dots, N\}, \lambda \in [0, 1]$$

are admissible and form an essentially complete class, thus it is sufficient to consider only monotone rules. The regret of a monotone rule $\delta_{T,\lambda}$ equals

$$R(\delta_{T,\lambda}, \gamma) = \begin{cases} \theta_\gamma \cdot \left\{ \lambda B(T, N, p_\gamma) + \sum_{T < n \leq N} B(n, N, p_\gamma) \right\} & \text{if } \theta_\gamma > 0, \\ -\theta_\gamma \cdot \left\{ 1 - \left(\lambda B(T, N, p_\gamma) + \sum_{T < n \leq N} B(n, N, p_\gamma) \right) \right\} & \text{if } \theta_\gamma \leq 0, \end{cases}$$

where $B(n, N, p_\gamma)$ denotes the binomial probability density function with parameters N and p_γ and $\theta_\gamma \equiv p_\gamma - p_0$.

It will be convenient to use a one-dimensional index for monotone rules $D(\delta_{T,\lambda}) \equiv T + (1 - \lambda)$. There is a one to one correspondence between index values $D \in [0, N + 1]$ and the set of all distinct monotone decision rules. $D = 0$ corresponds to the decision rule that assigns all population members to the innovation, no matter what the experimental outcomes are. $D = N + 1$ corresponds to the most conservative decision rule that always assigns the status quo treatment.

Lemma 3 establishes properties of maximum Type I and Type II regret of monotone statistical treatment rules for binomially distributed X that lead to simple characterisations of minimax regret, asymmetric minimax regret, and limited Type I regret rules. As before, maximum Type I regret is $\bar{R}_{Type\ I}(\delta) \equiv \max_{\gamma:p_\gamma \in [a,p_0]} R(\delta, \gamma)$ and maximum Type II regret is $\bar{R}_{Type\ II}(\delta) \equiv \max_{\gamma:p_\gamma \in (p_0,b]} R(\delta, \gamma)$.

Lemma 3 *If X has a binomial distribution, then*

- a) $\bar{R}_{Type\ I}(\delta)$ is a continuous and strictly decreasing function of $D(\delta)$ with $\bar{R}_{Type\ I}(\delta) = 0$ for $D(\delta) = N + 1$.
- b) $\bar{R}_{Type\ II}(\delta)$ is a continuous and strictly increasing function of $D(\delta)$ with $\bar{R}_{Type\ II}(\delta) = 0$ for $D(\delta) = 0$.

It follows from lemma 3 that there is a unique value of $D(\delta_M)$ such that $\bar{R}_{Type\ I}(\delta_M) = \bar{R}_{Type\ II}(\delta_M)$. δ_M is the minimax regret treatment rule. While its characterisation is implicit, monotonicity and continuity of the maximum Type I and Type II regret as functions of $D(\delta)$ makes computation very straightforward. The same characterisation of the minimax regret treatment rule for $p_\gamma \in [0, 1]$ was derived by Stoye (2006, Proposition 1(iii)) using game theoretic methods.

Likewise, there is a unique value $D(\delta_{A(K)})$ such that $K \cdot \bar{R}_{Type\ I}(\delta_{A(K)}) = \bar{R}_{Type\ II}(\delta_{A(K)})$. $\delta_{A(K)}$ is the minimax regret statistical treatment rule for asymmetric reference dependent welfare function $W_{A(K)}$. Limited Type I regret statistical treatment rule with Type I regret threshold \bar{r} is also easily characterized. If $\bar{r} \geq \max_{\gamma \in \Gamma} R(\delta_M, \gamma)$, then there is a unique value $D(\delta_{L(\bar{r})})$ such that $\bar{R}_{Type\ I}(\delta_{L(\bar{r})}) = \bar{r}$, then $\delta_{L(\bar{r})}$ is the limited Type I regret treatment rule. If $\bar{r} < \max_{\gamma \in \Gamma} R(\delta_M, \gamma)$, then the Type I regret constraint is not binding and the limited Type I regret treatment rule is the same as the minimax regret treatment rule δ_M .

The following proposition derives the exact large sample limit of maximum regret of minimax-

regret statistical treatment rules. Unlike in the normal case covered in Section 3, the minimax-regret rule in the Bernoulli case does not generally coincide with the plug-in rule:

$$\delta_P \equiv 1 \left| \frac{X}{N} > p_0 \right|.$$

In large samples, however, the difference between δ_M and δ_P has little effect on maximum regret. Proposition 4 shows that as sample size grows, the maximum regret of minimax regret rules and plug-in rules (normalized by \sqrt{N}) converge to the same limit. That limit is the same as minimax regret in a problem with N normally distributed outcomes with fixed variance $p_0(1-p_0)$.

Proposition 4

$$\lim_{N \rightarrow \infty} \sqrt{\frac{N}{p_0(1-p_0)}} \max_{\gamma \in \Gamma} R(\delta_P, \gamma) = \lim_{N \rightarrow \infty} \sqrt{\frac{N}{p_0(1-p_0)}} \max_{\gamma \in \Gamma} R(\delta_M, \gamma) = \max_{h > 0} [h\Phi(-h)] \approx .17.$$

4.2 Bounded outcomes

Now consider a more general setting. Let the outcomes of treatment $t = 1$ be bounded variables $y(1) \in [0, 1]$. Let $p_0 \equiv E[y(0)] \in (0, 1)$ denote the known average treatment outcome of the status quo treatment $t = 0$. Let $\{P_\gamma, \gamma \in \Gamma\}$ be the set of probability distributions $P[y(1)]$ that the planner considers feasible. Assume that $E_\gamma[y(1)] \in [a, b], 0 \leq a < p_0 < b \leq 1$. Also, let $\{P_\gamma, \gamma \in \Gamma_B\}$ denote the set of all Bernoulli distributions with $E_\gamma[y(1)] \in [a, b]$ and assume that $\Gamma_B \subset \Gamma$. The technique outlined below relies on including all the Bernoulli distributions in the feasible set.

Schlag (2007) proposed an elegant technique, which he calls the *binomial average*, that extends statistical treatment rules defined for samples of Bernoulli outcomes to samples of bounded outcomes. The resulting statistical treatment rules inherit important properties of their Bernoulli ancestors. Let $\delta : \{0, \dots, N\} \rightarrow [0, 1]$ be a statistical treatment rule defined for the sum of N i.i.d. Bernoulli distributed outcomes (as in the previous subsection). Let $X = (x_1, \dots, x_N)$ be an i.i.d. sample of bounded random variables with unknown distribution $P_\gamma[y(1)]$ and let $Z = (z_1, \dots, z_N)$ be a sample of i.i.d. uniform $(0, 1)$ random variables independent of X . Then the binomial average

extension of δ is defined as

$$\bar{\delta}(X) \equiv E_Z \delta \left(\sum_{k=0}^N 1[z_k \leq x_k] \right).$$

Verbally, this extension can be described as a simple process:

- a) randomly replace each bounded observation $x_i \in [0, 1]$ with a Bernoulli observation $\tilde{x}_i = 1$ with probability x_i and with $\tilde{x}_i = 0$ with probability $1 - x_i$,
- b) apply statistical treatment rule δ to $(\tilde{x}_1, \dots, \tilde{x}_N)$.

