

Manski, Charles; Tetenov, Aleksey

Working Paper

The quantile performance of statistical treatment rules using hypothesis tests to allocate a population to two treatments

cemmap working paper, No. CWP44/14

Provided in Cooperation with:

Institute for Fiscal Studies (IFS), London

Suggested Citation: Manski, Charles; Tetenov, Aleksey (2014) : The quantile performance of statistical treatment rules using hypothesis tests to allocate a population to two treatments, cemmap working paper, No. CWP44/14, Centre for Microdata Methods and Practice (cemmap), London, <https://doi.org/10.1920/wp.cem.2014.4414>

This Version is available at:

<https://hdl.handle.net/10419/130006>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

The Quantile Performance of Statistical Treatment Rules Using Hypothesis Tests to Allocate a Population to Two Treatments

Charles Manski
Aleksey Tetenov

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP44/14



An ESRC Research Centre

THE QUANTILE PERFORMANCE OF STATISTICAL TREATMENT RULES USING HYPOTHESIS TESTS TO ALLOCATE A POPULATION TO TWO TREATMENTS

Charles F. Manski

Department of Economics and Institute for Policy Research, Northwestern University

and

Aleksey Tetenov

Collegio Carlo Alberto

November 2014

Abstract

This paper modifies the Wald development of statistical decision theory to offer new perspective on the performance of certain statistical treatment rules. We study the quantile performance of *test rules*, ones that use the outcomes of hypothesis tests to allocate a population to two treatments. Let λ denote the quantile used to evaluate performance. Define a test rule to be λ -quantile optimal if it maximizes λ -quantile welfare in every state of nature. We show that a test rule is λ -quantile optimal if and only if its error probabilities are less than λ in all states where the two treatments yield different welfare. We give conditions under which λ -quantile optimal test rules do and do not exist. A sufficient condition for existence of optimal rules is that the state space be finite and the data enable sufficiently precise estimation of the true state. Optimal rules do not exist when the state space is connected and other regularity conditions hold, but near-optimal rules may exist. These nuanced findings differ sharply from measurement of mean performance, as mean optimal test rules generically do not exist. We present further analysis that holds when the data are real-valued and generated by a sampling distribution which satisfies the monotone-likelihood ratio (MLR) property with respect to the average treatment effect. We use the MLR property to characterize the stochastic-dominance admissibility of STRs when the data have a continuous distribution and then generate findings on the quantile admissibility of test rules.

We have benefitted from the opportunity to present parts of this work at the January 2014 Workshop on Likelihood and Simplicity at Bar-Ilan University, the November 2014 Cemmap Conference on Microdata Methods and Practice, and in seminars at the Center for the Study of Rationality of the Hebrew University of Jerusalem, Northwestern University, the University of California at Berkeley, and University College London.

1. Introduction

This paper modifies one feature of the Wald (1950) development of finite-sample statistical decision theory to offer new perspective on the performance of certain statistical treatment rules. We study the quantile performance of rules that use the outcomes of hypothesis tests to allocate a population of observationally identical persons to two treatments.

Wald considered the broad problem of using informative sample data to make decisions under uncertainty. He posed the task as choice of a statistical decision function, which maps potentially available data into a choice among the feasible actions. He recommended ex ante evaluation of statistical decision functions as procedures, chosen prior to realization of the data, specifying how a decision maker would use whatever data may be realized. Expressing the objective as minimization of a loss function, he proposed that the decision maker evaluate a statistical decision function by its mean performance across realizations of the sampling process, which he termed risk.

In the presence of uncertainty about the loss function and the sampling process yielding the data, Wald prescribed a three-step decision process. The first stage specifies the state space (parameter space), which indexes the loss functions and sampling distributions that the decision maker deems possible. The second stage eliminates inadmissible statistical decision functions. A decision function is inadmissible (weakly dominated) if there exists another one that yields at least as good mean sampling performance in every possible state of nature and strictly better mean performance in some state. The third stage uses some criterion to choose an admissible statistical decision function. Wald considered the minimax criterion, but researchers have also studied other criteria such as minimax regret and minimization of a subjective mean of the risk function (Bayes risk).

Manski (2004, 2005), Manski and Tetenov (2007), Hirano and Porter (2009), Stoye (2009, 2012), and Tetenov (2012) have used the Wald framework to study how a planner might use sample data on

treatment response to choose treatments for the members of a population. In this setting, a statistical decision function uses the data to choose a treatment allocation, so such a function has been called a statistical treatment rule (STR). The planner's objective has been expressed as maximization of a social welfare function that sums treatment outcomes across the population.

An important special form of STR uses the outcome of an hypothesis test to allocate a population of observationally identical persons to two treatments. For example, data from a randomized clinical trial (RCT) may be used to inform allocation of a group of similar medical patients between a status quo treatment and an innovation. A common suggestion is to use a hypothesis test to choose between the two treatments. The null hypothesis is that the innovation is no better than the status quo and the alternative is that the innovation is better. Using the data to test the null, a health planner assigns all patients to the innovation if the null is rejected and all to the status quo treatment if the null is not rejected.

The U. S. Food and Drug Administration (FDA) uses such a test to decide on drug approval. A pharmaceutical firm wanting approval of a new drug (the innovation) performs RCTs that compare the new drug with an approved drug or placebo (the status quo). Approval of the new drug normally requires rejection of the null hypothesis in two independent trials (Fisher and Moyé, 1999).

How well do rules based on hypothesis tests perform? This question has been difficult to address within the Wald framework. Consider the basic matter of admissibility. The only findings to date concern a special setting in which (a) treatment outcomes are binary (success or failure), (b) the welfare achieved by a treatment allocation is a monotone and weakly concave function of its population-wide success rate, (c) the planner knows the population success rate that would occur if all persons were to receive the status quo treatment, and (d) the available data on response to the innovation is the empirical success rate in an RCT.

In this context, Manski (2005) observed that a theorem of Karlin and Rubin (1956) characterizes the set of admissible STRs when welfare is linear in the population success rate. Then the admissible

rules are the monotone step functions that assign everyone respectively to the innovation or status quo if the empirical success rate of the innovation in the RCT is above or below some specified threshold. This implies that a test-based STR, henceforth called a *test rule*, is admissible if and only if the test respectively rejects or does not reject the null hypothesis when the empirical success rate in the RCT is above or below a specified threshold. Subsequent work of Manski and Tetenov (2007) considers situations in which welfare is a concave-monotone function of the population success rate and finds that STRs that are monotone step-functions remain admissible if the welfare function has sufficiently weak curvature.

To provide new perspective on test rules, this paper modifies the criterion used by Wald to evaluate the sampling performance of statistical decision functions. Whereas Wald proposed measurement of mean performance across potential samples, we instead measure quantile performance. Econometricians and statisticians have long known that the mean and quantiles of a probability distribution provide interesting alternative point predictors of the realization of a random variable. Analogously, maximization of expected and quantile utility provide interesting alternative criteria for decision making under uncertainty.

Decision making using a quantile-utility criterion was proposed in Manski (1988) in a setting without sample data. It was observed there that maximization of expected and quantile utility differ in important respects. Whereas the ranking of actions by expected utility is invariant only to cardinal transformations of the objective function, the ranking by quantile utility is invariant to ordinal transformations. Whereas expected utility conveys risk preferences through the shape of the utility function, quantile utility does so through the specified quantile, with higher values conveying more risk preference. Whereas expected utility is not well-defined when the distribution of utility has unbounded support with fat tails, quantile utility is always well-defined.

It turns out to be much simpler to analyze the quantile performance of test rules than their mean performance. Whereas the mean performance of a test rule depends jointly on the probabilities and

magnitudes of errors in treatment assignment, quantile performance depends only on error probabilities. In this respect, quantile performance is similar to classical hypothesis testing, which also focuses exclusively on error probabilities. However, the quantile performance of test rules differs from classical testing in the way that it uses error probabilities.

After formally introducing concepts and notation in Section 2, we develop the primary analysis in Section 3. We present two polar findings, determined by the error probabilities of an hypothesis test. Let $\lambda \in (0, 1)$ denote the quantile used to evaluate performance. Define a test rule to respectively be λ -quantile optimal or minimal if it respectively maximizes or minimizes λ -quantile welfare in every state of nature. Proposition 1 shows that a test-based rule is λ -quantile optimal if and only if its error probabilities are less than λ in all states where the two treatments yield different welfare. Contrariwise, a test rule is λ -quantile minimal if and only if its error probabilities are greater than or equal to λ in all states where the treatments yield different welfare. The properties of optimality and minimality are much stronger than admissibility and inadmissibility.

We give conditions under which λ -quantile optimal test rules do and do not exist. We show that optimal rules always exist when $\lambda > 1/2$ and sometimes exist when $\lambda \leq 1/2$. The main positive finding in the latter setting (Proposition 2) implies that optimal rules exist when the state space is finite and the data enable sufficiently precise estimation of the true state. The main negative finding (Proposition 3) is that optimal rules do not exist when the state space is connected and other regularity conditions hold. In this setting near-optimal rules may exist. These nuanced findings differ sharply from measurement of mean performance, as mean optimal test rules generically do not exist.

A fundamental feature of the quantile performance of test rules is that all error probabilities symmetrically determine the results. In contrast, the classical theory of hypothesis testing differentiates between null and alternative hypotheses, and correspondingly between Type I and Type II errors. It restricts attention to tests that yield a predetermined probability of a Type I error (conventionally 0.05 or 0.01) and seeks a test of this type that yields an adequately small probability of a Type II error, typically

0.20. For example, an FDA document providing guidance for the design of RCTs evaluating new medical devices states that the probability of a Type I error is conventionally set to 0.05 and that the probability of a Type II error depends on the claim for the device but should not exceed 0.20 (U. S. Food and Drug Administration, 2014). The International Conference on Harmonisation (ICH) has provided similar guidance for the design of RCTs evaluating pharmaceuticals. The ICH document states the following (International Conference on Harmonization, 1999, p. 1923):

“Conventionally the probability of type I error is set at 5% or less or as dictated by any adjustments made necessary for multiplicity considerations; the precise choice may be influenced by the prior plausibility of the hypothesis under test and the desired impact of the results. The probability of type II error is conventionally set at 10% to 20%; it is in the sponsor’s interest to keep this figure as low as feasible especially in the case of trials that are difficult or impossible to repeat. Alternative values to the conventional levels of type I and type II error may be acceptable or even preferable in some cases.”

Such asymmetric treatment of the two hypotheses is illogical from the perspective of statistical decision theory. Proposition 4 shows that, given a test with predetermined size less than λ , a decision maker concerned with λ -quantile performance may be able to do better by shrinking the acceptance region for the null hypothesis to some extent and correspondingly enlarging the acceptance region for the alternative. For example, the FDA view that 0.05 and 0.20 are acceptable probabilities of Type I and Type II error suggests that if the agency were to assess new medical devices by quantile performance in RCTs, it might set $\lambda = 0.20$ and want to adopt a test rule that yields equal probabilities of Type I and Type II errors that are less than 0.20.

