



Tetenov, A. (2012). Statistical treatment choice based on asymmetric minimax regret criteria. *Journal of Econometrics*, 166(1), 157-165. DOI: 10.1016/j.jeconom.2011.06.013

Peer reviewed version

License (if available):
CC BY-NC-ND

Link to published version (if available):
[10.1016/j.jeconom.2011.06.013](https://doi.org/10.1016/j.jeconom.2011.06.013)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via Elsevier at <http://www.sciencedirect.com/science/article/pii/S0304407611001266>. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/pure/about/ebr-terms>

Statistical Treatment Choice Based on Asymmetric Minimax Regret Criteria*

Aleksey Tetenov[†]

This version: January 2011

(First draft: October 2007)

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. 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 started by Manski (2004) 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) and derive exact finite sample solutions to these problems for experiments with normal, Bernoulli and bounded distributions of outcomes. The paper also evaluates the properties of treatment choice and sample size selection based on classical hypothesis tests and power calculations in terms of regret.

JEL classification: C12, C21, C44.

Keywords: treatment effects, loss aversion, statistical decision theory, hypothesis testing.

*I am grateful to Chuck Manski and Elie Tamer, as well as the guest co-editor and two anonymous referees for their very helpful comments. I have also benefited from the opportunity to present this work in seminars at Caltech, Collegio Carlo Alberto, Cornell, Georgetown, and University of Arizona.

[†]Collegio Carlo Alberto, Moncalieri (Torino), Italy. aleksey.tetenov@carloalberto.org

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 an 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 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 how much data should be gathered to get a sufficiently accurate estimate of the average 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 decision rules 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 investigating treatment choice procedures that minimize maximum regret by Manski (2004, 2005, 2007, 2009), Hirano and Porter (2009)¹, Stoye (2007a, 2007b, 2009) and Schlag (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, which is necessarily lower.

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

¹It should be noted that Hirano and Porter (2009) consider local asymptotic minimax regret criterion, which differs from the "global" minimax regret considered by other authors listed here and does not generally imply results about asymptotics of minimax regret over the whole parameter space.

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.

The asymmetric minimax regret criterion proposed here combines minimax regret with a kinked linear welfare function that is intended to capture a policy maker's loss aversion. Maximum Type II regret of asymmetric minimax regret solutions is larger than their maximum Type I regret by a given factor. I show that when treatment effect estimates are normally distributed, hypothesis testing rules with a given level α correspond to asymmetric minimax regret solutions for some asymmetry factor $K(\alpha)$ for any sample size and variance. It turns out, however, that extreme degrees of loss aversion are needed to obtain treatment choice rules corresponding to hypothesis tests with standard significance levels.

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, update it using the data and maximize expected welfare with regard to the 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, rather than a subjective, prior distribution does not have a strong economic justification. Decision making based on maximum regret is a conservative approach to dealing with the lack of substantiated prior beliefs, since maximum regret is the sharp upper bound on expected regret for decision makers with any prior distributions.

Some decision makers will find asymmetric minimax regret criterion proposed here a simple alternative to the conventional hypothesis testing method for treatment choice problems, which has

the advantage of taking into account the magnitude of potential welfare losses in addition to raw error probabilities. Even for those who aren't interested in using the criterion for decision making, the paper could provide useful insights for understanding the welfare implications of applying conventional statistical methods to treatment choice problems.

The paper proceeds in the following order. Section 2 exposit the decision-theoretic formulation of the problem and introduces the asymmetric minimax regret criterion. In section 3, I consider a simple but instructive case where the experiment generates a normally distributed random variable with known variance. There I derive asymmetric minimax regret decision rules and establish their correspondence with hypothesis testing rules. I analyze conventional treatment choice rules based on hypothesis testing and sample size choice based on power analysis in light of their regret properties. Section 4 analyzes treatment choice in a more practically applicable setting with binary or bounded random treatment outcomes. Exact minimax regret results were obtained for these problems by Stoye (2009) and Schlag (2007). I extend their results to derive asymmetric minimax regret solutions using a different technique and demonstrate that minimax regret solutions proposed by these authors for bounded outcomes could be suboptimal if the decision maker can place an informative upper bound on the variance of the outcome distribution, which is the case in many applications. All proofs are collected in an appendix.

2 Statistical treatment rules, welfare and regret

The basic setting is the same as in Manski (2004, 2005). The planner's problem is to assign members of a large population to one of two available treatments after observing some statistical evidence of their effectiveness. Let $t = 0$ denote the status quo treatment and $t = 1$ the innovation. Each member $j \in J$ of the population has a response vector $y_j(t)$ describing potential outcomes under both treatments t . The population is a probability space (J, Ω, P) and the distribution $P[y(\cdot)]$ of the random vector $y(\cdot)$ describes treatment response across the population. The population is "large," in the sense that J is uncountable and $P(j) = 0$. The distribution P is ex ante known to belong to a feasible set $\{P_\gamma, \gamma \in \Gamma\}$. Parameter γ is called the *state of nature*.

