

EUI WORKING PAPERS

ECO No. 2006/26



Designing Non-Parametric Estimates
and Tests for Means

KARL H. SCHLAG



EUROPEAN UNIVERSITY INSTITUTE

Department of Economics

EUROPEAN UNIVERSITY INSTITUTE
DEPARTMENT OF ECONOMICS

Designing Non-Parametric Estimates and Tests for Means

KARL H. SCHLAG

EUI Working Paper **ECO** No. 2006/26

This text may be downloaded for personal research purposes only. Any additional reproduction for such purposes, whether in hard copy or electronically, requires the consent of the author(s), editor(s). If cited or quoted, reference should be made to the full name of the author(s), editor(s), the title, the working paper or other series, the year, and the publisher.

The author(s)/editor(s) should inform the Economics Department of the EUI if the paper is to be published elsewhere, and should also assume responsibility for any consequent obligation(s).

ISSN 1725-6704

© 2006 Karl H. Schlag

Printed in Italy
European University Institute
Badia Fiesolana
I – 50016 San Domenico di Fiesole (FI)
Italy

<http://www.iue.it/>
<http://cadmus.iue.it/dspace/index.jsp>

Designing Non-Parametric Estimates and Tests for Means¹

Karl H. Schlag²

June 18, 2006 (first version May 12, 2006)

¹The author would like to thank Dean Foster, Markku Lanne, Charles Manski, Fortunato Pesarin and Richard Spady.

²Economics Department, European University Institute, Via della Piazzuola 43, 50133 Florence, Italy, Tel: 0039-055-4685951, email: schlag@iue.it

Abstract

We show how to derive nonparametric estimates from results for Bernoulli distributions, provided the means are the only parameters of interest. The only information is that the support of each random variable is contained in a known bounded set. Examples include presenting minimax risk properties of the sample mean and a minimax regret estimate for costly treatment.

With the same method we are able to design nonparametric exact statistical inference tests for means using existing uniformly most powerful (unbiased) tests for Bernoulli distributions. These tests are parameter most powerful in the sense that there is no alternative test with the same size that yields higher power over any set of alternatives that only depends on the means. As examples we present for the first time an exact unbiased nonparametric test for a single mean and for the equality of two means (both for independent samples and for paired experiments).

We also show how to improve performance of Hannan consistent rules.

Keywords: exact, distribution-free, nonparametric inference, binomial average, finite sample theory, Hannan consistency, universal consistent.

JEL classification: C12, C44, C14.

1 Introduction

In this paper we investigate estimation and inference for means in the nonparametric setting where only a bounded set containing the support of the underlying distributions is known. Criteria for the selection of an estimate is minimax risk, a special case being minimax regret. Hypothesis tests are evaluated according to size and power in a given finite sample.

When the underlying distribution is entirely unknown no “sensible” estimate nor hypothesis test for the mean exists following Bahadur and Savage (1956). Due to the possibility of fat tails, the expected loss of any estimate is necessarily unbounded. Any unbiased level α test of the mean has constant power equal to α .

We assume minimal additional knowledge of the underlying distribution in the form of being given a bounded set that is known to contain its support. Such bounds arise very natural when data has been generated by evaluating outcomes on a finite scale such as a grade scale or percentages. Imposing such bounds means that there is a nominally best and worst outcome. With this additional assumption estimates and hypothesis tests can be designed. By choosing an appropriate affine transformation there is no loss of further generality by assuming that the set known to contain the support is equal to $[0, 1]$.

Only few papers can be found in the literature that fall in this nonparametric setting under our objectives stated above. Hodges and Lehmann (1950) derive a point estimate for the mean and for the difference between two means under the quadratic loss function. Schlag (2003) solves for minimax regret in a two-armed bandit problem with discounting. Schlag (2006) designs a sequential testing rule that attains minimax regret. While there are some papers that have designed hypothesis tests that are exact in terms of an upper bound specified on the size (e.g. Stringer, 1963, Anderson, 1967, Bickel et al., 1989, Romano and Wolf, 2000), their performance in terms of power is not known for a given finite sample size.

This paper presents a new method for deriving nonparametric estimates and hypothesis tests by expanding on a simple technique used in Schlag (2003, 2006) to solve for minimax regret. Accordingly, exact results for Bernoulli distributions can be immediately extended to exact results for the set of distributions with support in $[0, 1]$. As the class of Bernoulli distributions is often easy handle, many new results emerge. We just recently discovered that Cucconi (1968, in Italian, cited by Pesarin, 1984) already suggested this use of auxiliary randomization to extend Wald’s (1947)

sequential probability ratio test to a nonparametric test for the mean of a random variable with support in $[0, 1]$. However that was the only application and these are the only known citations.

