NORTHWESTERN UNIVERSITY

# Essays in Statistical Decision Theory of Treatment Choice

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Economics

By

Aleksey Tetenov

EVANSTON, ILLINOIS

June 2008

# ABSTRACT

Essays in Statistical Decision Theory of Treatment Choice

Aleksey Tetenov

I study the problem of choice between two treatments for a population of observationally identical individuals based on statistical evidence about average treatment effects that does not reveal the best treatment with certainty. I approach the problem from the perspective of statistical decision theory, derive treatment rules that minimize maximum regret and contrast them with inference and decision making methods of classical statistics.

In Chapter 1, the choice is between a status quo treatment with a known outcome distribution and an innovation whose outcomes are observed only in a randomized experiment. I introduce criteria that asymmetrically treat Type I regret (from adopting an inferior innovation) and Type II regret (from rejecting a superior innovation). I derive exact finite sample solutions for experiments with normal, Bernoulli, or bounded distributions of individual outcomes and discuss approaches for other sampling distributions. For normal outcomes, asymmetric minimax regret rules coincide with

classical hypothesis testing rules, but conventional test levels imply unrealistic degrees of asymmetry.

In Chapter 2, written with Charles Manski, the treatments have binary outcomes and the objective is to maximize a concave-monotone function of the success rate. We show that admissibility of statistical treatment rules depends on that function's curvature. We establish a general complete class for concave and strictly monotone functions and a more specific result for functions with weak curvature, including the logarithmic function often used to model risk aversion. We compute minimax regret rules for specific welfare functions to demonstrate how they depend on the functions' curvature.

Chapter 3 studies the measurement of the precision of inference on partially identified parameters. Planners of surveys and experiments that partially identify parameters of interest can choose between using resources to reduce sampling error or to reduce the extent of partial identification. Previous research unanimously measured precision of inference by the length of 95% confidence intervals for the identification region. In a problem with normally distributed data, I show that other measures of precision (maximum mean squared error and maximum regret for treatment choice) yield qualitatively different conclusions about the relative value of reducing sampling error and the extent of partial identification.

## Acknowledgements

I am grateful to my dissertation committee – Charles F. Manski, Elie Tamer, and Benjamin F. Jones – for valuable feedback and suggestions on this dissertation and other matters. I am particularly indebted to Chuck Manski, who has been an incredibly dedicated advisor and teacher, and an endless source of wisdom, inspiration, and support throughout my graduate studies at Northwestern. I owe to him my interest in the problems of statistical treatment choice and partial identification, which led to this dissertation, and I've benefited enormously from our many valuable discussions on these topics. It has also been my pleasure to collaborate with Chuck, and I thank him for letting me include our joint paper as the second chapter of this dissertation.

In large part, I owe the motivation to press on with the research contained in the third chapter to Elie Tamer's interest and comments on a much earlier version of the paper, and to a very helpful discussion with Joerg Stoye.

While this dissertation is entirely theoretical, I have also been very fortunate to collaborate on applied empirical projects at Northwestern with Greg Duncan, Ben Jones, and Brian Uzzi.

The first two years of coursework would have been unimaginable without Paul Gao, Wallace Mok, Ben Passty, Zahra Siddique, and Kristina Steffenson, whose company made the endless hours of studying much more enjoyable and rewarding.

This dissertation is dedicated to my late grandmother and to my mother. From a young age they nurtured in me a sense of curiosity about the world, without which I would not have chosen to pursue academic research. My mother made enormous efforts and sacrifices to ensure that I received the best possible education and arranged a mix of public, private, and home schooling that gave me the best of each. Furthermore, I entirely owe my knowledge of elementary mathematics, which served as a solid foundation for my further studies, to her talented teaching.

# Table of Contents

ABSTRACT 3

Acknowledgements 5

List of Tables 9

List of Figures 10

Chapter 1. Statistical Treatment Choice Based on Asymmetric Minimax Regret Criteria 11

1.1. Introduction 11

1.2. Statistical Treatment Rules, Welfare and Regret 15

1.3. Simple Normal Experiment 23

1.4. Exact Statistical Treatment Rules for Binary and Bounded Outcomes 35

1.5. Evaluating Regret Using Approximations and Bounds 42

1.6. Proofs 49

Chapter 2. Admissible Treatment Rules for a Risk-Averse Planner with Experimental Data on an Innovation 60

2.1. Introduction 60

2.2. Background 64

2.3. Admissible Treatment Rules for Risk-Averse Planners 67

2.4. Bayes and Minimax-Regret Rules 79
2.5. Implications for Treatment Choice in Practice 84
Chapter 3. Measuring Precision of Statistical Inference on Partially Identified Parameters 92
3.1. Introduction 92
3.2. Statistical Setting and the Confidence Interval Approach 95
3.3. Minimax Mean Squared Error Approach 98
3.4. Minimax Regret Approach 100
3.5. Conclusion 106
3.6. Proofs 107
References 115

## List of Tables

1.1 Maximum Type I and Type II regret of statistical treatment rules induced by hypothesis tests based on a normally distributed estimate with known variance. 28

# List of Figures


1.1 Maximum Type I and Type II regret as functions of the decision rule threshold. 26

1.2 Regret functions of minimax regret and hypothesis test based decision rules. 29

1.3 Evaluation of maximum regret of the plug-in ($T = .5$) statistical tretment rule. 46

1.4 Maximum Type I and Type II regret approximations for a range of threshold statistical treatment rules. 47

2.1 Maximum regret, $N = 1...100$. (a) $\alpha = 0.25, f(x) = x$; (b) $\alpha = 0.75, f(x) = x$; (c) $\alpha = 0.25, f(x) = \log x$; (d) $\alpha = 0.75, f(x) = \log x$. 87

2.2 Threshold sample size, $N = 1...20$: (a) $\alpha = 0.25$; (b) $\alpha = 0.75$. 89

3.1 Bounds on $q(\delta, \theta_O)$ that guarantee attaining the lower bound on maximum regret $(P/2)$. 105

CHAPTER 1

# Statistical Treatment Choice Based on Asymmetric Minimax Regret Criteria

## 1.1. 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, 2008a, 2008b), Hirano and Porter (2006), Stoye (2007a, 2007b, 2007c), Eozenou, Rivas, and Schlag (2006) 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. 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 chapter'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 chapter 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 $\alpha$ 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 chapter proceeds in the following order. Section 1.2 exposits the decision-theoretic formulation of the problem and the criteria used to address it. In section 1.3, 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 1.4 analyzes treatment choice when treatment outcomes are either binary or bounded random variables. Exact mimimax regret results were obtained for these problems by Stoye (2007b) 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 1.5, 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. Section 1.6 collects all proofs.

## 1.2. 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 $(J, \Omega, P)$ 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 $\{P_\gamma, \gamma \in \Gamma\}$. $\gamma$ will be called the *state of the world*. I assume that average treatment outcomes $E_\gamma[y(t)]$ are finite for all $t$ and $\gamma$.

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 $\gamma$ 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 $\gamma$ and will be denoted by $Q_\gamma$. A (random) function $\delta$ 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 $\delta$ 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 $\gamma$. If the planner's welfare function is $U(a, \gamma)$, then the expected welfare from using statistical treatment rule $\delta$ in state of the world $\gamma$ equals

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

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

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

The analysis of this chapter is based on a normalization of the expected welfare called *regret*. The regret of statistical treatment rule $\delta$ is the difference between the highest expected welfare achievable by any feasible statistical treatment rule in state of the world $\gamma$ and the expected welfare of statistical treatment rule $\delta$ :

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

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

$$R(\delta,\gamma) = W\left(\delta^*_\gamma,\gamma\right) - W(\delta,\gamma) = \begin{cases} \theta_\gamma \cdot (1 - E_\gamma[\delta(X)]) & \text{if } \theta_\gamma > 0 \\ -\theta_\gamma \cdot E_\gamma[\delta(X)] & \text{if } \theta_\gamma \leq 0. \end{cases} \tag{1.2}$$

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.

### 1.2.1. 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 $\gamma$ that fall under the null hypothesis). The significance level (also called $\alpha$-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 $\bar{\theta}_\gamma$.

### 1.2.2. 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 $\delta'$ is preferred to $\delta$ 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

$$\delta_M \in \arg\min_{\delta \in \mathcal{D}} \max_{\gamma \in \Gamma} R(\delta, \gamma) \tag{1.3}$$

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

### 1.2.3. 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:

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

The expected welfare for this kinked linear welfare function equals

$$\begin{aligned} W_{A(K)}(\delta,\gamma) &\equiv \int_{X\in\mathcal{X}} U_{A(K)}(\delta(X),\gamma)\,dQ_\gamma \\ &= E_\gamma[y(0)] + \begin{cases} \theta_\gamma E_\gamma[\delta(X)] & \text{if } \theta_\gamma > 0, \\ K\cdot\theta_\gamma E_\gamma[\delta(X)] & \text{if } \theta_\gamma \le 0. \end{cases} \end{aligned} \tag{1.4}$$

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 $\gamma\in\Gamma$ :

$$W(\delta_2,\gamma) \gtreqless W(\delta_1,\gamma) \Longleftrightarrow W_{A(K)}(\delta_2,\gamma) \gtreqless W_{A(K)}(\delta_1,\gamma).$$

Thus, the set of admissible statistical treatment rules is the same for all asymmetrical linear welfare functions (1.4) and for the standard linear welfare (1.1).

The regret function for expected welfare (1.4) equals

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

The only difference between this regret function and the regret function for standard linear welfare (1.2) is the factor $K$ for $\theta_\gamma \le 0$. Maximum regret under the asymmetrical welfare function can be expressed through the regret function for linear welfare as

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

where

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

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

$$\min_{\delta \in \mathcal{D}} \max \left( K \cdot \bar{R}_{Type\ I}(\delta), \bar{R}_{Type\ II}(\delta) \right) \tag{1.5}$$

for the linear expected welfare (1.1). In problem (1.5) the planner gives $K$ times greater weight to regret from Type I errors.

### 1.2.4. 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:

$$\begin{aligned} \delta_{L(\bar{r})} \in \quad & \arg\min_{\delta \in \mathcal{D}} \max_{\gamma \in \Gamma} R(\delta, \gamma), \\ & s.t.\ \bar{R}_{Type\ I}(\delta) \leq \bar{r}. \end{aligned} \tag{1.6}$$

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.

### 1.3. 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 $\sigma^2$:

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

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 $(x_1, ..., x_N)$ all have a normal distribution $\mathcal{N}(\theta_\gamma, \sigma_0^2)$, then $X$ is a sufficient statistic for $(x_1, ..., x_N)$ with variance $\sigma^2 = \frac{\sigma_0^2}{N}$. Comparing single normal draw experiments with different values of $\sigma$, 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 $\sqrt{N}\left(\hat{\theta} - \theta_\gamma\right) \xrightarrow{D} \mathcal{N}(0, \sigma_0^2)$. 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 $\sigma$ 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 (1.1), 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 $\delta'$ 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*

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

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 $\delta_T$ in state of the world $\gamma$ equals

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

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

Maximum Type I and Type II regret equal

$$\text{(1.7)} \quad \begin{aligned} \bar{R}_{Type\ I}\left(\delta_T\right) &= \max_{\gamma:\theta_\gamma \leq 0}\left\{-\theta_\gamma \cdot \Phi\left(\frac{\theta_\gamma - T}{\sigma}\right)\right\} = \sigma \cdot \max_{h \leq 0}\left\{-h\Phi\left(h - \frac{T}{\sigma}\right)\right\}, \\ \bar{R}_{Type\ II}\left(\delta_T\right) &= \max_{\gamma:\theta_\gamma > 0}\left\{\theta_\gamma \cdot \Phi\left(\frac{T - \theta_\gamma}{\sigma}\right)\right\} = \sigma \cdot \max_{h > 0}\left\{h\Phi\left(\frac{T}{\sigma} - h\right)\right\}. \end{aligned}$$

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 $R\left(\delta_T, \theta_\gamma\right) = R\left(\delta_{-T}, -\theta_\gamma\right)$, it follows that $\bar{R}_{Type\ II}\left(\delta_T\right) = \bar{R}_{Type\ I}\left(\delta_{-T}\right)$. Lemma 1.1 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.

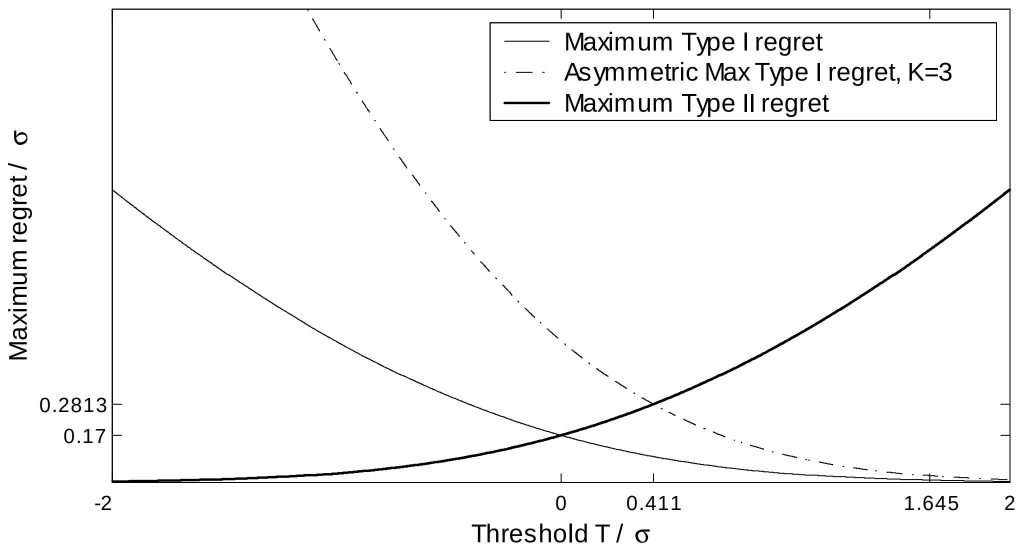


Figure 1.1: Maximum Type I and Type II regret as functions of the decision rule threshold.

**Lemma 1.1.** *a)* $\bar{R}_{Type\ I}(\delta_T)$ *is a continuous, strictly decreasing function of* $T$*,*

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

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

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

Figure 1.1 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 $\sigma$. The maximum regret $\max_{\gamma \in \Gamma} R(\delta_T, \theta_\gamma) = \max\left(\bar{R}_{Type\ I}(\delta_T), \bar{R}_{Type\ II}(\delta_T)\right)$ 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 $\delta_0$. 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

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

By substituting right-hand expressions from (1.7), 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 $\sigma$.

A conventional one-sided hypothesis test with significance level $\alpha$ 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 $\alpha$ for any $\theta_\gamma \leq 0$. Since $\frac{X-\theta_\gamma}{\sigma}$ has a standard normal distribution,

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

The statistical treatment rule based on results of a hypothesis test with level $\alpha$ is a threshold rule $\delta_{H(\alpha)}$with threshold $H(\alpha) \equiv \sigma\Phi^{-1}(1-\alpha)$. For a given test level $\alpha$, the threshold $T$ is proportional to the standard error $\sigma$. Thus a hypothesis test based

| 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 |
| $\alpha = .025$ | $T = 1.96\sigma$ | $.003665\sigma$ | $1.026\sigma$ | 279.9 |
| $\alpha = .01$ | $T = 2.326\sigma$ | $.001304\sigma$ | $1.264\sigma$ | 969.6 |

Table 1.1: Maximum Type I and Type II regret of statistical treatment rules induced by hypothesis tests based on a normally distributed estimate with known variance.

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 $\alpha$. In this normal model, the correspondence between a hypothesis test based rule with level $\alpha$ and an asymmetric minimax regret rule with level $K(\alpha)$ does not depend on the standard error of $\sigma$, 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 1.1 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

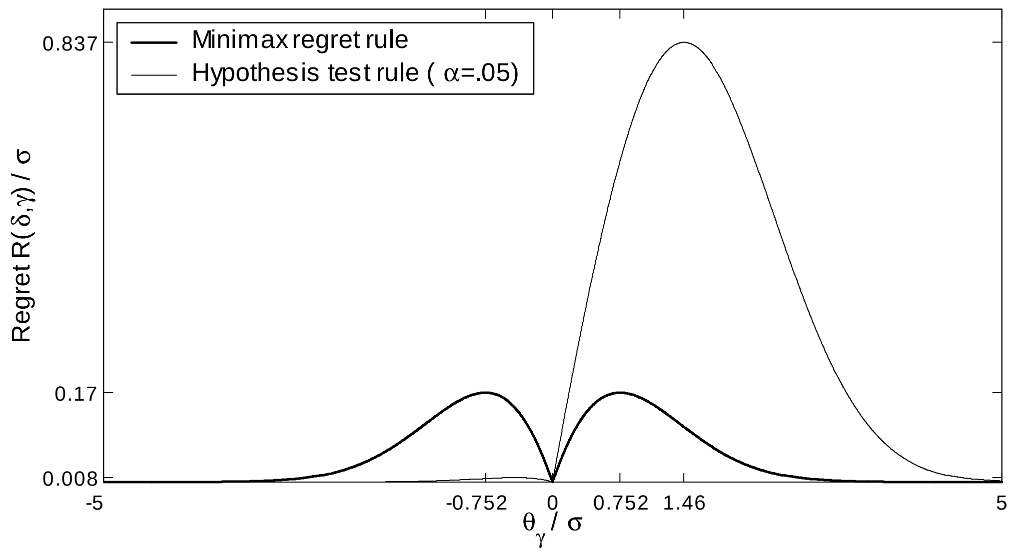


Figure 1.2: Regret functions of minimax regret and hypothesis test based decision rules.

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 1.2 compares the regret functions of the minimax regret treatment rule $\delta_0$ and the treatment rule $\delta_{H(.05)}$ induced by a hypothesis test with significance level $\alpha = .05$ over a range of feasible values of $\theta_\gamma$. The scale of both axes is normalized by $\sigma$. 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.

### 1.3.1. 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 $\sigma$.

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 1.1 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 (1.6). 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 $\sigma^2$, 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.