Population members are observationally identical, thus the treatment assignment decision can be fully characterized by an action $a \in [0, 1]$, which is the proportion of population that is assigned

by the planner to the innovative treatment $t = 1$ independently of potential outcomes. Proportion $1 - a$, then, is assigned to the status quo treatment $t = 0$.

The planner observes an outcome of an experiment $X \in \mathcal{X}$ that is informative about γ . The distribution of X depends on the unknown state of nature γ and is denoted by Q_γ . A function $\delta : \mathcal{X} \rightarrow [0, 1]$ mapping experimental outcomes into treatment assignment proportions is a *statistical treatment rule* (or simply a *decision rule*). Then $\delta(X)$ is the action chosen when X is observed. The set of all feasible statistical treatment rules is labeled \mathcal{D} .

First, consider a utilitarian planner whose payoff from taking action a in state of nature γ is the average treatment outcome across the population:

$$U(a, \gamma) \equiv (1 - a) E_\gamma [y(0)] + a E_\gamma [y(1)] = E_\gamma [y(0)] + \theta_\gamma \cdot a.$$

Then the primary statistic of interest to the planner is the average treatment effect

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

Average treatment outcomes $E_\gamma [y(t)]$ are assumed to be finite for all t and γ .

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

$$W(\delta, \gamma) \equiv \int U(\delta(X), \gamma) dQ_\gamma = E_\gamma [y(0)] + \theta_\gamma E_\gamma [\delta(X)], \quad (1)$$

where $E_\gamma [\delta(X)] = \int \delta(X) dQ_\gamma$.

In addition to the welfare function to evaluate outcomes the planner needs to specify a criterion to deal with the uncertainty of γ . These generally fall into two categories: Bayesian and uniform. A Bayesian criterion maximizes $\int W(\delta, \gamma) d\mu(\gamma)$ for some measure μ on the space Γ , which could be a probability measure (a proper prior) or infinite (improper prior). Uniform criteria treat all states of nature symmetrically and are not affected by reparametrization of Γ . Stoye (2010) provides an extensive comparison of such criteria.

The simplest uniform criterion is: $\max_{\delta \in \mathcal{D}} \min_{\gamma \in \Gamma} W(\delta, \gamma)$, called minimax in statistics when the for-

ulated objective is to minimize loss. Manski (2004) shows that in a problem of treatment choice between status quo and an innovation the minimax decision rule is very conservative: the status quo treatment has to be chosen regardless of experimental data, any other decision rule could yield welfare lower than $E_\gamma[y(0)]$ if $\theta_\gamma < 0$.

2.1 Minimax regret

The conservativeness of the minimax criterion has led Manski (2004) and other authors to focus instead on the minimax regret criterion introduced by Savage (1951). The regret of a statistical treatment rule δ is the difference between the highest expected welfare achievable by any feasible statistical treatment rule in state of nature γ and the expected welfare of δ :

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

The highest welfare in state of nature γ is achieved by a decision rule that selects the optimal (in state γ) treatment regardless of experimental outcomes. The regret function, then, equals

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

Regret, then, is the product of the average probability of suboptimal treatment assignment and the magnitude of the resulting utilitarian welfare loss.

The *minimax regret* criterion select a statistical treatment rule that minimizes maximum regret

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

Minimax regret does not have the problem of excessive conservativeness highlighted by Manski (2004), yet remains a very simple uniform decision criterion. Regret (as opposed to the negative of welfare) embodies the intuitive notion of a loss function, since it normalizes the loss to zero when an optimal decision is made given the knowledge of the underlying parameters. In estimation problems with standard statistical loss functions for which $L(|\hat{\theta} - \theta|) \geq 0$ and $L(0) = 0$, minimax and minimax regret criteria are identical. Axiomatic properties of minimax regret were first studied

by Milnor (1954) and recently by Hayashi (2008) and Stoye (2010); they apply directly to minimax regret with reference-dependent utility function considered below when the reference point $E_\gamma[y(0)]$ is known.