The underlying idea is simple. Randomly transform any observation from $[0, 1]$ into $\{0, 1\}$ as follows. If outcome y is observed then replace it with outcome 1 with probability y and replace it with outcome 0 with probability $1 - y$. We call this the *Bernoulli transformation*. After independently transforming all observations, continue as if the underlying distribution were Bernoulli. The insight is to understand that it is as if each of the underlying random variables is Bernoulli distributed according to the original mean.

We start by applying this technique to derive unbiased minimax risk estimates of the mean. We find that the randomized estimate resulting from calculating the sample mean of the Bernoulli transformed data attains minimax risk among the unbiased estimates for any underlying loss function. If loss is convex, then combining this with a result of Hodges and Lehmann (1950), we obtain that the sample mean of the original data attains minimax risk among the set of all unbiased estimates.

For the context of regret (Savage, 1951, Milnor, 1954) we immediately obtain minimax regret estimates for the choice between a known and an unknown treatment and for the case of choice between two unknown treatments when testing is costly. In each case we simply extend the respective existing results for Bernoulli distributions (Manski, 2004, Canner, 1970). New results are also easily derived with this technique as demonstrated by Schlag (2006) who obtains the first nonparametric result on minimax regret under sequential sampling.

An entirely new field for applications demonstrated in this paper concerns nonparametric hypothesis testing. Properties of a test having level α or being unbiased carry over from Bernoulli distributions to our nonparametric setting via the Bernoulli transformation. The power at any distribution of the transformed test equals the power of the original test at the Bernoulli distribution with the same mean. In particular, uniformly most powerful tests for Bernoulli distributions translate into “parameter most powerful” tests. These are tests that cannot be improved on for any set of alternatives that is only described in terms of the underlying means. These tests are randomized and require careful implementation (see appendix). However, regardless of their random nature, they are most powerful and hence present a benchmark for comparison with alternative exact deterministic tests.

For instance we present the first unbiased nonparametric test for a mean and this is also the first nonparametric most powerful test for a mean. Based on the results in this paper, for the first time statements of the following type can be proven. For the test of mean equal to 0.5 against the mean being larger than 0.5 there is no test with size 5% and type II error of 20% assigned to means larger than 0.7 that is based on only 36 independent observations but there is for 37. This specific result follows from simple calculations using the uniformly most powerful one sided test of a mean (see appendix of this paper). We also present the first nonparametric two sided tests for a mean and tests for the comparison between two means both for independent samples as well as for a paired experiment. In the latter two cases the randomized version of the Fisher test (Fisher, 1935, Tocher, 1950) and the McNemar (1947) test are extended to randomized tests for the nonparametrical setting. Other exact nonparametric tests are easily designed from existing tests for the Bernoulli case.

A similarly new area of application concerns minimax “Hannan regret” in the adversarial multi-armed bandit (Auer et al., 1995) that allows for arbitrary sequences of payoffs. Hannan regret compares own payoffs to a measure of performance of constant action based on realized payoffs. We show how to equivalently define Hannan regret in terms of the underlying distributions to then be able to derive the same type of results. The binomial transformation reduces the maximal Hannan regret of any rule. Consequently, Hannan consistent rules (Hannan, 1957, Hart and Mas-Colell, 2000, 2001) need only be defined in terms of sequences of binary payoffs.

The paper is organized as follows. In Section 2 we describe all results for the simplest case where the data consists of a given number of independent observations of a single random variable. The binomial transformation of the data is described followed by the main result that shows how minimax risk properties for Bernoulli distributions carry over to general distributions with support in $[0, 1]$. This also means that minimax risk can be derived by only considering Bernoulli distributions which allows us in Subsection 2.2 to prove existence of minimax risk for the non-parametric case. In Subsection 2.3 the implications for hypothesis testing are illustrated. Each result is illustrated with examples. In Section 3 we extend the above to allow for any finite number of random variables, separately presenting the scenarios where data collection is exogenous and endogenous, differentiating the latter according to whether the sample sizes of each random variable are fixed at the outset or whether they may

depend on earlier observations. In Section 4 we comment on extensions. We consider an adversarial multi-armed bandit where nature chooses the entire sequence of payoffs and consider what happens when each data point consists of several observations as in a paired experiment. We show how one can model bounds on the degree of independence of the random variables and briefly present the scenario with covariate information. Section 5 contains the conclusions. The appendix illustrates how to apply the PMP one-sided test for a mean.