Section 4 presents further analysis that holds when the data are real-valued and generated by a sampling distribution which satisfies the monotone-likelihood ratio (MLR) property with respect to the average treatment effect. Statisticians and econometricians have long appreciated that the MLR property is mathematically benign. Karlin and Rubin (1956) and Manski and Tetenov (2007) have previously used it in their studies of the mean admissibility of STRs. In this paper we use it to characterize

admissibility of STRs when the data have a continuous distribution. Proposition 5 shows that any STR is weakly stochastically dominated by a STR that varies treatment assignment shares monotonically with the data. Proposition 6 then fully characterizes which fractional monotone STRs are quantile admissible.

2. Concepts and Notation

2.1. The Planning Problem

The setup is as in Manski (2004) and Manski and Tetenov (2007). A planner must assign one of two treatments to each member of a treatment population, denoted J . The feasible treatments are $T = \{a, b\}$. Each $j \in J$ has a response function $u_j(\cdot): T \rightarrow Y$ mapping treatments $t \in T$ into individual welfare outcomes $u_j(t) \in R$. Treatment is individualistic; that is, a person's outcome may depend on the treatment he is assigned but not on the treatments assigned to others. The population is a probability space (J, Ω, P) , and the probability distribution $P[u(\cdot)]$ of the random function $u(\cdot): T \rightarrow R$ describes treatment response across the population. The population is “large,” in the formal sense that J is uncountable and $P(j) = 0, j \in J$.

While treatment response may be heterogeneous, the members of the population are observationally identical to the planner. That is, the planner does not observe person-specific covariates that would enable systematic differentiation of treatment of different persons. However, the planner can randomly allocate persons to the two treatments with specified allocation probabilities.

A statistical treatment rule maps sample data into a treatment allocation. Let Q denote the sampling distribution generating the available data and let Ψ denote the sample space; that is, Ψ is the set of data samples that may be drawn under Q . Let Δ denote the space of functions that map $T \times \Psi$ into the

unit interval and that satisfy the adding-up conditions: $\delta \in \Delta \Rightarrow \delta(a, \psi) + \delta(b, \psi) = 1, \forall \psi \in \Psi$. Each function $\delta \in \Delta$ defines a statistical treatment rule, $\delta(a, \psi)$ and $\delta(b, \psi)$ being the fractions of the population assigned to treatments a and b when the data are ψ . Observe that this definition of an STR does not specify which persons receive each treatment, only the assignment shares. Designation of the particular persons receiving each treatment is immaterial because assignment is random, the population is large, and the planner has an additive welfare function. As $\delta(a, \psi) + \delta(b, \psi) = 1$, we use the shorthand $\delta(\psi)$ to denote the fraction assigned to treatment b. The fraction assigned to treatment a is $1 - \delta(\psi)$.

The planner wants to maximize population welfare, which adds welfare outcomes across persons. Given data ψ , the population welfare that would be realized if the planner were to choose rule δ is

$$(1) \quad U(\delta, P, \psi) \equiv E[u(a)] \cdot [1 - \delta(\psi)] + E[u(b)] \cdot \delta(\psi) \equiv \alpha \cdot [1 - \delta(\psi)] + \beta \cdot \delta(\psi),$$

where $\alpha \equiv E[u(a)] \equiv \int_j u_j(a) dP(j)$ and $\beta \equiv E[u(b)] \equiv \int_j u_j(b) dP(j)$ are assumed to be finite. Inspection of (1) shows that, whatever value ψ may take, it is optimal to set $\delta(\psi) = 0$ if $\alpha > \beta$ and $\delta(\psi) = 1$ if $\alpha < \beta$. All allocations are optimal if $\alpha = \beta$.

The problem of interest is treatment choice when knowledge of P and Q does not suffice to determine the ordering of α and β . Hence, the planner does not know the optimal treatment.

2.2. Evaluating STRs by their State-Dependent Welfare Distributions

The starting point for development of implementable criteria for treatment choice under uncertainty is specification of a state space, say S. Thus, let $\{(P_s, Q_s), s \in S\}$ be the set of (P, Q) pairs that the planner deems possible. The planner does not know the optimal treatment if S contains at least one state such that $\alpha_s > \beta_s$ and another such that $\alpha_s < \beta_s$. We assume this throughout.

Considered as a function of ψ , $U(\delta, P_s, \psi)$ is a random variable with state-dependent sampling distribution $Q_s[U(\delta, P_s, \psi)]$. Following Wald's view of statistical decision functions as procedures, we use the vector $\{Q_s[U(\delta, P_s, \psi)], s \in S\}$ of state-dependent welfare distributions to evaluate rule δ . In principle this vector is computable, whatever the state space and sampling process may be. Hence, in principle, a planner can compare the vectors of state-dependent welfare distributions yielded by different STRs and base treatment choice on this comparison.

How might a planner compare the state-dependent welfare distributions yielded by different STRs? The planner wants to maximize welfare, so it seems self-evident that he should weakly prefer rule δ to an alternative rule δ' if, in every $s \in S$, $Q_s[U(\delta, P_s, \psi)]$ equals or stochastically dominates $Q_s[U(\delta', P_s, \psi)]$. It is less obvious how one should compare rules whose state-dependent welfare distributions are not uniformly ordered in this manner.

Writing in the mid-twentieth century, Wald adopted the then ubiquitous practice of using expected utility (equivalently, expected loss) to evaluate actions that yield probability distributions of outcomes. Following Wald, recent research on treatment choice with sample data has evaluated STRs by their mean performance across realizations of the sampling process. The mean performance of rule δ in state s , denoted $W(\delta, P_s, Q_s)$, is

$$(2) \quad W(\delta, P_s, Q_s) = \alpha_s \cdot \{1 - E_s[\delta(\psi)]\} + \beta_s \cdot E_s[\delta(\psi)],$$

where $E_s[\delta(\psi)] \equiv \int_{\Psi} \delta(\psi) dQ_s(\psi)$ is the mean (across potential samples) fraction of persons who are assigned to treatment b .

Should a planner seek to maximize mean performance across potential samples rather than some other functional of the welfare distribution? A common view among economists has been that the von Neumann and Morgenstern (1944) and Savage (1954) axiomatic derivations of expected utility

maximization provide rationales to favor this criterion over others. However, subsequent developments in decision theory have called into question whether the axioms that yield expected utility maximization are as compelling as they once seemed. See, for example, Binmore (2009). Among other developments, Rostek (2010) has provided an axiomatic derivation of quantile utility maximization.

Eschewing axiomatic thinking, Manski (1988, 2011) argues that one may reasonably seek to maximize any functional of the welfare distribution that at least weakly respects stochastic dominance. One may, for example, measure performance by a quantile of the welfare distribution. To evaluate quantile performance, observe that welfare in state s may be written

$$(1') \quad U(\delta, P_s, \psi) = \alpha_s + (\beta_s - \alpha_s)\delta(\psi) = \beta_s + (\alpha_s - \beta_s)[1 - \delta(\psi)].$$

The invariance of quantiles to increasing continuous transformations implies that, for any $\lambda \in (0, 1)$, the λ -quantile of welfare in state s , denoted $V_\lambda(\delta, P_s, Q_s)$, is

$$(3) \quad \begin{aligned} V_\lambda(\delta, P_s, Q_s) &= \alpha_s + (\beta_s - \alpha_s) \cdot v_{\lambda s}[\delta(\psi)] && \text{if } \alpha_s \leq \beta_s, \\ &= \beta_s + (\alpha_s - \beta_s) \cdot v_{\lambda s}[1 - \delta(\psi)] && \text{if } \alpha_s \geq \beta_s, \end{aligned}$$

where $v_{\lambda s}[\delta(\psi)]$ and $v_{\lambda s}[1 - \delta(\psi)]$ denote the λ -quantile of $\delta(\psi)$ and $1 - \delta(\psi)$ in state s . Thus, welfare quantiles are one of two linear functions of quantiles of the treatment allocation.

Mean and quantile performance are obviously identical in states where $\alpha_s = \beta_s$. However, they generically differ in states where $\alpha_s \neq \beta_s$. Expressions (2) and (3) simplify when the STR is uniformly singleton. Rule δ is uniformly singleton if, for every possible data realization, δ assigns the entire population to one of the two treatments. Thus, for each $\psi \in \Psi$, either $\delta(\psi) = 0$ or $\delta(\psi) = 1$.

Consider a state with $\alpha_s \neq \beta_s$. Let $R_s(\delta)$ be the state-dependent probability that δ yields an error,

choosing the inferior treatment over the superior one. That is,

$$(4) \quad \begin{aligned} R_s(\delta) &= Q_s[\delta(\psi) = 0] && \text{if } \alpha_s < \beta_s, \\ &= Q_s[\delta(\psi) = 1] && \text{if } \alpha_s > \beta_s. \end{aligned}$$

It is the case that

$$(5) \quad W(\delta, P_s, Q_s) = \min(\alpha_s, \beta_s) \cdot R_s(\delta) + \max(\alpha_s, \beta_s) \cdot [1 - R_s(\delta)],$$

$$(6) \quad \begin{aligned} V_\lambda(\delta, P_s, Q_s) &= \min(\alpha_s, \beta_s) && \text{if } R_s(\delta) \geq \lambda, \\ &= \max(\alpha_s, \beta_s) && \text{if } R_s(\delta) < \lambda. \end{aligned}$$

Observe that mean and quantile performance share the property of being monotonically decreasing in the error probability, falling from $\max(\alpha_s, \beta_s)$ to $\min(\alpha_s, \beta_s)$ as $R_s(\delta)$ increases from 0 to 1. However, they differ substantially in the pattern of decrease. Whereas mean performance varies linearly with the error probability, quantile performance is a step function.

We may now define mean and quantile admissibility and inadmissibility. Rule δ is *mean inadmissible (admissible)* if there exists (does not exist) another rule δ' such that $W(\delta', P_s, Q_s) \geq W(\delta, P_s, Q_s)$ for all $s \in S$ and $W(\delta', P_s, Q_s) > W(\delta, P_s, Q_s)$ for some s . Analogously, δ is *λ -quantile inadmissible (admissible)* if there exists (does not exist) a δ' such that $V_\lambda(\delta', P_s, Q_s) \geq V_\lambda(\delta, P_s, Q_s)$ for all $s \in S$ and $V_\lambda(\delta', P_s, Q_s) > V_\lambda(\delta, P_s, Q_s)$ for some s .

Going further, we define quantile optimality and minimality. Rule δ is *λ -quantile optimal* if $V_\lambda(\delta, P_s, Q_s) = \sup_{\delta' \in \Delta} V_\lambda(\delta', P_s, Q_s) = \max(\alpha_s, \beta_s)$ for all $s \in S$. δ is *λ -quantile minimal* if $V_\lambda(\delta, P_s, Q_s) = \inf_{\delta' \in \Delta} V_\lambda(\delta', P_s, Q_s) = \min(\alpha_s, \beta_s)$ for all $s \in S$. λ -quantile optimality and minimality are stronger

properties than λ -quantile admissibility and inadmissibility. Every λ -quantile optimal rule is λ -quantile admissible. Every λ -quantile minimal rule is dominated by a data-invariant rule that sets either $\delta(\psi) = 0$ or $\delta(\psi) = 1$ for all $\psi \in \Psi$. Section 3 studies the circumstances in which λ -quantile optimal and minimal rules do and do not exist.