2.2 Asymmetric minimax regret

To allow for asymmetric concern about Type I and Type II errors evident in hypothesis testing, I consider minimax regret treatment choice with asymmetric reference-dependent welfare functions exhibiting loss aversion (Kahneman and Tversky, 1979). For an asymmetry coefficient $K > 0$, let the welfare function U_K be linear² in the average treatment outcomes with the same slope as U above the reference point $E_\gamma[y(0)]$ and a K times steeper below the reference point. Formally, define U_K as:

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

The expected welfare of a statistical treatment rule δ then equals

$$W_K(\delta, \gamma) = \int U_K(\delta(X), \gamma) dQ_\gamma = E_\gamma[y(0)] + \begin{cases} \theta_\gamma E_\gamma[\delta(X)] & \text{if } \theta_\gamma > 0, \\ K \cdot \theta_\gamma E_\gamma[\delta(X)] & \text{if } \theta_\gamma \leq 0. \end{cases} \quad (4)$$

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 $\gamma \in \Gamma$:

$$W(\delta_2, \gamma) \succcurlyeq W(\delta_1, \gamma) \iff W_K(\delta_2, \gamma) \succcurlyeq W_K(\delta_1, \gamma).$$

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

²A kinked-linear utility is used for tractability. In Section 3.1 I discuss why this restriction does not have a substantial impact.

The regret function for expected welfare (4) equals

$$R_K(\delta, \gamma) = \begin{cases} \theta_\gamma \cdot (1 - E_\gamma[\delta(X)]) = R(\delta, \gamma) & \text{if } \theta_\gamma > 0, \\ -K\theta_\gamma \cdot E_\gamma[\delta(X)] = KR(\delta, \gamma) & \text{if } \theta_\gamma \leq 0. \end{cases}$$

The only difference from the regret function for standard linear welfare (2) is the factor K for $\theta_\gamma \leq 0$ which evaluates losses relative to the status at a higher rate. Maximum regret under the asymmetric welfare function can be expressed through the regret function for linear welfare as

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

where $\bar{R}_{Type\ I}(\delta) \equiv \max_{\gamma: \theta_\gamma \leq 0} R(\delta, \gamma)$ is the *maximum Type I regret* (across states of nature in which the innovation is inferior) and $\bar{R}_{Type\ II}(\delta) \equiv \max_{\gamma: \theta_\gamma > 0} R(\delta, \gamma)$ is the *maximum Type II regret*. 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, asymmetric welfare functions could be considered indirectly by solving the weighted minimax regret problem

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

for the linear expected welfare (1).

3 Normal experiment

I first consider an experiment whose outcome $\hat{\theta} \sim \mathcal{N}(\theta_\gamma, \sigma^2)$ is a scalar normally distributed random variable with unknown mean $\theta_\gamma \in \mathbb{R}$ (the average treatment effect of interest) and known variance σ^2 . If we let $\sigma^2 = N^{-1}\sigma_0^2$, normal experiments with different variances are informative for comparing experiments with different sample sizes.

It follows from the results of Karlin and Rubin (1956, Theorem 1) that if the distribution of $\hat{\theta}$ exhibits the monotone likelihood ratio property (which holds for normal and binomial distributions)

and the welfare function is (1), then the class of *monotone decision rules*

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

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 nature). Since $P_\gamma(\hat{\theta} = T) = 0$ for normally distributed $\hat{\theta}$, a smaller class of *threshold decision rules*

$$\delta_T(\hat{\theta}) \equiv \mathbb{I}[\hat{\theta} > T], T \in \mathbb{R}$$

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

Given that $\hat{\theta}$ is normally distributed, the regret of a threshold decision rule δ_T in state of nature γ equals

$$R(\delta_T, \gamma) = \begin{cases} -\theta_\gamma P_\gamma(\hat{\theta} > T) = -\theta_\gamma \Phi(\sigma^{-1}(\theta_\gamma - T)) & \text{if } \theta_\gamma \leq 0, \\ \theta_\gamma P_\gamma(\hat{\theta} \leq T) = \theta_\gamma \Phi(\sigma^{-1}(T - \theta_\gamma)) & \text{if } \theta_\gamma > 0, \end{cases}$$

the probability of an incorrect treatment choice multiplied by the magnitude of the loss $|\theta_\gamma|$.

Substituting $h = \sigma^{-1}\theta_\gamma$, maximum Type I and Type II regret could be expressed as

$$\begin{aligned} \bar{R}_{Type\ I}(\delta_T) &= \sigma \max_{h \leq 0} \{-h\Phi(h - \sigma^{-1}T)\}, \\ \bar{R}_{Type\ II}(\delta_T) &= \sigma \max_{h > 0} \{h\Phi(\sigma^{-1}T - h)\}. \end{aligned} \tag{6}$$

These functions have finite positive values for every $T \in \mathbb{R}$. 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 , $\bar{R}_{Type\ I}(\delta_T) \rightarrow +\infty$ as $T \rightarrow -\infty$ and $\bar{R}_{Type\ I}(\delta_T) \rightarrow 0$ as $T \rightarrow +\infty$.

b) $\bar{R}_{Type\ II}(\delta_T)$ is a continuous, strictly increasing function of T , $\bar{R}_{Type\ II}(\delta_T) \rightarrow 0$ as $T \rightarrow -\infty$ and $\bar{R}_{Type\ II}(\delta_T) \rightarrow +\infty$ as $T \rightarrow +\infty$.

