



Manski, C. F., & Tetenov, A. (2007). Admissible treatment rules for a risk-averse planner with experimental data on an innovation. *Journal of statistical planning and inference*, 137(6), 1998-2010.
<https://doi.org/10.1016/j.jspi.2006.06.037>

Peer reviewed version

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

Link to published version (if available):
[10.1016/j.jspi.2006.06.037](https://doi.org/10.1016/j.jspi.2006.06.037)

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

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via Elsevier at <http://www.sciencedirect.com/science/article/pii/S0378375806001820>. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

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

**ADMISSIBLE TREATMENT RULES FOR A RISK-AVERSE PLANNER
WITH EXPERIMENTAL DATA ON AN INNOVATION**

Charles F. Manski

Department of Economics and Institute for Policy Research, Northwestern University

and

Aleksey Tetenov

Department of Economics, Northwestern University

Revised: May 2006

The research presented here was supported in part by National Science Foundation grant SES-0314312. We have benefitted from the comments of two referees and from the opportunity to present this work in seminars at University College London and Vanderbilt University.

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 paper 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 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 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 $[x(1-x)^{-1}]^M g(x)$ 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 $[x(1-x)^{-1}]^M g(x)$ is strictly increasing. We also show that the class of KR-monotone rules is minimal complete if $f(\cdot)$ is strictly concave or $[x(1-x)^{-1}]g(x)$ is strictly increasing.

Section 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, Su, and Berry (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 yields 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 paper 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 5, which considers the implications of our analysis for treatment choice in practice.

2. Background

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 $y_j(\cdot): T \rightarrow Y$ mapping treatments $t \in T$ into outcomes $y_j(t) \in Y$.

The population is a probability space (J, Ω, P) , and the probability distribution $P[y(\cdot)]$ of the random function $y(\cdot): T \rightarrow Y$ 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 paper, outcomes are binary with $y_j(t) = 1$ denoting success and $y_j(t) = 0$ failure should person j receive 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 ζ of the population to treatment b and the remaining $1 - \zeta$ to treatment a . The population success rate under this fractional rule is

$$(1) \quad \alpha \cdot (1 - \zeta) + \beta \cdot \zeta = \alpha + (\beta - \alpha) \cdot \zeta.$$

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

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

The Empirical Evidence and Admissible Treatment Rules

Suppose that historical experience reveals α but not β . The available evidence on β 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 $z(\cdot): [0, \dots, N] \rightarrow [0, 1]$ 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

$$(2) \quad W(z; \beta) \equiv \sum_{n=0}^N p(n; \beta) \cdot f[\alpha + (\beta - \alpha) \cdot z(n)].$$

Expected welfare is a function of β , which is unknown. Let B index the values of β that the planner deems feasible. We assume that $0 < \alpha < 1$ and B includes at least one value smaller than α and at least one value greater than α . 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.

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

$$(3) \quad W(z; \beta) = \alpha + (\beta - \alpha) \cdot E_{\beta}[z(n)],$$

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

$$(4) \quad \begin{aligned} z(n) &= 0 && \text{for } n < k, \\ z(n) &= \lambda && \text{for } n = k, \\ z(n) &= 1 && \text{for } n > k, \end{aligned}$$

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 λ 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 \mathcal{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 \mathcal{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 rule. 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.

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 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 3.2 shows that the class of M -step monotone rules is essentially complete for all differentiable concave-monotone $f(\cdot)$ such that $[x(1-x)^{-1}]^M g(x)$ is weakly increasing in x . This class is

complete if $f(\cdot)$ is also strictly concave or $[x(1-x)^{-1}]^M g(x)$ is strictly increasing in x . Section 3.3 shows that the class of KR-monotone rules is minimal complete if $f(\cdot)$ is also strictly concave or if $[x(1-x)^{-1}]g(x)$ is strictly increasing in x . However, we show that KR-monotone rules can be inadmissible if $f(\cdot)$ is sufficiently curved.