The random variables $1[z_k \leq x_k], k = 0, \dots, N$ are i.i.d. Bernoulli with expectation $E_\gamma[y(1)]$, thus $\sum_{k=0}^N 1[z_k \leq x_k]$ has a Binomial distribution with parameters N and $E_\gamma[y(1)]$. For any state of the world γ , let $\bar{\gamma}$ be the state of the world in which $P_{\bar{\gamma}}[y(1)]$ is a Bernoulli distribution with the same mean $E_\gamma[y(1)]$. Then $E_\gamma(\tilde{\delta}) = E_{\bar{\gamma}}(\delta)$ and $R(\tilde{\delta}, \gamma) = R(\delta, \bar{\gamma})$. The regret of a binomial average treatment rule $\tilde{\delta}$ in state of the world γ is the same as the regret of δ in a Bernoulli state of the world $\bar{\gamma}$ with the same mean treatment outcomes. It follows that maximum Type I (II) regret of $\tilde{\delta}$ in the problem with bounded outcomes ($\gamma \in \Gamma$) is equal to maximum Type I (II) regret of δ in the problem with Bernoulli outcomes ($\gamma \in \Gamma_B$).

If statistical treatment rule δ satisfies some decision criterion based on maximum Type I and maximum Type II regret for the feasible set of Bernoulli outcome distributions, then its binomial average extension $\tilde{\delta}$ satisfies the same criterion for the feasible set of bounded outcome distributions. Suppose, for example, that δ_M minimizes maximum regret for Bernoulli distributions. Suppose there was a treatment rule δ' for bounded distributions that had lower maximum regret than $\tilde{\delta}_M$. Then δ' would have to have lower maximum regret over Γ_B than δ_M , which would imply that δ_M does not minimize maximum regret for the problem with Bernoulli distributions.

Binomial average extension yields exact minimax regret, asymmetric minimax regret and limited Type I regret statistical treatment rules if the set of feasible outcome distributions Γ includes the set of Bernoulli outcome distributions with the same means Γ_B . In many applications, however, the planner knows that Bernoulli outcome distributions are not feasible. If the outcome variable is annual income of a participant in a job training program, the planner may assume not only that the variable is bounded, but also that its variance is much smaller than the variance of a Bernoulli distribution with the same mean. If Bernoulli outcome distributions are excluded, then binomial average based treatment rules may not be optimal. The following proposition shows that a plug-in

statistical treatment rule

$$\delta_P \equiv 1 \left| \frac{1}{N} \sum_{i=1}^N x_i > p_0 \right|$$

has lower asymptotic maximum regret than a binomial average extension of δ_M , a minimax regret statistical treatment rule in the Bernoulli case.

Proposition 5 *Let $p_0 = E[y(0)]$ and let $\{P_\gamma, \gamma \in \Gamma\}$ be the set of feasible probability distributions of $y(1)$ such that $E_\gamma(y(1) - E_\gamma[y(1)])^2 < \sigma_h^2$, where $\sigma_h^2 < p_0(1 - p_0)$. Let (x_1, \dots, x_N) be i.i.d. random outcomes of treatment $t = 1$. Then*

$$\sqrt{N} \sup_{\gamma \in \Gamma} R(\delta_P, \gamma) \leq \sigma_h \cdot \max_{h>0} [h\Phi(-h)] + o(1).$$

Maximum regret of binomial average extension $\tilde{\delta}_M$ is by design the same as the maximum regret of the minimax regret treatment rule δ_M in the Bernoulli case. As long as for some $\Delta > 0$, Γ contains distributions with all possible means in a Δ -neighborhood of p_0

$$\forall p \in [p_0 - \Delta, p_0 + \Delta], \exists \gamma : E_\gamma[y(1)] = p,$$

the results of proposition 4 apply and

$$\lim_{N \rightarrow \infty} \sqrt{N} \max_{\gamma \in \Gamma} R(\tilde{\delta}_M, \gamma) = \sqrt{p_0(1 - p_0)} \cdot \max_{h>0} [h\Phi(-h)] > \sigma_h \cdot \max_{h>0} [h\Phi(-h)].$$

Thus, for large enough N , $\max_{\gamma \in \Gamma} R(\tilde{\delta}_M, \gamma) > \sup_{\gamma \in \Gamma} R(\delta_P, \gamma)$. This underscores the importance of placing appropriate restrictions on the set of feasible treatment outcome distributions before looking for minimax regret or asymmetric maximum regret based treatment rules.

5 Evaluating regret using approximations and bounds

In conclusion, I would like to discuss methods for dealing with statistical problems which do not have neat finite sample solutions such as described in the previous sections and give an example illustrating their properties. I will restrict attention to the case when the data consists of N i.i.d. observations (x_1, \dots, x_N) such that $E[x_i] = \theta_\gamma$, where $\theta_\gamma \equiv E[y(1)] - E[y(0)]$ is the average

treatment effect. For many sets of feasible distributions of $\{x_i\}$, there aren't proven complete class theorems that justify restricting attention to a small class of decision rules. Considering all feasible statistical treatment rules that are functions of (x_1, \dots, x_N) can be prohibitively difficult, but progress can be made by considering a suitable subset of feasible decision rules. Based on their sufficiency in an idealized problem with normally distributed outcomes considered in Section 3, the class of threshold decision rules $\delta_T \equiv 1[\bar{X} > T]$ based on the sample mean $\bar{X} \equiv \frac{1}{N} \sum_{i=1}^N x_i$ is a reasonable and tractable candidate class of statistical treatment rules to consider.

The regret of a threshold decision rule δ_T equals

$$R(\delta_T, \gamma) = \begin{cases} \theta_\gamma \cdot P_\gamma(\bar{X} \leq T) & \text{if } \theta_\gamma > 0, \\ -\theta_\gamma \cdot P_\gamma(\bar{X} > T) & \text{if } \theta_\gamma \leq 0. \end{cases}$$

To evaluate maximum Type I and Type II regret of δ_T ,

$$\begin{aligned} \bar{R}_{Type\ I}(\delta_T) &= \sup_{\theta \leq 0} \left\{ -\theta \cdot \sup_{\gamma: \theta_\gamma = \theta} P_\gamma(\bar{X} > T) \right\}, \\ \bar{R}_{Type\ II}(\delta_T) &= \sup_{\theta > 0} \left\{ \theta \cdot \sup_{\gamma: \theta_\gamma = \theta} P_\gamma(\bar{X} \leq T) \right\}, \end{aligned}$$

the planner needs to know, for each value of θ , the range of feasible probabilities that the sample mean \bar{X} exceeds the threshold T . Note that for each γ , $P_\gamma(\bar{X} > T)$ is a non-increasing function of T and $P_\gamma(\bar{X} \leq T)$ is non-decreasing. It follows that $\bar{R}_{Type\ I}(\delta_T)$ is non-increasing and $\bar{R}_{Type\ II}(\delta_T)$ is non-decreasing in T , thus solutions to minimax regret, asymmetric minimax regret, and limited Type I regret problems can be easily found if the researcher has a way to evaluate $\bar{R}_{Type\ I}(\delta_T)$ and $\bar{R}_{Type\ II}(\delta_T)$. The problem of evaluating $P_\gamma(\bar{X} \lesseqgtr T)$ for distributions of x_i that do not yield a convenient closed-form expression is well studied in statistics. I will consider three main approaches: brute force calculation or simulation, asymptotic approximation, and large deviation bounds.