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 $\max_{\gamma \in \Gamma} R(\delta_T, \theta_\gamma) = \max(\bar{R}_{T_{\text{type I}}}(\delta_T), \bar{R}_{T_{\text{type II}}}(\delta_T))$ is minimized when $\bar{R}_{T_{\text{type I}}}(\delta_T) = \bar{R}_{T_{\text{type II}}}(\delta_T)$, which happens at $T = 0$ since $R(\delta_T, \theta_\gamma) = R(\delta_{-T}, -\theta_\gamma)$ and $\bar{R}_{T_{\text{type II}}}(\delta_T) = \bar{R}_{T_{\text{type I}}}(\delta_{-T})$. The minimax regret treatment rule in this problem coincides with the *plug-in rule* $\delta_0 = \mathbb{I}[\hat{\theta} > 0]$ which assigns treatments as if the estimate $\hat{\theta}$ of the average treatment effect was the true value.

The asymmetric minimax regret statistical treatment rule δ_K^A under welfare function W_K is uniquely characterized by the equation

$$K \cdot \bar{R}_{T_{\text{type I}}}(\delta_K^A) = \bar{R}_{T_{\text{type II}}}(\delta_K^A).$$

It is evident from (6) that if T_K is the asymmetric minimax regret threshold for $\sigma = 1$, then the threshold for other values of σ equals σT_K . The implicit characterization of T_K can be rewritten as

$$K \max_{h \leq 0} \{-h\Phi(h - T_K)\} = \max_{h > 0} \{h\Phi(T_K - h)\}. \quad (7)$$

The threshold value T_K is an increasing function of the asymmetry factor K .

3.1 Comparison with hypothesis testing

Hypothesis tests and power calculations are a general device not explicitly designed to take into account the magnitude of welfare losses specific to treatment choice problems, hence cannot be expected, for example, to minimize maximum regret. However, since these techniques are ubiquitously used in treatment choice problems, it is of interest to evaluate their performance in terms of the regret function, which captures the welfare consequences more fully (for this class of problems) than the statistical power function.

A conventional one-sided hypothesis test with significance level α rejects the null hypothesis ($\theta_\gamma \leq 0$) and assigns the innovative treatment if $\hat{\theta} > \sigma\Phi^{-1}(1 - \alpha)$. This critical value guarantees that the probability of a Type I error does not exceed α for any $\theta_\gamma \leq 0$. The statistical treatment rule based on results of a hypothesis test with level α is a threshold rule $\delta_{H(\alpha)}$ with threshold $H(\alpha) \equiv \sigma\Phi^{-1}(1 - \alpha)$ proportional to the standard error σ . Thus a hypothesis test based treatment rule can also be rationalized as a solution to an asymmetric minimax regret problem with asymmetry




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 .

factor

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

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

Table 1 provides maximum Type I and Type 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 would also 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 Type II 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 does have much lower regret over $\theta_\gamma \leq 0$ ($.0082\sigma$ vs. $.17\sigma$ for the minimax regret rule). The hypothesis test rule generates the largest expected welfare losses relative to the mimiax regret rule when the innovative treatment is moderately effective.

Since it is a product of the magnitude of error $|\theta_\gamma|$ and its probability, regret of any threshold decision rule converges to zero as $\theta_\gamma \rightarrow 0$. Regret of all threshold decision rules also goes to zero for large values of θ_γ because probability of error declines exponentially, while θ_γ grows linearly. This provides an intuition for finding in Manski and Tetenov's (2007) finding that many nonlinear (e.g. logarithmic) transformations of the planner's utility have little effect on minimax regret treatment rules. Similarly, minimax regret rules computed for a nonlinear loss averse utility function estimated by Tversky and Kahneman (1992) do not differ greatly from those obtained here for kinked linear utility with the same asymmetry parameter.