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 β . It is therefore reasonable to conjecture that good treatment rules are 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 1 proves the result.

Proposition 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. \square

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^* :

$$z^*(n) = z^*(n') \equiv \frac{p(n; \alpha)}{p(n; \alpha) + p(n'; \alpha)} z(n) + \frac{p(n'; \alpha)}{p(n; \alpha) + p(n'; \alpha)} z(n'),$$

$$z^*(m) \equiv z(m), \forall m \notin \{n, n'\}.$$

For any value of β ,

$$(5) \quad W(z^*; \beta) - W(z; \beta) = p(n; \beta) \cdot \{f[\alpha + (\beta - \alpha) \cdot z^*(n)] - f[\alpha + (\beta - \alpha) \cdot z(n)]\} + \\ + p(n'; \beta) \cdot \{f[\alpha + (\beta - \alpha) \cdot z^*(n')] - f[\alpha + (\beta - \alpha) \cdot z(n')]\}.$$

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

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

The same inequality holds for $f[\alpha + (\beta - \alpha) \cdot z^*(n')]$. Substituting these inequalities into (5) and rearranging terms yields

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

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

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

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

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 $\beta \in B \setminus \{\alpha\}$.

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: $(n', n) = (1, 2), (1, 3), \dots, (1, N), (2, 3), (2, 4), \dots, (N-1, N)$. The result is a fractional monotone treatment rule that performs at least as well as z for all values of β and that dominates z if $f(\cdot)$ is strictly increasing.

Q. E. D.

Proposition 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.

3.2. M-step Monotone Rules

It appears that no result stronger than Proposition 1 can be proved without placing restrictions on the shape of $f(\cdot)$ beyond monotonicity and concavity. This section shows that Proposition 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 $[x(1-x)^{-1}]^M g(x)$ 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 $[x(1-x)^{-1}]^M g(x)$ strictly increases with x . Proposition 2 proves the result.

Proposition 2: Let $f(\cdot)$ be weakly increasing, concave and differentiable on $(\inf \{B\}, \sup \{B\})$, with derivative $g(\cdot)$. If $[x(1-x)^{-1}]^M g(x)$ weakly increases with x , then the M -step monotone rules are an essentially complete class. If $f(\cdot)$ is also strictly concave or if $[x(1-x)^{-1}]^M g(x)$ strictly increases with x , then the M -step monotone rules are a complete class. \square

Proof: Proposition 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 $[x(1-x)^{-1}]^M g(x)$ 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

$$(6) \quad \begin{aligned} z'(n) &= z(n) - p(n'; \alpha) \cdot p^{-1}(n; \alpha) \cdot [z'(n') - z(n)], \\ z'(n') &= \min [1, z(n') + p^{-1}(n'; \alpha) \cdot p(n; \alpha) \cdot z(n)]. \end{aligned}$$

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 $[x(1-x)^{-1}]^M g(x)$ strictly increases with x .

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

$$(7) \quad \begin{aligned} 0 \leq x < y \leq 1, \beta \neq \alpha &\Rightarrow f[\alpha + (\beta - \alpha) \cdot x] - f[\alpha + (\beta - \alpha) \cdot y] \geq g(\alpha) \cdot (\beta - \alpha) \cdot (x - y), \\ 0 \leq y < x \leq 1, \beta \neq \alpha &\Rightarrow f[\alpha + (\beta - \alpha) \cdot x] - f[\alpha + (\beta - \alpha) \cdot y] \geq g(\beta) \cdot (\beta - \alpha) \cdot (x - y). \end{aligned}$$

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

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

$$(8) \quad p(n'; x)p^{-1}(n; x)g(x) = [C(N, n')/C(N, n)] \cdot [x(1-x)^{-1}]^{n'-n-M} \cdot [x(1-x)^{-1}]^M g(x).$$

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 $[x(1-x)^{-1}]^M g(x)$ is strictly increasing, then so is $p(n'; x)p^{-1}(n; x)g(x)$.