Brute force calculation or simulation The main challenge for calculation or simulation is in selecting a finite set of feasible distributions that reliably approximates $\sup_{\gamma: \theta_\gamma = \theta} P_\gamma(\bar{X} \leq T)$ or $\sup_{\gamma: \theta_\gamma = \theta} P_\gamma(\bar{X} > T)$ for different values of θ . For some distributions (e.g. for discrete distributions with small finite support) such a set is easily constructed by creating a "grid" of distributions

with different parameter values. In nonparametric problems, however, it may be difficult to construct a finite set of distributions that will be certain to reliably approximate $\sup_{\gamma:\theta_\gamma=\theta} P_\gamma(\bar{X} \leq T)$ or $\sup_{\gamma:\theta_\gamma=\theta} P_\gamma(\bar{X} > T)$ for each θ . If an insufficiently rich set of distributions is chosen, the approximation will be lower than actual maximum regret.

Asymptotic approximation With the knowledge of $\theta_\gamma \equiv E_\gamma[x_i] \in \mathbb{R}$ and $\sigma_\gamma^2 \equiv V_\gamma[x_i] \in \mathbb{R}$, the planner can use the asymptotic normal approximation

$$P_\gamma(\bar{X} \leq T) \approx \Phi\left(\frac{\sqrt{N}}{\sigma_\gamma}(T - \theta_\gamma)\right).$$

To evaluate maximum Type I and Type II regret of a threshold decision rule it is sufficient to know minimum and maximum feasible variance σ_γ^2 for each feasible value of θ_γ . Normal approximations of tail probabilities of \bar{X} could be either higher or lower than the actual values, thus approximate values of maximum Type I/II regret could also be either above or below actual values.

Large deviation bounds There are a number of inequalities for tail probabilities of the distribution of sample mean \bar{X} . Using these inequalities allows the statistician to construct finite sample upper bounds on maximum Type I and Type II regret. Unlike asymptotic approximations, bounds constructed using large deviation inequalities are guaranteed not to be lower than actual maximum Type I/II regret values, which may be useful for conservative decision making.

The simplest large deviation bound is given by the one-sided Chebyshev's inequality, which requires only that x'_i s have bounded variance:

$$\begin{aligned} T < \theta_\gamma &\Rightarrow P_\gamma(\bar{X} \leq T) \leq \frac{1}{1 + \left(\frac{\sqrt{N}}{\sigma_\gamma}(T - \theta_\gamma)\right)^2}, \\ T > \theta_\gamma &\Rightarrow P_\gamma(\bar{X} > T) \leq \frac{1}{1 + \left(\frac{\sqrt{N}}{\sigma_\gamma}(T - \theta_\gamma)\right)^2}. \end{aligned}$$

If outcome variables are bounded $x_i \in [a, b]$, then Hoeffding's exponential inequality (1963,

Theorem 2) applies to the tail probabilities of \bar{X} :

$$T < \theta_\gamma \Rightarrow P_\gamma(\bar{X} \leq T) \leq \exp \left\{ -2 \left(\frac{\sqrt{N}}{b-a} (T - \theta_\gamma) \right)^2 \right\},$$

$$T > \theta_\gamma \Rightarrow P_\gamma(\bar{X} > T) \leq \exp \left\{ -2 \left(\frac{\sqrt{N}}{b-a} (T - \theta_\gamma) \right)^2 \right\}.$$

Hoeffding's inequality was used by Manski (2004) to compute bounds on maximum regret of plug-in (empirical success) treatment rules.

If a feasible distribution has finite absolute third moment $\beta_\gamma \equiv E |x_i - \theta_\gamma|^3 \in \mathbb{R}$, then bounds on $P_\gamma(\bar{X} \leq T)$ could be derived from the Berry-Esseen inequality:

$$|P_\gamma(\bar{X} \leq T) - \Phi(z)| \leq \min \left(C_0, C_1 \frac{1}{(1+|z|)^3} \right) \cdot \frac{\beta_\gamma}{\sigma_\gamma^3 \sqrt{N}}, \text{ where } z \equiv \frac{\sqrt{N}}{\sigma_\gamma} (T - \theta_\gamma).$$

Lowest proven values for the constants C_0 and C_1 are $C_0 \leq 0.7975$ (van Beek, 1972) and $C_1 \leq 32$ (Paditz, 1989). For large enough sample sizes, the Berry-Esseen inequality could show that the tail probabilities are arbitrarily close to their normal approximation, which is significantly smaller than the Chebyshev's and Hoeffding's bounds.

5.1 A numerical example

I will illustrate how the different methods of evaluating maximum regret of threshold rules may perform in practice on a simple example inspired by the problem of rare side effects in clinical trials. Let the average outcome of the status quo treatment $t = 0$ be $E[y(0)] = .5$ (outcome values refer to individual welfare of clinical outcomes). Suppose that a new treatment has been assigned to $N = 1000$ randomly selected patients. The treatment has three potential outcomes: $y(1) = 1$ and $y(1) = 0$ correspond to the positive and negative outcomes of the treatment on the condition that it is intended to treat, while $y(1) = -100$ corresponds to a rare, dangerous side effect. The set of feasible treatment outcome distributions Γ includes all probability distributions with the support $\{-100, 0, 1\}$ that have a limited probability of the rare side effect $P_\gamma[y(1) = -100] \leq \frac{1}{1000}$. Let \bar{X} be the sample average of the 1000 trial outcomes of the new treatment.

First, let's consider how well the different methods approximate the regret of a plug-in statistical

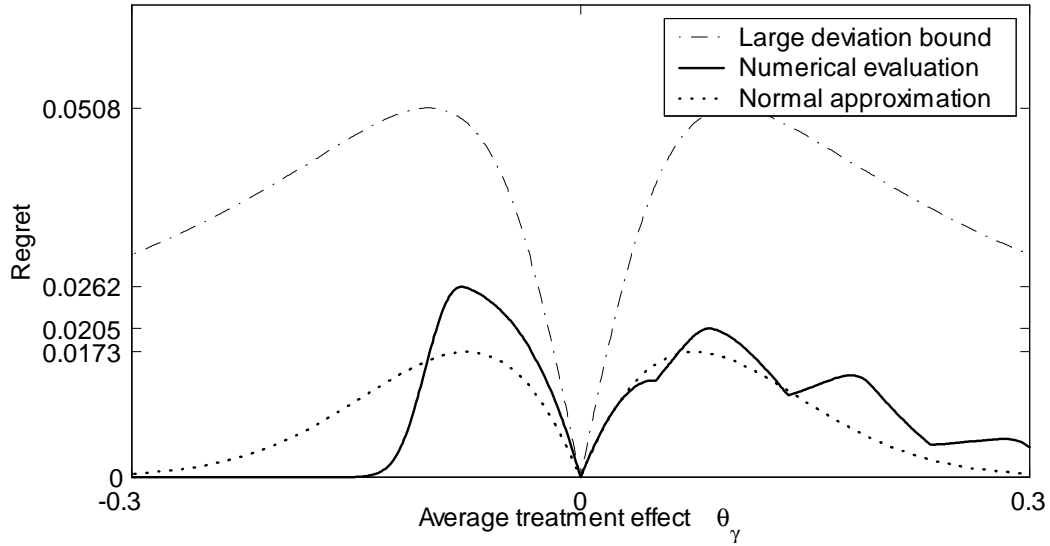


Figure 3: Evaluation of maximum regret of the plug-in ($T = .5$) statistical treatment rule.