We can analogously define mean optimality and minimality, but these concepts are not useful in practice. Rule δ is *mean optimal* if $W(\delta, P_s, Q_s) = \max(\alpha_s, \beta_s)$ for all $s \in S$ and is *mean minimal* if $W(\delta, P_s, Q_s) = \min(\alpha_s, \beta_s)$ for all $s \in S$. These properties are achievable only if, for all $s \in S$, the sample data reveal the sign of $\beta_s - \alpha_s$ with probability one.

3. Quantile Optimality and Minimality of Test Rules

Construction of a test rule begins by partitioning the state space into disjoint subsets S_a and S_b , where S_a contains all states in which treatment a is optimal and S_b contains all states in which b is optimal. Thus, $\alpha_s > \beta_s \Rightarrow s \in S_a$, $\alpha_s < \beta_s \Rightarrow s \in S_b$, and the states with $\alpha_s = \beta_s$ are somehow split between the two sets. Let s^* denote the unknown true state. The two hypotheses are $[s^* \in S_a]$ and $[s^* \in S_b]$.

A test rule δ partitions the sample space Ψ into disjoint acceptance regions $\Psi_{\delta a}$ and $\Psi_{\delta b}$. When the data ψ lie in $\Psi_{\delta a}$, the rule accepts hypothesis $[s^* \in S_a]$ by setting $\delta(\psi) = 0$. When ψ lies in $\Psi_{\delta b}$, the rule accepts $[s^* \in S_b]$ by setting $\delta(\psi) = 1$. We use the word "accepts" rather than the traditional term "does not reject" because treatment choice is an affirmative action.

The above shows that test-based rules are uniformly singleton. Indeed, the converse holds as well. If δ is uniformly singleton, one can collect all of the data values for which the rule assigns everyone to treatment a, call this subset of the sample space the acceptance region $\Psi_{\delta a}$, and do likewise for $\Psi_{\delta b}$.

3.1. Basic Finding

Given the equivalence of uniformly singleton and test rules, we can study the quantile performance of uniformly singleton rules and apply the findings immediately to test rules. The basic finding is this elementary but powerful proposition:

Proposition 1: Let rule δ be uniformly singleton. δ is λ -quantile optimal if and only if $R_s(\delta) < \lambda$ for all $s \in S$ s. t. $\alpha_s \neq \beta_s$. δ is λ -quantile minimal if and only if $R_s(\delta) \geq \lambda$ for all $s \in S$ s. t. $\alpha_s \neq \beta_s$. \square

Proof: $V_\lambda(\delta, P_s, Q_s) = \max(\alpha_s, \beta_s) = \min(\alpha_s, \beta_s)$ when $\alpha_s = \beta_s$. Hence, it suffices to consider states where $\alpha_s \neq \beta_s$. Application of equation (6) to these states yields the results.

Q. E. D.

The proposition has strikingly contrasting implications for evaluation of test rules, dependent on the error probabilities. A rule with error probabilities uniformly less than λ is the best STR possible from the perspective of λ -quantile welfare. One with error probabilities uniformly greater than or equal to λ is the worst possible. Observe that all error probabilities symmetrically determine the results. The proposition does not distinguish Type I and Type II errors as in the classical theory of hypothesis testing.

A special but important class of hypothesis tests juxtaposes two simple hypotheses. Then the Neyman-Pearson Lemma shows that, among all tests with a specified probability of a Type I error, the likelihood-ratio test minimizes the probability of a Type II error, and vice versa. In the context of treatment choice, having two simple hypotheses means that S contains two states, with treatment a better in one state and b better in the other. Then the Neyman-Pearson Lemma implies that a planner considering use of a test rule need not look beyond the class of likelihood-ratio tests. Applying

Proposition 1 to likelihood ratio tests yields this corollary, which makes explicit the form of error probabilities for likelihood-ratio tests.

Corollary: Let $S = \{0, 1\}$, with $\alpha_0 > \beta_0$ and $\alpha_1 < \beta_1$. Let the sample data have distinct state-dependent sampling distributions Q_0 and Q_1 with either Lebesgue density or probability mass functions $q_0(\cdot)$ and $q_1(\cdot)$. Let $\eta \geq 0$ and let δ be the likelihood-ratio rule with threshold η ; thus, $\Psi_{\delta a} = [\psi \in \Psi: q_1(\psi) \leq \eta q_0(\psi)]$ and $\Psi_{\delta b} = [\psi \in \Psi: q_1(\psi) > \eta q_0(\psi)]$. Then δ is λ -quantile optimal if and only if $\max \{Q_0[q_1(\psi) > \eta q_0(\psi)], Q_1[q_1(\psi) \leq \eta q_0(\psi)]\} < \lambda$. δ is λ -quantile minimal if and only if $\min \{Q_0[q_1(\psi) > \eta q_0(\psi)], Q_1[q_1(\psi) \leq \eta q_0(\psi)]\} \geq \lambda$. \square

Proof: Rule δ has error probabilities $R_0(\delta) = Q_0[q_1(\psi) > \eta q_0(\psi)]$ and $R_1(\delta) = Q_1[q_1(\psi) \leq \eta q_0(\psi)]$.

Hence, the result is an immediate application of the proposition.

Q. E. D.

In practice, the sets S_a and S_b typically contain multiple elements; that is, they are composite rather than simple hypotheses. The Neyman-Pearson Lemma generically does not extend to tests of composite hypotheses, so a planner considering test rules may not want to restrict attention to rules based on likelihood-ratio tests. Nevertheless, a planner measuring λ -quantile performance will still want to determine if there exist test-based rules that are λ -quantile optimal and, if so, to choose such a test. The planner will not want to choose a rule that is λ -quantile minimal.

An obvious way to improve on a minimal rule is to choose a data-invariant rule, one that makes either $\Psi_{\delta a}$ or $\Psi_{\delta b}$ the entirety of the sample space. A planner who lets $\Psi_{\delta a} = \Psi$ always assigns the entire population to treatment a, regardless of the data. This rule has zero probability of error for all $s \in S_a$ and probability one of error for all $s \in S_b$. Hence, the rule is neither λ -quantile optimal nor minimal,

whatever value the planner specifies for λ . Letting $\Psi_{\delta b} = \Psi$ yields the analogous result.

A potentially more interesting way to improve on a minimal rule is to choose a uniformly fractional rule, one such that $0 < \delta(\psi) < 1$ for all $\psi \in \Psi$. In every state such that $\alpha_s \neq \beta_s$, such a rule always yields welfare that is strictly greater than $\min(\alpha_s, \beta_s)$ and strictly less than $\max(\alpha_s, \beta_s)$. Hence, the rule is neither λ -quantile optimal nor minimal.

When do there exist quantile optimal test rules? Sections 3.2 through 3.5 report some findings.

3.2. Existence of Quantile Optimal Test Rules

There exists an obvious λ -quantile optimal rule whenever $\lambda > \frac{1}{2}$. While one ordinarily thinks of ψ as data that are informative about treatment response, statistical decision theory also encompasses study of STRs that make treatment choice vary with uninformative data. That is, δ may make the treatment allocation depend on data generated by a randomizing device. Suppose in particular that $\Psi = \{0, 1\}$, $Q_s(\psi = 0) = Q_s(\psi = 1) = \frac{1}{2}$ for all $s \in S$, and δ is the rule that lets $\Psi_{\delta a} = \{0\}$ and $\Psi_{\delta b} = \{1\}$. The error probabilities for this test rule are $R_s(\delta) = \frac{1}{2}$ for all $s \in S$. Hence, δ is λ -quantile minimal for all $\lambda \leq \frac{1}{2}$ and λ -quantile optimal for all $\lambda > \frac{1}{2}$. See Sections 3.4 and 3.5 for further discussion of randomized test rules.

To the best of our knowledge, there exists no similarly obvious way to form a λ -quantile optimal rule when $\lambda \leq \frac{1}{2}$. In this domain, optimality becomes a more stringent condition as λ decreases and as the state space expands. It appears infeasible to perform an elementary general analysis, but we can make progress by examining particular contexts.

A preliminary finding is that there exists no λ -quantile optimal rule with $\lambda \leq \frac{1}{2}$ if the data are uninformative about the sign of the average treatment effect, in the sense that the sampling distribution does not vary with the sign of $\beta - \alpha$. Suppose that there exist two states 0 and 1 such that $\alpha_0 > \beta_0$, $\alpha_1 <$

β_1 , and $Q_0 = Q_1$. Then, for any test rule δ , the error probabilities in the two states satisfy $R_0(\delta) + R_1(\delta) = 1$. Hence, the error probability in one state is necessarily at least equal to $1/2$. Hence, quantile optimality is impossible for $\lambda \leq 1/2$.

Suppose that the sampling distribution of the data varies with the sign of the average treatment effect. Proposition 2 demonstrates that quantile optimal rules exist if there is a positive distance between the sets S_a and S_b (for example, if S is finite), and the data enable sufficiently precise estimation of the true state. In contrast, Proposition 3 shows that for $\lambda < 1/2$, no λ -quantile optimal test rule exists if the set S is connected and other regularity conditions hold. In combination, the two propositions show that quantile optimality is neither an empty concept nor ubiquitous. It is attainable by a test rule in some settings but not in others. In the setting of Proposition 3, we show by example that a test rule may be nearly optimal even though not exactly so.

Proposition 2: Let S be a subset of a metric space (Θ, d) with distance $d(\cdot, \cdot)$. Let

$$(7) \quad \varepsilon = \frac{1}{2} \cdot \min_{s \in S_a, s' \in S_b} d(s, s') > 0.$$

Suppose that an estimator $\check{s}(\cdot): \Psi \rightarrow \Theta$ is ε -far from the true state s with probability below λ :

$$(8) \quad Q_s[d(\check{s}(\psi), s) \geq \varepsilon] < \lambda \quad \text{for every } s \in S.$$

Then the minimum-distance test rule

$$(9) \quad \delta_{\text{md}}(\psi) \equiv 1[\min_{s \in S_b} d(\check{s}(\psi), s) < \min_{s \in S_a} d(\check{s}(\psi), s)]$$

is λ -quantile optimal. \square

Proof: It follows from the definition of ε that in every state $s \in S_a$,

$$d(\check{s}(\psi), s) < \varepsilon \Rightarrow d(\check{s}(\psi), s) < d(\check{s}(\psi), s') \text{ for all } s' \in S_b.$$

The same is true for $s \in S_b$ and $s' \in S_a$. Hence, a necessary condition for rule (8) to yield an error when s is the true state is that $d(\check{s}(\psi), s) \geq \varepsilon$. Hence,

$$R_s(\delta_{md}) \leq Q_s[d(\check{s}(\psi), s) \geq \varepsilon].$$

It follows from this and (8) that $R_s(\delta_{md}) < \lambda$ for every $s \in S$. Hence, the rule is optimal.

Q. E. D.

Remark: A sufficient condition for (7) to hold is that the average treatment effect $\beta_s - \alpha_s$ be uniformly continuous in s and bounded away from zero. If the state space S is finite, condition (8) is satisfied whenever $\check{s}(\cdot)$ is a weakly consistent estimator of the true state and the sample size is sufficiently large.