### 1.3.2. 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 $\theta_\gamma = \bar{\theta}$:

$$\begin{aligned} P\left(X \leq H(.05) | \theta_\gamma = \bar{\theta}\right) &= \Phi\left(\Phi^{-1}(.95) - \frac{\bar{\theta}}{\bar{\sigma}}\right) = .1, \\ \bar{\sigma} &= \frac{\bar{\theta}}{\Phi^{-1}(.95) - \Phi^{-1}(.1)} = \frac{\bar{\theta}}{2.926}. \end{aligned}$$

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 $.837\sigma$, 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 eight 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 $\delta_0$. 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 three times smaller than the one selected by power calculations.

### 1.3.3. 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 $(x_1, ..., x_N)$ consists of $N$ independent normally distributed observations with unknown mean $\theta_\gamma$ and unknown variance $\sigma_\gamma^2$. Let the set of feasible states of the world be

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

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

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

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 $S^2 \equiv \frac{1}{N}\sum_{i=1}^N \left(x_i - \bar{X}\right)^2$ the sample variance. It is well known (cf. Berger, 1985) that the pair $\left(\bar{X}, S^2\right)$ is a sufficient statistic for $(x_1, ..., x_N)$, 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 1.2.** *Let* $\delta_T \equiv 1\left|\bar{X} > T\right|$ *be a threshold statistical treatment rule such that* $T^* \equiv \frac{\sqrt{N}}{\bar{\sigma}}\left|T\right|$ *satisfies the condition*

$$\max_{h\in(0,T^*)}\left\{h\cdot\Phi\left(\frac{\bar{\sigma}}{\underline{\sigma}}\left(T^*-h\right)\right)\right\} \leq \max_{h\geq T^*}\left\{h\cdot\Phi\left(T^*-h\right)\right\}, \tag{1.8}$$

*then*

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

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

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

Condition (1.8) 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 $\delta_T$ could be higher on the set $\Gamma$ 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 $\delta_T$.

It follows from Proposition 1.2 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 $\Gamma$) if their threshold values satisfy condition (1.8). The range of thresholds for which condition (1.8) 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 $\frac{\sqrt{N}}{\bar{\sigma}}|T|\leq .22$.

## 1.4. 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.

### 1.4.1. 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 $(x_1,...,x_N)$, each having a Bernoulli distribution with mean $p_\gamma$. The sum of outcomes $X=\sum_{i=1}^{n}x_i$ has a binomial distribution with parameters $N$ and $p_\gamma$. $X$ is a sufficient statistic for $(x_1,...,x_N)$, 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

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

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 $p_\gamma$ 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 1.3 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 1.3.** *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 1.3 that there is a unique value of $D(\delta_M)$ such that

$$\bar{R}_{Type\ I}(\delta_M) = \bar{R}_{Type\ II}(\delta_M).$$

$\delta_M$ 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 (2007b, Proposition 1(iii)) using game theoretic methods.

Likewise, there is a unique value $D\left(\delta_{A(K)}\right)$ such that

$$K \cdot \bar{R}_{Type\ I}\left(\delta_{A(K)}\right) = \bar{R}_{Type\ II}\left(\delta_{A(K)}\right).$$

$\delta_{A(K)}$ 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 $D\left(\delta_{L(\bar{r})}\right)$ such that $\bar{R}_{Type\ I}\left(\delta_{L(\bar{r})}\right) = \bar{r}$, 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 $\delta_M$.

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 1.3, 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 $\delta_M$ and $\delta_P$ has little effect on maximum regret. Proposition 1.4 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 1.4.**

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

*which approximately equals* .17.

### 1.4.2. 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 $\delta : \{0,...,N\} \to [0,1]$ be a statistical treatment rule defined for the sum of $N$ i.i.d. Bernoulli distributed outcomes (as in the previous subsection). Let $X = (x_1,...,x_N)$ be an i.i.d. sample of bounded random variables with unknown distribution $P_\gamma[y(1)]$ and let $Z = (z_1,...,z_N)$ be a sample of i.i.d. uniform $(0,1)$ random variables independent of $X$. Then the binomial average extension of $\delta$ 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 $\delta$ to $(\tilde{x}_1,...,\tilde{x}_N)$.

The random variables $1[z_k \leq x_k], k = 0,...,N$ 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 $\gamma$, 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 $\gamma$ is the same as the regret of $\delta$ 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 $\delta$ in the problem with Bernoulli outcomes ($\gamma \in \Gamma_B$).

If statistical treatment rule $\delta$ 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 $\delta_M$ 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 $\delta'$ would have to have lower maximum regret over $\Gamma_B$ than $\delta_M$, which would imply that $\delta_M$ 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 $\Gamma$ includes the set of Bernoulli outcome distributions with the same means $\Gamma_B$. 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 it's 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 $\delta_M$, a minimax regret statistical treatment rule in the Bernoulli case.

**Proposition 1.5.** *Let $p_0 = E\left[y\left(0\right)\right]$ and let $\left\{P_\gamma, \gamma \in \Gamma\right\}$ be the set of feasible probability distributions of $y\left(1\right)$ such that $E_\gamma \left(y\left(1\right) - E_\gamma\left[y\left(1\right)\right]\right)^2 < \sigma_h^2$, where $\sigma_h^2 < p_0\left(1 - p_0\right)$. Let $\left(x_1, ..., x_N\right)$ be i.i.d. random outcomes of treatment $t = 1$. Then*

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

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

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

the results of proposition 1.4 apply and

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

Thus, for large enough $N$, $\max_{\gamma \in \Gamma} R(\tilde{\delta}_M, \gamma) > \sup_{\gamma \in \Gamma} R\left(\delta_P, \gamma\right)$. 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.

### 1.5. 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 $(x_1, ..., x_N)$ 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 $(x_1, ..., x_N)$ 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 1.3, the class of threshold decision rules $\delta_T \equiv 1\left|\bar{X} > T\right|$ 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 $\delta_T$ 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 $\delta_T$,

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

the planner needs to know, for each value of $\theta$, the range of feasible probabilities that the sample mean $\bar{X}$ exceeds the threshold $T$. Note that for each $\gamma$, $P_\gamma(\bar{X} > T)$ 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 $P_\gamma\left(\bar{X} \lesseqgtr T\right)$ 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 $\theta$. 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 $\theta$. If an insufficiently rich set of distributions is chosen, the approximation will be lower than actual maximum regret.

**Asymptotic approximation**

With the knowledge of $\theta_\gamma \equiv E_\gamma [x_i]$ and $\sigma_\gamma^2 \equiv V_\gamma [x_i]$, 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 $\sigma_\gamma^2$ for each feasible value of $\theta_\gamma$. 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 $\bar{X}$:

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

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 $\beta_\gamma \equiv E\left|x_i - \theta_\gamma\right|^3 \in \mathbb{R}$, then bounds on $P_\gamma\left(\bar{X} \leq T\right)$ could be derived from the Berry-Esseen inequality:

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

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.

### 1.5.1. 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 $E\left[y\left(0\right)\right] = .5$ (outcome values refer to individual welfare of clinical outcomes). Suppose

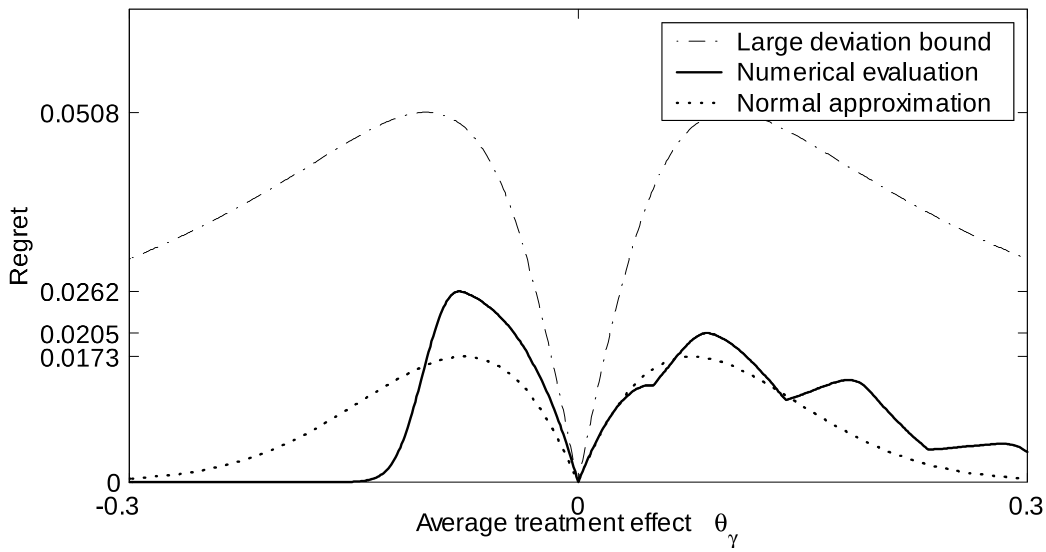


Figure 1.3: Evaluation of maximum regret of the plug-in ($T = .5$) statistical tretment rule.

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 $\Gamma$ 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 treatment rule $\delta_P \equiv 1 \left|\bar{X} > .5\right|$, which assigns the population to new treatment if it outperforms the status quo treatment in the trial by any margin. Figure 1.3 displays the maximum regret of $\delta_P$ for a range of feasible values of the average

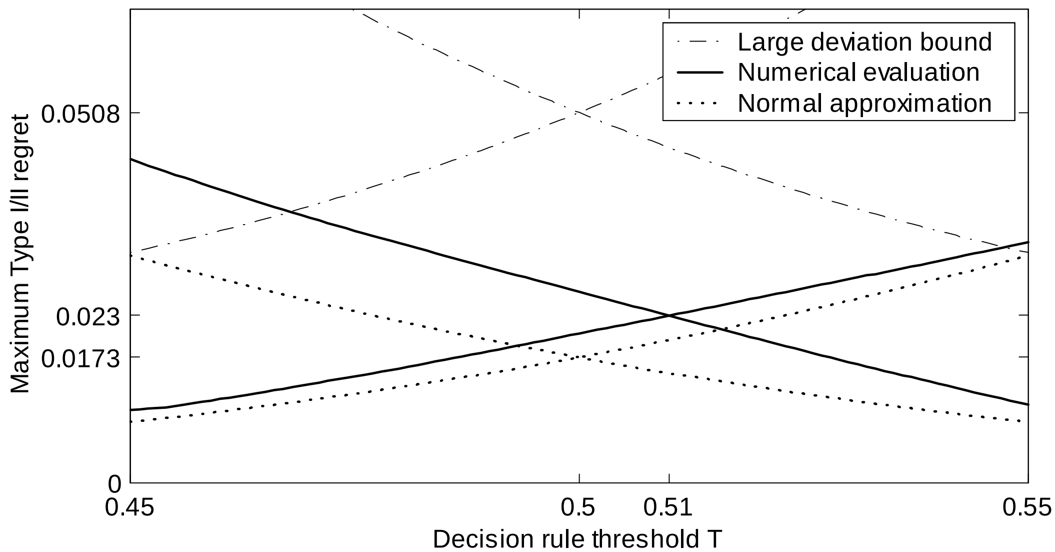


Figure 1.4: Maximum Type I and Type II regret approximations for a range of threshold statistical treatment rules.

treatment effect $\theta_\gamma$. There are multiple feasible outcome distributions with the same $\theta_\gamma$, so the lines represent the maximum aprroximated regret among those distributions. Figure 1.4 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 $\theta$. 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 $\delta_P$ equal to .0173.

The thick solid lines in Figures 1.3 and 1.4 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 $P_\gamma\left(\bar{X} \leq .5\right)$ 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 1.4 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.

## 1.6. Proofs

**Lemma 1.1**

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

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

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

To show that $\max_{h \leq 0} \left\{-h\Phi\left(h - T\right)\right\}$ 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 $h\left(h - T\right) > 1$ and $h - T < 0$

for all $h < H$ and for all $T \in [T_0 - \Delta, T_0 + \Delta]$. Then for such $h$ and $T$ :

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

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

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

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

For any $T < 0$

$$\max_{h\le 0}\left\{-h\Phi\left(h-T\right)\right\} \ge -T\Phi\left(0\right) = -\frac{T}{2}$$

(by substituting $h = T$), thus $\max_{h\le 0}\left\{-h\Phi\left(h-T\right)\right\} \to \infty$ as $T \to -\infty$.

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

$h = -T$ and equals $\frac{1}{4T}$. Then

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

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

**Proposition 1.2**

a) If $T > 0$, then the maximum Type II regret of threshold decision rule $\delta_T$ over the set $\Gamma$ equals

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

The third line uses substitutions $h \equiv \frac{\sqrt{N}}{\bar{\sigma}} \theta$ 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 $\max_{\sigma \in [\underline{\sigma}, \bar{\sigma}]} \Phi \left[ \frac{\bar{\sigma}}{\sigma} (T^* - h) \right] = \Phi \left[ T^* - h \right]$ for $h \ge T^*$. The last equality follows from the condition (1.8).

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 (1.8) 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 $\delta'\left(\bar{X},S^2\right)$ has both lower maximum Type I regret and lower maximum Type II regret than $\delta_T$ over the set $\Gamma$. Since $\delta_T$ 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 $\delta_{T'} \equiv 1\left|\bar{X} > T'\right|$ 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 $\delta_T$ is a solution to the minimax regret, asymmetric minimax regret, or limited type I regret problem over the feasible set $\bar{\Gamma}$. Thus $\delta'$ cannot have both lower maximum Type I and lower maximum Type II regret than $\delta_T$. $\square$

**Lemma 1.3**

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 $p_\gamma$ 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 $D(\delta) \geq 1$. If 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 $p_\gamma \in (0, p_0)$. $\square$

**Proposition 1.4**

It follows from lemma 1.3 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:

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

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 $\sqrt{\frac{N}{p_0(1-p_0)}}\bar{R}_{Type\ II}(\delta_P) \to \max_{h>0}[h\Phi(-h)]$, 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 $\frac{\sqrt{N}}{\sigma_0} R\left(\delta_P, p_0 + \frac{\sigma_0}{\sqrt{N}}h\right)$ uniformly converges to $h\Phi(-h)$ for $h \in \left(0, \frac{3}{2}\sigma_0^{-2}\right]$ as $N \to \infty$. Since the function $h\Phi(-h)$ reaches its maximum at $h \approx .752$ and $\frac{3}{2}\sigma_0^{-2} \geq 6$, $\max_{h>0}[h\Phi(-h)] = \max_{h\in\left(0,\frac{3}{2}\sigma_0^{-2}\right]}[h\Phi(-h)]$. 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 $N > N_1$:

$$\sup_{h \in \left(0, \frac{3}{2}\sigma_0^{-2}\right]} \left| \Phi\left(-h\frac{\sigma_0}{\sigma}\right) - \Phi\left(-h\right) \right| < \frac{\sigma_0^2}{3}\varepsilon, \tag{1.9}$$

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

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

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

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

There exists $N_2$ such that for $N > N_2$, $h \in \left(0, \frac{3}{2}\sigma_0^{-2}\right]$ 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 $h \in \left(0, \frac{3}{2}\sigma_0^{-2}\right]$

$$\left| P_\gamma\left(\frac{\sqrt{N}}{\sigma}\left(\frac{X}{N} - p_\gamma\right) \leq z\right) - \Phi\left(z\right) \right| \leq \frac{B}{\sqrt{N}}. \tag{1.10}$$

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

Since

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

letting $z = -\frac{\sigma_0}{\sigma} h$ in (1.10) yields for $N > N_3$

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

Combining this result wtih (1.9) shows that for $N > \max(N_1, N_3)$

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

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 $h > 0$,

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

The one-sided Chebyshev's inequality shows that

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

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

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

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

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

□

**Proposition 1.5**

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

For all $\gamma$ such that $V_\gamma\left[y\left(1\right)\right] < \frac{\sigma_h^2}{9}$ (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

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

Applying the result to the formula for regret of the plug-in rule $\delta_P$ yields

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

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

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

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

Applying the result to the formula for regret of the plug-in rule $\delta_P$ yields

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

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

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

$$\left|P_\gamma\left(\frac{\sqrt{N}}{\sqrt{V_\gamma\left[y\left(1\right)\right]}}\left(\frac{1}{N}\sum_{i=1}^N x_i - p_\gamma\right) \leq z\right) - \Phi\left(z\right)\right| \leq \frac{E_\gamma\left|y\left(1\right) - p_\gamma\right|^3}{V_\gamma^{3/2}\left[y\left(1\right)\right]\sqrt{N}}. \tag{1.11}$$

Substitute $z = \frac{\sqrt{N}}{\sqrt{V_\gamma[y(1)]}}\left(p_0 - p_\gamma\right)$ into (1.11) and it becomes

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

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

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

thus

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

Applying the result to the regret formula for $\delta_P$ yields

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

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 $\gamma$ with $p_\gamma > 0$. If $V_\gamma [y(1)] < \frac{\sigma_h^2}{9}$, or $V_\gamma [y(1)] \ge \frac{\sigma_h^2}{9}$ and $p_\gamma - p_0 \ge 6\frac{\sigma_h}{\sqrt{N}}$, then

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

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

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

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

CHAPTER 2

# Admissible Treatment Rules for a Risk-Averse Planner with Experimental Data on an Innovation

This chapter was cowritten with Charles F. Manski and originally published in *Journal of Statistical Planning and Inference* © 2007 Elsevier B.V.

## 2.1. Introduction

Problems of choice between a status quo treatment and an innovation occur often in practice. In the medical arena, the status quo may be a standard treatment for a health condition and the innovation may be a new treatment proposed by researchers. Historical experience administering the status quo treatment to populations of patients may have made its properties well understood. In contrast, the properties of the innovation may be uncertain, the only available information deriving from a randomized clinical trial. Then choice between the status quo treatment and the innovation is a statistical decision problem.

This chapter studies the admissibility of treatment rules when the decision maker is a planner (e.g., a physician) who must choose treatments for a population of persons who are observationally identical but who may vary in their response to treatment. We focus on the relatively simple case where treatments have binary outcomes, which we label success and failure. Then the feasible treatment rules are the functions that map the number of experimental successes into a treatment allocation specifying the fraction of the population who receive each treatment.