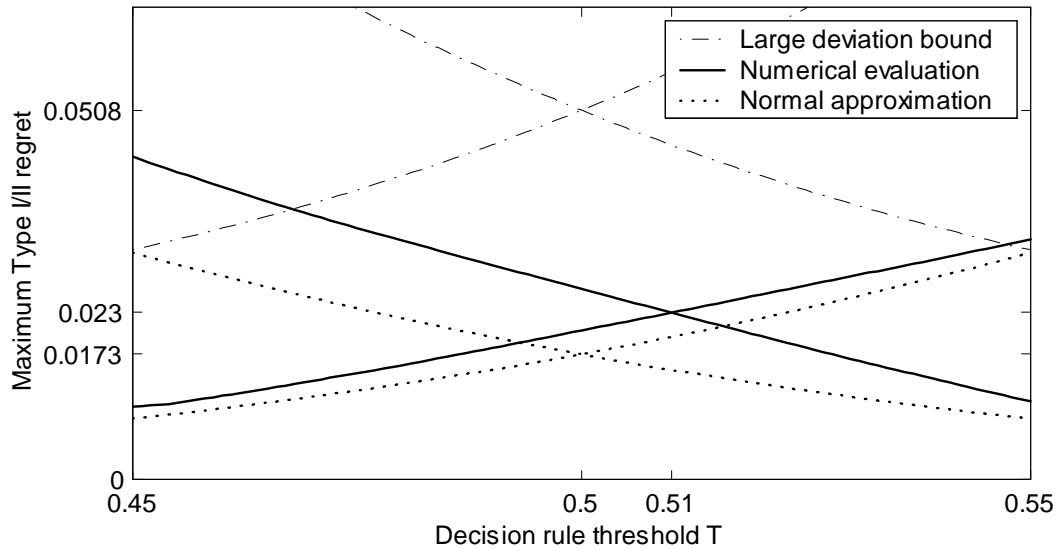


Figure 4: Maximum Type I and Type II regret approximations for a range of threshold statistical treatment rules.

treatment rule $\delta_P \equiv 1[\bar{X} > .5]$, which assigns the population to new treatment if it outperforms the status quo treatment in the trial by any margin. Figure 3 displays the maximum regret of δ_P for a range of feasible values of the average treatment effect $\theta_\gamma \equiv E_\gamma[y(1)] - E[y(0)]$. There are multiple feasible outcome distributions with the same θ_γ , so the lines represent the maximum approximated regret among those distributions. Figure 4 shows maximum Type I and Type II regret approximations for threshold decision rules with thresholds ranging from $T = .45$ to $T = .55$. The top lines show the best upper bounds on maximum regret derived from large deviation bounds. That is, the best of the bounds derived from Chebyshev's, Hoeffding's, or Berry-Esseen inequalities. Each inequality is applied to all feasible values of distribution moments for a given θ . Chebyshev's inequality provides the smallest bounds in this example despite fairly large sample size because some of the feasible outcome distributions have large range $[-100, 1]$ and large third moments. It provides an upper bound of .0508 for both maximum Type I ($\theta \leq 0$) and maximum Type II ($\theta > 0$) regret.

The lower dotted lines show maximum regret computed using the normal approximation to the distribution of \bar{X} based on the feasible values for the variance of outcome distributions. The normal approximation suggests that both maximum Type I and maximum Type II regret of δ_P equal to .0173.

The thick solid lines in Figures 3 and 4 show the maximum regret evaluated numerically. The set of feasible distributions in this problem is simple enough (two-dimensional and continuous) to be reliably approximated by a finite set of distributions. For this example, the probabilities $P_\gamma(\bar{X} \leq .5)$ and corresponding regret values were evaluated on a grid of 60,000 distributions. These calculations show that maximum Type I regret of the plug-in rule equals .0262, while the maximum Type II regret equals .0205. Figure 4 shows that among threshold decision rules, minimax regret is attained by the decision rule with threshold $T = .51$, rather than by the plug-in rule, and its maximum regret equals .0230.

In this example, the large deviation bounds on maximum regret are much higher than its actual values, while normal approximations are significantly lower. Both of them suggest that the plug-in decision rule minimizes maximum regret, even though its maximum regret is 12% higher than the minimum attainable by a different threshold decision rule. The difference between these approximations and actual maximum regret presents a bigger problem for the selection of

trial sample size. Using the normal approximation to evaluate maximum regret could lead the statistician to choose sample size about 40% smaller than is necessary to make decisions with the desired maximum regret. Using the large deviation bounds, on the other hand, could lead her to choose a sample size almost five times larger than necessary.

Asymptotic approximations and large deviation bounds provide convenient and tractable methods for evaluating maximum regret of threshold decision rules. This example shows, however, that even in realistic problems with fairly large sample size, they could significantly misrepresent the maximum regret of decision rules. Whenever possible, such results should be verified by direct computation or simulation.

A Proofs

Lemma 1 I will prove the results in part a), the proof of part b) is analogous. Note that it is w.l.o.g. to set $\sigma = 1$ to simplify notation, then

$$\bar{R}_{Type\ I}(\delta_T) = \max_{h \leq 0} \{-h\Phi(h-T)\}.$$

For every fixed $h < 0$, $-h\Phi(h-T)$ is a strictly decreasing function of T . Furthermore, for any fixed T , $-h\Phi(h-T)$ is a continuous function of h , with $\lim_{h \rightarrow -\infty} \{-h\Phi(h-T)\} = 0$, and $-h\Phi(h-T) > 0$ for $-\infty < h < 0$, thus $-h\Phi(h-T)$ attains its maximum on $h \in (-\infty, 0)$. Therefore $\max_{h \leq 0} \{-h\Phi(h-T)\}$ is a strictly decreasing function of T .

To show that $\max_{h \leq 0} \{-h\Phi(h-T)\}$ is continuous in T for all $T \in \mathbb{R}$, let's fix $T = T_0$ and pick some $\Delta > 0$. Then there exists $H < 0$ such that $h(h-T) > 1$ and $h-T < 0$ for all $h < H$ and for all $T \in [T_0 - \Delta, T_0 + \Delta]$. Then for such h and T :

$$\begin{aligned} \frac{d}{dh} \{-h\Phi(h-T)\} &= -\Phi(h-T) - h\phi(h-T) > \\ &> \frac{\phi(h-T)}{h-T} - h\phi(h-T) \\ &= \phi(h-T) \frac{1-h(h-T)}{h-T} > 0. \end{aligned}$$

The second line follows from an well known inequality for the normal distribution:

$$\begin{aligned} 1 - \Phi(\eta) &< \frac{\phi(\eta)}{\eta} \text{ for } \eta > 0 \\ \Rightarrow \Phi(\eta) &< -\frac{\phi(\eta)}{\eta} \text{ for } \eta < 0. \end{aligned}$$

Since $\frac{d}{dh} \{-h\Phi(h-T)\} > 0$ for all $h < H$ and all $T \in [T_0 - \Delta, T_0 + \Delta]$, the maximum of $-h\Phi(h-T)$ over h for each T is achieved on the closed interval $h \in [H, 0]$. The derivative of $-h\Phi(h-T)$ with respect to T is bounded on the rectangle $(h, T) \in [H, 0] \times [T_0 - \Delta, T_0 + \Delta]$, thus $\max_{h \leq 0} \{-h\Phi(h-T)\} = \max_{h \in [H, 0]} \{-h\Phi(h-T)\}$ is continuous in T at T_0 .

For any $T < 0$

$$\max_{h \leq 0} \{-h\Phi(h-T)\} \geq -T\Phi(0) = -\frac{T}{2}$$

(by substituting $h = T$), thus $\max_{h \leq 0} \{-h\Phi(h - T)\} \rightarrow \infty$ as $T \rightarrow -\infty$.