3.2 Sample size selection

I will illustrate sample size selection based on maximum regret by comparing it with 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 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 selects $\bar{\theta} > 0$ as the clinically relevant size of positive treatment effect. Following the guidelines, she selects the sample size for which the variance of $\hat{\theta}$ equals $\bar{\sigma}^2$, where $\bar{\sigma}^2$ satisfies the condition that $\hat{\theta}$ will fall under the 5% hypothesis test threshold $H(.05) = \bar{\sigma}\Phi^{-1}(.95)$ with probability 10% if $\theta_\gamma = \bar{\theta}$:

$$\begin{aligned} P_{\bar{\theta}}(\hat{\theta} \leq H(.05)) &= \Phi(\Phi^{-1}(.95) - \bar{\theta}/\bar{\sigma}) = .1, \\ \bar{\sigma} &= \bar{\theta}/(\Phi^{-1}(.95) - \Phi^{-1}(.1)) = .342\bar{\theta}. \end{aligned}$$

At the chosen value $\theta_\gamma = \bar{\theta}$ the regret then equals $.1\bar{\theta}$. The regret over $\theta_\gamma > 0$, however, is not maximized at $\theta_\gamma = \bar{\theta}$, but at a smaller value $\theta_\gamma = 1.46\bar{\sigma} = .5\bar{\theta}$ for which the probability of a Type II error is 57% and the Type II regret achieves its maximum of $.837\bar{\sigma} = .286\bar{\theta}$.

The exact rationale for trying to limit the probability of Type II errors to a particular value of 10% at a particular value $\bar{\theta}$ is unclear to me, but suppose that the researcher didn't want Type II regret to exceed $.1\bar{\theta}$. If the hypothesis test decision rule was taken for granted, researcher planning the experiment would then want to select σ large enough so that maximum Type II regret (equal to $.837\sigma$) would not exceed $.1\bar{\theta}$. This would require $\sigma \leq .12\bar{\theta}$, implying a sample size over eight times larger than the one selected by conventional power calculations described above.

If a researcher instead planned to use a minimax regret decision rule δ_0 and wanted a sample size sufficient to limit the maximum regret by $.1\bar{\theta}$, she would select a sample size such that $.17\sigma = .1\bar{\theta}$, which would be almost three times smaller than the one derived by power calculations.

4 Exact statistical treatment rules for binary and bounded outcomes

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

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 we need to consider only 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

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

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

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

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

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 2 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 and asymmetric minimax regret treatment 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 2 *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 2 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 in Stoye (2009) using game theoretic methods.

Likewise, there is a unique value $D(\delta_K^A)$ such that $K\bar{R}_{Type I}(\delta_K^A) = \bar{R}_{Type II}(\delta_K^A)$. δ_K^A is the minimax regret statistical treatment rule for asymmetric welfare function W_K .

The following proposition derives the exact large sample limit of maximum regret of asymmetric minimax regret decision rules δ_K^A and shows that the same large sample limit is attained by decision rules derived from the normal approximation

$$\delta_K^N \equiv \mathbb{I} \left[X/N - p_0 > N^{-1/2} \sigma_0 T_K \right]. \quad (8)$$

with the threshold value T_K implicitly defined by (7) and $\sigma_0 = \sqrt{p_0(1-p_0)}$. For symmetric minimax regret, this approximation yields the plug-in rule $\delta^P = \delta_1^N = \mathbb{I}[X/N > p_0]$. The limit is equal to the maximum regret in a problem with N i.i.d. normally distributed outcomes with variance $\sigma_0^2 = p_0(1-p_0)$.

Proposition 3 *Asymptotic maximum regret of both minimax regret and plug-in statistical treatment rules is equal to*

$$\lim_{N \rightarrow \infty} \sqrt{N} \max_{\gamma \in \Gamma} R(\delta_K^A, \gamma) = \lim_{N \rightarrow \infty} \sqrt{N} \max_{\gamma \in \Gamma} R(\delta_K^N, \gamma) = \sigma_0 \max_{h > 0} \{h\Phi(T_K - h)\}.$$

The result in Proposition 3 establishes global asymptotic properties of decision rules over the whole parameter space Γ , which do not generally follow from the local asymptotic minimax results in Hirano and Porter (2009). The result will be used in the following section to highlight one of the limitations of extending minimax regret decision rules from binomial to bounded outcomes.

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 feasible probability distributions $P[y(1)]$. 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

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

This extension could be described algorithmically (as in Schlag (2007)):

- 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) take the treatment assignment probability prescribed by decision rule δ to $(\tilde{x}_1, \dots, \tilde{x}_N)$,
- c) average the assignment probabilities derived in step (b) over repeated independent draws (this is captured in the formula above by the expectation over the distribution of Z).

The random variables $\mathbf{I}[z_k \leq x_k], k = 0, \dots, N$ are i.i.d. Bernoulli with expectation $E_\gamma[y(1)]$, thus $\sum_{k=0}^N \mathbf{I}[z_k \leq x_k]$ has a Binomial distribution with parameters N and $E_\gamma[y(1)]$. For any state of nature γ , let $\bar{\gamma}$ be the state of nature 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 nature γ is the same as the regret of δ in a Bernoulli state of nature $\bar{\gamma}$ 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. For example, δ^M minimizes maximum regret for Bernoulli distributions. Suppose there was a

treatment rule $\tilde{\delta}'$ 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 and asymmetric minimax 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 because binomial averaging artificially inflates the variance of data from bounded outcomes to match the variance of binomial outcomes with the same mean. The following proposition shows that in this case a simple plug-in decision rule $\delta^P = \mathbb{I} \left[N^{-1} \sum_{i=1}^N x_i > p_0 \right]$ has a lower maximum regret than a binomial average extension of an exact minimax regret rule δ^M in the Bernoulli case.