Section 2.2 formalizes the planner's problem and reviews the case where the objective of the planner is to maximize the population rate of treatment success. In this setting, a theorem of Karlin and Rubin (1956) shows that the admissible rules are ones which assign all members of the population to the status quo treatment if the number of experimental successes is below a specified threshold and all to the innovation if the number of successes is above the threshold. An interior fractional allocation of the population is possible in an admissible rule only when the number of experimental successes exactly equals the threshold. Karlin and Rubin called this class of treatment rules *monotone,* but we will refer to them as *KR-monotone.*

In Section 2.3, we suppose that the objective of the planner is to maximize a concave-monotone function $f(\cdot)$ of the rate of treatment success. We show that this seemingly modest generalization of the welfare function is consequential. Now admissible treatment rules need not be KR-monotone; in fact, KR-monotone rules may be inadmissible. However, a weaker notion of monotonicity remains relevant. Define a *fractional monotone* rule to be one in which the fraction of the population assigned to the innovation weakly increases with the experimental success rate.We show that the class of fractional monotone rules is essentially complete. That is, given any rule which is not fractional monotone, there exists a fractional monotone rule that performs at least as well in all feasible states of nature. If $f(\cdot)$ is concave and strictly monotone, the class of fractional monotone rules is complete. That is, given any rule which is not fractional monotone, there exists a fractional monotone rule that performs at least as well in all feasible states of nature and better in some state of nature.

Further findings emerge when the welfare function has weak curvature. Let $f(\cdot)$ be differentiable with derivative function $g(\cdot)$. Suppose that, for a given positive integer $M$, the quantity $\left[x(1-x)^{-1}\right]^{M} g(x)$ weakly increases with $x$. Define an *M-step monotone* rule to be a fractional monotone rule that gives an interior fractional treatment assignment for no more than $M$ consecutive values of the number of experimental successes. This definition extends the class of KR-monotone rules, which is the special case with $M = 1$. We show that the class of $M$-step monotone rules is essentially complete if the above conditions hold. This class is complete if $f(\cdot)$ is strictly concave or if $\left[x(1-x)^{-1}\right]^{M} g(x)$ is strictly increasing. We also show that the class of KR-monotone rules is minimal complete if $M = 1$ and if $f(\cdot)$ is strictly concave or $\left[x(1-x)^{-1}\right]^{M} g(x)$ is strictly increasing.

Section 2.4 investigates particular treatment rules.We find that Bayes rules and the minimax-regret rule depend on the curvature of the welfare function. These rules are KR-monotone if the curvature is sufficiently weak. However, they deliver interior fractional treatment allocations if the curvature is sufficiently strong. Computation of Bayes rules is typically straightforward. Computation of a minimax-regret rule is simple when this rule is KR-monotone but is more challenging otherwise.

Our consideration of planning problems where welfare is a nonlinear function of the rate of treatment success appears to be new to research studying treatment choice using experimental data. Previous research has examined planning problems in which experimental findings are used to inform treatment choice; see, for example, Canner

(1970), Cheng et al. (2003), and Manski (2004, 2005). However, these and (as far as we are aware) other studies have invariably assumed without comment that welfare is the rate of treatment success.

From a decision theoretic perspective, concave-monotone welfare functions are intriguing because they sometimes yield the conclusion that planners should fractionally allocate observationally identical persons across different treatments. It has been common to presume that a planner should treat observationally identical persons identically. The analysis in this chapter shows that this presumption sometimes is inappropriate when a risk averse planner uses experimental data to inform treatment choice.

From a substantive perspective, consideration of concave-monotone functions of the success rate is interesting because, in expected utility theory, such functions imply distaste for mean-preserving spreads of gambles and thus express risk aversion. Public discourse on health matters, although not entirely coherent, suggests strong risk aversion. This is evident in the ancient admonition of the Hippocratic Oath that a physician should "First, do no harm." It is also evident in the drug approval process of the U.S. Food and Drug Administration, which requires that the manufacturers of pharmaceuticals demonstrate "substantial evidence of effect" for their products (see Gould, 2002). We discuss this matter further in the concluding Section 2.5, which considers the implications of our analysis for treatment choice in practice.

## 2.2. Background

### 2.2.1. The Planning Problem

The basic concepts are as in Manski (2004, 2005). The planner's problem is to choose treatments from a finite set $T$ of mutually exclusive and exhaustive treatments. Each member $j$ of the treatment population, denoted $J$, has a response function $y_j(\cdot) : T \to Y$ mapping treatments $t \in T$ into outcomes $y_j(t) \in Y$. The population is a probability space $(J, \Omega, P)$, and the probability distribution $P[y(\cdot)]$ of the random function $y(\cdot) : T \to Y$ describes treatment response across the population. The population is "large," in the sense that $J$ is uncountable and $P(j) = 0, j \in J$.

In this chapter, outcomes are binary with $y_j(t) = 1$ denoting success and $y_j(t) = 0$ failure if person $j$ receives treatment $t$. There are two treatments, $t = a$ denoting the status quo and $t = b$ the innovation. The population success rates if everyone were to receive the same treatment are $\alpha \equiv P[y(a) = 1]$ and $\beta \equiv P[y(b) = 1]$, respectively. Consider a rule that assigns a fraction $\zeta$ of the population to treatment $b$ and the remaining $1 - \zeta$ to treatment $a$. The population success rate under this fractional rule is

$$\alpha \cdot (1 - \zeta) + \beta \cdot \zeta = \alpha + (\beta - \alpha)\zeta. \tag{2.1}$$

Welfare is $f[\alpha + (\beta - \alpha)\zeta]$, where $f(\cdot)$ is an increasing, concave transformation of the success rate.

The optimal treatment rule is obvious if $(\alpha, \beta)$ are known. The planner should choose $\zeta = 1$ if $\beta > \alpha$ and $\zeta = 0$ if $\beta < \alpha$; all values of $\zeta$ yield the same welfare if $\beta = \alpha$. The problem of interest is treatment choice when $(\alpha, \beta)$ are only partially known.

### 2.2.2. The Empirical Evidence and Admissible Treatment Rules

Suppose that historical experience reveals $\alpha$ but not $\beta$. The available evidence on $\beta$ comes from a randomized experiment, where $N$ subjects are drawn at random and assigned to treatment $b$. Of these subjects, a number $n$ experience outcome $y(b) = 1$ and the remaining $N - n$ experience $y(b) = 0$. The outcomes of all subjects are observed.

In this setting, the sample size $N$ indexes the sampling process and the number $n$ of experimental successes is a sufficient statistic for the sample data. The feasible statistical treatment rules are the functions $z(\cdot) : [0, ..., N] \rightarrow [0, 1]$ that map the number of experimental successes into a treatment allocation. Thus, for each value of $n$, rule $z$ allocates a fraction $z(n)$ of the population to treatment $b$ and the remaining $1 - z(n)$ to treatment $a$.

Following Wald (1950), we evaluate a statistical treatment rule by its expected performance across repeated samples. Let $p(n; \beta) \equiv C(N, n) \cdot \beta^n (1 - \beta)^{N-n}$ denote the binomial probability of $n$ successes in $N$ trials, where $C(N, n) \equiv N! [n! \cdot (N - n)!]^{-1}$. Then the expected welfare yielded by rule $z(\cdot)$ across repeated samples is

$$W(z; \beta) \equiv \sum_{n=0}^{N} p(n; \beta) \cdot f[\alpha + (\beta - \alpha) z(n)]. \tag{2.2}$$

Expected welfare is a function of $\beta$, which is unknown. Let $B$ index the values of that the planner deems feasible. We assume that $0 < \alpha < 1$ and $B$ includes at least one value smaller than $\alpha$ and at least one value greater than $\alpha$. Rule $z'$ (weakly) dominates rule $z$ if $W(z; \beta) \leq W(z'; \beta)$ for all $\beta \in B$ and $W(z; \beta) < W(z'; \beta)$ for some $\beta \in B$. Rule $z$ is admissible if there exists no other rule $z'$ that dominates $z$; if a dominating rule exists, $z$ is inadmissible.

### 2.2.3. Admissible Rules for a Risk-Neutral Planner

Manski (2005, Chapter 3) considers the case in which welfare is the population rate of treatment success; thus, $f(\cdot)$ is the identity function. Then the expected welfare of rule $z$ is

$$W(z;\beta) = \alpha + (\beta - \alpha) \cdot E_\beta [z(n)], \tag{2.3}$$

where $E_\beta [z(n)] = \sum_n p(n;\beta) \cdot z(n)$. Rule $z$ is admissible if there exists no $z'$ such that $(\beta - \alpha) \cdot E_\beta [z(n) - z'(n)] \leq 0$ for all $\beta \in B$ and $(\beta - \alpha) \cdot E_\beta [z(n) - z'(n)] < 0$ for some $\beta \in B$.

A KR-monotone treatment rule, defined in Karlin and Rubin (1956), has the form

$$\begin{aligned} z(n) &= 0 \quad \text{for } n < k, \\ z(n) &= \lambda \quad \text{for } n = k, \\ z(n) &= 1 \quad \text{for } n > k, \end{aligned} \tag{2.4}$$

where $0 \leq k \leq N$ and $0 \leq \lambda \leq 1$. Thus, a KR-monotone rule allocates all persons to treatment $a$ if $n$ is smaller than the specified threshold $k$, a fraction $\lambda$ to treatment $b$ if $n = k$, and all to treatment $b$ if $n$ is larger than $k$.

Manski (2005, Proposition 3.1) applies Karlin and Rubin (1956, Theorem 4) to show that the class of KR-monotone rules is minimal complete if $B$ excludes the extreme values 0 and 1. That is, the admissible rules and the KR-monotone rules coincide.

Everywhere fractional treatment rules are inadmissible even when the only empirical evidence about the innovation is the outcome of a single experiment. Let $B$ exclude the extreme values $\{0, 1\}$ and consider a sample of size 1. If $N = 1$, there are two possible values for the threshold $k$ and, hence, two types of KR-monotone rules. Setting $k = 0$ yields rules in which $z(0)$ can take any value and $z(1) = 1$. Setting $k = 1$ yields rules in which $z(0) = 0$ and $z(1)$ can take any value.

### 2.3. Admissible Treatment Rules for Risk-Averse Planners

Determination of the admissible treatment rules when the function $f(\cdot)$ is nontrivially concave is a challenging problem. However, there are ways to make progress. This section presents findings that shed some light on the matter. Section 2.3.1 shows that the class of fractional monotone rules is essentially complete for all concave-monotone $f(\cdot)$. This class is complete if $f(\cdot)$ is concave and strictly monotone.

Section 2.3.2 shows that the class of $M$-step monotone rules is essentially complete for all differentiable concave-monotone $f(\cdot)$ such that $\left[x(1-x)^{-1}\right]^M g(x)$ is weakly increasing in $x$. This class is complete if $f(\cdot)$ is also strictly concave or $\left[x(1-x)^{-1}\right]^M g(x)$ is strictly increasing in $x$. Section 2.3.3 shows that the class of KR-monotone rules is minimal complete if $M = 1$ and $f(\cdot)$ is also strictly concave or if $\left[x(1-x)^{-1}\right]^M g(x)$ is strictly increasing in $x$. However, we show that KR-monotone rules can be inadmissible if $f(\cdot)$ is sufficiently curved.

### 2.3.1. The Fractional Monotone Rules are an Essentially Complete Class

The binomial density function possesses the strict form of the monotone-likelihood ratio property: $(n > n', \beta > \beta') \Rightarrow p(n;\beta)/p(n;\beta') > p(n';\beta)/p(n';\beta')$. Thus, larger values of $n$ are unambiguously evidence for larger values of $\beta$. It is therefore reasonable to conjecture that good treatment rules are the ones that make the fraction of the population allocated to treatment $b$ increase with $n$.

The results of Karlin and Rubin (1956) show that a strong form of this conjecture is correct if $f(\cdot)$ is linear in the population success rate. The Karlin and Rubin theorems do not apply to nonlinear $f(\cdot)$. Nevertheless, the conjecture remains correct in the weaker sense that the class of fractional monotone treatment rules is essentially complete for all concave-monotone welfare functions and complete when $f(\cdot)$ is concave and strictly monotone. Formally, we say that a treatment rule $z$ is fractional monotone if $n < n' \Rightarrow z(n) \leq z(n')$. Proposition 2.1 proves the result.

**Proposition 2.1.** *If $f(\cdot)$ is weakly increasing and concave, the class of fractional monotone rules is essentially complete. If $f(\cdot)$ is also strictly increasing, the class of fractional monotone rules is complete.*

**Proof.** Suppose that $z$ is not fractional monotone, so $z(n) < z(n')$ for some $n > n'$. Consider replacing $z$ with the following treatment rule $z^*$:

$$\begin{aligned} z^*(n) &\equiv z^*(n') \equiv \frac{p(n;\alpha)}{p(n;\alpha)+p(n';\alpha)} z(n) + \frac{p(n';\alpha)}{p(n;\alpha)+p(n';\alpha)} z(n'), \\ z^*(m) &\equiv z(m), \ \forall m \notin \{n, n'\}. \end{aligned}$$

For any value of $\beta$,

$$W(z^*;\beta) - W(z;\beta) = p(n;\beta)\cdot\{f[\alpha+(\beta-\alpha)z^*(n)] - f[\alpha+(\beta-\alpha)z(n)]\}$$
$$+p(n';\beta)\cdot\{f[\alpha+(\beta-\alpha)z^*(n')] - f[\alpha+(\beta-\alpha)z(n')]\}. \tag{2.5}$$

The function $f(\cdot)$ is concave and $z^*(n)$ is a convex combination of $z(n)$ and $z(n')$. Hence,

$$f[\alpha+(\beta-\alpha)z^*(n)] \geq$$
$$\frac{p(n;\alpha)}{p(n;\alpha)+p(n';\alpha)} f[\alpha+(\beta-\alpha)z(n)] + \frac{p(n';\alpha)}{p(n;\alpha)+p(n';\alpha)} f[\alpha+(\beta-\alpha)z(n')].$$

The same inequality holds for $f[\alpha+(\beta-\alpha)z^*(n')]$. Substituting these inequalities into (2.5) and rearranging terms yields

$$W(z^*;\beta) - W(z;\beta) \geq$$
$$\frac{p(n;\beta)\cdot p(n';\alpha) - p(n';\beta)\cdot p(n;\alpha)}{p(n;\alpha)+p(n';\alpha)}\cdot\{f[\alpha+(\beta-\alpha)z(n')] - f[\alpha+(\beta-\alpha)z(n)]\}.$$

The following inequalities use the monotone-likelihood ratio property and the fact that $z(n) < z(n')$:

$$\beta < \alpha \Rightarrow \quad p(n;\beta)\cdot p(n';\alpha) - p(n';\beta)\cdot p(n;\alpha) < 0,$$
$$f[\alpha+(\beta-\alpha)z(n')] - f[\alpha+(\beta-\alpha)z(n)] \leq 0,$$
$$\beta > \alpha \Rightarrow \quad p(n;\beta)\cdot p(n';\alpha) - p(n';\beta)\cdot p(n;\alpha) > 0,$$
$$f[\alpha+(\beta-\alpha)z(n')] - f[\alpha+(\beta-\alpha)z(n)] \geq 0.$$

It follows that $W(z^*;\beta) \geq W(z;\beta)$ for all $\beta \in B$. If $f(\cdot)$ is strictly increasing, the right-hand side inequalities are strict and $W(z^*;\beta) > W(z;\beta)$ for all $\beta \in B\backslash\{\alpha\}$.

Given any rule $z$ that is not fractional monotone, we can iteratively apply the transformation described above to all pairs $(n', n)$ for which $z(n') > z(n)$, in the following order: $(n', n) = (1,2), (1,3), ..., (1,N), (2,3), (2,4), ..., (N-1,N)$. The result is a fractional monotone treatment rule that performs at least as well as $z$ for all values of $\beta$ and that dominates $z$ if $f(\cdot)$ is strictly increasing. □

Proposition 2.1 implies that a risk-neutral or risk-averse planner can restrict attention to fractional monotone treatment rules; there is no reason to contemplate other rules. The proposition does not imply that all fractional monotone rules are worthy of consideration. Indeed, we already know that a risk-neutral planner can restrict attention to rules that are KR-monotone.

### 2.3.2. M-Step Monotone Rules

It appears that no result stronger than Proposition 2.1 can be proved without placing restrictions on the shape of $f(\cdot)$ beyond monotonicity and concavity. This section shows that Proposition 2.1 can be strengthened considerably if $f(\cdot)$ is restricted to be differentiable with derivative function $g(\cdot)$ that does not decrease too rapidly.

We define a treatment rule to be $M$-step monotone if $n < n' \Rightarrow z(n) \leq z(n')$ and, for a given positive integer $M$, $n + M \leq n' \Rightarrow z(n) = 0$ or $z(n') = 1$. Suppose that $\left[x(1-x)^{-1}\right]^M g(x)$ weakly increases with $x$. With this restriction on the curvature of $f(\cdot)$, the class of $M$-step monotone rules is essentially complete, whatever the sample size $N$ may be. Moreover, this class is complete if $f(\cdot)$ is strictly concave or if $\left[x\left(1-x\right)^{-1}\right]^{M} g\left(x\right)$ strictly increases with $x$. Proposition 2.2 proves the result.

**Proposition 2.2.** *Let $f(\cdot)$ be weakly increasing, concave and differentiable on $(\inf\left\{B\right\},\sup\left\{B\right\})$, with derivative $g(\cdot)$. If $\left[x\left(1-x\right)^{-1}\right]^{M} g\left(x\right)$ weakly increases with $x$, then the $M$-step monotone rules are an essentially complete class. If $f(\cdot)$ is also strictly concave or if $\left[x\left(1-x\right)^{-1}\right]^{M} g\left(x\right)$ strictly increases with $x$, then the $M$-step monotone rules are a complete class.*

**Proof.** Proposition 2.1 showed that the class of fractional monotone treatment rules is essentially complete. Suppose that $z$ is fractional monotone but not $M$-step monotone. We will iteratively construct an $M$-step monotone rule that performs at least as well as $z$. Part A of the proof describes the content of each step of the iteration. Part B gives the iteration. Parts A and B show that the class of $M$-step monotone rules is essentially complete. Part C of the proof shows that this class is complete if $f(\cdot)$ is strictly concave or if $\left[x\left(1-x\right)^{-1}\right]^{M} g\left(x\right)$ strictly increases with $x$.