For any $T > 0$ and $h < 0$, $\Phi(h - T) \leq \frac{1}{(h-T)^2}$ by Chebyshev's inequality. Also, differentiation of $-\frac{h}{(h-T)^2}$ with respect to h shows that $\max_{h \leq 0} \left\{-\frac{h}{(h-T)^2}\right\}$ is achieved at $h = -T$ and equals $\frac{1}{4T}$. Then

$$\max_{h \leq 0} \{-h\Phi(h - T)\} \leq \max_{h \leq 0} \left\{-\frac{h}{(h-T)^2}\right\} = \frac{1}{4T}$$

and $\frac{1}{4T} \rightarrow 0$, thus $\max_{h \leq 0} \{-h\Phi(h - T)\} \rightarrow 0$ as $T \rightarrow \infty$.

Proposition 2 a) If $T > 0$, then the maximum Type II regret of threshold decision rule δ_T over the set Γ equals

$$\begin{aligned} \max_{\gamma \in \Gamma: \theta_\gamma > 0} R(\delta_T, \gamma) &= \max_{\gamma \in \Gamma: \theta_\gamma > 0} \{\theta_\gamma P_\gamma(\bar{X} \leq T)\} = \max_{\gamma \in \Gamma: \theta_\gamma > 0} \left\{ \theta_\gamma \Phi \left[\frac{\sqrt{N}}{\sigma_\gamma} (T - \theta_\gamma) \right] \right\} \\ &= \max_{\theta > 0} \left\{ \theta \cdot \max_{\sigma \in [\underline{\sigma}, \bar{\sigma}]} \Phi \left[\frac{\sqrt{N}}{\sigma} (T - \theta) \right] \right\} \\ &= \frac{\bar{\sigma}}{\sqrt{N}} \max_{h > 0} \left\{ h \cdot \max_{\sigma \in [\underline{\sigma}, \bar{\sigma}]} \Phi \left[\frac{\bar{\sigma}}{\sigma} (T^* - h) \right] \right\} \\ &= \frac{\bar{\sigma}}{\sqrt{N}} \max \left\{ \max_{h \in (0, T^*)} \left(h \Phi \left[\frac{\bar{\sigma}}{\underline{\sigma}} (T^* - h) \right] \right), \max_{h \geq T^*} (h \Phi [T^* - h]) \right\} \\ &= \frac{\bar{\sigma}}{\sqrt{N}} \max_{h \geq T^*} \{h \Phi [T^* - h]\} \end{aligned}$$

The third line uses substitutions $h \equiv \frac{\sqrt{N}}{\sigma} \theta$ and $T^* \equiv \frac{\sqrt{N}}{\bar{\sigma}} T$. The fourth line uses the fact that $\max_{\sigma \in [\underline{\sigma}, \bar{\sigma}]} \Phi \left[\frac{\bar{\sigma}}{\sigma} (T^* - h) \right] = \Phi \left[\frac{\bar{\sigma}}{\underline{\sigma}} (T^* - h) \right]$ for $h < T^*$ and $\max_{\sigma \in [\underline{\sigma}, \bar{\sigma}]} \Phi \left[\frac{\bar{\sigma}}{\sigma} (T^* - h) \right] = \Phi [T^* - h]$ for $h \geq T^*$. The last equality follows from the condition (8).

Similar derivation for maximum Type II regret over the set $\bar{\Gamma}$ shows that for $T > 0$:

$$\begin{aligned} \max_{\gamma \in \bar{\Gamma}: \theta_\gamma > 0} R(\delta_T, \gamma) &= \max_{\theta > 0} \left\{ \theta \cdot \Phi \left[\frac{\sqrt{N}}{\bar{\sigma}} (T - \theta) \right] \right\} \\ &= \frac{\bar{\sigma}}{\sqrt{N}} \max_{h > 0} \{h \Phi [T^* - h]\} \\ &= \frac{\bar{\sigma}}{\sqrt{N}} \max \left\{ \max_{h \in (0, T^*)} (h \Phi [T^* - h]), \max_{h \geq T^*} (h \Phi [T^* - h]) \right\} \\ &= \frac{\bar{\sigma}}{\sqrt{N}} \max_{h \geq T^*} \{h \Phi [T^* - h]\}. \end{aligned}$$

The last equality holds because $\Phi [T^* - h] \leq \Phi \left[\frac{\bar{\sigma}}{\underline{\sigma}} (T^* - h) \right]$ for $h < T^*$, thus condition (8) implies

$$\max_{h \in (0, T^*)} (h\Phi [T^* - h]) \leq \max_{h \geq T^*} (h\Phi [T^* - h]).$$

If $T \leq 0$, then $T - \theta < 0$ for all $\theta > 0$, thus $\max_{\sigma \in [\underline{\sigma}, \bar{\sigma}]} \Phi \left[\frac{\sqrt{N}}{\sigma} (T - \theta) \right] = \Phi \left[\frac{\sqrt{N}}{\bar{\sigma}} (T - \theta) \right]$ and

$$\begin{aligned} \max_{\gamma \in \Gamma: \theta_\gamma > 0} R(\delta_T, \gamma) &= \max_{\theta > 0} \left\{ \theta \cdot \max_{\sigma \in [\underline{\sigma}, \bar{\sigma}]} \Phi \left[\frac{\sqrt{N}}{\sigma} (T - \theta) \right] \right\} \\ &= \max_{\theta > 0} \left\{ \theta \cdot \Phi \left[\frac{\sqrt{N}}{\bar{\sigma}} (T - \theta) \right] \right\} = \max_{\gamma \in \bar{\Gamma}: \theta_\gamma > 0} R(\delta_T, \gamma). \end{aligned}$$

The proof is analogous for Type I regret.

b) Suppose that $\delta'(\bar{X}, S^2)$ has both lower maximum Type I regret and lower maximum Type II regret than δ_T over the set Γ . Since δ_T achieves maximum Type I and II regret over the subset $\bar{\Gamma}$, it follows that

$$\begin{aligned} \sup_{\gamma \in \bar{\Gamma}: \theta_\gamma \leq 0} R(\delta', \gamma) &\leq \sup_{\gamma \in \Gamma: \theta_\gamma \leq 0} R(\delta', \gamma) < \max_{\gamma \in \Gamma: \theta_\gamma \leq 0} R(\delta_T, \gamma) = \max_{\gamma \in \bar{\Gamma}: \theta_\gamma \leq 0} R(\delta_T, \gamma), \\ \sup_{\gamma \in \bar{\Gamma}: \theta_\gamma > 0} R(\delta', \gamma) &\leq \sup_{\gamma \in \Gamma: \theta_\gamma > 0} R(\delta', \gamma) < \max_{\gamma \in \Gamma: \theta_\gamma > 0} R(\delta_T, \gamma) = \max_{\gamma \in \bar{\Gamma}: \theta_\gamma > 0} R(\delta_T, \gamma). \end{aligned}$$

Since the class of threshold decision rules is essentially complete for the problem with fixed variance, there must be a threshold decision rule $\delta_{T'} \equiv 1_{|\bar{X} > T'|}$ such that $R(\delta', \gamma) \leq R(\delta_{T'}, \gamma)$ for all $\gamma \in \bar{\Gamma}$. Then

$$\begin{aligned} \sup_{\gamma \in \bar{\Gamma}: \theta_\gamma \leq 0} R(\delta_{T'}, \gamma) &\leq \sup_{\gamma \in \bar{\Gamma}: \theta_\gamma \leq 0} R(\delta', \gamma) < \max_{\gamma \in \bar{\Gamma}: \theta_\gamma \leq 0} R(\delta_T, \gamma), \\ \sup_{\gamma \in \bar{\Gamma}: \theta_\gamma > 0} R(\delta_{T'}, \gamma) &\leq \sup_{\gamma \in \bar{\Gamma}: \theta_\gamma > 0} R(\delta', \gamma) < \max_{\gamma \in \bar{\Gamma}: \theta_\gamma > 0} R(\delta_T, \gamma), \end{aligned}$$

which contradicts the conjecture that δ_T is a solution to the minimax regret, asymmetric minimax regret, or limited type I regret problem over the feasible set $\bar{\Gamma}$. Thus δ' cannot have both lower maximum Type I and lower maximum Type II regret than δ_T .