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

$$(9) \quad \begin{aligned} W(z'; \beta) - W(z; \beta) &= p(n'; \beta) \cdot \{f[\alpha + (\beta - \alpha) \cdot z'(n')] - f[\alpha + (\beta - \alpha) \cdot z(n')]\} \\ &\quad + p(n; \beta) \cdot \{f[\alpha + (\beta - \alpha) \cdot z'(n)] - f[\alpha + (\beta - \alpha) \cdot z(n)]\} \\ &\geq p(n'; \beta) \cdot g(\beta) \cdot (\beta - \alpha) \cdot [z'(n') - z(n')] + p(n; \beta) \cdot g(\alpha) \cdot (\beta - \alpha) \cdot [z'(n) - z(n)] \\ &= (\beta - \alpha) \cdot [p(n'; \beta)p^{-1}(n; \beta)g(\beta) - p(n'; \alpha)p^{-1}(n; \alpha)g(\alpha)] \cdot p(n; \beta) \cdot [z'(n') - z(n')] \geq 0. \end{aligned}$$

The first inequality follows from (7). The second equality follows from (6). The last inequality holds for all β 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 $[x(1-x)^{-1}]^M g(x)$ 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: $(n, n') = (0, N), (0, N-1), \dots, (0, M), (1, N), (1, N-1), \dots, (1, 1+M), \dots, (N-M, N)$. 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 $n = 2, 3, \dots$, 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 β . Hence, the class of M -step monotone rules is essentially complete.

C. If $f(\cdot)$ is strictly concave or $[x(1-x)^{-1}]^M g(x)$ 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.

Q. E. D.

3.3. KR-Monotone Rules Revisited

The KR-monotone rules are the M-step monotone rules with $M = 1$. Thus, Proposition 2 shows that the class of KR-monotone rules is essentially complete if $[x(1-x)^{-1}]g(x)$ weakly increases with x . This class is complete if $f(\cdot)$ is strictly concave or if $[x(1-x)^{-1}]g(x)$ strictly increases with x . For example, Proposition 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 is $\log(x) - x$, whose derivative is $x^{-1}(1-x)$.

This section develops further properties of the KR-monotone rules. To begin, Proposition 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 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. \square

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:

$$(10) \quad \forall n: z(n) \geq z'(n) \text{ with strict inequality for some } n \Rightarrow W(z; \beta) > W(z'; \beta), \beta \in (\alpha, 1),$$

$$\forall n: z(n) \leq z'(n) \text{ with strict inequality for some } n \Rightarrow W(z; \beta) > W(z'; \beta), \beta \in (0, \alpha).$$

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.

Q. E. D.

Combining Propositions 2 and 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 3, Corollary: Let $f(\cdot)$ be strictly increasing and differentiable on $(\inf \{B\}, \sup \{B\})$, with $[x(1-x)^{-1}]g(x)$ weakly increasing in x . Then the class of KR-monotone rules is minimal complete if $f(\cdot)$ is strictly concave or if $[x(1-x)^{-1}]g(x)$ is strictly increasing. \square

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 α 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

$$(11) \quad \begin{aligned} z_{vw}(n) &= v \quad \text{for } n \leq k, \\ z_{vw}(n) &= w \quad \text{for } n > k. \end{aligned}$$

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

Proposition 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 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

$$(12a) \quad g(\alpha)/g(\beta_L) > [d_L/(1 - d_L)][(1 - w)/v],$$

$$(12b) \quad g(\alpha)/g(\beta_H) < [d_H/(1 - d_H)][(1 - w)/v].$$

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

$$(13a) \quad g[(1 - v)\alpha + v\beta_L]/g[(1 - w)\alpha + w\beta_L] < [d_L/(1 - d_L)][(1 - w)/v],$$

$$(13b) \quad g[(1 - v)\alpha + v\beta_H]/g[(1 - w)\alpha + w\beta_H] > [d_H/(1 - d_H)][(1 - w)/v]. \quad \square$$