Proposition 4 *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 $E_\gamma(y(1) - E_\gamma[y(1)])^2 < \sigma_u^2$, where $\sigma_u^2 < p_0(1 - p_0)$. Let (x_1, \dots, x_N) be i.i.d. random outcomes of treatment $t = 1$. Then*

$$\sqrt{N} \sup_{\gamma \in \Gamma} R(\delta^P, \gamma) \leq \sigma_u \max_{h>0} \{h\Phi(-h)\} + o(1).$$

Maximum regret of the 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 3 apply and

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

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 outcome distributions before looking for minimax regret based treatment rules, since rules that are optimal (in minimax regret sense) for a larger set of feasible distribution need not be optimal for a smaller feasible set.

5 Appendix: 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:

$$\Phi(\eta) < -\frac{\phi(\eta)}{\eta} \text{ for } \eta < 0.$$

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 $T \rightarrow \infty$.

Lemma 2 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 $p_\gamma \in (0, p_0)$,

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

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 3 It follows from lemma 2 that

$$\min(K\bar{R}_{Type I}(\delta_K^N), \bar{R}_{Type II}(\delta_K^N)) \leq \max_{\gamma \in \Gamma} R(\delta_K^A, \gamma) \leq \max(K\bar{R}_{Type I}(\delta_K^N), \bar{R}_{Type II}(\delta_K^N)).$$

If $\sqrt{N}\bar{R}_{Type I}(\delta_K^N)$ and $\sqrt{N}\bar{R}_{Type II}(\delta_K^N)$ both converge to $\sigma_0 \max_{h>0} [h\Phi(T_K - h)]$, then it follows that $\max_{\gamma \in \Gamma} R(\delta_K^A, \gamma)$ converges to the same limit. I will establish it for $\sqrt{N}\bar{R}_{Type II}(\delta_K^N)$, the proof for $\sqrt{N}\bar{R}_{Type I}(\delta_K^N)$ is analogous.

First, I will show using Berry-Esseen inequality that

$$\sup_{p_\gamma \in [p_0, p_0 + N^{-1/2}B]} \sqrt{N}R(\delta_K^A, \gamma) \rightarrow \sigma_0 \max_{h>0} [h\Phi(T_K - h)], \quad (9)$$

and then, using Chebyshev's inequality that for all N

$$\sup_{p_\gamma \in (p_0 + N^{-1/2}B, 1]} \sqrt{N}R(\delta_K^A, \gamma) < \sigma_0 \max_{h>0} [h\Phi(T_K - h)], \quad (10)$$

where

$$B = \max \left(2\sigma_0 T_K, \left(\sigma_0 \max_{h>0} [h\Phi(T_K - h)] \right)^{-1}, \sigma_0 \arg \max_{h>0} [h\Phi(T_K - h)] \right).$$

Let $h = \sqrt{N}\sigma_0^{-1}(p_\gamma - p_0)$ and $\sigma_\gamma = \sqrt{p_\gamma(1-p_\gamma)}$, then $h \in [0, \sigma_0^{-1}B]$ for $p_\gamma \in [p_0, p_0 + N^{-1/2}B]$

and

$$\begin{aligned} \sqrt{N}R(\delta_K^A, \gamma) &= \sqrt{N}(p_\gamma - p_0) P_\gamma \left(X/N - p_0 \leq N^{-1/2}\sigma_0 T_K \right) = \\ &= \sigma_0 h P_\gamma \left(X/N - p_\gamma \leq N^{-1/2}\sigma_0 T_K - (p_\gamma - p_0) \right) \\ &= \sigma_0 h P_\gamma \left(\sqrt{N}\sigma_\gamma^{-1} (X/N - p_\gamma) \leq \frac{\sigma_0}{\sigma_\gamma} (T_K - h) \right). \end{aligned} \quad (11)$$

The Berry-Esseen inequality with $C = 1$ (cf. Shiryaev (1995, p. 63,374)) applied to X , which is a sum of N i.i.d. Bernoulli variables with mean p_γ , yields for every $z \in \mathbb{R}$