(A) Suppose that $z$ is not $M$-step monotone, with $z(n) > 0$ and $z(n') < 1$ for some $(n, n')$ such that $n + M \leq n'$. We will compare $z$ to an alternative treatment rule $z'$ in which $z(n)$ and $z(n')$ are replaced by

$$\begin{aligned} z'(n) &= z(n) - p(n';\alpha)\cdot p^{-1}(n;\alpha)\cdot\left[z'(n') - z(n')\right], \\ z'\left(n'\right) &= \min\left[1, z\left(n'\right) + p^{-1}\left(n';\alpha\right)\cdot p\left(n;\alpha\right)\cdot z\left(n\right)\right]. \end{aligned} \tag{2.6}$$

Observe that rule $z'$ has either $z'(n) = 0$ or $z'(n') = 1$. We will show that $z'$ performs at least as well as $z$. It dominates $z$ if $f(\cdot)$ is strictly concave or if $\left[x\left(1-x\right)^{-1}\right]^{M} g\left(x\right)$ strictly increases with $x$.

To show this requires two preliminary steps. First, weak concavity of $f(\cdot)$ implies that

$$\begin{aligned} 0 &\leq x < y \leq 1, \beta \neq \alpha \Rightarrow \\ &f\left[\alpha + \left(\beta - \alpha\right) x\right] - f\left[\alpha + \left(\beta - \alpha\right) y\right] \geq g\left(\alpha\right) \cdot \left(\beta - \alpha\right) \cdot \left(x - y\right), \\ 0 &\leq y < x \leq 1, \beta \neq \alpha \Rightarrow \\ &f\left[\alpha + \left(\beta - \alpha\right) x\right] - f\left[\alpha + \left(\beta - \alpha\right) y\right] \geq g\left(\beta\right) \cdot \left(\beta - \alpha\right) \cdot \left(x - y\right). \end{aligned} \tag{2.7}$$

These inequalities are strict if $f(\cdot)$ is strictly concave.

Second, $p(n'; x)p^{-1}(n; x) \cdot g(x)$ weakly increases with $x$ for $n + M \leq n'$. This holds because

(2.8)

$$p(n'; x)p^{-1}(n; x) \cdot g(x) = \left[C\left(N, n'\right) / C\left(N, n\right)\right] \cdot \left[x\left(1-x\right)^{-1}\right]^{n'-n-M} \cdot \left[x\left(1-x\right)^{-1}\right]^{M} g\left(x\right).$$

The first term on the right-hand side is a positive constant. The second term is a positive and weakly increasing function on $(0, 1)$. The last term is positive and weakly increasing by assumption. If $\left[x\left(1-x\right)^{-1}\right]^{M} g\left(x\right)$ is strictly increasing, then so is $p(n'; x)p^{-1}(n; x) \cdot g(x)$.

Now consider the difference in welfare between rules $z'$ and $z$. All rules yield the same welfare if $\beta = \alpha$. For $\beta \neq \alpha$,

$$\begin{aligned}
W(z';\beta) - W(z;\beta) &= p(n';\beta) \cdot \{f[\alpha + (\beta-\alpha) z'(n')] - f[\alpha + (\beta-\alpha) z(n')]\} \\
&\quad + p(n;\beta) \cdot \{f[\alpha + (\beta-\alpha) z'(n)] - f[\alpha + (\beta-\alpha) z(n)]\} \\
&\geq p(n';\beta) \cdot g(\beta) \cdot (\beta-\alpha) \cdot [z'(n') - z(n')] \\
&\quad + p(n;\beta) \cdot g(\alpha) \cdot (\beta-\alpha) \cdot [z'(n) - z(n)] \\
&= (\beta-\alpha) \cdot \left[p(n';\beta)p^{-1}(n;\beta) \cdot g(\beta) - p(n';\alpha)p^{-1}(n;\alpha) \cdot g(\alpha)\right] \\
&\quad \cdot p(n;\beta) \cdot [z'(n') - z(n')] \\
&\geq 0 \qquad (2.9)
\end{aligned}$$

The first inequality follows from (2.7). The second equality follows from (2.6). The last inequality holds for all $\beta$ because the first two terms have the same sign when they do not equal zero and the last two terms are strictly positive. If $f(\cdot)$ is strictly concave or if $\left[x(1-x)^{-1}\right]^M g(x)$ is strictly increasing, then $W(z';\beta) - W(z;\beta) > 0$ for all $\beta \neq \alpha$.

(B) We iteratively apply the transformation described above to all pairs $(n, n')$ for which $n + M \leq n'$, $z(n) > 0$ and $z(n') < 1$, in this order:
$(n, n') = (0, N), (0, N-1), ..., (0, M), (1, N), (1, N-1), ..., (1, 1+M), ..., (N-M, N)$.
We show that, performed in this order, each iteration preserves fractional monotonicity of the treatment rule and that the outcome is an $M$-step rule.

Consider the first iteration, which has $n = 0$ and $n' = N$. If $z(0) > 0$ and $z(N) < 1$, the iteration either reduces $z(0)$ to zero or increases $z(N)$ to one. Both results preserve fractional monotonicity.

Next let $n = 0$ and $n' < N$. The preceding iteration has considered the pair $(0, n' + 1)$; thus, $z(0) > 0$ implies that $z(n' + 1) = 1$. The current iteration either reduces $z(0)$ to zero or increases $z(n')$ to one. Again, both results preserve fractional monotonicity.

After completion of all iterations with $n = 0$, we either have that $z(0) > 0$ or $z(0) = 0$. If $z(0) > 0$, then $z(n') = 1$ for all $n' \geq M$. Hence, an $M$-step rule has been achieved and no further iteration is necessary.

If $z(0) = 0$, we perform the iterations for $n = 1$. As in the first round, these iterations preserve fractional monotonicity and deliver an $M$-step rule if $z(1) > 0$ at their completion. If $z(1) = 0$, we perform the iterations for $n = 2, 3, ...$, continuing through further rounds of iteration until an $M$-step rule is achieved.

The ultimate result of the iterative process is an $M$-step monotone rule that performs at least as well as $z$ for all values of $\beta$. Hence, the class of $M$-step monotone rules is essentially complete.

(C) If $f(\cdot)$ is strictly concave or $\left[x\left(1-x\right)^{-1}\right]^{M} g\left(x\right)$ is strictly increasing, the treatment rule obtained through the iterated modification dominates the original rule $z$. Thus, the $M$-step monotone rules form a complete class. □

### 2.3.3. KR-Monotone Rules Revisited

The KR-monotone rules are the $M$-step monotone rules with $M = 1$. Thus, Proposition 2.2 shows that the class of KR-monotone rules is essentially complete if $\left[x\left(1-x\right)^{-1}\right]g\left(x\right)$ weakly increases with $x$. This class is complete if $f(\cdot)$ is strictly concave or if $\left[x\left(1-x\right)^{-1}\right]g\left(x\right)$ strictly increases with $x$. For example, Proposition 2.2 shows that the class of KR-monotone rules is complete if $f(x) = \log(x)$. A welfare function that just barely satisfies the conditions of Proposition 2.2 is $\log(x) - x$, whose derivative is $x^{-1}(1-x)$.

This section develops further properties of the KR-monotone rules. To begin, Proposition 2.3 shows that if $f(\cdot)$ is strictly increasing and the class of KR-monotone rules is complete, then this class is minimal complete.

**Proposition 2.3.** *Let $f(\cdot)$ be strictly increasing. If the class of KR-monotone rules is essentially complete, then every KR-monotone rule is admissible. If the class of KR-monotone rules is complete, then it is minimal complete.*

**Proof.** Suppose that $z$ is an inadmissible KR-monotone rule. Then there exists a rule $z'$ that dominates $z$. By assumption, the KR-monotone rules are an essentially complete class. So there exists a KR-monotone rule $z' \neq z$ that dominates $z$. However, one of the following two conditions must hold if $f(\cdot)$ is strictly increasing:

$$\begin{aligned}
\forall n: z(n) \geq z'(n), \text{ with strict inequality for some } n &\Rightarrow \\
W(z;\beta) &> W(z';\beta), \beta \in (\alpha, 1), \\
\forall n: z(n) \leq z'(n), \text{ with strict inequality for some } n &\Rightarrow \\
W(z;\beta) &> W(z';\beta), \beta \in (0, \alpha).
\end{aligned} \tag{2.10}$$

Therefore, $z'$ cannot dominate $z$. Thus, all KR-monotone rules are admissible.

If the class of KR-monotone rules is complete, there exist no admissible rules outside of this class. Hence, the class of KR-monotone rules is minimal complete. □

Combining Propositions 2.2 and 2.3 shows that Theorem 4 of Karlin and Rubin (1956) extends to welfare functions that are concave-monotone with sufficiently weak curvature. We state this result as:

**Proposition 2.3, Corollary.** *Let $f(\cdot)$ be strictly increasing and differentiable on $(\inf\{B\}, \sup\{B\})$, with $\left[x(1-x)^{-1}\right] g(x)$ weakly increasing in $x$. Then the class of KR-monotone rules is minimal complete if $f(\cdot)$ is strictly concave or if $\left[x(1-x)^{-1}\right] g(x)$ is strictly increasing.*

The corollary shows that all KR-monotone rules are admissible when the welfare function has sufficiently weak curvature. However, we can show that some KR-monotone rules are inadmissible when $f(\cdot)$ has sufficiently strong curvature. Again let $f(\cdot)$ be strictly increasing and differentiable, with $g(\cdot)$ denoting the derivative function. Let the space $B$ contain only two values, one lower than $\alpha$ and the other higher; thus, $B = \{\beta_L, \beta_H\}$, where $\beta_L < \alpha < \beta_H$. For a specified $k$ with $0 \leq k \leq N$ and a specified pair $(v, w)$ with $0 \leq v \leq w \leq 1$, define the treatment rule

$$\begin{aligned} z_{vw}(n) &= v \text{ for } n \leq k, \\ z_{vw}(n) &= w \text{ for } n > k. \end{aligned} \tag{2.11}$$

A special case is the KR-monotone rule $z_{01}$.

Proposition 2.4 compares rule $z_{01}$ with a non-extreme fractional rule $z_{vw}$; that is, one with $0 < v \leq w < 1$. We find that rule $z_{01}$ strictly dominates $z_{vw}$ if the derivative function $g(\cdot)$ decreases sufficiently slowly and vice versa if $g(\cdot)$ decreases sufficiently rapidly.

**Proposition 2.4.** *Let $f(\cdot)$ be strictly increasing and differentiable. Fix $k$. Let $d_L \equiv \sum_{n>k} p(n; \beta_L)$ and $d_H \equiv \sum_{n>k} p(n; \beta_H)$. Let $0 < v \leq w < 1$. Rule $z_{01}$ strictly dominates $z_{vw}$ if*

$$g(\alpha)/g(\beta_L) > [d_L/(1-d_L)] \cdot [(1-w)/v], \tag{2.12}$$

$$g(\alpha)/g(\beta_H) < [d_H/(1-d_H)] \cdot [(1-w)/v]. \tag{2.13}$$

*Rule $z_{vw}$ strictly dominates $z_{01}$ if*

$$g[(1-v)\alpha + v\beta_L]/g[(1-w)\alpha + w\beta_L] < [d_L/(1-d_L)] \cdot [(1-w)/v], \tag{2.14}$$

$$g[(1-v)\alpha + v\beta_H]/g[(1-w)\alpha + w\beta_H] > [d_H/(1-d_H)] \cdot [(1-w)/v]. \tag{2.15}$$

**Proof.** By (2.2), the expected welfare of rules $z_{01}$ and $z_{vw}$ in the two feasible states of nature are as follows:

$$W(z_{01};\beta_L) = (1-d_L)\, f(\alpha) + d_L f(\beta_L), \tag{2.16}$$

$$W(z_{01};\beta_H) = (1-d_H)\, f(\alpha) + d_H f(\beta_H), \tag{2.17}$$

$$W(z_{vw};\beta_L) = (1-d_L)\, f((1-v)\,\alpha + v\beta_L) + d_L f((1-w)\,\alpha + w\beta_L), \tag{2.18}$$

$$W(z_{vw};\beta_H) = (1-d_H)\, f((1-v)\,\alpha + v\beta_H) + d_H f((1-w)\,\alpha + w\beta_H). \tag{2.19}$$

Rule $z_{01}$ strictly dominates $z_{vw}$ if $W(z_{01};\beta_L) > W(z_{vw};\beta_L)$ and
$W(z_{01};\beta_H) > W(z_{vw};\beta_H)$. Rule $z_{vw}$ strictly dominates if these inequalities are reversed.

Ceteris paribus, the direction of the inequalities depends on the curvature of $f(\cdot)$. By assumption, $f(\cdot)$ is concave and strictly increasing. Hence, its derivative $g(\cdot)$ is weakly decreasing and everywhere positive. Use the mean-value theorem to rewrite (2.18) and (2.19) as

$$\begin{aligned} W(z_{vw};\beta_L) &= (1-d_L)\, f(\alpha) + d_L f(\beta_L) \\ &\quad + (1-d_L)\,(\beta_L-\alpha)\, g\left[(1-v_L)\,\alpha + v_L\beta_L\right] v \\ &\quad + d_L\,(\beta_L-\alpha)\, g\left[(1-w_L)\,\alpha + w_L\beta_L\right](w-1), \end{aligned} \tag{2.20}$$

$$\begin{aligned} W(z_{vw};\beta_H) &= (1-d_H)\, f(\alpha) + d_H f(\beta_H) \\ &\quad + (1-d_H)\,(\beta_H-\alpha)\, g\left[(1-v_H)\,\alpha + v_H\beta_H\right] v \\ &\quad + d_H\,(\beta_H-\alpha)\, g\left[(1-w_H)\,\alpha + w_H\beta_H\right](w-1), \end{aligned} \tag{2.21}$$

where $v_L \in [0, v], w_L \in [w, 1], v_H \in [0, v]$, and $w_H \in [w, 1]$. Recall that $\beta_L < \alpha < \beta_H$. Comparison of (2.16) and (2.17) with (2.20) and (2.21) shows that rule $z_{01}$ strictly dominates $z_{vw}$ if and only if

$$(2.22) \qquad (1 - d_L)\, g\left[(1 - v_L)\, \alpha + v_L \beta_L\right] v + d_L g\left[(1 - w_L)\, \alpha + w_L \beta_L\right] (w - 1) > 0$$

and

$$(2.23) \qquad (1 - d_H)\, g\left[(1 - v_H)\, \alpha + v_H \beta_H\right] v + d_H g\left[(1 - w_H)\, \alpha + w_H \beta_H\right] (w - 1) < 0.$$

Rule $z_{vw}$ strictly dominates $z_{01}$ if and only if these inequalities are reversed.

Whether (2.22)-(2.23) hold, or the reverse inequalities, depends on how rapidly the derivative function $g(\cdot)$ decreases with its argument. Direct analysis of the inequalities is complicated by the fact that the intermediate values $(v_L, w_L, v_H, w_H)$ used in the mean-value theorem are themselves determined by $g(\cdot)$. However, the fact that $g(\cdot)$ is a decreasing function implies that simpler sufficient conditions for dominance can be obtained by letting the intermediate values vary over their feasible ranges. Inequalities (2.12)-(2.13) are the sufficient condition for rule $z_{01}$ to strictly dominate $z_{vw}$ and inequalities (2.14)-(2.15) are the sufficient condition for $z_{vw}$ to strictly dominate $z_{01}$. □

### 2.4. Bayes and Minimax-Regret Rules

To learn more about how the shape of the welfare function affects treatment choice, we next study the behavior of Bayes rules and the minimax-regret rule. Sections 2.4.1 and 2.4.2 present some analytical findings for Bayes rules and the minimax-regret rule, respectively. Section 2.5 will report some numerical findings for the minimax-regret and other rules.

### 2.4.1. Bayesian Planning

A Bayesian planner places a prior subjective probability distribution, say $\Pi$, on the set $B$. Observing the number $n$ of experimental successes in the randomized trial, he forms a posterior distribution, say $\Pi(\beta|n)$. Treating $\beta$ as a random variable with distribution $\Pi(\beta|n)$, the planner then solves the problem

$$\max_{\zeta\in[0,1]} \int f\left[\alpha+(\beta-\alpha)\zeta\right] d\Pi\left(\beta|n\right). \tag{2.24}$$

Proposition 2.5 shows that, given a regularity condition, the Bayes rule assigns the entire population to treatment $a(\zeta=0)$ if the posterior mean of $\beta$ does not exceed $\alpha$ and assigns a positive fraction to treatment $b(\zeta>0)$ otherwise. The proposition also gives a sufficient condition for the Bayes rule to be fractional $(0<\zeta<1)$.