Lemma 3 I will provide the proof for $\bar{R}_{Type I}(\delta)$, the proof for $\bar{R}_{Type II}(\delta)$ is analogous.

For a fixed $\bar{\delta}$, $R(\bar{\delta}, \gamma)$ is a bounded continuous function of p_γ on the closed interval $[a, p_0]$, thus

attains its maximum. Also,

$$|D(\delta_1) - D(\delta_2)| < \varepsilon \Rightarrow \sup_{\gamma: p_\gamma \in [a, p_0]} |R(\delta_1, \gamma) - R(\delta_2, \gamma)| < \varepsilon,$$

thus $\max_{\gamma: p_\gamma \in [a, p_0]} R(\delta, \gamma)$ is a continuous function of $D(\delta)$.

For any fixed $p_\gamma \in (0, p_0)$,

$$R(\delta_{T, \lambda}, \gamma) = -\theta \cdot \left\{ 1 - \lambda \cdot B(T, N, p_\gamma) - \sum_{T < n \leq N} B(n, N, p_\gamma) \right\}$$

is a strictly decreasing function of $D(\delta) = T + (1 - \lambda)$. For $p_\gamma = 0$, $R(\delta_{T, \lambda}, \gamma)$ is also a strictly decreasing function of $D(\delta)$ for $D(\delta) \in [0, 1]$ and $R(\delta_{T, \lambda}, 0) = 0$ for $D(\delta) \geq 1$. It follows that

$\max_{\gamma: p_\gamma \in [a, p_0]} R(\delta, \gamma)$ is a strictly decreasing function of $D(\delta)$.

If $D(\delta) = N + 1$, then $T = N, \lambda = 0$, thus $R(\delta_{T, \lambda}, \gamma) = 0$ for any $p_\gamma \in (0, p_0)$.

Proposition 4 It follows from lemma 3 that maximum regret of the minimax regret treatment rule lies between maximum Type I and maximum Type II regret of the plug-in treatment rule:

$$\min(\bar{R}_{Type\ I}(\delta_P), \bar{R}_{Type\ II}(\delta_P)) \leq \max_{\gamma \in \Gamma} R(\delta_M, \gamma) \leq \max(\bar{R}_{Type\ I}(\delta_P), \bar{R}_{Type\ II}(\delta_P)).$$

If $\sqrt{\frac{N}{p_0(1-p_0)}} \bar{R}_{Type\ I}(\delta_P)$ and $\sqrt{\frac{N}{p_0(1-p_0)}} \bar{R}_{Type\ II}(\delta_P)$ both converge to $\max_{h>0} [h\Phi(-h)]$, then it follows that $\max_{\gamma \in \Gamma} R(\delta_M, \gamma)$ converges to the same limit. I will establish that $\sqrt{\frac{N}{p_0(1-p_0)}} \bar{R}_{Type\ II}(\delta_P) \rightarrow \max_{h>0} [h\Phi(-h)]$, the proof for $\sqrt{\frac{N}{p_0(1-p_0)}} \bar{R}_{Type\ I}(\delta_P)$ is analogous.

To simplify notation, I will use the following substitutions:

$$\begin{aligned} \sigma &= \sqrt{p_\gamma(1-p_\gamma)}, \\ \sigma_0 &= \sqrt{p_0(1-p_0)}, \text{ and} \\ h &= \frac{\sqrt{N}}{\sigma_0} (p_\gamma - p_0). \end{aligned}$$

I will use the Berry-Esseen inequality to show that $\frac{\sqrt{N}}{\sigma_0} R(\delta_P, p_0 + \frac{\sigma_0}{\sqrt{N}} h)$ uniformly converges to $h\Phi(-h)$ for $h \in (0, \frac{3}{2}\sigma_0^{-2}]$ as $N \rightarrow \infty$. Since the function $h\Phi(-h)$ reaches its maximum at

$h \approx .752$ and $\frac{3}{2}\sigma_0^{-2} \geq 6$, $\max_{h>0} [h\Phi(-h)] = \max_{h \in (0, \frac{3}{2}\sigma_0^{-2}]} [h\Phi(-h)]$. I use Chebyshev's inequality to show that $\frac{\sqrt{N}}{\sigma_0} R\left(\delta_P, p_0 + \frac{\sigma_0}{\sqrt{N}}h\right) < \max_{h>0} [h\Phi(-h)]$ for $h > \frac{3}{2}\sigma_0^{-2}$.

For any $\varepsilon > 0$, there is N_1 such that for all $N > N_1$:

$$(9) \quad \sup_{h \in (0, \frac{3}{2}\sigma_0^{-2}]} \left| \Phi\left(-h\frac{\sigma_0}{\sigma}\right) - \Phi(-h) \right| < \frac{\sigma_0^2}{3}\varepsilon,$$

because the standard normal c.d.f. Φ has a bounded derivative, $\sigma_0 \neq 0$, and $\sup_{h \in (0, \frac{3}{2}\sigma_0^{-2}]} \left(1 - \frac{\sigma_0}{\sigma}\right) \rightarrow 0$ as $N \rightarrow \infty$.

Application of the uniform Berry-Esseen inequality to X , which is a sum of N independent Bernoulli random variables with mean p_γ (cf. Shiryaev (1995, p. 63)), yields

$$\left| P_\gamma\left(\frac{\sqrt{N}}{\sigma}\left(\frac{X}{N} - p_\gamma\right) \leq z\right) - \Phi(z) \right| \leq \frac{p_\gamma^2 + (1-p_\gamma)^2}{\sqrt{p_\gamma(1-p_\gamma)}\sqrt{N}},$$

for any $z \in \mathbb{R}$, where Φ is the standard normal c.d.f.

There exists N_2 such that for $N > N_2$, $h \in (0, \frac{3}{2}\sigma_0^{-2}]$ implies $p_\gamma \in \left[p_0, \frac{1+p_0}{2}\right]$. The function $\frac{p_\gamma^2 + (1-p_\gamma)^2}{\sqrt{p_\gamma(1-p_\gamma)}}$ is continuous and bounded on $p_\gamma \in \left[p_0, \frac{1+p_0}{2}\right]$ since $p_0 > 0$ and $\frac{1+p_0}{2} < 1$. Let $B \equiv \max_{p_\gamma \in \left[p_0, \frac{1+p_0}{2}\right]} \frac{p_\gamma^2 + (1-p_\gamma)^2}{\sqrt{p_\gamma(1-p_\gamma)}}$. Then for $N > N_2$ and $h \in (0, \frac{3}{2}\sigma_0^{-2}]$

$$(10) \quad \left| P_\gamma\left(\frac{\sqrt{N}}{\sigma}\left(\frac{X}{N} - p_\gamma\right) \leq z\right) - \Phi(z) \right| \leq \frac{B}{\sqrt{N}}.$$

There exists $N_3 > N_2$ such that $\frac{B}{\sqrt{N}} < \frac{\sigma_0^2}{3}\varepsilon$ for $N > N_3$.