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

For sufficiently large N , $p_0 + N^{-1/2}B < 1$, then for some finite M , $\frac{p_\gamma^2 + (1-p_\gamma)^2}{\sqrt{p_\gamma(1-p_\gamma)}} \leq M$ for all $p_\gamma \in [p_0, p_0 + N^{-1/2}B]$. Setting $z = \frac{\sigma_0}{\sigma_\gamma} (T_K - h)$ and applying the inequality to (11) yields

$$\left| \sqrt{N}R(\delta_K^A, \gamma) - \sigma_0 h \Phi \left(\frac{\sigma_0}{\sigma_\gamma} (T_K - h) \right) \right| \leq \sigma_0 h N^{-1/2} M \leq N^{-1/2} M B.$$

Since Φ has a bounded derivative,

$$\sup_{h \in [0, \sigma_0^{-1}B]} \left| \Phi \left(\frac{\sigma_0}{\sigma_\gamma} (T_K - h) \right) - \Phi(T_K - h) \right| \leq \phi(0) \sup_{h \in [0, \sigma_0^{-1}B]} \left| \left(1 - \frac{\sigma_0}{\sigma_\gamma} \right) (T_K - h) \right| \rightarrow 0$$

as $N \rightarrow \infty$, since $\sigma_\gamma \rightarrow \sigma_0$ in the shrinking interval $[p_0, p_0 + N^{-1/2}B]$.

Combining the results,

$$\begin{aligned} \sup_{h \in [0, \sigma_0^{-1}B]} \left| \sqrt{N}R(\delta_K^A, \gamma) - \sigma_0 h \Phi(T_K - h) \right| &\leq \sup_{h \in [0, \sigma_0^{-1}B]} \left| \sqrt{N}R(\delta_K^A, \gamma) - \sigma_0 h \Phi\left(\frac{\sigma_0}{\sigma_\gamma}(T_K - h)\right) \right| + \\ &+ \sup_{h \in [0, \sigma_0^{-1}B]} \left| \Phi\left(\frac{\sigma_0}{\sigma_\gamma}(T_K - h)\right) - \Phi(T_K - h) \right|. \end{aligned}$$

(9) follows since both terms converge to zero as $N \rightarrow \infty$ and $\sigma_0 \max_{h>0} [h\Phi(T_K - h)] = \max_{h \in [0, \sigma_0^{-1}B]} [\sigma_0 h \Phi(T_K - h)]$ due to the choice of B .

For $p_\gamma - p_0 > N^{-1/2}B$, $(p_0 - p_\gamma) + N^{-1/2}\sigma_0 T_K < (p_0 - p_\gamma)/2 < 0$, since $B \geq 2\sigma_0 T_K$. Applying Chebyshev's inequality to $X/N - p_\gamma$, which has variance $\sigma_\gamma^2/N \leq 1/(4N)$ yields

$$\begin{aligned} P_\gamma \left(X/N - p_0 \leq N^{-1/2}\sigma_0 T_K \right) &= P_\gamma \left(X/N - p_\gamma \leq (p_0 - p_\gamma) + N^{-1/2}\sigma_0 T_K \right) \leq \\ &\leq P_\gamma \left(X/N - p_\gamma \leq (p_0 - p_\gamma)/2 \right) \leq \\ &\leq \frac{\sigma_\gamma^2/N}{(p_\gamma - p_0)^2/4} \leq \frac{1}{N(p_\gamma - p_0)^2}. \end{aligned}$$

Then (10) follows from

$$\begin{aligned} \sqrt{N}R(\delta_K^A, \gamma) &= \sqrt{N}(p_\gamma - p_0) P_\gamma \left(X/N - p_0 \leq N^{-1/2}\sigma_0 T_K \right) \leq \\ &\leq \frac{1}{\sqrt{N}(p_\gamma - p_0)} < \frac{1}{B} \leq \sigma_0 \max_{h>0} [h\Phi(T_K - h)]. \end{aligned}$$

Proposition 4 Let V_γ denote the variance of a random variable in state of nature 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 $V_\gamma[y(1)] < \sigma_u^2/9$ (thus $V_\gamma\left[\frac{1}{N}\sum_{i=1}^N x_i - p_\gamma\right] < N^{-1}\sigma_u^2/9$) the one-sided Chebyshev's inequality yields

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

Applying the result to the formula for regret of the plug-in rule δ^P yields

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

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

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

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

Applying the inequality to the regret of plug-in rule δ^P yields

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

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

For all γ that do not fall into one of the two cases considered above, $p_\gamma - p_0 < 6N^{-1/2}\sigma_u$ and $V_\gamma [y(1)] \in [\sigma_u^2/9, \sigma_u^2]$. The Berry-Esseen inequality (cf. Shiryaev (1995, p. 374)), applied to the sum of N i.i.d. random variables $(x_i - p_\gamma)$, yields for any $z \in \mathbb{R}$

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

Substituting $z = \sqrt{N}V_\gamma^{-1/2} [y(1)] \cdot (p_0 - p_\gamma)$ gives us

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

The second inequality hold because given $y(1) - p_\gamma \in [0, 1]$, $E_\gamma |y(1) - p_\gamma|^3 \leq V_\gamma [y(1)]$. The third inequality follows from $V_\gamma [y(1)] \geq \sigma_u^2/9$.