Proposition 3: Let S be a connected subset of a metric space (Θ, d) with distance $d(\cdot, \cdot)$. Let $S_{a>} \equiv \{s \in S: \alpha_s > \beta_s\}$ and $S_{b>} \equiv \{s \in S: \alpha_s < \beta_s\}$. Assume that the closure of the set $S_{a>} \cup S_{b>}$ is S ; that is, for any $s \in S$ and any $r > 0$, there exists $s' \in S_{a>} \cup S_{b>}$ such that $d(s, s') < r$. Let the probability $Q_s(\Psi_0)$ be continuous in s for every measurable subset of the sample space $\Psi_0 \subset \Psi$. Then no λ -quantile optimal test rule exists for $\lambda < 1/2$. \square

Proof: Let $\lambda < 1/2$ and suppose that test-rule δ is λ -quantile optimal. Let $s_a \in S_{a>}$ and $s_b \in S_{b>}$. Optimality requires that $Q_{s_a}(\Psi_{\delta b}) < \lambda$ and $Q_{s_b}(\Psi_{\delta b}) > 1 - \lambda$.

Given that $Q_s(\Psi_{\delta b})$ is continuous in s and that S is connected, there exists s^* such that $Q_{s^*}(\Psi_{\delta b}) = 1/2$. (See Rudin, 1976, Theorem 4.22). Continuity of $Q_s(\Psi_{\delta b})$ in s implies that there exists $r > 0$ such that

$d(s, s^*) < r \Rightarrow |Q_s(\Psi_{\delta b}) - Q_{s^*}(\Psi_{\delta b})| = |Q_s(\Psi_{\delta b}) - 1/2| < 1/2 - \lambda$. Given that $S = \text{cl}(S_{a>} \cup S_{b>})$, there exists either $s' \in S_{a>}$ or $s' \in S_{b>}$ such that $d(s', s^*) < r$. Hence, $|Q_{s'}(\Psi_{\delta b}) - 1/2| < 1/2 - \lambda$.

If $s' \in S_{a>}$, the optimality condition $Q_{s'}(\Psi_{\delta b}) < \lambda$ implies that $|Q_{s'}(\Psi_{\delta b}) - 1/2| = 1/2 - Q_{s'}(\Psi_{\delta b}) > 1/2 - \lambda$, which contradicts the conclusion reached above. If $s' \in S_{b>}$, the optimality condition $Q_{s'}(\Psi_{\delta b}) > 1 - \lambda$ implies that $|Q_{s'}(\Psi_{\delta b}) - 1/2| = Q_{s'}(\Psi_{\delta b}) - 1/2 > 1/2 - \lambda$, which again contradicts the conclusion reached above. Hence, δ is not λ -quantile optimal.

Q. E. D.

Remark: The state space has the required structure if a is a status-quo treatment with known mean outcome $\alpha^* \in (0, 1)$ and b is an innovation with mean outcome known to lie in the interval $(0, 1)$. Then $S = (0, 1)$, with $\alpha_s = \alpha^*$ and $\beta_s = s$ for $s \in S$. It is the case that $S_{a>} = (0, \alpha^*)$, $S_{b>} = (\alpha^*, 1)$, and $\text{cl}(S_{a>} \cup S_{b>}) = (0, 1)$. The sampling distribution has the required continuity if, for example, Q_s is Normal(s, k) for some fixed $k > 0$ or if Q_s is Binomial (n, s) for some integer n .

3.3. Existence of Near-Optimal Test Rules

Quantile optimality is a very strong property, so it should not be surprising that it is sometimes unattainable. In settings such as Proposition 3 where no test rule is exactly quantile optimal, there may nevertheless exist STRs that a planner deems acceptably close to optimal. In particular, there may exist a rule δ that is ε -optimal in the sense that $V_\lambda(\delta, P_s, Q_s) \geq \max(\alpha_s, \beta_s) - \varepsilon$ for all $s \in S$ and a specified $\varepsilon > 0$.

Consideration of ε -optimality may be particularly relevant to medical practice, which has long distinguished between the statistical and clinical significance of treatment effects. While the idea of clinical significance has been interpreted in various ways, many writers call an average treatment effect β

– α clinically significant if $|\beta - \alpha| > \varepsilon$ for a specified value of ε deemed minimally consequential in clinical practice. The ICH put it this way (International Conference on Harmonisation, 1999, p. 1923): “The treatment difference to be detected may be based on a judgement concerning the minimal effect which has clinical relevance in the management of patients.” A numerical example is given by Sedgwick (2014). Evaluating a topical ointment to prevent infection after minor surgery, he states (p. 1): “The smallest effect of clinical interest was an absolute decrease in the incidence of infection of 5%.” Thus, he specifies $\varepsilon = 0.05$ in a context where the outcome could take the value 0 (no infection) or 1 (infection).

The following Lemma shows that empirical success rules based on random samples of outcomes are ε -optimal when the sample sizes are sufficiently large.

Lemma: Let (n_a, n_b) be the sample sizes for each treatment and denote the independent treatment outcomes by $\psi = (y_{a,1}, \dots, y_{a,n_a}, y_{b,1}, \dots, y_{b,n_b})$. Let $E_s(y_a) = \alpha_s$ and $E_s(y_b) = \beta_s$. Denote the sample means by $m_a = (\sum_{i=1..n_a} y_{a,i})/n_a$ and $m_b = (\sum_{i=1..n_b} y_{b,i})/n_b$. Define the *empirical success rule* $\delta_{es}(\psi) \equiv 1[m_b > m_a]$. Consider any $\varepsilon > 0$ and $\lambda \in (0, 1)$.

(A). Assume that the variance of outcomes is uniformly bounded in all states s : $\text{Var}_s(y_t) \leq \nu$ for some finite ν . Then δ_{es} is λ -quantile ε -optimal if $4\nu\varepsilon^{-2}(1/n_a + 1/n_b) < \lambda$. If the design is balanced, with $n = n_a = n_b$, letting $n > 8\nu/(\varepsilon^2\lambda)$ ensures ε -optimality.

(B). Assume that outcomes are uniformly bounded in all states: wlog, let $0 \leq y_t \leq 1$. Then δ_{es} is λ -quantile ε -optimal if $2\exp(-1/2\varepsilon^2n_a) + 2\exp(-1/2\varepsilon^2n_b) < \lambda$. If the design is balanced, letting $n > -2\log(\lambda/4)/\varepsilon^2$ ensures ε -optimality. \square

Proof: (A) The variances of m_a and m_b are bounded from above by $\text{Var}_s(m_a) \leq \nu/n_a$ and $\text{Var}_s(m_b) \leq \nu/n_b$. By Chebyshev's inequality, $Q_s(|m_a - \alpha| \geq \varepsilon/2) \leq \text{Var}_s(m_a)/(\varepsilon/2)^2 \leq 4\nu/(\varepsilon^2n_a)$. Analogously, $Q_s(|m_b - \beta| \geq \varepsilon/2) \leq 4\nu/(\varepsilon^2n_b)$.

Consider states for which $|\beta_s - \alpha_s| \geq \varepsilon$. If $|(m_b - m_a) - (\beta - \alpha)| < \varepsilon$ for data realization ψ , then $U(\delta_{es}, P_s, \psi) = \max(\alpha_s, \beta_s)$. Hence, the error probability $R_s(\delta_{es})$ is bounded above as follows:

$$\begin{aligned} R_s(\delta_{es}) &\leq 1 - Q_s[|(m_b - m_a) - (\beta - \alpha)| < \varepsilon] \leq 1 - Q_s[|m_b - \beta| < \varepsilon/2 \text{ and } |m_a - \alpha| < \varepsilon/2] \\ &= Q_s[|m_b - \beta| \geq \varepsilon/2 \text{ or } |m_a - \alpha| \geq \varepsilon/2] \leq Q_s[|m_b - \beta| \geq \varepsilon/2] + Q_s[|m_a - \alpha| \geq \varepsilon/2] \\ &\leq 4v\varepsilon^{-2}(1/n_a + 1/n_b). \end{aligned}$$

It follows that $R_s(\delta_{es}) < \lambda$ for sufficiently large sample sizes (n_a, n_b) . Then $V_\lambda(\delta_{es}, P_s, Q_s) = \max(\alpha_s, \beta_s)$ for all s such that $|\beta_s - \alpha_s| \geq \varepsilon$. Moreover, regardless of sample size, $V_\lambda(\delta_{es}, P_s, Q_s) \geq \max(\alpha_s, \beta_s) - \varepsilon$ for all s such that $|\beta_s - \alpha_s| < \varepsilon$. Hence, δ_{es} is ε -optimal. With a balanced design, the threshold value of n that ensures ε -optimality is obtained by solving the equation $8v\varepsilon^{-2}/n = \lambda$.

(B) The large deviations inequality of Hoeffding (1963, Theorem 2) shows that $Q_s(|m_a - \alpha| \geq \varepsilon/2) \leq 2\exp(-1/2\varepsilon^2 n_a)$ and $Q_s(|m_b - \beta| \geq \varepsilon/2) \leq 2\exp(-1/2\varepsilon^2 n_b)$. Applying the same argument as part (A) with these inequalities yields the ε -optimality result. With a balanced design, the threshold value of n that ensures ε -optimality is obtained by solving the equation $4\exp(-1/2\varepsilon^2 n) = \lambda$.

Q. E. D.

To illustrate application of the lemma, consider the setting of Sedgwick (2014). He specifies $\varepsilon = 0.05$. The data are from an RCT with a balanced design. He considers a test with the conventional probabilities of Type I and Type II errors, namely 0.05 and 0.20. This suggest consideration of quantile performance with $\lambda = 0.20$.

Parts A and B of the lemma both apply in this setting. Either may yield the lower threshold sample size, depending on the value of v that one is able to specify. In the absence of prior knowledge of α and β , the uniform bound on variance is $v = 1/4$, so the Chebychev inequality yields the threshold sample size $2/(0.05^2 \times 0.2) = 4000$. The Hoeffding inequality yields a lower threshold sample size, this being $-2\log(0.05)/(0.05)^2 = 2397$. Thus, the lemma shows that a sufficient condition for the empirical success rule to be 0.2-quantile 0.05-optimal is that the trial have balanced sample sizes $n > 2397$.

Neither the Chebychev nor the Hoeffding inequality use the information that the outcome is binary in the Sedgwick setting. Bringing this information to bear can be shown to yield substantially smaller threshold sample sizes. An exact calculation using knowledge that the outcome is binary shows that the empirical success rule is 0.2-quantile 0.05-optimal if the trial has balanced sample sizes $n > 161$.

3.4. Randomized Test Rules

In Section 3.2 we observed that statistical decision theory encompasses study of STRs that make the treatment allocation depend on data generated by a randomizing device. Letting $\Psi = \{0, 1\}$ and $Q_s(\psi = 0) = Q_s(\psi = 1) = \frac{1}{2}$ for all $s \in S$, we showed that the randomized test rule with $\Psi_{\delta a} = \{0\}$ and $\Psi_{\delta b} = \{1\}$ is λ -quantile optimal for all $\lambda > \frac{1}{2}$. The only non-randomized rules available in this setting are the data-invariant rules that always choose one treatment or the other. These rules are not optimal.