Proof: By (2), the expected welfare of rules z_{01} and z_{vw} in the two feasible states of nature are as follows:

$$(14a) \quad W(z_{01}; \beta_L) = (1 - d_L)f(\alpha) + d_Lf(\beta_L),$$

$$(14b) \quad W(z_{01}; \beta_H) = (1 - d_H)f(\alpha) + d_Hf(\beta_H),$$

$$(14c) \quad W(z_{vw}; \beta_L) = (1 - d_L)f[(1 - v)\alpha + v\beta_L] + d_Lf[(1 - w)\alpha + w\beta_L],$$

$$(14d) \quad W(z_{vw}; \beta_H) = (1 - d_H)f[(1 - v)\alpha + v\beta_H] + d_Hf[(1 - w)\alpha + w\beta_H].$$

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 z_{01} 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 (14c) and (14d) as

$$(14c)' \quad W(z_{vw}; \beta_L) = (1 - d_L)f(\alpha) + d_L f(\beta_L) \\ + (1 - d_L)(\beta_L - \alpha)g[(1 - v_L)\alpha + v_L\beta_L]v + d_L(\beta_L - \alpha)g[(1 - w_L)\alpha + w_L\beta_L](w - 1),$$

$$(14d)' \quad W(z_{vw}; \beta_H) = (1 - d_H)f(\alpha) + d_H f(\beta_H) \\ + (1 - d_H)(\beta_H - \alpha)g[(1 - v_H)\alpha + v_H\beta_H]v + d_H(\beta_H - \alpha)g[(1 - w_H)\alpha + w_H\beta_H](w - 1),$$

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 (14a) and (14b) with (14c)' and (14d)' shows that rule z_{01} strictly dominates z_{vw} if and only if

$$(15a) \quad (1 - d_L)g[(1 - v_L)\alpha + v_L\beta_L]v + d_Lg[(1 - w_L)\alpha + w_L\beta_L](w - 1) > 0$$

and

$$(15b) \quad (1 - d_H)g[(1 - v_H)\alpha + v_H\beta_H]v + d_Hg[(1 - w_H)\alpha + w_H\beta_H](w - 1) < 0.$$

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

Whether (15a)-(15b) 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 (12a)-(12b) are the sufficient condition for rule z_{01} to strictly dominate z_{vw} and inequalities (13a)-(13b) are the sufficient condition for z_{vw} to strictly dominate z_{01} .

Q. E. D.

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 4.1 and 4.2 present some analytical findings for Bayes rules and the minimax-regret rule respectively. Section 5 will report some numerical findings for the minimax-regret and other rules.

4.1. Bayesian Planning

A Bayesian planner places a prior subjective probability distribution, say Π , 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 β as a random variable with distribution $\Pi(\beta|n)$, the planner then solves the problem

$$(16) \quad \max_{\zeta \in [0, 1]} \int f[\alpha + (\beta - \alpha)\zeta] d\Pi(\beta|n).$$

Proposition 5 shows that, given a regularity condition, the Bayes rule assigns the entire population to treatment a ($\zeta = 0$) if the posterior mean of β does not exceed α 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 5: Consider problem (16). Let $\Pi(\beta|n)$ be non-degenerate. Let $E_{\Pi(\beta|n)}[\beta]$ denote the posterior mean of β .

(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

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

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 $\int f(\beta)d\Pi(\beta|n) < f(\alpha)$. \square

Proof: (a) Strict concavity of $f(\cdot)$ implies that $\int f[\alpha + (\beta - \alpha)\zeta]d\Pi(\beta|n)$ is strictly concave in ζ . Hence, problem (16) 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 β . Hence, $\int f[\alpha + (\beta - \alpha)\zeta]d\Pi(\beta|n) < f[\alpha + (E_{\Pi(\beta|n)}[\beta] - \alpha)\zeta]$. Hence, $E_{\Pi(\beta|n)}[\beta] \leq \alpha \Rightarrow \int f[\alpha + (\beta - \alpha)\zeta]d\Pi(\beta|n) < f(\alpha)$.