Applying this inequality to the regret formula for δ^P yields

$$\begin{aligned}
\sqrt{N}R(\delta^P, \gamma) &= \sqrt{N}(p_\gamma - p_0) P_\gamma \left(N^{-1} \sum_{i=1}^N x_i \leq p_0 \right) \leq \\
&\leq \sqrt{N}(p_\gamma - p_0) \left(\Phi \left(\sqrt{N}V_\gamma^{-1/2} [y(1)] \cdot (p_0 - p_\gamma) \right) + \frac{3}{\sqrt{N}\sigma_u} \right) \leq \\
&\leq V_\gamma^{1/2} [y(1)] \max_{h>0} h\Phi(-h) + \frac{3(p_\gamma - p_0)}{\sigma_u} \leq \sigma_u \max_{h>0} h\Phi(-h) + 18N^{-1/2}.
\end{aligned}$$

The second inequality uses substitution $h = \sqrt{N}V_\gamma^{-1/2} [y(1)] \cdot (p_0 - p_\gamma)$. The last inequality uses $p_\gamma - p_0 < 6N^{-1/2}\sigma_u$.

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

$$\sqrt{N}R(\delta^P, \gamma) \leq \frac{\sigma_u}{6} < \sigma_u \max_{h>0} [h\Phi(-h)].$$

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

$$\sqrt{N}R(\delta^P, \gamma) \leq \sigma_u \max_{h>0} h\Phi(-h) + \frac{18}{\sqrt{N}},$$

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

References

- [1] Berger, J., 1985. *Statistical Decision Theory and Bayesian Analysis*, 2nd ed. Springer Verlag, New York.
- [2] Hayashi, T., 2008. Regret Aversion and Opportunity-Dependence. *Journal of Economic Theory* 139, 242-268.
- [3] Hirano, K., Porter, J., 2009. Asymptotics for Statistical Treatment Rules. *Econometrica* 77, 1683 - 1701.
- [4] International Committee on Harmonization, 1998. Guideline E9: Statistical Principles for Clinical Trials.
- [5] Kahneman, D., Tversky, A., 1979. Prospect Theory: An Analysis of Decision Under Risk. *Econometrica* 47, 263-291.
- [6] Karlin, S., Rubin, H., 1956. The Theory of Decision Procedures for Distributions with Monotone Likelihood Ratio. *The Annals of Mathematical Statistics* 27, 272-299.
- [7] Manski, C., 2004. Statistical Treatment Rules for Heterogeneous Populations. *Econometrica* 72, 1221-1246.
- [8] Manski, C., 2005. *Social Choice with Partial Knowledge of Treatment Response*. Princeton UP.
- [9] Manski, C., 2007. Minimax-Regret Treatment Choice with Missing Outcome Data. *Journal of Econometrics* 139, 105-115.
- [10] Manski, C., 2009. Adaptive Partial Drug Approval. *The Economists' Voice* 6, Iss. 4, Art. 9.
- [11] 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.
- [12] Milnor, J., 1954. Games Against Nature. In R.M. Thrall, C.H. Coombs, R.L. Davis (eds.): *Decision Processes*. Wiley.

- [13] Paditz, L., 1989. On the Analytical Structure of the Constant in the Nonuniform Version of the Esseen Inequality. *Statistics* 20, 453-464.
- [14] Savage, L., 1951. The Theory of Statistical Decision. *Journal of the American Statistical Association* 46, 55-67.
- [15] Schlag, K., 2007. Eleven - Designing Randomized Experiments under Minimax Regret. Working paper, European University Institute.
- [16] Shiryaev, A., 1995. *Probability* (2nd edition). Springer Verlag, New York.
- [17] Stoye, J., 2007a. Minimax Regret Treatment Choice with Incomplete Data and Many Treatments. *Econometric Theory* 23, 190-199.
- [18] Stoye, J., 2007b. Minimax Regret Treatment Choice with Finite Samples and Missing Outcome Data. *Proceedings of the 5th International Symposium on Imprecise Probability, Prague*.
- [19] Stoye, J., 2009. Minimax Regret Treatment Choice with Finite Samples. *Journal of Econometrics* 151, 70-81.
- [20] Stoye, J., 2010. *Statistical Decisions under Ambiguity. Theory and Decision*, forthcoming.
- [21] Tversky, A., Kahneman, D., 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5(4), 297–323.
- [22] 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.
- [23] Wald, A., 1950. *Statistical Decision Functions*. Wiley, New York.