In this section we consider randomized testing when $\lambda \leq \frac{1}{2}$. The simple randomized rule described earlier is minimal when $\lambda \leq \frac{1}{2}$, hence uninteresting. However, data generated by a randomizing device can still be useful when combined with informative data. An example suffices to demonstrate this.

To formalize randomized testing, we now let Ψ denote the sample space for informative data. We introduce a random variable v distributed $\text{Uniform}(0, 1)$, independent of ψ , to serve as a randomization device. Then the joint sample space is $\Psi \times [0, 1]$ and Q denotes the sampling distribution generating realizations of (ψ, v) . The shorthand $\delta(\psi, v)$ denotes the fraction of the population that rule δ assigns to treatment b. With this background, we present an example in which there exists a randomized λ -quantile optimal test rule but no non-randomized λ -quantile optimal test rule.

Example: Let $S = \{s_a, s_b\}$ consist of two states with $\beta_{s_a} < \alpha_{s_a}$ and $\beta_{s_b} > \alpha_{s_b}$. Let $\Psi = \{0, 1\}$ and let the sampling distributions Q_{s_a} and Q_{s_b} of the informative data be:

$$Q_{s_a}(\psi = 0) = 0.96, \quad Q_{s_a}(\psi = 1) = 0.04,$$

$$Q_{s_b}(\psi = 0) = 0.32, \quad Q_{s_b}(\psi = 1) = 0.68.$$

Consider a randomized test rule $\delta(\psi, v)$ with

$$\delta(0, v) = 1[v > 0.75], \quad \delta(1, v) = 1 \text{ for all } v.$$

The distribution of welfare $U(\delta, P_{s_a}, \psi)$ in state s_a is

$$Q_{s_a}[U(\delta, P_{s_a}, \psi) = \beta_{s_a}] = Q_{s_a}[\delta(\psi, v) = 1] = 0.96 \cdot 0.25 + 0.04 = 0.28,$$

$$Q_{s_a}[U(\delta, P_{s_a}, \psi) = \alpha_{s_a}] = Q_{s_a}[\delta(\psi, v) = 0] = 0.96 \cdot 0.75 = 0.72.$$

The distribution of welfare in state s_b is

$$Q_{s_b}[U(\delta, P_{s_b}, \psi) = \beta_{s_b}] = Q_{s_b}[\delta(\psi, v) = 1] = 0.32 \cdot 0.25 + 0.68 = 0.76,$$

$$Q_{s_b}[U(\delta, P_{s_b}, \psi) = \alpha_{s_b}] = 0.24.$$

For $\lambda = 0.3$, rule δ is λ -quantile optimal.

Now consider the space of non-randomized test rules; that is, rules in which $\delta(\psi, v)$ does not vary with the realization v . Given that $Q_{s_a}(\psi = 0) = 0.96$, a non-randomized rule can be 0.3-quantile optimal only if $\delta(0, v) = 0$ for all v , but it then follows that $U(\delta, P_{s_b}, \psi) = \alpha_{s_b}$ with probability 0.32. Hence, δ is not 0.3-quantile optimal. Thus, there does not exist any 0.3-quantile optimal non-randomized test rule. \square

The above discussion shows that randomization may be useful when evaluating STRs by their quantile performance. One may ask whether it is useful from the perspective of mean performance. The answer is negative. The mean welfare of a randomized STR $\delta(\psi, v)$ in state s equals

$$(2') \quad W(\delta, P_s, Q_s) = \alpha_s \cdot \{1 - E_s[\delta(\psi, v)]\} + \beta_s \cdot E_s[\delta(\psi, v)],$$

where $E_s[\delta(\psi, v)] \equiv \int_{\Psi} \int_{[0,1]} \delta(\psi, v) dv dQ_s(\psi)$. The same mean welfare is obtained by a non-randomized rule $\delta'(\psi) \equiv \int_{[0,1]} \delta(\psi, v) dv$ that averages treatment assignment fractions over v . Thus the class of non-randomized rules is essentially complete with regard to mean welfare.

3.5. Classical Testing with Predetermined Size

Whereas decision-theoretic evaluation of test rules considers all error probabilities symmetrically, classical hypothesis testing asymmetrically differentiates between null and alternative hypotheses, setting a predetermined probability of a Type I error (size). A necessary condition for a classical test to yield a λ -quantile optimal test rule is that the size γ of the test be set to a value less than λ . This done, the test rule is λ -quantile optimal if and only if the resulting probability of a Type II error is also less than λ .

An heuristic argument suggests that, given a test rule δ based on a test with size $\gamma < \lambda$, a decision maker concerned with λ -quantile performance can do better by shrinking the acceptance region for the null hypothesis to some extent and correspondingly enlarging the acceptance region for the alternative. Let S_a and S_b be the null and alternative hypotheses, with $\Psi_{\delta a}$ and $\Psi_{\delta b}$ being the corresponding acceptance regions for δ . As in Proposition 3, let $S_{a>} \equiv (s \in S: \alpha_s > \beta_s)$ and $S_{b>} \equiv (s \in S: \alpha_s < \beta_s)$. These are the only states relevant to decision making so we henceforth restrict attention to them.

A test with size γ has Type I error probability $(\sup R_s(\delta), s \in S_{a>}) = \gamma$ and Type II error probability $(\sup R_s(\delta), s \in S_{b>})$. Let δ' be an alternative rule based on a test where $\Psi_{\delta'a}$ is a subset of $\Psi_{\delta a}$ such that $\gamma < (\sup R_s(\delta'), s \in S_{a>}) < \lambda$. Then rules δ' and δ have the same optimal λ -quantile performance on $S_{a>}$, but rule δ' may outperform δ in terms of Type II error probability.

A simple way to formalize the heuristic argument is to let δ' be a randomized version of rule δ that uses the randomization device to appropriately shrink $\Psi_{\delta a}$ and enlarge $\Psi_{\delta b}$. Proposition 4 gives conditions under which a randomized rule dominates δ or even is λ -quantile optimal.

Proposition 4: Let $\delta(\psi) = 1[\psi \in \Psi_{\delta b}]$ be a test rule with size $\gamma < \lambda$. There exists randomized test rules that dominate δ if $R_s(\delta) \in [\lambda, \lambda \cdot (1-\gamma)/(1-\lambda))$ for some $s \in S_{b>}$. There exists randomized test rules that are λ -quantile optimal if $R_s(\delta) < \lambda \cdot (1-\gamma)/(1-\lambda)$ for all $s \in S_{b>}$. \square

Proof: Let v be distributed Uniform(0, 1), independent of ψ . Consider a randomized test rule

$$\delta'(\psi, v) \equiv 1[v \leq \psi] + \delta(\psi) \cdot 1[v > \psi],$$

where $v \in (0, 1)$. For $s \in S_{a>}$,

$$R_s(\delta') = Q_s[\delta'(\psi, v) = 1] = v + (1-v)R_s(\delta) \leq v + (1-v)\gamma = \gamma + v \cdot (1-\gamma).$$

Hence, $R_s(\delta') < \lambda$ for any $v < (\lambda - \gamma)/(1 - \gamma)$. Hence, $V_\lambda(\delta', P_s, Q_s) = V_\lambda(\delta, P_s, Q_s) = \alpha_s$.

For $s \in S_{b>}$,

$$R_s(\delta') = Q_s[\delta'(\psi, v) = 0] = (1-v)R_s(\delta) \leq R_s(\delta).$$

Hence, $V_\lambda(\delta', P_s, Q_s) \geq V_\lambda(\delta, P_s, Q_s)$.

The above shows that δ' has the same λ -quantile performance as δ on $S_{a>}$ and weakly outperforms δ on $S_{b>}$. Suppose that $R_s(\delta) \in [\lambda, \lambda \cdot (1-\gamma)/(1-\lambda))$ for some $s \in S_{b>}$. Then $R_s(\delta') = (1-v)R_s(\delta) < \lambda$ for any $v > 1 - \lambda/R_s(\delta)$. Hence, $V_\lambda(\delta', P_s, Q_s) = \beta_s > \alpha_s = V_\lambda(\delta, P_s, Q_s)$. Given the assumption that $R_s(\delta) < \lambda \cdot (1-\gamma)/(1-\lambda)$, it follows that $1 - \lambda/R_s(\delta) < (\lambda - \gamma)/(1 - \gamma)$. Hence, there exists v such that $1 - \lambda/R_s(\delta) < v < (\lambda - \gamma)/(1 - \gamma)$. Hence, rule δ' dominates δ in λ -quantile.

Finally, suppose that $R_s(\delta) \in [\lambda, \lambda \cdot (1-\gamma)/(1-\lambda))$ for all $s \in S_{b>}$. Then the above argument shows that $V_\lambda(\delta', P_s, Q_s) = \beta_s$ for all $s \in S_{b>}$. Hence, δ' is λ -quantile optimal.

Q. E. D.

4. Treatment Choice When Continuous Real Data Satisfy the Monotone-Likelihood Ratio Property in the Average Treatment Effect

The analysis of Section 3 placed little structure on the data used to inform treatment choice. Hence, the findings should be widely applicable. This section presents further analysis that holds when the data have considerable structure. Specifically, the data are assumed to be real-valued and generated by a sampling distribution that satisfies the monotone-likelihood ratio (MLR) property with respect to the average treatment effect ($\beta - \alpha$).

These assumptions are quite special, so it should not be expected that the findings of this section will have broad application. Nevertheless, treatment choice in this setting is intellectually worthy of study because the MLR property is mathematically benign. Karlin and Rubin (1956) and Manski and Tetenov (2007) have previously used it to study the mean admissibility of STRs. Here we use it to study admissibility in stochastic dominance, a strong property with immediate implications for quantile admissibility.

The analysis begins with Proposition 5, which shows that the fractional monotone treatment rules form an essentially complete class with respect to stochastic dominance when the data satisfy the maintained assumptions. A *fractional monotone* rule is one in which $\delta(\psi)$ is weakly increasing in ψ . *Essential completeness* means that any randomized decision rule $\delta(\psi, v)$ can be replaced by a fractional monotone rule $\delta'(\psi)$ that weakly stochastically dominates $\delta(\psi, v)$ in each state s . The planner then does not need to consider any other types of STRs. Manski and Tetenov (2007, Proposition 1) show that fractional monotone rules form an essentially complete class when the planner wants to maximize the expectation $E_s[f(U(\delta, P_s, \psi, v))]$ of a concave-monotone function $f(\cdot)$ of the population welfare and ψ is binomial. Here we find that a planner with any decision criterion that respects stochastic dominance can restrict attention to fractional monotone rules. Criteria that respect stochastic dominance include maximizing a quantile, the mean, or the mean of an increasing function of $U(\delta, P_s, \psi, v)$.

The choice problem posed in Proposition 5 is more general than that of choice between two

treatments, this because we found that the proposition holds more generally. After stating the proposition, we show that it applies to the treatment-choice problem of interest, to test rules among others. The proofs of all propositions and lemmas in this section are collected in an Appendix.

Proposition 5: Let $u(a, s)$ be the payoff function from action $a \in [a_l, a_h] \subset \mathbb{R}$ in state $s \in \mathcal{S}$. Assume that $u(a, s)$ is weakly monotonic in a for each s . For $s \in \mathcal{S}$, let the data $\psi \in \mathbb{R}$ have a continuous distribution $Q_s(\psi)$ and density $q_s(\psi)$ with respect to Lebesgue measure. Let $v \sim \text{Uniform}[0, 1]$ be a randomization variable independent of ψ . Assume that there exists a state s_0 for which $u(a, s_0)$ is constant in a .