**Proposition 2.5.** *Consider problem (2.24). Let $\Pi(\beta|n)$ be non-degenerate. Let $E_{\Pi(\beta|n)}[\beta]$ denote the posterior mean of $\beta$.*

*a) Let $f(\cdot)$ be strictly concave. Then the Bayes rule is unique, therefore admissible. The solution is $\zeta=0$ if $E_{\Pi(\beta|n)}[\beta]\leq\alpha$.*

*b) Let $f(\cdot)$ be continuously differentiable. Let $f(\cdot)$ and $\Pi(\beta|n)$ be sufficiently regular that*

$$\frac{\partial}{\partial\zeta}\left\{\int f\left[\alpha+(\beta-\alpha)\zeta\right]d\Pi\left(\beta|n\right)\right\}=\int\left\{\frac{\partial}{\partial\zeta}f\left[\alpha+(\beta-\alpha)\zeta\right]\right\}d\Pi\left(\beta|n\right)$$

*in a neighborhood of* $\zeta = 0$*. Then all solutions satisfy* $\zeta > 0$ *if* $E_{\Pi(\beta|n)}[\beta] > \alpha$*. All solutions satisfy* $\zeta \in (0, 1)$ *if* $E_{\Pi(\beta|n)}[\beta] > \alpha$ *and* $\int f(\beta)\, d\Pi(\beta|n) < f(\alpha)$*.*

**Proof.** (a) Strict concavity of $f(\cdot)$ implies that $\int f[\alpha + (\beta - \alpha)\zeta]\, d\Pi(\beta|n)$ is strictly concave in $\zeta$. Hence, problem (2.24) has a unique solution. If $\zeta = 0$, then $\int f[\alpha + (\beta - \alpha)\zeta]\, d\Pi(\beta|n) = f(\alpha)$. For each $\zeta > 0$, $f[\alpha + (\beta - \alpha)\zeta]$ is strictly concave as a function of $\beta$. Hence, $\int f[\alpha + (\beta - \alpha)\zeta]\, d\Pi(\beta|n) < f\left[\alpha + \left(E_{\Pi(\beta|n)}[\beta] - \alpha\right)\zeta\right]$. Hence, $E_{\Pi(\beta|n)}[\beta] \leq \alpha \Rightarrow \int f[\alpha + (\beta - \alpha)\zeta]\, d\Pi(\beta|n) < f(\alpha)$.

(b) $\frac{\partial}{\partial\zeta} f[\alpha + (\beta - \alpha)\zeta]\Big|_{\zeta=0} = (\beta - \alpha) \cdot \frac{\partial}{\partial x} f(x)\Big|_{x=\alpha}$. Hence, $\frac{\partial}{\partial\zeta}\left\{\int f[\alpha + (\beta - \alpha)\zeta]\, d\Pi(\beta|n)\right\}\Big|_{\zeta=0} = \left(E_{\Pi(\beta|n)}[\beta] - \alpha\right) \cdot \frac{\partial}{\partial x} f(x)\Big|_{x=\alpha}$. By assumption, $E_{\Pi(\beta|n)}[\beta] > \alpha$ and $\frac{\partial}{\partial x} f(x)\Big|_{x=\alpha} > 0$. Hence, $\int f[\alpha + (\beta - \alpha)\zeta]\, d\Pi(\beta|n)$ strictly increases with $\zeta$ in a neighborhood of $\zeta = 0$, implying that solutions to (2.24) are positive. If $\int f(\beta)\, d\Pi(\beta|n) < f(\alpha)$, then $\zeta = 1$ does not solve (2.24). Hence, solutions are fractional. □

Observe that the concavity and differentiability restrictions placed on $f(\cdot)$ are used in different parts of the proposition. The proof of part (a) only uses the assumption that $f(\cdot)$ is strictly concave. The proof of part (b) only uses the assumption that $f(\cdot)$ is continuously differentiable and the stated regularity condition.

### 2.4.2. Minimax-Regret Planning

The minimax-regret criterion for treatment choice uses no prior information beyond the planner's knowledge that $\beta$ lies in the set $B$. Let $Z$ denote the space of all functions that map $[0, ..., N] \to [0, 1]$. For each $\beta \in B$, $\max[f(\alpha), f(\beta)]$ is the maximum welfare achievable given knowledge of $\beta$, $W(z; \beta)$ is the expected welfare achieved by rule $z(\cdot)$, and the difference between these quantities is regret $R(z;\beta)$:

$$
\begin{aligned}
R(z;\beta) &\equiv \max\left[f(\alpha), f(\beta)\right] - W(z;\beta) \\
&= \sum_{n=0}^{N} p(n;\beta)\cdot\left\{\max\left[f(\alpha), f(\beta)\right] - f\left[\alpha + (\beta-\alpha)\cdot z(n)\right]\right\}.
\end{aligned} \tag{2.25}
$$

A minimax-regret rule $z_{\text{mmr}}$ solves the problem

$$\inf_{z\in Z}\sup_{\beta\in B} R(z;\beta). \tag{2.26}$$

(Another criterion that uses no information beyond knowledge of $B$ is the maximin rule. We do not consider it because it is ultra-conservative, entirely ignoring the sample data. If $B$ contains any value smaller than $\alpha$, the maximin rule assigns the entire population to the status quo, whatever the sample size may be.)

Stoye (2007b) has shown that when $f(\cdot)$ is linear and $B = [0,1]$, there exists an easily computable KR-monotone minimax-regret rule that satisfies the condition

$$\max_{\beta<\alpha}\left[(\alpha-\beta)\sum_{n=0}^{N} p(n;\beta)\, z_{\text{mmr}}(n)\right] = \max_{\beta>\alpha}\left[(\beta-\alpha)\sum_{n=0}^{N} p(n;\beta)\left(1 - z_{\text{mmr}}(n)\right)\right].$$

We show here that minimax-regret rules have a similar characterization for nonlinear $f(\cdot)$ if the class of KR-monotone rules is essentially complete. To simplify exposition, let $f(\cdot)$ be strictly increasing and continuous on $(0,1)$. Let $B$ be a closed subset of $(0,1)$.

Each KR-monotone rule is defined by two numbers: the integer $k$ specifying the location of the step and the fraction $\lambda$ specifying the fraction of the population assigned to treatment $b$ when there are $k$ experimental successes. The sum $k + (1-\lambda) = \sum_{n=0\ldots N}\left[1 - z(n)\right]$ uniquely indexes each KR-monotone rule. That is, there is a one-to-one correspondence between the set of all KR-monotone rules and the interval $[0, N]$ through this index.

For each sample outcome, the proportion of population assigned to treatment $b$ and the regret in each state of nature change monotonically with the value of the index. If $z'$ and $z$ are KR-monotone treatment rules and $\sum_{n=0\ldots N}[1 - z'(n)] < \sum_{n=0\ldots N}[1 - z(n)]$, then $z'(n) \geq z(n)$ for all $n$, with strict inequality for one value of $n$. Moreover, $\beta < \alpha \Rightarrow R(z'; \beta) > R(z; \beta)$ and $\beta > \alpha \Rightarrow R(z'; \beta) < R(z; \beta)$.

The quantity $\max_{\beta \in B \cap (0,\alpha]} R(\cdot; \beta)$ is a strictly decreasing and continuous function of the index $\sum_{n=0\ldots N}[1 - z(n)]$, while $\max_{\beta \in B \cap [\alpha,1)} R(\cdot; \beta)$ is strictly increasing and continuous. Hence, there is a unique rule that minimizes maximum regret among KR-monotone rules. It satisfies the condition

$$\max_{\beta \in B \cap (0,\alpha]} R(z_{\text{mmr}}; \beta) = \max_{\beta \in B \cap [\alpha,1)} R(z_{\text{mmr}}; \beta).$$

If the class of KR-monotone rules is essentially complete, this treatment rule solves problem (2.26). The same results hold for KR-monotone rules if $B = (0, 1)$ and $f(\cdot)$ satisfies the conditions of Proposition 4.2 for $M = 1$.

The situation is different if $f(\cdot)$ has strong curvature and $B$ contains positive values arbitrarily close to zero. Then a minimax-regret rule never assigns the entire population to treatment $b$. Proposition 2.6 gives the result.

**Proposition 2.6.** *Let $\alpha > 0$. Let $B$ contain a sequence of positive values that converges to zero. If*

$$\lim_{\beta \to 0_+} \beta^M f(\beta) = -\infty \tag{2.27}$$

*for some $M \geq 0$, then $z_{\text{mmr}}(n) < 1$ for all $n \leq M$ regardless of sample size $N$.*

**Proof.** Let $z_0$ denote the treatment rule that always assigns everyone to treatment $a$; thus, $z_0(n) = 0$ for all values of $n$. This rule has finite maximum regret $\sup_{\beta > \alpha} f(\beta) - f(\alpha) \leq f(1) - f(\alpha)$. Hence, any treatment rule with infinite maximum regret cannot be minimax regret. Suppose $z(n) = 1$, then

$$
\begin{aligned}
\sup_{\beta \in B} R(z;\beta) &\geq \lim_{\beta \to 0_+} R(z;\beta) \geq \lim_{\beta \to 0_+} \left(p(n;\beta) \cdot \{f(\alpha) - f[\alpha + (\beta - \alpha) z(n)]\}\right) \\
&= C(N,n) \cdot \lim_{\beta \to 0_+} \left(\beta^n (1-\beta)^{N-n} \cdot [f(\alpha) - f(\beta)]\right).
\end{aligned}
$$

This quantity is infinite because $\lim_{\beta \to 0} \left[\beta^n (1-\beta)^{N-n}\right] \cdot f(\alpha)$ is finite, $\lim_{\beta \to 0} (1-\beta)^{N-n} = 1$ and $\lim_{\beta \to 0} [\beta^n \cdot f(\beta)] = -\infty$ follows from (2.27). Hence, the minimax-regret rule must have $z_{\text{mmr}}(n) < 1$. □

To illustrate, consider the welfare function $f(x) = -x^{-K}$, where $K > 1$. Then (2.27) holds for $M < K$ and $z_{\text{mmr}}(n) < 1$ for $n < K$. Consider the function $f(x) = -\exp(x^{-1})$. Then (2.27) holds for all values of $M$ and $z_{\text{mmr}}(n) < 1$ for any $n$.

## 2.5. Implications for Treatment Choice in Practice

This concluding section explores some implications of our analysis for the practice of treatment choice. In the course of doing so, we present numerical findings that add texture to the analysis.

### 2.5.1. Test-Based Rules in Medicine

Although problems of choice between a status quo treatment and an innovation occur often in practice, explicit use of statistical decision theory to make such choices is rare. In the medical arena, the branch of statistics that has strongly influenced practice has been hypothesis testing rather than decision theory. Indeed, testing the hypothesis of zero average treatment effect is institutionalized in the U.S. food and drug administration (FDA) drug approval process, which calls for comparison of a new treatment under study ($t = b$) with a placebo or an approved treatment ($t = a$). FDA approval of the new treatment normally requires one-sided rejection of the null hypothesis of zero average treatment effect $\{H_0\colon E[y(b)] = E[y(a)]\}$ in two independent clinical trials (Fisher and Moyé, 1999). In the context of treatments with binary outcomes, this means performance of a test with null hypothesis $\{H_0 : \beta = \alpha\}$ and alternative $\{H_1 : \beta > \alpha\}$.

The use of an hypothesis test to choose between the status quo treatment and an innovation gives the status quo a privileged position and, thus, might be loosely construed as an expression of risk aversion. However, the classical practice of handling the null and alternative hypotheses asymmetrically, fixing the probability of a type I error and seeking to minimize the probability of a type II error, is difficult to motivate from the perspective of treatment choice. Moreover, error probabilities at most measure the chance of choosing a sub-optimal rule. They do not measure the loss in welfare resulting from a sub-optimal choice.

Even if statistical decision theory does not motivate treatment rules based on hypothesis tests, we can productively use decision theory to evaluate such rules. In the setting of this chapter, a conventional test-based rule assigns treatment $b$ to the entire population if the number of experimental successes is large enough to reject $H_0$ and assigns treatment $a$ otherwise. Thus, a test-based rule has the form

$$
\begin{aligned}
z(n) &= 0 \quad \text{for } n \leq d(s, \alpha), \\
z(n) &= 1 \quad \text{for } n > d(s, \alpha),
\end{aligned}
\tag{2.28}
$$

where $s$ is the specified size of the test and $d(s, \alpha)$ is the associated critical value. Given that $n$ is binomial, $d(s, \alpha) = \min i : p(n > i; \alpha) \leq s$.

Test-based rules are KR-monotone. Hence, by the Corollary to Proposition 2.3, these rules are admissible if the welfare function has sufficiently weak curvature. This fact gives some grounding for the application of test-based rules, but admissibility is only a necessary condition for a treatment rule to be attractive. To obtain further understanding, we next compare the maximum regret of test-based rules with that of other treatment rules.

### 2.5.2. Comparing a Test-Based Rule with the Minimax-Regret Rule and the Plug-in Rule

Figure 2.1 shows the maximum regret of the treatment rule based on the exact binomial test with conventional size $s = 0.05$. The four panels of the figure consider two welfare functions (linear and log) and two values of $\alpha$ (0.25 and 0.75). In each panel, the $x$-axis gives the sample size $N$, ranging from 1 to 100. The $y$-axis gives maximum regret multiplied by $N^{1/2}$. For comparison, Figure 2.1 also shows the maximum regret of the minimax-regret rule and of the empirical plug-in rule. The latter rule assigns the entire

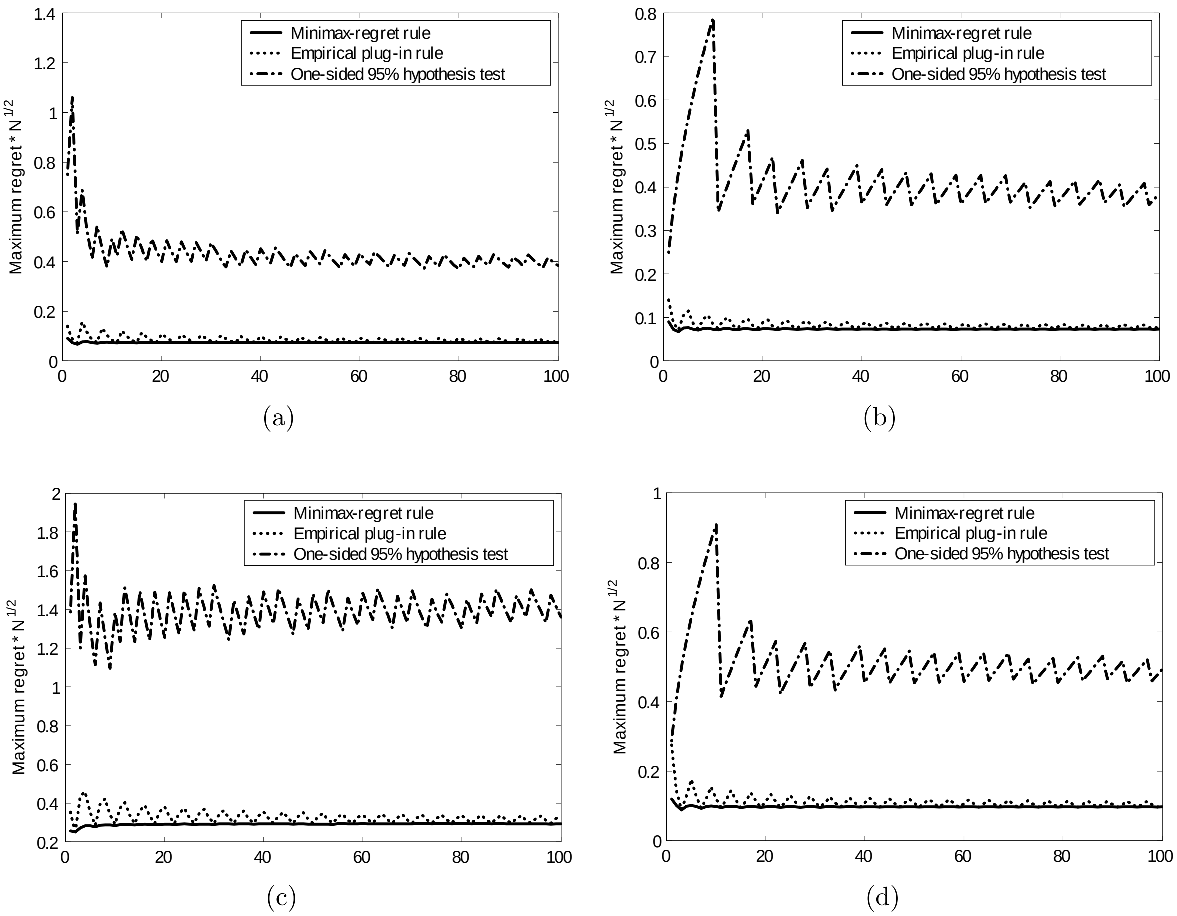


Figure 2.1: Maximum regret, $N = 1...100$. (a) $\alpha = 0.25$, $f(x) = x$; (b) $\alpha = 0.75$, $f(x) = x$; (c) $\alpha = 0.25$, $f(x) = \log x$; (d) $\alpha = 0.75$, $f(x) = \log x$.

population to treatment $b$ if the empirical rate of treatment success exceeds $\alpha$, and assigns everyone to treatment $a$ otherwise.

Consider first the behavior of maximum regret as a function of sample size. In every case, the primary large-scale feature of the plot is its invariance with $N$. This shows that maximum regret converges to zero at rate $N^{1/2}$ as sample size increases. A curious small-scale feature of the plots for the test-based and plug-in rules is jaggedness as the sample size varies across adjacent values of $N$. This occurs because these rules are step functions that remain constant over multiple values of $N$; for example, the plug-in rule is $z(n) = 1\,|n > \alpha N|$. The plots for the minimax-regret rule show no such jaggedness because the minimax-regret rule is fractional at the threshold and changes more smoothly with $N$.

Now compare the maximum regret of the three treatment rules. In every case, the maximum regret of the test-based rule is much larger than that of the minimax-regret rule. When the sample size is larger than ten, the ratio of the former maximum regret to the latter is typically about 5 to 1. These ratios quantify the inferiority of the test-based rule when viewed from the vantage of maximum regret.

One should not conclude that the test-based rule is inferior in all states of nature. Being admissible, this rule must yield smaller regret in some states of nature. The test-based rule, which "stacks the deck" in favor of the status quo treatment, delivers smaller regret than the minimax-regret rule in states of nature with $\beta < \alpha$ and larger regret in states with $\beta > \alpha$. The clear inferiority of the rule in terms of maximum regret arises because, under both the linear and log welfare functions, the latter losses are much larger than the former gains.

Observe that the maximum regret of the plug-in rule is close to that of the minimax-regret rule. Indeed, the two are nearly the same at the bottom of each jag of the plug-in rule. Although the minimax-regret rule is relatively easy to compute, the plug-in rule is simpler yet. Hence, the plots indicate that a practitioner who is not equipped to compute the minimax-regret rule would suffer little by using the plug-in rule as an approximation.