(b) $\{\partial f[\alpha + (\beta - \alpha)\zeta]/\partial\zeta\}_{\zeta=0} = (\beta - \alpha) \cdot [df(x)/dx]_{x=\alpha}$. Hence, $\{\partial\{\int f[\alpha + (\beta - \alpha)\zeta]d\Pi(\beta|n)\}/\partial\zeta\}_{\zeta=0} = (E_{\Pi(\beta|n)}[\beta] - \alpha) \cdot [df(x)/dx]_{x=\alpha}$. By assumption, $E_{\Pi(\beta|n)}[\beta] > \alpha$ and $[df(x)/dx]_{x=\alpha} > 0$. Hence, $\int f[\alpha + (\beta - \alpha)\zeta]d\Pi(\beta|n)$ strictly increases with ζ in a neighborhood of $\zeta = 0$, implying that solutions to (16) are positive. If $\int f(\beta)d\Pi(\beta|n) < f(\alpha)$, then $\zeta = 1$ does not solve (16). Hence, solutions are fractional.

Q. E. D.

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.

4.2. Minimax-Regret Planning

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

$$(17) \quad R(z; \beta) \equiv \max [f(\alpha), f(\beta)] - W(z; \beta) = \sum_{n=0}^N p(n; \beta) \cdot \{\max [f(\alpha), f(\beta)] - f[\alpha + (\beta - \alpha) \cdot z(n)]\}.$$

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

$$(18) \quad \inf_{z \in Z} \sup_{\beta \in B} R(z; \beta).$$

(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 α , the maximin rule assigns the entire population to the status quo, whatever the sample size may be.)

Stoye (2005) 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

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

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 λ specifying the fraction of the population assigned to treatment b when there are k experimental successes. The sum $k + (1 - \lambda) = \sum_{n=0..N} [1 - z(n)]$ 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 $\sum_{n=0..N} [1 - z'(n)] < \sum_{n=0..N} [1 - z(n)]$, 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 $\sum_{n=0..N} [1 - z(n)]$, 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 (18). The same results hold for KR-monotone rules if $B = (0, 1)$ and $f(\cdot)$ satisfies the conditions of Proposition 2 for $M = 1$.

The situation is different if $f(\cdot)$ has strong curvature and B contains positive values arbitrarily close to 0. Then a minimax-regret rule never assigns the entire population to treatment b . Proposition 6 gives the result.

Proposition 6: Let $\alpha > 0$. Let B contain a sequence of positive values that converges to zero. If

$$(19) \quad \lim_{\beta \rightarrow 0_+} \beta^M \cdot f(\beta) = -\infty$$

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

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

This quantity is infinite because $\lim_{\beta \rightarrow 0} \beta^n (1 - \beta)^{N-n} \cdot f(\alpha)$ is finite, $\lim_{\beta \rightarrow 0} (1 - \beta)^{N-n} = 1$ and $\lim_{\beta \rightarrow 0} \beta^n \cdot f(\beta) = -\infty$ follows from (19). Hence, the minimax-regret rule must have $z_{\text{mmr}}(n) < 1$.

Q. E. D.

To illustrate, consider the welfare function $f(x) = -x^{-K}$, where $K > 1$. Then (19) holds for $M < K$ and $z_{\text{mmr}}(n) < 1$ for $n < K$. Consider the function $f(x) = -\exp(1/x)$. Then (19) holds for all values of M and $z_{\text{mmr}}(n) < 1$ for any n .

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.

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 $\{H_0: E[y(b)] = E[y(a)]\}$ 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 paper, 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

$$(20) \quad \begin{aligned} z(n) &= 0 & \text{for } n \leq d(s, \alpha), \\ z(n) &= 1 & \text{for } n > d(s, \alpha), \end{aligned}$$

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 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.