Let $Q_s(\psi)$ possess the *monotone likelihood ratio property* for all pairs (s, s_0) such that $u(a, s)$ is not constant in a . That is,

$$(10a) \quad \text{if } u(a, s) \text{ is non-increasing in } a, \text{ then } q_s(\psi)/q_{s_0}(\psi) \geq q_s(\psi')/q_{s_0}(\psi') \text{ for all } \psi < \psi',$$

$$(10b) \quad \text{if } u(a, s) \text{ is non-decreasing in } a, \text{ then } q_s(\psi)/q_{s_0}(\psi) \leq q_s(\psi')/q_{s_0}(\psi') \text{ for all } \psi < \psi'.$$

Then for any randomized strategy $\delta(\psi, v): \Psi \times [0, 1] \rightarrow [a_l, a_h]$, there exists a monotone non-randomized strategy $\delta'(\psi): \Psi \rightarrow [a_l, a_h]$ whose distribution of payoffs $Q_s[u(\delta'(\psi), s)]$ weakly first order stochastically dominates the distribution of payoffs $Q_s[u(\delta(\psi, v), s)]$ of δ in each state s . δ' could be constructed by monotonically rearranging the values taken by $\delta(\psi, v)$ in state s_0 :

$$(11) \quad \delta'(\psi) \equiv G^{-1}_{\delta, s_0}(F_0(\psi)),$$

where $G^{-1}_{\delta, s_0}(\cdot)$ is the quantile function of the distribution $Q_{s_0}[\delta(\psi, v)]$ of the action $\delta(\psi, v)$ in state s_0 and $F_0(t) \equiv Q_{s_0}(\psi \leq t)$ is the c.d.f. of ψ in state s_0 . \square

Proposition 5 applies to the treatment-choice problem with action $a \in [0, 1]$ denoting the fraction of the population assigned to treatment b . Payoff function (1) is decreasing in a when $\beta_s - \alpha_s < 0$ and is increasing in a when $\beta_s - \alpha_s > 0$. If $Q_s(\psi)$ possesses the monotone likelihood ratio property in $\beta - \alpha$, then (10a) and (10b) hold. In this setting, $\delta'(\psi)$ defined in (11) is a fractional monotone treatment rule. The Corollary below applies the proposition to this setting.

Corollary: If $Q_s(\psi)$ is continuous and possesses the monotone likelihood ratio property in $(\beta - \alpha)$, then the class of fractional monotone STRs is essentially complete under any decision criterion that respects stochastic dominance. (A criterion respects stochastic dominance if rule δ' is weakly preferred to δ when the distribution of outcomes of δ' weakly stochastically dominates the distribution of outcomes of δ .) \square

Quantiles respect stochastic dominance, so the corollary applies. It implies that, for any rule δ , $V_\lambda(\delta', P_s, Q_s) \geq V_\lambda(\delta, P_s, Q_s)$ for all $s \in S$.

Thus far, the quantiles of fractional monotone treatment rules are abstract objects. Proposition 6 will provide an explicit characterization of these quantiles, facilitating analysis. As a prelude, we give a lemma that characterizes quantiles of monotone, possibly discontinuous, functions of continuously distributed random variables. The lemma is followed by a corollary applying it to the treatment-choice problem.

Lemma 1: Let X be a random variable with a continuous probability distribution P and $f(\cdot)$ a monotonic function. Denote the λ -quantile of $f(X)$ by $Q_\lambda[f(X)] = \inf \{g: P[f(X) \leq g] \geq \lambda\}$. If $f(\cdot)$ is non-decreasing, it equals the limit of $f(x)$ from the left at the λ -quantile of X :

$$(12a) \quad Q_\lambda[f(X)] = \lim_{x \rightarrow q^-} f(x),$$

where $q \equiv \inf [x: P(X \leq x) \geq \lambda]$. If $f(\cdot)$ is non-increasing, it equals the limit of $f(x)$ from the right at the upper $(1 - \lambda)$ -quantile of X :

$$(12b) \quad Q_\lambda[f(X)] = \lim_{x \rightarrow r^+} f(x),$$

where $r \equiv \sup [x: P(X \leq x) \leq 1 - \lambda]$. \square

Remark: In general, neither (12a) nor (12b) can be simplified by replacing the one-sided limits with the value of $f(x)$ at any x . For example, let X have a Uniform(0, 1) distribution and let $f(x) = x \cdot 1[x < 1/2] + 1[x \geq 1/2]$. X is continuous and $f(\cdot)$ is non-decreasing. The median $Q_{1/2}[f(X)]$ equals $1/2$, but there is no x at which $f(x) = 1/2$.

Corollary: If ψ has a continuous distribution in state s and $\delta(\psi)$ is non-decreasing in ψ , then for $\alpha_s \leq \beta_s$:

$$(13a) \quad V_\lambda(\delta, P_s, Q_s) = \alpha_s + (\beta_s - \alpha_s) \cdot \lim_{\psi \rightarrow q_{\lambda s}^-} \delta(\psi),$$

where $q_{\lambda s} \equiv \inf [x : Q_s(\psi \leq x) \geq \lambda]$ is the λ -quantile of ψ in state s . For states with $\alpha_s \geq \beta_s$:

$$(13b) \quad V_\lambda(\delta, P_s, Q_s) = \beta_s + (\alpha_s - \beta_s) \cdot [1 - \lim_{\psi \rightarrow r_{\lambda s}^+} \delta(\psi)],$$

where $r_{\lambda s} \equiv \sup [x : Q_s(\psi \leq x) \leq 1 - \lambda]$ is the upper $(1 - \lambda)$ -quantile of ψ in state s . \square

Remark: If $\alpha_s \leq \beta_s$ and $\delta(\cdot)$ is continuous at $q_{\lambda s}$, then

$$(14a) \quad V_\lambda(\delta, P_s, Q_s) = \alpha_s + (\beta_s - \alpha_s) \cdot \delta(q_{\lambda s}).$$

Similarly, if $\alpha_s \geq \beta_s$ and $\delta(\cdot)$ is continuous at $r_{\lambda s}$, then

$$(14b) \quad V_\lambda(\delta, P_s, Q_s) = \beta_s + (\alpha_s - \beta_s) \cdot (1 - \delta(r_{\lambda s})).$$

The corollary shows that, in each state s , the λ -quantile of welfare of every fractional monotone rule is determined by the one-sided limit of $\delta(\psi)$ at a particular point ($q_{\lambda s}$ or $r_{\lambda s}$) that is invariant across rules. This simplifies comparison of alternative monotone rules. Let $\psi^+ \equiv \{q_{\lambda s}, s: \beta_s - \alpha_s > 0\}$ and $\psi^- \equiv \{r_{\lambda s}, s: \beta_s - \alpha_s < 0\}$ denote the sets of comparison points. Proposition 6 characterizes the class of λ -quantile admissible monotone rules, first for the case in which the sets ψ^- and ψ^+ overlap or touch and then for the case in which they are distant from each other.

Proposition 6: Let $Q_s(\psi)$ be continuous and possess the monotone likelihood ratio property in $(\beta - \alpha)$.

(A). Suppose that $\inf \psi^- < \inf \psi^+ \leq \sup \psi^- < \sup \psi^+$ and that the sets ψ^- and ψ^+ are intervals on \mathbb{R} . Then a fractional monotone rule δ is λ -quantile admissible if and only if:

$$(15a) \quad \delta(\psi) = 0 \quad \text{for all } \psi < \inf \psi^+,$$

$$(15b) \quad \delta(\psi) = 1 \quad \text{for all } \psi > \sup \psi^-.$$

(B). Suppose that $\inf \psi^- \leq \sup \psi^- < \inf \psi^+ \leq \sup \psi^+$. Then a fractional monotone rule δ is λ -quantile admissible if and only if:

- (16a) $\delta(\psi) = 0$ for all $\psi < \sup \psi^-$,
- (16b) if $\sup \psi^- \in \psi^-$, then $\lim \delta(\psi) = 0$ for $\psi \rightarrow (\sup \psi^-)^+$,
- (16c) $\delta(\psi) = 1$ for all $\psi > \inf \psi^+$,
- (16d) if $\inf \psi^+ \in \psi^+$, then $\lim \delta(\psi) = 1$ for $\psi \rightarrow (\inf \psi^+)^-$. \square

Proposition 6 applies both to test rules and to ones that allocate positive fractions of the population to both treatments. It shows that an admissible rule allocates everyone to treatment a if the realized data value ψ is sufficiently small and everyone to b if ψ is sufficiently large. Allocations for intermediate values of ψ can be singleton or fractional, but the fraction allocated to b must weakly increase with ψ .

The reasoning underlying the proposition can be understood most easily when considering the subset of continuous decision rules $\delta(\cdot)$ for which characterization (14a, 14b) holds. If $\psi \in \psi^-$, then $V_\lambda(\delta, P_s, Q_s) = \beta_s + (\alpha_s - \beta_s) \cdot [1 - \delta(\psi)]$ in some state of nature with $\beta_s - \alpha_s < 0$. Hence, minimizing $\delta(\psi)$ increases λ -quantile welfare. Similarly, if $\psi \in \psi^+$, then $V_\lambda(\delta, P_s, Q_s) = \alpha_s + (\beta_s - \alpha_s) \cdot \delta(\psi)$ in some state of nature with $\beta_s - \alpha_s > 0$. Here maximizing $\delta(\psi)$ increases λ -quantile welfare. When seeking a rule $\delta'(\cdot)$ that dominates $\delta(\cdot)$, we have to take into account two constraints. First we have to preserve monotonicity. Second, if the same point ψ belongs to both sets $\psi \in \psi^-$ and $\psi \in \psi^+$, changing the value of $\delta(\psi)$ necessarily reduces λ -quantile welfare in at least one state of nature. Hence δ' can dominate δ only if $\delta'(\psi) = \delta(\psi)$ at such ψ .

Here is an example in which the conditions of Proposition 6 are satisfied.

Example: Let $[(\beta_s - \alpha_s), s \in S] = \mathbb{R}$ and let Q_s be $\text{Normal}(\beta_s - \alpha_s, 1)$. In this case, $q_{\lambda s} = (\beta_s - \alpha_s) + \Phi^{-1}(\lambda)$ and $r_{\lambda s} = (\beta_s - \alpha_s) + \Phi^{-1}(1 - \lambda)$. $\psi^+ = (\Phi^{-1}(\lambda), \infty)$ and $\psi^- = (-\infty, \Phi^{-1}(1 - \lambda))$.