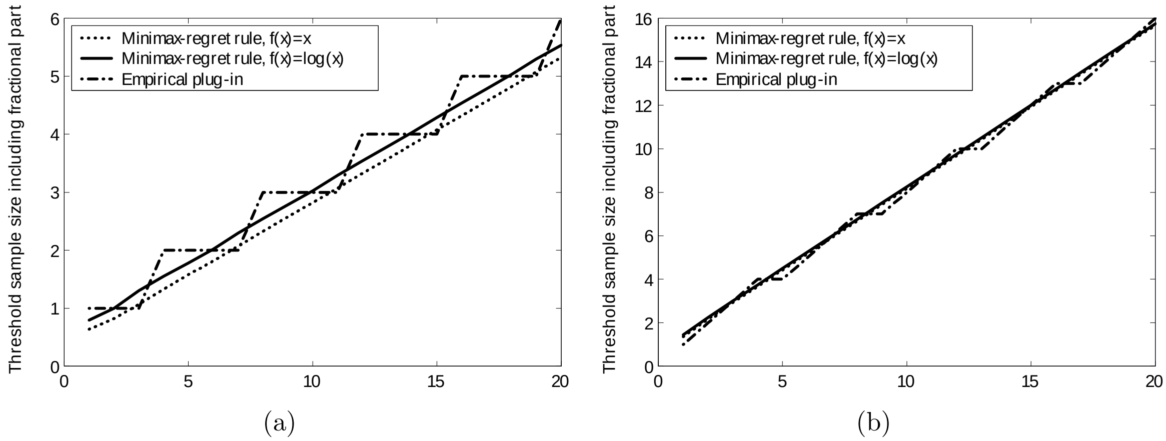


Figure 2.2: Threshold sample size, $N = 1...20$: (a) $\alpha = 0.25$; (b) $\alpha = 0.75$.

### 2.5.3. Variation of the Minimax-Regret Rule with the Welfare Function

Finally, we return to the question that most motivates this chapter, namely how the welfare function affects treatment choice. The analysis of Sections 2.3 and 2.4 has made clear that moving from a linear welfare function to one that is strongly curved can have important consequences. If $f(\cdot)$ has strong curvature, KR-monotone rules may not be admissible (Proposition 2.4) and the minimax-regret rule never assigns the entire population to treatment $b$ (Proposition 2.6). However, our analysis has not explored the consequences of moving from a linear welfare function to one with sufficiently weak curvature that the KR-monotone rules continue to form the minimal complete class (Corollary to Proposition 2.3).

To shed some light on this, we compare the minimax-regret rule for the linear and log welfare functions. Figure 2.2 presents this comparison for the same values of $\alpha$ as in Figure 2.1. In each panel, the $x$-axis gives the sample size $N$, ranging from 1 to 20. The findings to be discussed here are similar for the larger sample sizes shown in Figure

2.1.We do not present larger sample sizes in Figure 2.1 because making the $x$-axis run from 1 to 100 seriously diminishes one's ability to see important features of the plots.

The $y$-axis of each panel gives a one-dimensional representation of the minimax-regret rule. Having the KR-monotone form, this rule is defined by two numbers: the integer $k$ specifying the location of the step and the fraction $\lambda$ specifying the fraction of the population assigned to treatment $b$ when there are $k$ experimental successes. A one-dimensional representation of the rule is achieved by computing $k + (1 - \lambda)$. For example, the value 2.7 on the $y$-axis of Figure 2.2 means that $k = 2$ and $\lambda = 0.3$.

Figure 2.2 shows that moving from the linear to log welfare function has very little effect on the minimax-regret rule. The KR-threshold $k + (1 - \lambda)$ is nearly the same under both welfare functions. When $\alpha = 0.75$, the quantitative change in $k + (1 - \lambda)$ is so small as barely to be visible. When $\alpha = 0.25$, the change is more noticeable but its magnitude is still small. In both cases, the plots with respect to sample size are close to parallel to one another. We have computed analogs to Figure 2.2 for $\alpha$ as small as 0.01, and found that the variation in the rule across welfare functions is still small and that the plots remain close to parallel.

Figure 2.2 also shows the empirical plug-in rule. It is very similar to the two minimax-regret rules, the primary difference being that it is a step function rather than one that varies smoothly with $N$. This similarity explains why, in Figure 2.1, we found that the maximum regret of the plug-in rule is close to that of the minimax-regret rule.

Taken in combination, our analytical findings and the numerical findings in Figure 2.2 indicate that concavity of the welfare function does not per se have important consequences for treatment choice. What matters is the degree of curvature of the welfare function. We cannot say how curved a welfare function should be in practice. The answer to this question is necessarily context specific.

CHAPTER 3

# Measuring Precision of Statistical Inference on Partially Identified Parameters

## 3.1. Introduction

It has become widely recognized that many types of statistical data only partially identify the parameters of interest as simple as population means, meaning that the parameters cannot be estimated with arbitrary precision simply by increasing the sample size. Statisticians designing surveys and experiments which generate such data could use limited resources either to reduce the extent of partial identification or to reduce sampling error. The former can be accomplished, for example, by putting more effort into pursuing sampled population members who did not respond to a survey. The latter by increasing sample size or improving measurement precision. To inform these choices, it is useful to analytically derive the relative effects of both margins of planning on the precision of inference, which the planner could then compare to their relative costs.

The problem was first considered in the Cochran-Mosteller-Tukey report on the Kinsey study published in 1954. Concerned with nonrandom nonresponse to the study's questions, CMT advocated a conservative approach to inference that sets limits on population parameters by allowing for any values of the variable in the part of the population that was not sampled or refused to respond. A variety of applications of the same approach, now known as *partial identification*, has been developed by Manski

(1995, 2007a) and other researchers. CMT calculated for different sample sizes and refusal rates the relative effects of reducing nonresponse or increasing the sample size on the precision of inference about the population means. They judged the precision of inference by the length of a 95% confidence interval around the estimated identification region. The same measure of precision has been used to illustrate the effects of missing data on the precision of inference by Horowitz and Manski (1998) and McFadden (2006).

Length of a confidence interval for the identification region is not the only reasonable way to measure the precision of inference on the parameter of interest. In this chapter I consider two other measures of precision and show that they yield qualitatively different conclusions about the relative merits of reducing sampling error and reducing the extent of partial identification. First, I consider the maximum mean squared error (MSE) of the point estimate around the true value of the parameter, which has often been used by statisticians to measure the precision of estimators of point identified parameters.

The second measure considered in this chapter is the maximum regret of a statistical treatment rule. It is applicable when the parameter of interest is the difference in average returns of two mutually exclusive policies or treatments for a population of interest and the goal of inference is to decide which one should be used. Regret, then, is the average welfare loss incurred from choosing an inferior treatment for the population based on the observed statistical data. In recent years, econometricians started studying statistical treatment rules that minimize maximum regret both when the average treatment effect of interest is point identified (Manski 2004, 2005; Hirano and Porter 2006; Stoye 2007b; Schlag 2007; Manski and Tetenov 2007) and when it is partially identified (Manski 2007a, 2007b, 2008a, 2008b; Stoye 2007a, 2007c).

I apply these measures of precision to the following partial identification problem. Let the real-valued parameter of interest $\theta = \theta_O + \theta_U$ be the sum of a point identified component $\theta_O$ and a partially identified component $\theta_U$. For the point identified component $\theta_O$, the statistician observes an unbiased normally distributed estimate with known standard error $\sigma$. The partially identified component $\theta_U$ is only known to lie in a given bounded interval of length $2P$. The problem is deliberately simplified to demostrate in an analytically tractable setting the qualitative differences between the conclusions about the relative benefits of reducing sampling error vs. narrowing the identification region drawn based on alternative measures of precision. I derive the minimax estimator of $\theta$ under the maximum MSE criterion and a minimax regret statistical treatment rule under the maximum regret criterion. I show that as $\sigma \to 0$, both of these measures of precision imply greater relative importance of addressing the partial identificaton problem than measuring the length of confidence intervals suggests. For maximum regret, the result is particularly strong. If the standard error $\sigma$ falls below a certain proportion of the width of the identification region $2P$, then reducing it even further does not reduce maximum regret. Thus, more precise inference for treatment choice could be made only by reducing the width of the identification region. The same effect has been shown by Stoye (2007c) in a problem of treatment choice based on random samples of binary treatment outcomes with missing data.

The chapter proceeds as follows. Section 3.2 describes the statistical problem and reviews the results of measuring precision of inference by the length of confidence intervals. In section 3.3, I derive the estimator of $\theta$ that minimizes maximum MSE and evaluate the effect of changing the parameters of the problem on its minimax MSE. In section 3.4, I consider the problem from a statistical treatment choice perspective, derive a minimax regret statistical treatment rule and evaluate the effects of changing the parameters of the problem ($\sigma$ and $P$) on its minimax regret. Section 3.5 concludes and section 3.6 collects all proofs.

### 3.2. Statistical Setting and the Confidence Interval Approach

I will consider the following partial identification problem. The parameter of interest to the statistician is

$$\theta = \theta_O + \theta_U.$$

$\theta_O \in \mathbb{R}$ is a point identified (observable) component, for which the statistician could obtain an unbiased normally distributed estimate $X$ with standard error $\sigma$:

$$X \sim \mathcal{N}\left(\theta_O, \sigma^2\right).$$

$\theta_U$ is a partially identified (unobservable) component, which is only known to lie in a bounded interval of length $2P$:

$$\theta_U \in [-P, P].$$

The restriction that $\theta_U$ lies in a symmetric interval around zero is without loss of generality.

For example, $\theta$ could be the difference between average potential outcomes of two alternative treatments on a population of interest. Suppose that $\theta_O$ is the average outcome of one treatment, which is point identified by experimental data generated by assigning that treatment to a random sample of population members, while $-\theta_U$ is the average outcome of the second treatment, which is partially identified based on observational data.[1]

Alternatively, $\theta_O$ could be the average difference in potential outcomes of the two treatments point identified by an experiment that randomly assigned one of two treatments to members of the population of interest, while $-\theta_U$ is the difference between future costs of the two treatments that the randomized experiment did not seek to measure.

In this setting, the pair $(\sigma, P)$ describes the parameters of the experiment. The main question of this chapter is how do these parameters of the experiment affect the precision of inference on $\theta$ that the statistician could carry out based on its results (observation of $X$). Formally, let the function

$$M(\sigma, P) \geq 0$$

be a particular measure of maximum precision with which the statistician can carry out inference on $\theta$ based on the data from an experiment with parameters $(\sigma, P)$. Lower values of $M(\sigma, P)$ correspond to more precise inference and $M(\sigma, P) = 0$ corresponds to perfect precision (so, for example, $M(0, 0) = 0$). I will evaluate three such functions and the relative marginal effects of changes in $\sigma$ and $P$ on the precision they imply:

$$\frac{\partial M(\sigma, P)/\partial \sigma}{\partial M(\sigma, P)/\partial P}. \tag{3.1}$$

[1] I am grateful to Chuck Manski for this example.

When evaluating a proposed experiment, the planner could compare (3.1) to the ratio of marginal costs of decreasing $\sigma$ and $P$ and see whether the proposed allocation of resources maximizes the precison of inference (minimizes $M(\sigma, P)$).

First, let's consider using the length of a $100(1-\alpha)\%$ confidence interval for the identification interval as the measure of precision. In this model, the identification set for the parameter of interest $\theta$ is

$$\theta \in [\theta_O - P, \theta_O + P]. \tag{3.2}$$

Given that the random experimental outcome $X$ is normally distributed with mean $\theta_O$ and standard error $\sigma$, the confidence interval

$$\left[X - P - \Phi^{-1}(1-\alpha/2)\,\sigma, X + P + \Phi^{-1}(1-\alpha/2)\,\sigma\right] \tag{3.3}$$

contains the identification set (3.2) exactly with probability $1-\alpha$. $\Phi$ denotes the standard normal c.d.f., so for the conventional 95% confidence intervals, for example, $\Phi^{-1}(1-\alpha/2) \approx 1.96$. The precision of inference from an experiment with parameters $(\sigma, P)$, as measured by the length of a $100(1-\alpha)\%$ confidence interval then equals

$$M_{CI(\alpha)}(\sigma, P) \equiv 2\Phi^{-1}(1-\alpha/2)\,\sigma + 2P.$$

The marginal effects of changes in $\sigma$ and $P$ equal

$$\begin{aligned} \frac{\partial M_{CI(\alpha)}(\sigma, P)}{\partial \sigma} &= 2\Phi^{-1}(1-\alpha/2), \\ \text{and } \frac{\partial M_{CI(\alpha)}(\sigma, P)}{\partial P} &= 2. \end{aligned}$$

The ratio of these marginal effects equals

$$\frac{\partial M_{CI(\alpha)}(\sigma, P)/\partial\sigma}{\partial M_{CI(\alpha)}(\sigma, P)/\partial P} = \Phi^{-1}(1-\alpha/2). \tag{3.4}$$

Thus, if the length of conventional 95% confidence intervals is used as a measure of precision, then a reduction of the standard error $\sigma$ by $\varepsilon$ always brings the same improvement as a reduction of the half-length $P$ of the identification interval by $1.96\varepsilon$. Note that the evaluation of the relative effects of reducing the sampling error and the extent of partial identification depends on the chosen confidence level $100(1-\alpha)\%$. Thus, using a 99% confidence level instead of 95% would imply a relatively higher value of reducing the standard error instead of reducing the extent of partial identification.

### 3.3. Minimax Mean Squared Error Approach

Suppose, now, that instead of an interval the statistician is asked to provide a single point estimate of $\theta$. Let the estimator $\hat{\theta}(X)$ be a function mapping the observed experimental outcome $X$ into the estimate that the statistician provides upon observing $X$. There is a long tradition in statistics of measuring the precision of point estimators by their mean squared error

$$E_{(\theta_O,\theta_U)}\left(\hat{\theta}(X)-\theta\right)^2 = E_{(\theta_O,\theta_U)}\left(\hat{\theta}(X)-\theta_O-\theta_U\right)^2.$$

The expectation here is taken with respect to the probability distribution of $X$ for a given value of $\theta_O$. $\theta_U$ does not affect the probability distribution of the experimental outcome $X$, but affects the magnitude of error.

Proposition 3.1 shows that the estimator $\hat{\theta}(X) = X$ minimizes the maximum MSE

$$\max_{\substack{\theta_O \in \mathbb{R}, \\ \theta_U \in [-P,P]}} E_{(\theta_O,\theta_U)} \left(\hat{\theta}(X) - \theta_O - \theta_U\right)^2. \tag{3.5}$$

If we assumed that $\theta_U$ lies in an interval $\theta_U \in [L, U]$, then the corresponding minimax estimator would be $\hat{\theta}(X) = X + \frac{L+U}{2}$.

**Proposition 3.1.** *If $\theta_O \in \mathbb{R}$, $\theta_U \in [-P, P]$, and $X \sim \mathcal{N}(\theta_O, \sigma^2)$, then the estimator*

$$\hat{\theta}(X) \equiv X$$

*minimizes maximum MSE (3.5), which equals*

$$M_{MSE}(\sigma, P) \equiv \sigma^2 + P^2.$$

The marginal effects of changes in $\sigma$ and $P$ on the maximum mean squared error equal

$$\begin{aligned} \frac{\partial M_{MSE}(\sigma, P)}{\partial \sigma} &= 2\sigma, \\ \text{and } \frac{\partial M_{MSE}(\sigma, P)}{\partial P} &= 2P. \end{aligned}$$

And the ratio of these marginal effects equals

$$\frac{\partial M_{MSE}(\sigma, P)/\partial \sigma}{\partial M_{MSE}(\sigma, P)/\partial P} = \frac{\sigma}{P}. \tag{3.6}$$

This ratio yields qualitatively different conclusions than the ratio derived for confidence intervals about the relative benefits of reducing $\sigma$ vs. reducing $P$ when planning surveys or experiments. When $\frac{\sigma}{P} < \Phi^{-1}(1-\alpha/2)$,

$$\frac{\partial M_{MSE}(\sigma,P)/\partial\sigma}{\partial M_{MSE}(\sigma,P)/\partial P} < \frac{\partial M_{CI(\alpha)}(\sigma,P)/\partial\sigma}{\partial M_{CI(\alpha)}(\sigma,P)/\partial P}.$$

Thus, the maximum MSE measure of precision implies lower importance of further reducing standard errors than does the length of confidence interval measure. For the conventional 95% confidence intervals $\Phi^{-1}(1-\alpha/2) \approx 1.96$. Thus, in evaluating any proposed experiment or survey in which the standard error is going to be smaller than the length of the identification interval ($\sigma < 1.96P$) a planner using the maximum MSE measure of precision would allocate more resources to reducing the extent of partial identification than a planner measuring precision by the length of the confidence interval. The difference between the "marginal rates of substitution" produced by the two methods could be particularly large when considering large sample surveys and experiments in which the extent of partial identification could greatly exceed sampling error.

### 3.4. Minimax Regret Approach

The third measure of precision - minimax regret - is motivated by directly considering the economic loss resulting from incorrect inference about $\theta$ when $\theta$ is the difference in average returns of two alternative policy decisions and the ultimate aim of inference about $\theta$ is to choose which policy to implement. For example, the policies may be two proposed cancer therapies, with $\theta$ measuring the average difference in the welfare of cancer patients from a target population net of the average difference between the costs of these two therapies.

Let $\theta = r_2 - r_1$, where $r_1$ is the average return from implementing the first policy and $r_2$ the average return from implementing the second policy. Then the economic loss from choosing the second policy when, in fact, $r_1 > r_2$ $(\theta < 0)$ equals $r_1 - r_2 = -\theta$. The economic loss from choosing to implement the first policy when, in fact, $r_1 < r_2$ $(\theta > 0)$ equals $r_2 - r_1 = \theta$. The method by which the decision maker chooses which policy to implement based on experimental data $X$ could be summarized by a *statistical treatment rule* $\delta(X)$, which is a function mapping feasible realizations of $X \in \mathbb{R}$ into the $[0, 1]$ interval. $\delta\left(\bar{X}\right) = 0$ if the decision maker implements the first policy when outcome $\bar{X}$ is observed, $\delta\left(\bar{X}\right) = 1$ if she implements the second policy. $\delta\left(\bar{X}\right)$ could takes values between 0 and 1 if the decision maker could implement either policy with some probability upon observing outcome $\bar{X}$.