5.2. Comparing a Test-Based Rule with the Minimax-Regret Rule and the Plug-In Rule

Figure 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 α (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 1 also shows the maximum regret of the minimax-regret rule and of the empirical plug-in rule. The latter rule assigns the entire population to treatment b if the empirical rate of treatment success exceeds α , 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.

5.3. Variation of the Minimax-Regret Rule with the Welfare Function

Finally, we return to the question that most motivates this paper, namely how the welfare function affects treatment choice. The analysis of Sections 3 and 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 4) and the minimax-regret rule never assigns the entire population to treatment b (Proposition 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 3).

To shed some light on this, we compare the minimax-regret rule for the linear and log welfare functions. Figure 2 presents this comparison for the same values of α as in Figure 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 1. We do not present larger sample sizes in Figure 2 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 λ 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 means that $k = 2$ and $\lambda = 0.3$.

Figure 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 for α 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 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 one that varies smoothly with N . This similarity explains why, in Figure 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 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.

References

Canner, P. (1970), "Selecting One of Two Treatments When the Responses Are Dichotomous," *Journal of the American Statistical Association*, 65, 293-306.

Cheng, Y., F. Su, and D. Berry (2003), "Choosing Sample Size for a Clinical Trial Using Decision Analysis," *Biometrika*, 90, 923-936.

Fisher, L. and L. Moyé (1999), "Carvedilol and the Food and Drug Administration Approval Process: An Introduction," *Controlled Clinical Trials*, 20, 1-15.

Gould, A. (2002), "Substantial Evidence of Effect," *Journal of Biopharmaceutical Statistics*, 12, 53-77.

Karlin, S. and H. Rubin (1956), "The Theory of Decision Procedures for Distributions with Monotone Likelihood Ratio," *Annals of Mathematical Statistics*, 27, 272-299.

Manski, C. (2004), "Statistical Treatment Rules for Heterogeneous Populations," *Econometrica*, 72, 1221-1246.

Manski, C. (2005), *Social Choice with Partial Knowledge of Treatment Response*, Princeton: Princeton University Press.

Stoye, J. (2005), "Minimax-Regret Treatment Choice with Finite Samples," Department of Economics, New York University.

Wald, A. (1950), *Statistical Decision Functions*, New York: Wiley.