Since

$$P_\gamma\left(\frac{\sqrt{N}}{\sigma}\left(\frac{X}{N} - p_\gamma\right) \leq -\frac{\sigma_0}{\sigma}h\right) = P_\gamma\left(\frac{\sqrt{N}}{\sigma}\left(\frac{X}{N} - p_\gamma\right) \leq -\frac{\sqrt{N}}{\sigma}(p_\gamma - p_0)\right) = P_\gamma\left(\frac{X}{N} \leq p_0\right),$$

letting $z = -\frac{\sigma_0}{\sigma}h$ in (10) yields for $N > N_3$

$$\sup_{h \in (0, \frac{3}{2}\sigma_0^{-2}]} \left| P_\gamma\left(\frac{X}{N} \leq p_0\right) - \Phi\left(-\frac{\sigma_0}{\sigma}h\right) \right| \leq \frac{\sigma_0^2}{3}\varepsilon.$$

Combining this result with (9) shows that for $N > \max(N_1, N_3)$

$$\begin{aligned} & \sup_{h \in (0, \frac{3}{2}\sigma_0^{-2}]} \left| P_\gamma \left(\frac{X}{N} \leq p_0 \right) - \Phi(-h) \right| \leq \\ & \sup_{h \in (0, \frac{3}{2}\sigma_0^{-2}]} \left| \Phi \left(-\frac{\sigma_0}{\sigma} h \right) - \Phi(-h) \right| + \sup_{h \in (0, \frac{3}{2}\sigma_0^{-2}]} \left| P_\gamma \left(\frac{X}{N} \leq p_0 \right) - \Phi \left(-\frac{\sigma_0}{\sigma} h \right) \right| \leq \frac{2}{3}\sigma_0^2 \varepsilon, \end{aligned}$$

and, since $R \left(\delta_P, p_0 + \frac{\sigma_0}{\sqrt{N}} h \right) = \frac{\sigma_0}{\sqrt{N}} h \cdot P_\gamma \left(\frac{X}{N} \leq p_0 \right)$ for $h > 0$,

$$\begin{aligned} & \sup_{h \in (0, \frac{3}{2}\sigma_0^{-2}]} \left| \frac{\sqrt{N}}{\sigma_0} R \left(\delta_P, p_0 + \frac{\sigma_0}{\sqrt{N}} h \right) - h \Phi(-h) \right| = \\ & \sup_{h \in (0, \frac{3}{2}\sigma_0^{-2}]} \left| h \cdot P_\gamma \left(\frac{X}{N} \leq p_0 \right) - h \Phi(-h) \right| \leq \\ & \frac{3}{2}\sigma_0^{-2} \sup_{h \in (0, \frac{3}{2}\sigma_0^{-2}]} \left| P_\gamma \left(\frac{X}{N} \leq p_0 \right) - \Phi(-h) \right| \leq \varepsilon. \end{aligned}$$

The one-sided Chebyshev's inequality shows that

$$P_\gamma \left(\frac{X}{N} \leq p_0 \right) = P_\gamma \left(\frac{X}{N} - p_\gamma \leq p_0 - p_\gamma \right) \leq \frac{1}{1 + \frac{N}{\sigma^2} (p_\gamma - p_0)^2} \leq \frac{1}{4\sigma_0^2 h^2},$$

where the last inequality uses substitution $p_\gamma - p_0 = \frac{\sigma_0}{\sqrt{N}} h$ and the second one uses $\sigma^2 \leq \frac{1}{4}$. For $h > \frac{3}{2}\sigma_0^{-2}$ this implies

$$\begin{aligned} \frac{\sqrt{N}}{\sigma_0} R \left(\delta_P, p_0 + \frac{\sigma_0}{\sqrt{N}} h \right) &= h \cdot P_\gamma \left(\frac{X}{N} \leq p_0 \right) \\ &\leq h \cdot \frac{1}{4\sigma_0^2 h^2} \leq \frac{1}{6} < \max_{h>0} [h \Phi(-h)]. \end{aligned}$$

Thus, for $N > \max(N_1, N_3)$

$$\left| \max_{\gamma: p_\gamma > 0} \frac{\sqrt{N}}{\sigma_0} R(\delta_P, \gamma) - \max_{h>0} [h \Phi(-h)] \right| < \varepsilon.$$

Proposition 5 Let V_γ denote the variance of a random variable in state of the world γ and let $p_\gamma \equiv E_\gamma[y(1)]$. I will consider the case when $p_\gamma > p_0$, the proof for $p_\gamma \leq p_0$ is analogous.

For all γ such that $V_\gamma[y(1)] < \frac{\sigma_h^2}{9}$ (thus $V_\gamma \left[\frac{1}{N} \sum_{i=1}^N x_i - p_\gamma \right] < \frac{\sigma_h^2}{9N}$) the one-sided Chebyshev's

inequality shows that

$$P_\gamma \left(\frac{1}{N} \sum_{i=1}^N x_i \leq p_0 \right) = P_\gamma \left(\frac{1}{N} \sum_{i=1}^N x_i - p_\gamma \leq -(p_\gamma - p_0) \right) \leq \frac{1}{1 + \frac{9N}{\sigma_h^2} (p_\gamma - p_0)^2}.$$

Applying the result to the formula for regret of the plug-in rule δ_P yields

$$R(\delta_P, \gamma) = (p_\gamma - p_0) \cdot P_\gamma \left(\frac{1}{N} \sum_{i=1}^N x_i \leq p_0 \right) \leq \frac{\sigma_h}{3\sqrt{N}} \cdot \frac{\sqrt{\frac{9N}{\sigma_h^2}} (p_\gamma - p_0)}{1 + \frac{9N}{\sigma_h^2} (p_\gamma - p_0)^2} \leq \frac{\sigma_h}{6\sqrt{N}}.$$

To obtain the last inequality, observe that $\max_{h>0} \frac{h}{1+h^2} = \frac{1}{2}$.

For all γ such that $p_\gamma - p_0 \geq 6\frac{\sigma_h}{\sqrt{N}}$, also apply Chebyshev's inequality, using the fact that $V_\gamma \left[\frac{1}{N} \sum_{i=1}^N x_i - p_\gamma \right] < \frac{\sigma_h^2}{N}$:

$$P_\gamma \left(\frac{1}{N} \sum_{i=1}^N x_i \leq p_0 \right) = P_\gamma \left(\frac{1}{N} \sum_{i=1}^N x_i - p_\gamma \leq -(p_\gamma - p_0) \right) \leq \frac{1}{1 + \frac{N}{\sigma_h^2} (p_\gamma - p_0)^2}.$$

Applying the result to the formula for regret of the plug-in rule δ_P yields

$$R(\delta_P, \gamma) = (p_\gamma - p_0) \cdot P_\gamma \left(\frac{1}{N} \sum_{i=1}^N x_i \leq p_0 \right) \leq \frac{\sigma_h}{\sqrt{N}} \cdot \frac{\sqrt{\frac{N}{\sigma_h^2}} (p_\gamma - p_0)}{1 + \frac{N}{\sigma_h^2} (p_\gamma - p_0)^2} \leq \frac{\sigma_h}{6\sqrt{N}}.$$

The last inequality holds because $\sqrt{\frac{N}{\sigma_h^2}} (p_\gamma - p_0) \geq 6$ and $\max_{h>6} \frac{h}{1+h^2} < \frac{1}{6}$.