The regret of statistical treatment rule $\delta$ is the average (over the probability distribution of outcome $X$) economic loss incurred by the decision maker using $\delta$. It is a function of $\theta_O$ and $\theta_U$, and in this problem equals

$$R(\delta, (\theta_O, \theta_U)) \equiv \begin{cases} \theta \cdot [1 - E_{\theta_O}\delta(X)] & \text{if } \theta > 0, \\ -\theta \cdot E_{\theta_O}\delta(X) & \text{if } \theta \leq 0, \end{cases} \tag{3.7}$$

where $E_{\theta_O}\delta(X)$ denotes the average value of $\delta(X)$ given that $X \sim \mathcal{N}(\theta_O, \sigma^2)$. When $\theta > 0$, the first policy is inferior and $[1 - E_{\theta_O}\delta(X)]$ is the probability with which the decision maker would mistakenly choose it based on observation of the random experimental outcome $X$. When $\theta < 0$, the second policy is inferior and $E_{\theta_O}\delta(X)$ is the probability of choosing it.

Minimizing maximum regret was a criterion suggested by Savage (1951) as a clarification of Wald's *minimax principle* (1950). For a more detailed discussion on applying minimax regret criterion to statistical treatment choice problems see Manski (2004 or 2007a, Chapter 11).

Of course, if we will measure the precision of inference by the maximum regret of a statistical treatment rule, then we first ought to find statistical treatment rules that minimize maximum regret for given experimental parameters $(\sigma, P)$. Proposition 3.2 derives such rules and their maximum regret.

**Proposition 3.2.** *a) For $\sigma > 2P \cdot \phi(0)$, the unique minimax regret statistical treatment rule is*

$$\delta_{M(\sigma,P)}(X) \equiv 1\left|X > 0\right|. \tag{3.8}$$

*Its maximum regret equals*

$$\max_{\substack{\theta_O \in \mathbb{R}, \\ \theta_U \in [-P,P]}} R\left(\delta_{M(\sigma,P)}, (\theta_O, \theta_U)\right) = \max_{h>0}\left[h\Phi\left(\frac{P-h}{\sigma}\right)\right] > \frac{P}{2},$$

*which is a strictly increasing function of $\sigma$ for any given $P$.*

*b) For $\sigma \leq 2P \cdot \phi(0)$, statistical treatment rules*

$$\delta_{M(\sigma,P)}(X) \equiv \begin{cases} 1\left|X > 0\right| & \text{if } \sigma = 2P \cdot \phi(0), \\ \Phi\left(\left[(2P \cdot \phi(0))^2 - \sigma^2\right]^{-1/2} X\right) & \text{if } \sigma < 2P \cdot \phi(0), \end{cases} \tag{3.9}$$

*minimize maximum regret, which equals $\frac{P}{2}$.*

Two features of Proposition 3.2 are qualitatively similar to results obtained by Stoye (2007c), who studied minimax regret statistical treatment rules based on binary outcome data from an experiment with randomized treatment assignment in which the outcomes are missing with some probability.

First, when the extent of partial identification (in Stoye's problem, the maximum feasible proportion of missing outcomes) is below some threshold relative to the sampling error, the minimax regret statistical treatment rule is the same as it would be with point identification. In Proposition 3.2 (part a) the same result holds, the minimax regret statistical treatment rule (3.8) is identical for all values of $P \leq \frac{\sigma}{2\phi(0)}$, including the point identified case $P = 0$.

The second qualitative similarity is that maximum regret of the minimax regret statistical treatment rule becomes constant with respect to the sampling error once the sampling error falls below some threshold relative to the extent of partial identification. Thus, reducing the sampling error below that threshold (reducing $\sigma$ in this chapter, increasing sample size in Stoye's) could not further reduce minimax regret.

Since this second result could appear counterintuitive, it deserves further explanation. Let

$$q(\delta, \theta_O) \equiv E_{\theta_O}\delta(X)$$

denote the average probability (with respect to the distribution of $X$) with which the decision maker using statistical treatment $\delta$ will choose the second policy. Then for a given value of $P$, $\frac{P}{2}$ is the lower bound on maximum regret attainable by any statistical treatment rule for any value of $\sigma$. This could be seen by considering maximum regret over the subset $\{\theta_O = 0, \theta_U \in [-P, P]\}$

$$\max_{\substack{\theta_O \in \mathbb{R},\\ \theta_U \in [-P,P]}} R(\delta, (\theta_O, \theta_U)) \geq \max_{\theta_U \in [-P,P]} R(\delta, (0, \theta_U)) = \max(P \cdot q(\delta, 0), P \cdot (1 - q(\delta, 0))) \geq \frac{P}{2}.$$

In order to attain this lower bound, the statistical treatment rule $\delta$ must satisfy $q(\delta, 0) = \frac{1}{2}$. For values $\bar{\theta}_O \neq 0$, however, there is a range of values of $q\left(\delta, \bar{\theta}_O\right)$ for which

$$\max_{\theta_U \in [-P,P]} R\left(\delta, \left(\bar{\theta}_O, \theta_U\right)\right) \leq \frac{P}{2}.$$

This range is given by the inequalities

$$(3.10) \qquad \begin{aligned} q(\delta, \theta_O) &\geq 1 - \tfrac{P}{2(P+\theta_O)} \quad \text{for } \theta_O \geq -\tfrac{P}{2}, \\ q(\delta, \theta_O) &\leq \tfrac{P}{2(P-\theta_O)} \qquad \text{for } \theta_O \leq \tfrac{P}{2}. \end{aligned}$$

As shown in Proposition 3.2, for $\sigma \leq 2P \cdot \phi(0)$, it is possible to construct statistical treatment rules that satisfy these inequalities for all $\theta_O \in \mathbb{R}$. Figure 3.1 displays in bold lines the bounds (3.10) and shows that the function $q\left(\delta_{M(\sigma,P)}, \theta_O\right) = \Phi\left(\frac{\theta_O}{2P \cdot \phi(0)}\right)$, which is identical for all minimax regret rules defined by (3.9), fits within these bounds.

Statistical treatment rules derived in part b of Proposition 3.2 may not be the only ones that minimize maximum regret, but deriving one class of minimax regret rules is sufficient to make conclusions about the minimum value of maximum regret, and thus about the precision of inference from the data for treatment choice.

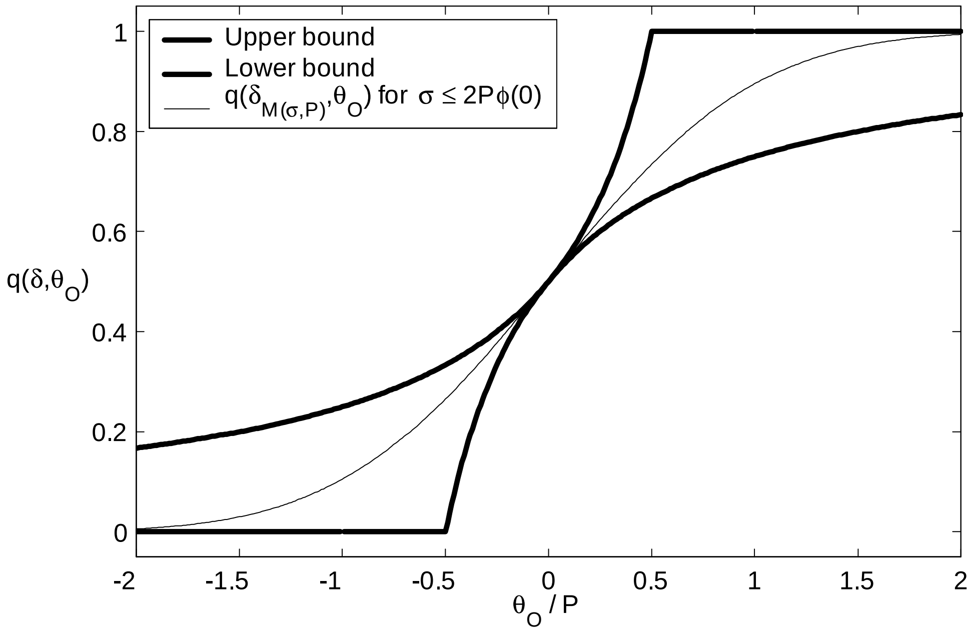


Figure 3.1: Bounds on $q(\delta, \theta_O)$ that guarantee attaining the lower bound on maximum regret $(P/2)$.

Suppose that the planner chooses minimax regret to measure inferential precision of the data generated by an experiment with parameters $(\sigma, P)$

$$M_{MMR}(\sigma, P) \equiv \begin{cases} \max\limits_{h>0} \left\{ h\Phi\left(\frac{P-h}{\sigma}\right)\right\} & \text{if } \sigma > 2P \cdot \phi(0), \\ \frac{P}{2} & \text{if } \sigma \leq 2P \cdot \phi(0). \end{cases}$$

Doing so could yield even more drastically different conclusions about the relative benefits of reducing the extent of partial identification and reducing sampling error than either the confidence interval length or the maximum MSE approach, since for

$\frac{\sigma}{P} \leq 2\phi(0) \approx 0.8$

$$\frac{\partial M_{MMR}(\sigma, P)}{\partial \sigma} = 0,$$
$$\frac{\partial M_{MMR}(\sigma, P)}{\partial P} = \frac{1}{2},$$

implying that reducing the extent of partial identification is not only relatively more important than reducing sampling error, it is the only way to reduce minimax regret and improve the inferential precision of experimental or survey data for treatment choice.

## 3.5. Conclusion

In this chapter, I considered two alternative measures of inferential precision for partially identified parameteres in addition to the length of 95% confidence interval, which is the primary measure previously considered by other researchers. All three measures yield qualitatively different conclusions about the relative merits of reducing sampling error and reducing the extent of partial identification in the data. In particular, both the maximum mean squared error and minimax regret (applicable when inference is carried out on the average treatment effect with the goal of choosing the best treatment) emphasize greater value of reducing the extent of partial identification compared to the confidence interval measure if the sampling error is relatively small compared to the width of the identification interval.

The statistical problem with a normal sampling distribution considered in the chapter is simple in comparison to many practical problems. However, it is sufficiently rich to capture some of the main features of partial identification problems and to concisely illustrate how choosing different criteria for measuring the precision of inference qualitatively impacts the conclusions about the relative value of reducing the extent of partial identification and reducing sampling error. The results could serve both as a rough practical approximation for partial identification problems with similar structure and as a useful indicator of potential findings for future research that considers more complex practical partial identification problems.

## 3.6. Proofs

**Proposition 3.1**

For an estimator $\hat{\theta}(X)$, define

$$\begin{aligned} b_{\hat{\theta}}(\theta_O) &\equiv E\hat{\theta}(X) - \theta_O, \\ v_{\hat{\theta}}(\theta_O) &\equiv E\left(\hat{\theta}(X) - E\hat{\theta}(X)\right)^2, \end{aligned}$$

with expectatations taken over the probability distribution of $X \sim \mathcal{N}(\theta_O, \sigma^2)$. It is well known (e.g., Berger 1985, p. 350) that $\hat{\theta}^*(X) = X$ is a minimax estimator of normal mean $\theta_O$ under squared error. That means for any other estimator $\hat{\theta}(X)$

$$\begin{aligned} \max_{\theta_O \in \mathbb{R}} E\left(\hat{\theta}(X) - \theta_O\right)^2 &\geq \max_{\theta_O \in \mathbb{R}} E\left(\hat{\theta}^*(X) - \theta_O\right)^2, \\ \max_{\theta_O \in \mathbb{R}} \left[v_{\hat{\theta}}(\theta_O) + b_{\hat{\theta}}^2(\theta_O)\right] &\geq \max_{\theta_O \in \mathbb{R}} \left[v_{\hat{\theta}^*}(\theta_O) + b_{\hat{\theta}^*}^2(\theta_O)\right] = \sigma^2. \end{aligned} \tag{3.11}$$

The maximum MSE of $\hat{\theta}(X)$ for estimating $\theta = \theta_O + \theta_U$ equals

$$\begin{aligned} \max_{\substack{\theta_O \in \mathbb{R}, \\ \theta_U \in [-P,P]}} E\left(\hat{\theta}(X) - \theta_O - \theta_U\right)^2 &= \max_{\substack{\theta_O \in \mathbb{R}, \\ \theta_U \in [-P,P]}} \left[v_{\hat{\theta}}(\theta_O) + \left(E\hat{\theta}(X) - \theta_O - \theta_U\right)^2\right] = \\ &= \max_{\theta_O \in \mathbb{R}} \left[v_{\hat{\theta}}(\theta_O) + \left(\left|b_{\hat{\theta}}(\theta_O)\right| + P\right)^2\right]. \end{aligned}$$

Since for $\hat{\theta}^*(X)$, $b_{\hat{\theta}^*}(\theta_O) = 0$ and $v_{\hat{\theta}^*}(\theta_O) = \sigma^2$, its maximum MSE equals $\sigma^2 + P^2$.

It follows from (3.11), that the maximum MSE of any other estimator $\hat{\theta}(X)$

$$\max_{\theta_O \in \mathbb{R}} \left[ v_{\hat{\theta}}(\theta_O) + \left( |b_{\hat{\theta}}(\theta_O)| + P \right)^2 \right] \geq \max_{\theta_O \in \mathbb{R}} \left[ v_{\hat{\theta}}(\theta_O) + b_{\hat{\theta}}^2(\theta_O) \right] + P^2 \geq \sigma^2 + P^2,$$

thus $\hat{\theta}^*(X) = X$ minimizes the maximum MSE. □

**Proposition 3.2(a)**

Let $(\theta_O, \theta_U) \in \Theta, \Theta = \mathbb{R} \times [-P, P]$. The proof of part a relies on a well known result (e.g., Berger 1985, p. 350) that if $\pi^*$ is a proper prior distribution on $\Theta$, the decision rule $\delta^*$ is Bayes with respect to $\pi^*$, and for all $(\theta_O, \theta_U) \in \Theta$

$$R(\delta^*, (\theta_O, \theta_U)) \leq \int R(\delta, (\theta_O, \theta_U)) \partial \pi^*(\theta_O, \theta_U),$$

then the decision rule $\delta^*$ is minimax. This result applies as well when $R$ denotes regret, then $\delta^*$ is a minimax-regret rule.

Decision rule

$$\delta^*(X) \equiv 1 |X > 0|$$

is Bayes with respect to any symmetric two-point prior distribution $\pi$ with $\pi(\theta_O^*, \theta_U^*) = .5$ and $\pi(-\theta_O^*, -\theta_U^*) = .5$, if $\theta_O^* > 0$ and $\theta_O^* + \theta_U^* > 0$.

When $\theta > 0$, for a given value of $\theta_O$, regret $R(\delta^*, (\theta_O, \theta_U)) = (\theta_O + \theta_U) \cdot [1 - E_{\theta_O} \delta^*(X)]$ is largest at $\theta_U = P$, since the first term is increasing in $\theta_U$ and the second term is positive and doesn't depend on $\theta_U$. Since $E_{\theta_O} \delta^*(X) = 1 - \Phi\left(-\frac{\theta_O}{\sigma}\right)$, maximum regret of $\delta^*$ over $\theta > 0$ then equals (with the substitution $h = \theta_O + P$)

$$\max_{\substack{\theta_O, \theta_U \in \Theta, \\ \theta_O + \theta_U > 0}} R\left(\delta^*, (\theta_O, \theta_U)\right) = \max_{\theta_O > -P} \left[ (\theta_O + P) \cdot \Phi\left(-\frac{\theta_O}{\sigma}\right)\right] = \max_{h>0} \left[ h\Phi\left(\frac{P-h}{\sigma}\right)\right].$$

The maximum is attained at

$$\theta_O^* = \arg\max_{h>0} \left[ h\Phi\left(\frac{P-h}{\sigma}\right)\right] - P.$$

When $\theta < 0$, regret $R\left(\delta^*, (\theta_O, \theta_U)\right) = -\left(\theta_O + \theta_U\right) \cdot E_{\theta_O} \delta^*\left(X\right)$ is maximized at $\theta_U = -P$ for a given $\theta_O$, and equals (with the substitution $h = -\left(\theta_O - P\right)$)

$$\max_{\substack{\theta_O, \theta_U \in \Theta, \\ \theta_O + \theta_U < 0}} R\left(\delta^*, (\theta_O, \theta_U)\right) = \max_{\theta_O < P} \left[ -\left(\theta_O - P\right) \cdot \Phi\left(\frac{\theta_O}{\sigma}\right)\right] = \max_{h>0} \left[ h\Phi\left(\frac{P-h}{\sigma}\right)\right].$$

Let's differentiate $h\Phi\left(\frac{P-h}{\sigma}\right)$ with respect to $h$

$$\frac{\partial}{\partial h}\left[ h\Phi\left(\frac{P-h}{\sigma}\right)\right] = \Phi\left(\frac{P-h}{\sigma}\right) - \frac{h}{\sigma}\phi\left(\frac{P-h}{\sigma}\right) = \Phi\left(\frac{P-h}{\sigma}\right)\left[1 - \frac{h}{\sigma}\frac{\phi\left(\frac{P-h}{\sigma}\right)}{\Phi\left(\frac{P-h}{\sigma}\right)}\right].$$

At $h = 0$, $\frac{\partial}{\partial h}\left[ h\Phi\left(\frac{P-h}{\sigma}\right)\right] = \Phi\left(\frac{P}{\sigma}\right) > 0$. The function $\frac{\phi(y)}{\Phi(y)} > 0$ is strictly decreasing with $\lim\limits_{y \to -\infty} \frac{\phi(y)}{\Phi(y)} = +\infty$, thus $\frac{h}{\sigma}\frac{\phi\left(\frac{P-h}{\sigma}\right)}{\Phi\left(\frac{P-h}{\sigma}\right)}$ is strictly increasing in $h$ over $h > 0$ and $\lim\limits_{h\to\infty}\left[1 - \frac{h}{\sigma}\frac{\phi\left(\frac{P-h}{\sigma}\right)}{\Phi\left(\frac{P-h}{\sigma}\right)}\right] = -\infty$. It follows that $\frac{\partial}{\partial h}\left[ h\Phi\left(\frac{P-h}{\sigma}\right)\right]$ changes sign once over $h > 0$ from positive to negative at $h^*$ given by $\frac{h^*}{\sigma}\frac{\phi\left(\frac{P-h^*}{\sigma}\right)}{\Phi\left(\frac{P-h^*}{\sigma}\right)} = 1$, thus $h\Phi\left(\frac{P-h}{\sigma}\right)$ attains its maximum over $h > 0$ at $h^*$.