Figure 1: Maximum regret, $N=1..100$, $\alpha=0.25$, $f(x)=x$

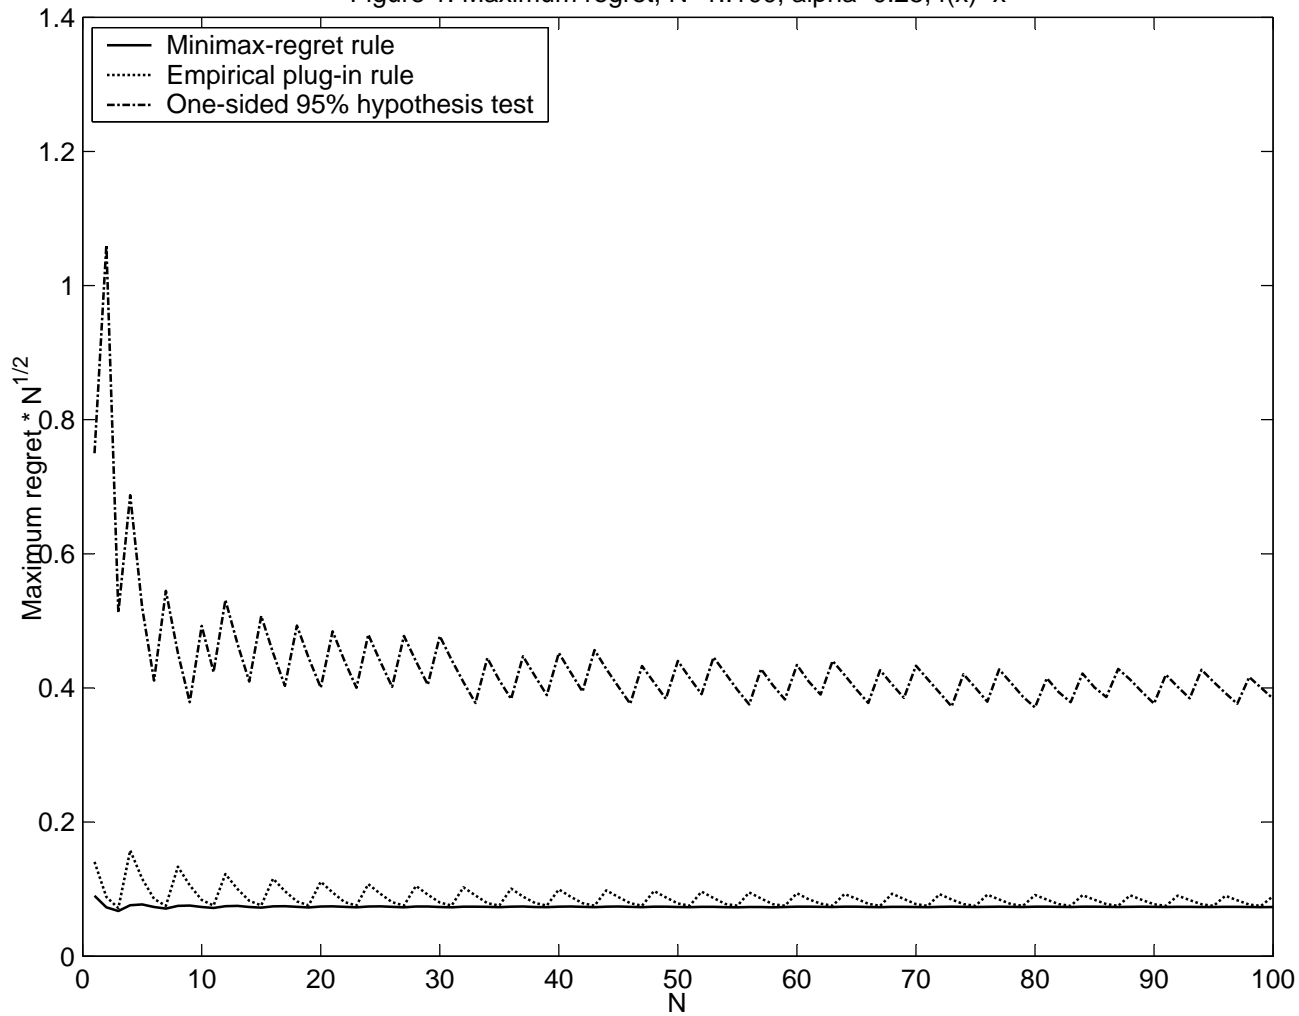


Figure 1: Maximum regret, $N=1..100$, $\alpha=0.75$, $f(x)=x$




Figure 1: Maximum regret, $N=1..100$, $\alpha=0.25$, $f(x)=\log(x)$