For all γ such that $p_\gamma - p_0 < 6\frac{\sigma_h}{\sqrt{N}}$ and $V_\gamma[y(1)] \in \left[\frac{\sigma_h^2}{9}, \sigma_h^2 \right]$, let's apply the Berry-Esseen inequality (cf. Shiryaev (1995, p. 374)) to the sum of N i.i.d. random variables $(x_i - p_\gamma)$, for any $z \in \mathbb{R}$:

$$(11) \quad \left| P_\gamma \left(\frac{\sqrt{N}}{\sqrt{V_\gamma[y(1)]}} \left(\frac{1}{N} \sum_{i=1}^N x_i - p_\gamma \right) \leq z \right) - \Phi(z) \right| \leq \frac{E_\gamma |y(1) - p_\gamma|^3}{V_\gamma^{3/2}[y(1)] \sqrt{N}}.$$

Substitute $z = \frac{\sqrt{N}}{\sqrt{V_\gamma[y(1)]}} (p_0 - p_\gamma)$ into (11) and it becomes

$$\left| P_\gamma \left(\frac{1}{N} \sum_{i=1}^N x_i \leq p_0 \right) - \Phi \left(\frac{\sqrt{N}}{\sqrt{V_\gamma[y(1)]}} (p_0 - p_\gamma) \right) \right| \leq \frac{E_\gamma |y(1) - p_\gamma|^3}{V_\gamma^{3/2}[y(1)] \sqrt{N}}.$$

Since $y(1) - p_\gamma \in [0, 1]$, $E_\gamma |y(1) - p_\gamma|^3 \leq V_\gamma [y(1)]$ and $V_\gamma [y(1)] \geq \frac{\sigma_h^2}{9}$

$$\frac{E_\gamma |y(1) - p_\gamma|^3}{V_\gamma^{3/2} [y(1)] \sqrt{N}} \leq \frac{1}{V_\gamma^{1/2} [y(1)] \sqrt{N}} \leq \frac{3}{\sigma_h \sqrt{N}},$$

thus

$$\left| P_\gamma \left(\frac{1}{N} \sum_{i=1}^N x_i \leq p_0 \right) - \Phi \left(\frac{\sqrt{N}}{\sqrt{V_\gamma [y(1)]}} (p_0 - p_\gamma) \right) \right| \leq \frac{3}{\sigma_h \sqrt{N}}.$$

Applying the result to the regret formula for δ_P yields

$$\begin{aligned} R(\delta_P, \gamma) &= (p_\gamma - p_0) \cdot P_\gamma \left(\frac{1}{N} \sum_{i=1}^N x_i \leq p_0 \right) \leq \\ &\leq (p_\gamma - p_0) \cdot \left(\Phi \left(\frac{\sqrt{N}}{\sqrt{V_\gamma [y(1)]}} (p_0 - p_\gamma) \right) + \frac{3}{\sigma_h \sqrt{N}} \right) \leq \\ &\leq \frac{\sqrt{V_\gamma [y(1)]}}{\sqrt{N}} \max_{h>0} h \Phi(-h) + \frac{3(p_\gamma - p_0)}{\sigma_h \sqrt{N}} \leq \\ &\leq \frac{\sigma_h}{\sqrt{N}} \max_{h>0} h \Phi(-h) + \frac{18}{N}. \end{aligned}$$

The last inequality uses $p_\gamma - p_0 < 6 \frac{\sigma_h}{\sqrt{N}}$.

The three cases considered are exhaustive of all states of the world γ with $p_\gamma > 0$. If $V_\gamma [y(1)] < \frac{\sigma_h^2}{9}$, or $V_\gamma [y(1)] \geq \frac{\sigma_h^2}{9}$ and $p_\gamma - p_0 \geq 6 \frac{\sigma_h}{\sqrt{N}}$, then

$$\sqrt{N} R(\delta_P, \gamma) \leq \frac{\sigma_h}{6} < \sigma_h \cdot \max_{h>0} [h \Phi(-h)].$$

If $V_\gamma [y(1)] \geq \frac{\sigma_h^2}{9}$ and $p_\gamma - p_0 < 6 \frac{\sigma_h}{\sqrt{N}}$, then

$$\sqrt{N} R(\delta_P, \gamma) \leq \sigma_h \cdot \max_{h>0} h \Phi(-h) + \frac{18}{\sqrt{N}},$$

thus $\sqrt{N} \sup_{\gamma \in \Gamma} R(\delta_P, \gamma) \leq \sigma_h \cdot \max_{h>0} [h \Phi(-h)] + o(1)$.

References

- [1] Berger, J., 1985. *Statistical Decision Theory and Bayesian Analysis*, 2nd ed. Springer Verlag, New York.
- [2] Hirano, K., Porter, J., 2006. *Asymptotics for Statistical Treatment Rules*. Unpublished manuscript, University of Arizona and University of Wisconsin-Madison.
- [3] International Committee on Harmonization, 1998. *Guideline E9: Statistical Principles for Clinical Trials*.
- [4] Karlin, S., Rubin, H., 1956. The Theory of Decision Procedures for Distributions with Monotone Likelihood Ratio. *The Annals of Mathematical Statistics* 27, 272-299.
- [5] Manski, C., 2004. Statistical Treatment Rules for Heterogeneous Populations. *Econometrica* 72, 1221-1246.
- [6] Manski, C., 2005. *Social Choice with Partial Knowledge of Treatment Response*. Princeton UP.
- [7] Manski, C., 2007. Minimax-Regret Treatment Choice with Missing Outcome Data. *Journal of Econometrics* 139, 105-115.
- [8] Manski, C., 2007. *Adaptive Minimax-Regret Treatment Choice, with Application to Drug Approval*. Unpublished manuscript, Northwestern University.
- [9] Manski, C., 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.
- [10] Paditz, L., 1989. On the Analytical Structure of the Constant in the Nonuniform Version of the Esseen Inequality. *Statistics* 20, 453-464.
- [11] Savage, L., 1951. The Theory of Statistical Decision. *Journal of the American Statistical Association* 46, 55-67.

- [12] Schlag, K., 2006. Minimax Regret in Practice - Four Examples on Treatment Choice. Unpublished manuscript, European University Institute.
- [13] Schlag, K., 2007. Eleven - Designing Randomized Experiments under Minimax Regret. Unpublished manuscript, European University Institute.
- [14] Shiryaev, A., 1995. Probability (2nd edition). Springer Verlag, New York.
- [15] Stoye, J., 2006. Minimax Regret Treatment Choice with Finite Samples. Unpublished manuscript, New York University.
- [16] Stoye, J., 2007. Minimax Regret Treatment Choice with Incomplete Data and Many Treatments. *Econometric Theory* 23, 190-199.
- [17] Stoye, J., 2007. Minimax Regret Treatment Choice with Finite Samples and Missing Outcome Data. *Proceedings of the 5th International Symposium on Imprecise Probability*, Prague.
- [18] van Beek, P., 1972. An Application of Fourier Methods to the Problem of Sharpening the Berry-Esseen Inequality. *Probability Theory and Related Fields* 23, 187-196.
- [19] Wald, A., 1950. *Statistical Decision Functions*. Wiley, New York.