When $\sigma > 2P \cdot \phi\left(0\right)$, $h^* > P$. To see this, evaluate $\frac{\partial}{\partial h}\left[ h\Phi\left(\frac{P-h}{\sigma}\right)\right]$ at $h = P$:

$$\left.\frac{\partial}{\partial h}\left[ h\Phi\left(\frac{P-h}{\sigma}\right)\right]\right|_{h=P} = \Phi\left(0\right) - \frac{P}{\sigma}\phi\left(0\right) = \frac{1}{2} - \frac{P \cdot \phi\left(0\right)}{\sigma} > 0.$$

Thus $\frac{\partial}{\partial h}\left[h\Phi\left(\frac{P-h}{\sigma}\right)\right]$ changes sign at $h^{*}>P$. Since $h^{*}>P$, maximum regret is attained at $\left(\theta_{O}^{*},P\right)$ and $\left(-\theta_{O}^{*},-P\right)$, where $\theta_{O}^{*}=h^{*}-P>0$.

Maximum regret of $\delta^{*}$ exceeds $\frac{P}{2}$ because $\frac{\partial}{\partial h}\left[h\Phi\left(\frac{P-h}{\sigma}\right)\right]>0$ for $P\leq h<h^{*}$, therefore

$$\max_{h>0}\left[h\Phi\left(\frac{P-h}{\sigma}\right)\right]=h^{*}\Phi\left(\frac{P-h^{*}}{\sigma}\right)>P\Phi\left(\frac{P-P}{\sigma}\right)=\frac{P}{2}.$$

Since $\delta^{*}$ is a Bayes rule with respect to prior $\pi^{*}$ with $\pi^{*}\left(\theta_{O}^{*},P\right)=.5$ and $\pi^{*}\left(-\theta_{O}^{*},-P\right)=.5$ and

$$\int R\left(\delta^{*},\left(\theta_{O},\theta_{U}\right)\right)\partial\pi^{*}\left(\theta_{O},\theta_{U}\right)=R\left(\delta^{*},\left(\theta_{O}^{*},P\right)\right)=\max_{\substack{\theta_{O}\in\mathbb{R},\\ \theta_{U}\in[-P,P]}}R\left(\delta^{*},\left(\theta_{O},\theta_{U}\right)\right),$$

$\delta^{*}$ minimizes maximum regret. Furthermore, since $\delta^{*}$ is a unique Bayes rule up to randomization at $X=0$, which does not affect $R\left(\delta,\left(\theta_{O},\theta_{U}\right)\right)$ for any values of $\left(\theta_{O},\theta_{U}\right)$, it is admissible.

To verify that minimax regret $\max_{h>0}\left[h\Phi\left(\frac{P-h}{\sigma}\right)\right]$ is a decreasing function of $\sigma$ for a given $P$ and $\sigma>2P\cdot\phi(0)$, observe that since $h^{*}>P$,

$$\max_{h>0}\left[h\Phi\left(\frac{P-h}{\sigma}\right)\right]=\max_{h>P}\left[h\Phi\left(\frac{P-h}{\sigma}\right)\right].$$

For any $h>P$, $h\Phi\left(\frac{P-h}{\sigma}\right)$ is strictly decreasing in $\sigma$ and has a unique maximum over $h>P$ for a given $\sigma$, thus $\max_{h>P}\left[h\Phi\left(\frac{P-h}{\sigma}\right)\right]$ is strictly decreasing in $\sigma$. $\square$

**Proposition 3.2(b)**

First, I will show that any rule $\delta$ for which $q(\delta, \theta_O) = E_{\theta_O}\delta(X)$ lies within the bounds (3.10) has maximum regret of $\frac{P}{2}$. The lower bound

$$q(\delta, \theta_O) \geq 1 - \tfrac{P}{2(P+\theta_O)} \quad \text{for } \theta_O \geq -\tfrac{P}{2},$$

guarantees that $R(\delta, (\theta_O, \theta_U)) \leq \frac{P}{2}$ over $\theta > 0$. Since $R(\delta, (\theta_O, \theta_U))$ is increasing in $\theta_U$ when $\theta > 0$,

$$\max_{\substack{\theta_O, \theta_U \in \Theta, \\ \theta_O + \theta_U > 0}} R(\delta, (\theta_O, \theta_U)) = \max_{\theta_O > -P} R(\delta, (\theta_O, P)) = \max_{\theta_O > -P} \left[(\theta_O + P) \cdot [1 - q(\delta, \theta_O)]\right].$$

For $\theta_O \geq -\frac{P}{2}$, if $q(\delta, \theta_O) \geq 1 - \frac{P}{2(P+\theta_O)} \geq 0$, then

$$(\theta_O + P) \cdot [1 - q(\delta, \theta_O)] \leq (\theta_O + P) \cdot \frac{P}{2(P + \theta_O)} = \frac{P}{2}.$$

For $\theta_O \in \left[-P, -\frac{P}{2}\right]$,

$$(\theta_O + P) \cdot [1 - q(\delta, \theta_O)] \leq \theta_O + P \leq \frac{P}{2}.$$

The proof for the upper bound, which ensures that $R(\delta, (\theta_O, \theta_U)) \leq \frac{P}{2}$ for $\theta < 0$, is analogous. Both the lower and the upper bound are equal to $\frac{1}{2}$ at $\theta_O = 0$, thus $q(\delta, 0) = \frac{1}{2}$ and

$$\max_{\theta_O, \theta_U \in \Theta} R(\delta, (\theta_O, \theta_U)) \geq \max_{\theta_U \in [-P, P]} R(\delta, (0, \theta_U)) = \frac{P}{2}.$$

Thus the maximum regret of $\delta$ equals $\frac{P}{2}$ if $q(\delta, \theta_O)$ satisfies inequalities (3.10).

Second, I will show that the function

$$q^*(\theta_O) \equiv \Phi\left(\frac{\theta_O}{2P\cdot\phi(0)}\right)$$

lies within the bounds (3.10). The proof will verify this for $\theta_O \geq 0$, it is analogous for $\theta_O < 0$.

When $\theta_O = 0$, $q^*(0) = \Phi(0) = \frac{1}{2}$, which coincides with both bounds. $q^*(\theta_O)$ satisfies the upper bound because for $\theta_O \in \left[0, \frac{P}{2}\right]$

$$\Phi\left(\frac{\theta_O}{2P\cdot\phi(0)}\right) \leq \frac{1}{2} + \frac{\theta_O}{2P\cdot\phi(0)}\cdot\phi(0) = \frac{P+\theta_O}{2P} \leq \frac{P}{2(P-\theta_O)}.$$

The first inequality follows from using $\phi(0)$ as an upper bound on the derivative of $\Phi$. The second one follows from $(P+\theta_O)(P-\theta_O) = P^2 - \theta_O^2 \leq P^2$.

The proof that $q^*(\theta_O) \geq 1 - \frac{P}{2(P+\theta_O)}$ for all $\theta_O \geq 0$ is split into two cases, $\theta_O \in [0, P]$ and $\theta_O \geq P$.

Case 1. For $\theta_O \in [0, P]$, I will prove that $q^*(\theta_O)$ increases faster than the lower bound, which guarantees that $q^*(\theta_O) \geq 1 - \frac{P}{2(P+\theta_O)}$, since both are equal at $\theta_O = 0$. It will be sufficient to consider $P = 1$, to simplify notation, and thus $\theta_O \in [0, 1]$. For $P = 1$, $q^*(\theta_O) = \Phi\left(\frac{\theta_O}{2\phi(0)}\right)$, $\phi(y) = \frac{1}{\sqrt{2\pi}}\exp\left(-\frac{1}{2}y^2\right)$, with $2\phi(0) = \sqrt{\frac{2}{\pi}}$, thus

$$\frac{\partial}{\partial\theta_O}q^*(\theta_O) = \frac{1}{2\phi(0)}\phi\left(\frac{\theta_O}{2\phi(0)}\right) = \sqrt{\frac{\pi}{2}}\cdot\frac{1}{\sqrt{2\pi}}\exp\left(-\frac{1}{2}\left(\sqrt{\frac{\pi}{2}}\theta_O\right)^2\right) = \frac{1}{2}\exp\left(-\frac{\pi}{4}\theta_O^2\right).$$

Since the function $e(y)$ is convex with $e(0) = 1$ and $e(1) < 3$, $e(y) \leq 1 + 2y$ for $y \in [0,1]$, therefore $e(y) \geq \frac{1}{1-2y}$ for $y \in [-1,0]$. Since $\frac{\pi}{4} < 1$ and $\theta_O^2 < 1$,

$$\frac{1}{2}\exp\left(-\frac{\pi}{4}\theta_O^2\right) \geq \frac{1}{2}\cdot\frac{1}{1+\frac{\pi}{2}\theta_O^2} = \frac{1}{2+\pi\theta_O^2}.$$

For $\theta \in [0,1]$, $\pi\theta_O < 4$, thus $\pi\theta_O^2 \leq 4\theta_O$ and $2 + \pi\theta_O^2 \leq 2 + 4\theta_O + 2\theta_O^2 = 2(1+\theta_O)^2$, therefore $\frac{1}{2+\pi\theta_O^2} \geq \frac{1}{2(1+\theta_O)^2}$, and

$$\frac{\partial}{\partial\theta_O}q^*(\theta_O) \geq \frac{1}{2+\pi\theta_O^2} \geq \frac{1}{2(1+\theta_O)^2} = \frac{\partial}{\partial\theta_O}\left[1 - \frac{1}{2(1+\theta_O)}\right].$$

Case 2. For $\theta_O \geq P$, I will also use $P = 1$ to simpify notation, so the aim is to prove that $q^*(\theta_O) = \Phi\left(\frac{\theta_O}{2\phi(0)}\right) \geq 1 - \frac{1}{2(1+\theta_O)}$. For $y > 0$, $1 - \Phi(y) < \frac{\phi(y)}{y}$, which implies

$$q^*(\theta_O) = \Phi\left(\frac{\theta_O}{2\phi(0)}\right) > 1 - \frac{2\phi(0)}{\theta_O}\phi\left(\frac{\theta_O}{2\phi(0)}\right) = 1 - \frac{1}{\pi\theta_O}\exp\left(-\frac{\pi}{4}\theta_O^2\right).$$

For $y \geq 0$, $e(y) \geq 1 + y$, thus for $y \leq 0$, $e(y) \leq \frac{1}{1-y}$. Using this inequality yields

$$q^*(\theta_O) > 1 - \frac{1}{\pi\theta_O}\cdot\frac{1}{1+\frac{\pi}{4}\theta_O^2} = 1 - \frac{1}{\pi\theta_O + \frac{\pi^2}{4}\theta_O^3} > 1 - \frac{1}{2(1+\theta_O)}$$

where the last inequality follows from observation that $\frac{\pi^2}{4} > 2$, and for $\theta_O \geq 1$, $\pi\theta_O > 2$ and $\theta_O^3 \geq \theta_O$.

Since $q^*(\theta_O)$ satisfies the inequalities (3.10), any statistical treatment rule with $q(\delta,\theta_O) = q^*(\theta_O)$ has maximum regret of $\frac{P}{2}$. It remains to show that this holds for statistical treatment rules (3.9) defined in part b of Proposition 3.2.

For $\sigma = 2P \cdot \phi(0)$, $\delta_{M(\sigma,P)}(X) = 1\,|X > 0|$, thus $q\left(\delta_{M(\sigma,P)}, \theta_O\right) = \Phi\left(\frac{\theta_O}{2P\cdot\phi(0)}\right) = q^*(\theta_O)$ and the rule minimizes maximum regret, which equals $\frac{P}{2}$.

For $\sigma < 2P \cdot \phi(0)$, it is simplest to derive $\delta_{M(\sigma,P)}(X)$ using the following construction[2]. Let $\sigma_0 = 2P \cdot \phi(0)$. Define an auxiliary random variable

$$Y \sim \mathcal{N}\left(0, \sigma_0^2 - \sigma^2\right),$$

independent of the observed outcome $X \sim \mathcal{N}(\theta_O, \sigma^2)$. Then $X + Y \sim \mathcal{N}(\theta_O, \sigma_0^2)$. Define the statistical treatment rule $\delta^*_{M(\sigma,P)}(X, Y)$ as

$$\delta^*_{M(\sigma,P)}(X, Y) \equiv 1\,|X + Y > 0|,$$

then clearly

$$q\left(\delta^*_{(\sigma,P)}, \theta_O\right) = \Phi\left(\frac{\theta_O}{2P \cdot \phi(0)}\right) = q^*(\theta_O).$$

Integrating $\delta^*_{M(\sigma,P)}(X, Y)$ with respect to the distribution of $Y$ yields

$$\delta_{M(\sigma,P)}(X) \equiv E\left(1\,|X + Y > 0|\right) = 1 - \Phi\left(-\left(\sigma_0^2 - \sigma^2\right)^{-1/2} X\right) = \Phi\left(\left(\sigma_0^2 - \sigma^2\right)^{-1/2} X\right),$$

which thus satisfies $q\left(\delta_{M(\sigma,P)}, \theta_O\right) = q^*(\theta_O)$ by construction and.minimizes maximum regret, which equals $\frac{P}{2}$. $\square$

[2]This proof technique is similar to Schlag's (2007) *binomial average*, in that both algebraically simplify the problem by adding some noise to the observed outcomes.

## References

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

[2] Canner, P. 1970. Selecting One of Two Treatments when the Responses are Dichotomous. *Journal of the American Statistical Association* 65: 293–306.

[3] Cheng, Y., F. Su, and D. Berry. 2003. Choosing Sample Size for a Clinical Trial Using Decision Analysis. *Biometrika* 90: 923–936.

[4] Cochran, W., F. Mosteller, and J. Tukey. 1954. *Statistical Problems of the Kinsey Report on Sexual Behavior in the Human Male.* Washington, DC: American Statistical Association.

[5] Eozenou P., J. Rivas, and K.H. Schlag. 2006. Minimax Regret in Practice - Four Examples on Treatment Choice. Working paper, European University Institute.

[6] Fisher, L., and L. Moyé. 1999. Carvedilol and the Food and Drug Administration Approval Process: an Introduction. *Controlled Clinical Trials* 20: 1–15.

[7] Gould, A. 2002. Substantial Evidence of Effect. *Journal of Biopharmaceutical Statistics* 12: 53–77.

[8] Hirano, K., and J.R. Porter. 2006. Asymptotics for Statistical Treatment rules. Working paper, University of Arizona and University of Wisconsin-Madison.

[9] Horowitz, J.L., and C.F. Manski. 1998. Censoring of Outcomes and Regressors Due to Survey Nonresponse: Identification and Estimation Using Weights and Imputations. *Journal of Econometrics* 84: 37–58.

[10] International Committee on Harmonization. 1998. *Guideline E9: Statistical Principles for Clinical Trials.*

[11] Karlin, S., and H. Rubin. 1956. The Theory of Decision Procedures for Distributions with Monotone Likelihood Ratio. *Annals of Mathematical Statistics* 27: 272–299.

[12] Manski, C.F. 1995. *Identification Problems in the Social Sciences.* Cambridge, Massachusetts: Harvard University Press.

[13] Manski, C.F. 2004. Statistical Treatment Rules for Heterogeneous Populations. *Econometrica* 72: 1221–1246.

[14] Manski, C.F. 2005. *Social Choice with Partial Knowledge of Treatment Response.* Princeton: Princeton University Press.

[15] Manski, C.F. 2007a. *Identification for Prediction and Decision.* Cambridge, Massachusetts: Harvard University Press.

[16] Manski, C.F. 2007b. Minimax-Regret Treatment Choice with Missing Outcome Data. *Journal of Econometrics* 139: 105–115.

[17] Manski, C.F. 2008a. Adaptive Partial Drug Approval. Working paper, Northwestern University.

[18] Manski, C.F. 2008b. Adaptive Partial Policy Innovation: Coping with Ambiguity through Diversification. Working paper, Northwestern University.

[19] Manski, C.F., and A. Tetenov. 2007. Admissible Treatment Rules for a Risk-Averse Planner with Experimental Data on an Innovation. *Journal of Statistical Planning and Inference* 137: 1998–2010.

[20] McFadden, D. 2006. How Consumers Respond to Incentives. *Jean-Jacques Laffont Lecture.* Toulouse, France. http://idei.fr/doc/conf/annual/paper_2006.pdf.

[21] Paditz, L. 1989. On the Analytical Structure of the Constant in the Nonuniform Version of the Esseen Inequality. *Statistics* 20: 453–464.

[22] Savage, L.J. 1951. The Theory of Statistical Decision. *Journal of the American Statistical Association* 46: 55–67.

[23] Schlag, K.H. 2007. Eleven - Designing Randomized Experiments under Minimax Regret. Working paper, European University Institute.

[24] Shiryaev, A.N. 1995. *Probability (2nd Edition).* New York: Springer Verlag.

[25] Stoye, J. 2007a. Minimax Regret Treatment Choice with Incomplete Data and Many Treatments. *Econometric Theory* 23: 190–199.

[26] Stoye, J. 2007b. Minimax Regret Treatment Choice with Finite Samples. Working paper, New York University.

[27] Stoye, J. 2007c. Minimax Regret Treatment Choice with Finite Samples and Missing Outcome Data. In *Proceedings of the Fifth International Symposium on Imprecise Probability: Theories and Applications*, ed. G. de Cooman, J. Veinarová, and M. Zaffalon. Prague, Czech Republic.

[28] 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.

[29] Wald, A. 1950. *Statistical Decision Functions.* New York: Wiley.