Consider $\lambda \leq 1/2$. Then $\inf \psi^+ \leq \sup \psi^-$ and monotone treatment rules are admissible if and only if they satisfy conditions (15a) and (15b):

$$\delta(\psi) = 0 \quad \text{for all } \psi < \inf \psi^+ = \Phi^{-1}(\lambda),$$

$$\delta(\psi) = 1 \quad \text{for all } \psi > \sup \psi^- = \Phi^{-1}(1 - \lambda).$$

Consider $\lambda > 1/2$. Then $\sup \psi^- < \inf \psi^+$ and monotone treatment rules are admissible if and only if they satisfy conditions (16a)–(16d). Conditions (16b) and (16d) are not binding because the sets ψ^- and ψ^+ are not closed. Conditions (16a) and (16c) are:

$$\delta(\psi) = 0 \quad \text{for all } \psi < \sup \psi^- = \Phi^{-1}(1 - \lambda),$$

$$\delta(\psi) = 1 \quad \text{for all } \psi > \inf \psi^+ = \Phi^{-1}(\lambda).$$

For the median, $\lambda = 1/2$, we obtain the stark result that there is essentially only one admissible monotone rule, with $\delta(\psi) = 0$ for all $\psi < 0$ and $\delta(\psi) = 1$ for all $\psi > 0$. This is also the only rule admissible for all quantiles.

Both mean admissible and quantile admissible rules are increasing in ψ , but the two criteria have somewhat different implications. The class of mean-admissible rules in this example is the class of all threshold rules $\delta(\psi) = 1[\psi > t]$, $t \in \mathbb{R}$. Threshold rules with $t \in [\min\{\Phi^{-1}(1 - \lambda), \Phi^{-1}(\lambda)\}, \max\{\Phi^{-1}(1 - \lambda), \Phi^{-1}(\lambda)\}]$ are λ -quantile admissible, but those with $t < \min\{\Phi^{-1}(1 - \lambda), \Phi^{-1}(\lambda)\}$ or $t > \max\{\Phi^{-1}(1 - \lambda), \Phi^{-1}(\lambda)\}$ are not. \square

We would like to determine whether Propositions 5 and 6 can be extended to settings in which the data have discrete support. A case of some applied interest is that in which outcomes are binary and the data, arising from a randomized trial, have a Binomial distribution. These extensions appear technically challenging.

5. Conclusion

Statisticians and econometricians have devoted considerable attention to development of the classical theory of hypothesis tests. This paper explores a distinct perspective on tests as a class of statistical treatment rules that use acceptance regions to assign a population to one of two treatments.

It has been difficult to study the mean performance of test rules, but we have found it productive to study their quantile performance. We view the paper as achieving two major findings. One, initiated in Propositions 1 and fleshed out in Propositions 2 through 4, is that the concept of quantile optimality has content. The other, developed in Proposition 5 and applied in Proposition 6, is that the concept of stochastic dominance admissibility has content.

We see considerable scope for further analysis in multiple directions. While exact quantile optimality is too strong a property to have general applicability, the concept of near optimality may be broadly useful. It may be particularly helpful to medical decision making, which already uses the idea of clinical significance. Hence, we see good reason to continue our study of near optimality, adding to the Lemma and initial application given in Section 3.3.

When multiple optimal or acceptable near-optimal STRs exist, a planner must choose among them. Contrariwise, when no acceptable near-optimal rule exists, a planner must choose among the admissible rules. Previous research has studied classical decision criteria (maximin, minimax-regret, Bayes) when STRs are evaluated by mean sampling performance. We think it important to learn the properties of these criteria when rules are evaluated by quantile performance.

Finally, we think it important to extend the study of quantile performance to settings more complex than allocation of observationally identical persons to two treatments. Suppose that members of the population can be treated differentially conditional on observable covariates and/or that the number of feasible treatments is larger than two. The concept of an STR extends immediately to these settings.

Let X be the covariate space and let Δ denote the space of functions that map $T \times X \times \Psi$ into the unit interval, with $\delta \in \Delta \Rightarrow \sum_{t \in T} \delta(t, x, \psi) = 1, \forall (x, \psi) \in X \times \Psi$. Then each $\delta \in \Delta$ defines an STR and δ is uniformly singleton if $\delta(t, x, \psi)$ only takes the values 0 and 1. However, the quantile welfare of an STR does not have a simple expression comparable to (3), which was the starting point for the analysis of this paper. Hence, study of quantile performance appears challenging.

Appendix: Proofs of Findings in Section 4

Proof of Proposition 5:

First we show that the non-randomized strategy $\delta'(\psi)$ defined in (11) re-arranges the values of $\delta(\psi, v)$ to be increasing in ψ and has the same probability distribution of actions (and hence payoffs) as $\delta(\psi, v)$ in state s_0 . Then we show that $Q_s[u(\delta'(\psi), s)]$ weakly stochastically dominates $Q_s[u(\delta(\psi, v), s)]$ in all other states of nature.

Given that ψ has a continuous distribution, random variable $F_0(\psi)$ has a Uniform(0, 1) distribution in state s_0 . Hence, random variable $\delta'(\psi) = G^{-1}_{\delta, s_0}(F_0(\psi))$ has c.d.f. G_{δ, s_0} in state s_0 .¹ Given that both $G^{-1}_{\delta, s_0}(\cdot)$ and $F_0(\cdot)$ are non-decreasing, $\delta'(\psi)$ is also non-decreasing in ψ . Given that F_0 is continuous and G^{-1}_{δ, s_0} is left-continuous, $\delta'(\psi)$ is also left-continuous.

In states where $u(a, s)$ is constant in a , the distributions of payoffs are identical for all strategies. Hence, weak stochastic dominance holds. Now suppose that state s satisfies (10a), so $u(a, s)$ is non-increasing in a . (The proof is analogous for states in which $u(a, s)$ is non-decreasing in a .)

We want to show that the distribution of $\delta'(\psi)$ is weakly stochastically dominated by the

¹ Let Q and F be the quantile and distribution functions of a random variable. It is the case that, for all u in $(0, 1)$ and all real t (see, for example, Pfeiffer, 1990, p.266) $Q(u) \leq t \Leftrightarrow u \leq F(t)$. If u is itself random with distribution P , it follows that $P[Q(u) \leq t] = P[u \leq F(t)]$. If u is uniform, $P[u \leq F(t)] = F(t)$.

distribution of $\delta(\psi, v)$. Denote the c.d.f. of action $\delta(\psi, v)$ in state s by

$$G_{\delta,s}(t) \equiv Q_s[u(\delta(\psi, v), s)] = \int q_s(\psi) \int 1[\delta(\psi, v) \leq t] dv d\psi.$$

Given any $t \in [a_l, a_h]$, consider the indicator functions $1[\delta(\psi, v) \leq t]$ and $1[\delta'(\psi) \leq t]$. These indicator functions generate rejection regions for classical hypothesis tests with null hypothesis s_0 and alternative hypothesis s . A randomized test with rejection region $\Omega \equiv \{(\psi, v): \delta(\psi, v) \leq t\}$ has power function $G_{\delta,s}(t)$ as a function of s . Similarly, $G_{\delta',s}(t) \equiv \int q_s(\psi) \cdot 1[\delta'(\psi) \leq t] d\psi$ is the power function (as a function of s) of a non-randomized test with rejection region $\Omega' \equiv \{\psi: \delta'(\psi) \leq t\}$. We have shown above that the two tests have equal power in state s_0 : $G_{\delta',s_0}(t) = G_{\delta,s_0}(t)$. Given that $\delta'(\psi)$ is non-decreasing in ψ , $1[\delta'(\psi) \leq t]$ is non-increasing in ψ and there exists ψ_t such that

$$1[\delta'(\psi) \leq t] = 1 \quad \text{for all } \psi < \psi_t,$$

$$1[\delta'(\psi) \leq t] = 0 \quad \text{for all } \psi > \psi_t.$$

Given that state s satisfies (10a), the test with rejection region $\Omega' = (\psi: \psi \leq \psi_t)$ is a likelihood-ratio test. The tests with rejection regions Ω and Ω' have the same size. It follows from the Neyman-Pearson lemma that test Ω' must be at least as powerful as Ω in state s .² That is, $G_{\delta',s}(t) \geq G_{\delta,s}(t)$.

We can thus establish that $G_{\delta',s}(t) \geq G_{\delta,s}(t)$ for all t . Hence, the distribution of $\delta(\psi, v)$ weakly stochastically dominates the distribution of $\delta'(\psi)$. Given that $u(a, s)$ is a weakly decreasing function of a , $Q_s[u(\delta'(\psi), s)]$ weakly stochastically dominates $Q_s[u(\delta(\psi, v), s)]$.

Q. E. D.

² For a version that covers randomized tests see, for example, Lehmann and Romano (2008), Theorem 3.2.1.

Proof of Lemma 1:

We show the proof for non-increasing $f(\cdot)$. The proof for non-decreasing functions is analogous.

First, we establish that $Q_\lambda[f(X)] \leq \lim_{x \rightarrow r^+} f(x)$. Consider any $x^* > r$. From monotonicity of $f(\cdot)$ it follows that $\lim_{x \rightarrow r^+} f(x) \geq f(x^*)$. Hence,

$$P[f(X) \leq \lim_{x \rightarrow r^+} f(x)] \geq P[f(X) \leq f(x^*)] \geq P(X \geq x^*).$$

The second inequality holds because $x \geq x^* \Rightarrow f(x) \leq f(x^*)$. Given that the above inequality holds for all $x^* > r$, it holds for the limit from the right:

$$P[f(X) \leq \lim_{x \rightarrow r^+} f(x)] \geq \lim_{x \rightarrow r^+} P(X \geq x).$$

Given that $P(X)$ is continuous, $\lim_{x \rightarrow r^+} P(X \geq x) = \lambda$. Therefore, $P[f(X) \leq \lim_{x \rightarrow r^+} f(x)] \geq \lambda$ and $Q_\lambda[f(X)] \leq \lim_{x \rightarrow r^+} f(x)$.

Now suppose that $\lim_{x \rightarrow r^+} f(x) > Q_\lambda[f(X)]$. Then there exists some $x^* > r$ for which $f(x^*) > Q_\lambda[f(X)]$. This implies, from the definition of $Q_\lambda[f(X)]$, that $P[f(X) \leq f^*] \geq \lambda$ for some $f^* < f(x^*)$. Then

$$\lambda \leq P[f(X) \leq f^*] \leq P[f(X) < f(x^*)] \leq P(X > x^*) \leq P(X \geq x^*) < \lambda.$$

The second inequality holds because $f(x) < f(x^*) \Rightarrow x > x^*$. The last inequality holds from the definition of r , continuity of X , and $x^* > r$. We have arrived at a contradiction. Hence, $Q_\lambda[f(X)] = \lim_{x \rightarrow r^+} f(x)$.

Q. E. D.

Proof of Proposition 6:

(A) Suppose that δ does not satisfy (15a) for some $\psi' < \inf \psi^+$. Then $\delta(\psi) > 0$ on the interval $(\psi', \inf \psi^+)$ because $\delta(\cdot)$ is non-decreasing. There is at most a countable set of points ψ at which $\delta(\cdot)$ is discontinuous, so there exists $\psi^* \in (\max(\psi', \inf \psi^-), \inf \psi^+)$ such that $\delta(\cdot)$ is continuous at ψ^* . Let ψ^{**} be any point in the open interval $(\psi^*, \inf \psi^+)$.