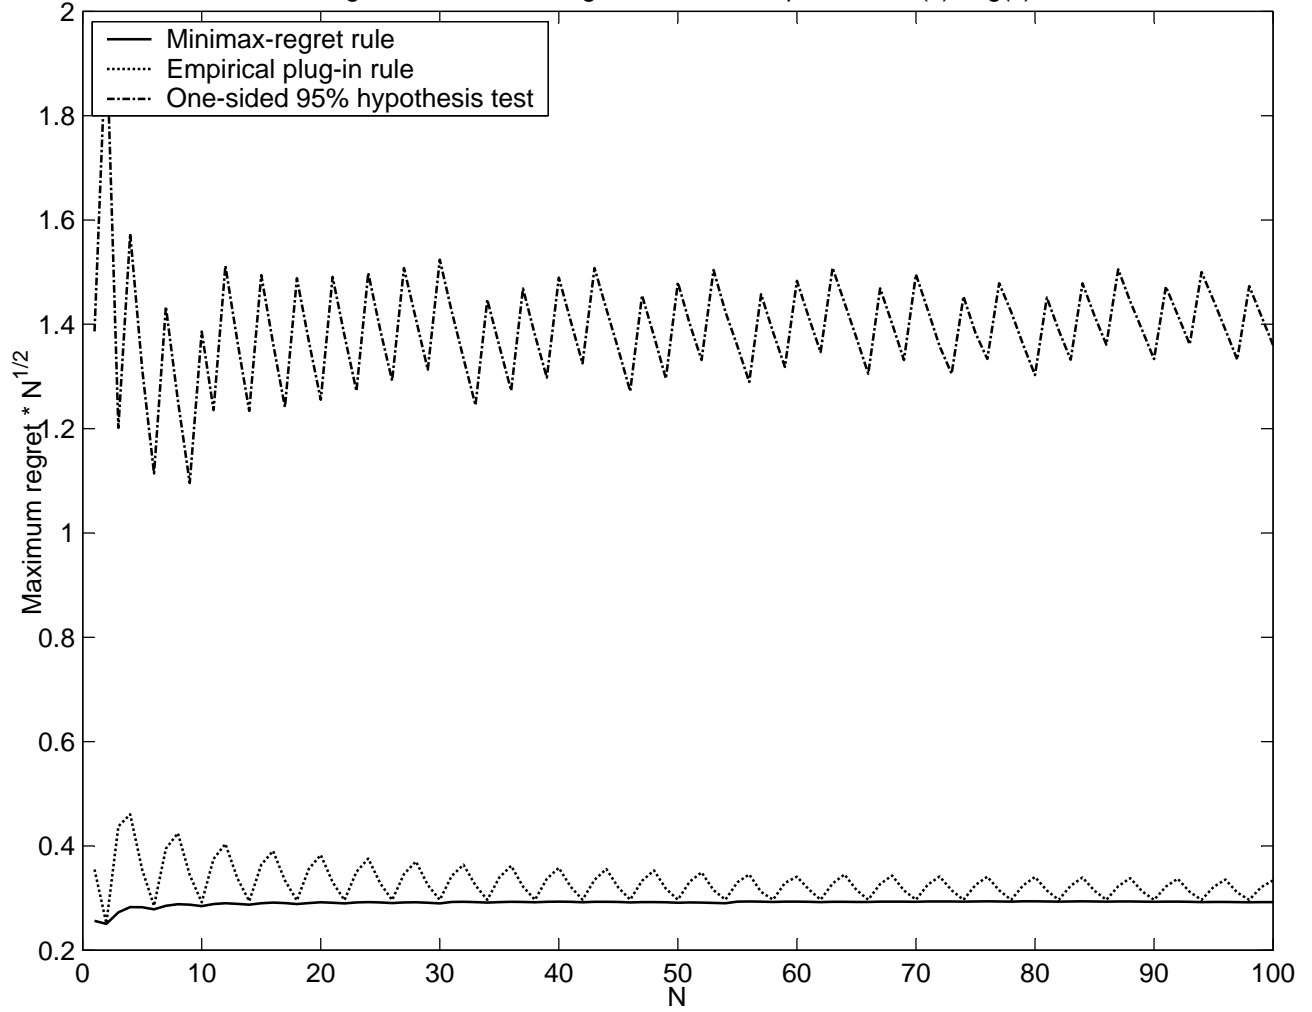


Figure 1: Maximum regret, $N=1..100$, $\alpha=0.75$, $f(x)=\log(x)$




Figure 2: Threshold sample size, $N=1..20$, $\alpha=0.25$




Figure 2: Threshold sample size, $N=1..20$, $\alpha=0.75$