Decision rule $\delta(\cdot)$ is inadmissible because it is λ -quantile dominated by

$$\begin{aligned} \delta'(\psi) &= 0 && \text{for } \psi \leq \psi^{**}, \\ &= \delta(\psi) && \text{for } \psi > \psi^{**}. \end{aligned}$$

Given that $\inf \psi^- < \psi^* < \inf \psi^+ \leq \sup \psi^-$ and ψ^- is an interval, there exists a state s such that $\psi^* = r_{\lambda s}$ and $\beta_s - \alpha_s < 0$. Then

$$\begin{aligned} V_\lambda(\delta', P_s, Q_s) &= \beta_s + (\alpha_s - \beta_s) \cdot (1 - \delta'(\psi^*)) = \alpha_s \\ &> \beta_s + (\alpha_s - \beta_s) \cdot (1 - \delta(\psi^*)) = V_\lambda(\delta, P_s, Q_s), \end{aligned}$$

because $\delta(\psi^*) > 0$, both $\delta(\cdot)$ and $\delta'(\cdot)$ are continuous at ψ^* and (14b) holds.

It remains to show that $V_\lambda(\delta', P_s, Q_s) \geq V_\lambda(\delta, P_s, Q_s)$ in every state s . For states with $\beta_s - \alpha_s < 0$, it is true because $\delta'(\psi) \leq \delta(\psi)$ for all ψ . Take any state of nature s with $\beta_s - \alpha_s > 0$. Since $\psi^{**} < \inf \psi^+ \leq q_{\lambda s}$, $\delta'(\psi) = \delta(\psi)$ in the neighborhood of $q_{\lambda s}$. Therefore, $\lim_{\psi \rightarrow q_{\lambda s}^-} \delta'(\psi) = \lim_{\psi \rightarrow q_{\lambda s}^-} \delta(\psi)$. It follows from (13a) that $V_\lambda(\delta', P_s, Q_s) = V_\lambda(\delta, P_s, Q_s)$.

The proof that condition (15b) is necessary is analogous.

Now we will establish the sufficiency of conditions (15a) and (15b) for admissibility. Suppose that δ satisfies them. For δ to be inadmissible, there must be a fractional monotone rule δ' such that $V_\lambda(\delta', P_s, Q_s) \geq V_\lambda(\delta, P_s, Q_s)$ in every state s and $V_\lambda(\delta', P_{s^*}, Q_{s^*}) > V_\lambda(\delta, P_{s^*}, Q_{s^*})$ in some state s^* . We will show

that if $V_\lambda(\delta', P_{s^*}, Q_{s^*}) > V_\lambda(\delta, P_{s^*}, Q_{s^*})$ in some state s^* , then necessarily $V_\lambda(\delta', P_{s'}, Q_{s'}) < V_\lambda(\delta, P_{s'}, Q_{s'})$ in some other state s' .

Suppose that $\beta_{s^*} - \alpha_{s^*} > 0$. Then it must be that $\lim_{\psi \rightarrow q_{\lambda s^*}^-} \delta'(\psi) > \lim_{\psi \rightarrow q_{\lambda s^*}^-} \delta(\psi)$. Either $q_{\lambda s^*} > \sup \psi^-$ or $q_{\lambda s^*} \leq \sup \psi^-$. If $q_{\lambda s^*} > \sup \psi^-$, it follows from (15b) that $\lim_{\psi \rightarrow q_{\lambda s^*}^-} \delta(\psi) = 1$ and it is impossible that $\lim_{\psi \rightarrow q_{\lambda s^*}^-} \delta'(\psi) > 1$, since $\delta'(\cdot) \leq 1$. Hence, it must be that $q_{\lambda s^*} \leq \sup \psi^-$.

Denote $d = \lim_{\psi \rightarrow q_{\lambda s^*}^-} \delta(\psi)$. Given that $\lim_{\psi \rightarrow q_{\lambda s^*}^-} \delta'(\psi) > d$, there exists some $\psi^* < q_{\lambda s^*}$ for which $\delta'(\psi^*) > d$. Given that $\delta'(\cdot)$ is non-decreasing, we can select $\psi^* > \inf \psi^-$ such that $\delta'(\psi^*) > d$. Given that the set ψ^- is an interval by assumption and $\inf \psi^- < \psi^* < q_{\lambda s^*} \leq \sup \psi^-$, there is some state s' with $\beta_{s'} - \alpha_{s'} < 0$ for which $r_{\lambda s'} = \psi^*$. We will show that δ' must perform worse than δ at s' .

Given that $\delta(\cdot)$ is non-decreasing, $\delta(\psi) \leq d$ for $\psi < q_{\lambda s^*}$ and $\lim_{\psi \rightarrow \psi^{*+}} \delta(\psi) \leq d$. On the other hand, $\lim_{\psi \rightarrow \psi^{*+}} \delta'(\psi) \geq \delta'(\psi^*) > d$. Then it follows from (13b) and $\lim_{\psi \rightarrow \psi^{*+}} \delta'(\psi) > \lim_{\psi \rightarrow \psi^{*+}} \delta(\psi)$ that $V_\lambda(\delta', P_{s'}, Q_{s'}) < V_\lambda(\delta, P_{s'}, Q_{s'})$. A similar argument applies if we assume $\beta_{s^*} - \alpha_{s^*} < 0$.

(B) Suppose that δ satisfies (16a) and (16b). For any s with $\beta_s - \alpha_s < 0$, either $r_{\lambda s} < \sup \psi^-$ or $r_{\lambda s} = \sup \psi^-$. If $r_{\lambda s} < \sup \psi^-$. It follows from (16a) that $\lim_{\psi \rightarrow r_{\lambda s}^+} \delta(\psi) = 0$. If $r_{\lambda s} = \sup \psi^-$, the same conclusion follows from (16b). Hence,

$$V_\lambda(\delta, P_s, Q_s) = \beta_s + (\alpha_s - \beta_s) \cdot (1 - \lim_{\psi \rightarrow r_{\lambda s}^+} \delta(\psi)) = \alpha_s$$

for all s with $\beta_s - \alpha_s < 0$. Similarly, if δ satisfies (16c) and (16d), then $V_\lambda(\delta, P_s, Q_s) = \beta_s$ for all s with $\beta_s - \alpha_s > 0$. If δ satisfies (16a)–(16d), it is λ -quantile optimal, therefore admissible.

To show that conditions (16a)–(16d) are necessary, first observe that the monotone rule

$$\delta^*(\psi) = 0 \quad \text{for } \psi \leq (\sup \psi^- + \inf \psi^+) / 2,$$

$$= 1 \quad \text{for } \psi > (\sup \psi^- + \inf \psi^+) / 2$$

is λ -quantile optimal. For all s with $\beta_s - \alpha_s < 0$, $\delta^*(\psi) = 0$ for ψ in the neighborhood of $r_{\lambda s}$. Hence,

$$V_\lambda(\delta^*, P_s, Q_s) = \beta_s + (\alpha_s - \beta_s) \cdot (1 - \lim_{\psi \rightarrow r_{\lambda s}^+} \delta^*(\psi)) = \alpha_s.$$

Similarly, $V_\lambda(\delta^*, P_s, Q_s) = \beta_s$ for all s with $\beta_s - \alpha_s > 0$. Given that there exists a λ -quantile optimal decision rule, only λ -quantile optimal decision rules are admissible.

If δ does not satisfy any of the conditions (16a)–(16d), it is not λ -quantile optimal. If (16a) is not satisfied, then $\delta(\psi') > 0$ for some $\psi' < \sup \psi^-$. Therefore there is a $\psi^* \in \psi^-$ such that $\psi' < \psi^*$. By monotonicity of $\delta(\cdot)$, $\delta(\psi) \geq \delta(\psi') > 0$ for all $\psi > \psi^*$. Therefore $\lim_{\psi \rightarrow (\psi^*)^+} \delta(\psi) > 0$, $V_\lambda(\delta, P_s, Q_s) < \alpha_s$, and $\delta(\cdot)$ is not λ -quantile optimal. If (16b) does not hold, there exists s with $\beta_s - \alpha_s < 0$, for which $r_{\lambda s} = \sup \psi^-$ and $\lim_{\psi \rightarrow r_{\lambda s}^+} \delta(\psi) > 0$. Hence $V_\lambda(\delta, P_s, Q_s) < \alpha_s$ and $\delta(\cdot)$ is not λ -quantile optimal. The necessity of (16c) and (16d) is proved analogously.

Q. E. D.

References

- Binmore, K. (2009), *Rational Decisions*, Princeton: Princeton University Press.
- Fisher, L. and L. Moyé (1999), "Carvedilol and the Food and Drug Administration Approval Process: An Introduction," *Controlled Clinical Trials*, 20, 1-15.
- Hirano, K. and J. Porter (2009), "Asymptotics for Statistical Treatment Rules," *Econometrica*, 77, 1683-1701.
- Hoeffding, W. (1963), "Probability Inequalities for Sums of Bounded Random Variables," *Journal of the American Statistical Association*, 58, 13-30.
- International Conference on Harmonisation (1999), "ICH E9 Expert Working Group. Statistical Principles for Clinical Trials: ICH Harmonized Tripartite Guideline," *Statistics in Medicine*, 18, 1905-1942.
- Karlin, S. and H. Rubin (1956), "The Theory of Decision Procedures for Distributions with Monotone Likelihood Ratio," *Annals of Mathematical Statistics*, 27, 272-299.
- Lehmann, E. and J. Romano (2005), *Testing Statistical Hypotheses*, New York: Springer.
- Manski, C. (1988), "Ordinal Utility Models of Decision Making Under Uncertainty," *Theory and Decision*, 25, 79-104.
- Manski, C. (2004), "Statistical Treatment Rules for Heterogeneous Populations," *Econometrica*, 72, 221-246.
- Manski, C. (2005), *Social Choice with Partial Knowledge of Treatment Response*, Princeton: Princeton University Press.
- Manski, C. (2011), "Actualist Rationality," *Theory and Decision*, 71, 195-210.
- Manski C. 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.
- Pfeiffer, P. (1990), *Probability for Applications*, New York: Springer.
- Rostek, M. (2010), "Quantile Maximization in Decision Theory," *Review of Economic Studies*, 77, 339-371.
- Rudin, W. (1976), *Principles of Mathematical Analysis*, New York: McGraw-Hill.
- Sedgwick, P. (2014), "Clinical Significance versus Statistical Significance," *BMJ*, 348:g2130. www.bmj.com/content/348/bmj.g2130, accessed October 11, 2014.
- Stoye J. (2009), "Minimax Regret Treatment Choice with Finite Samples," *Journal of Econometrics*, 151, 70-81.

Stoye, J. (2012), "Minimax Regret Treatment Choice with Covariates or with Limited Validity of Experiments," *Journal of Econometrics*, 166, 157-165.

Tetenov, A. (2012), "Statistical Treatment Choice Based on Asymmetric Minimax Regret Criteria," *Journal of Econometrics*, 166, 157-165.

Savage L. (1954), *The Foundations of Statistics*, New York: Wiley.

U. S. Food and Drug Administration (2014), "Statistical Guidance for Clinical Trials of Nondiagnostic Medical Devices,"

www.fda.gov/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/ucm106757.htm,

accessed October 11, 2014.

von Neumann, J. and O. Morgenstern (1944), *Theory of Games and Economic Behavior*, Princeton: Princeton University Press.

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