2 Single Sample Problem

Consider a random variable Y that yields outcomes y in $[0, 1]$ based on an unknown distribution P . The results herein directly also apply if the range of outcomes \mathcal{Y} is some subset of $[0, 1]$ containing $\{0, 1\}$. Moreover, the results can also be applied if the range of outcomes belongs to some fixed compact set $\mathcal{Y} \subset \mathbb{R}$ provided any outcome $y \in \mathcal{Y}$ is first transformed affinely into $[0, 1]$ using $\frac{y-d^0}{d^1-d^0}$ where $d^0 := \inf \mathcal{Y}$ and $d^1 := \sup \mathcal{Y}$.

Let $\mu = \mu(Y)$ denote the mean and σ_Y^2 denote the variance of Y . Let \mathcal{P} denote the set of all distributions with outcome range $[0, 1]$. Y is called *binary valued* if $Y \in \{0, 1\}$ in which case the underlying distribution is called *Bernoulli*. Let \mathcal{P}^b denote the set of all Bernoulli distributions.

Let Y^b be the binary valued random variable that results when transforming any outcome y realized by Y as follows. Outcome y is transformed into outcome 1 with probability y and into outcome 0 with probability $1 - y$. We refer to this random mapping as the *Bernoulli transformation of outcome y* . The Bernoulli distribution of Y^b is denoted by P^b . If Y is discrete then $\Pr(Y^b = 1) = \sum_{y \in [0, 1]} y \Pr(Y = y)$. $E(Y^b | Y = y) = y$, $\mu(Y^b) = \mu(Y)$ and $\sigma_Y^2 \leq \sigma_{Y^b}^2$. In fact, Y^b is the random variable with maximal variance among all random variables that have the same mean as Y .

Given $N \in \mathbb{N}$ consider a random sample of N independent observations of a realization of Y . Let $Y_{1,N}$ denote this random sequence and let $y_{1,N} = (y_1, \dots, y_N)$ denote a typical realization, the *sample mean* $\bar{y} = \frac{1}{N} \sum_{k=1}^N y_k$ being an unbiased estimator of the mean of Y . Let P^N denote the distribution of $Y_{1,N}$ induced by P .

Let $Y_{1,N}^b$ denote the random sequence that results when drawing a sequence $y_{1,N}$ from $Y_{1,N}$ and then independently Bernoulli transforming each element in $y_{1,N}$. Let $y_{1,N}^b = (y_n^b)_{n=1}^N$ be a typical realization when starting with $y_{1,N}$ and hence we also

call $y_{1,N}^b$ the *Bernoulli transformed sample*. Note that we can also interpret $Y_{1,N}^b$ as a random sample of N independent observations of Y^b so $Y_{1,N}^b = (Y^b)_{1,N}$. Consequently, $\sum_{n=1}^N y_n^b$ is binomially distributed and the sample mean $\bar{y}^b = \frac{1}{N} \sum_{n=1}^N y_n^b$ of the Bernoulli transformed sample $y_{1,N}^b$, also called the *binomial average*, is an alternative unbiased estimator of the mean of Y . Since $\sigma_Y^2 \leq \sigma_{Y^b}^2$, the variance of the binomial average \bar{y}^b is never strictly smaller than that of the sample mean \bar{y} where the variances of these two alternative estimates of the mean only coincide if Y is binary valued. Thus, a statistician would never prefer the binomial average over the sample mean as an unbiased estimate for the mean of Y . However below we find that the binomial average is very useful to derive estimates or to design tests relating to the unknown mean of Y .

Consider the general problem of finding a minimax risk estimate of $g(\mu)$ based on the observed sample y_1, \dots, y_N where the target $g(\mu)$ belongs to some space \mathcal{G} for all $\mu \in [0, 1]$. An *estimate* f maps the observed sample into an element of \mathcal{G} . The estimate f is called *deterministic* if $f : [0, 1]^N \rightarrow \mathcal{G}$.¹ We also allow for non deterministic estimates so in general $f : [0, 1]^N \rightarrow \Delta\mathcal{G}$ where ΔA denotes the set of distributions with support contained in the set A . f is called *binomial* if it first takes the Bernoulli transformation of the sample and then applies some mapping so if there is some $f_0 : \{0, 1\}^N \rightarrow \Delta\mathcal{G}$ such that $f(Y_{1,N}) = f_0(Y_{1,N}^b)$. So if f is binomial then $f(Y_{1,N}) = f(Y_{1,N}^b)$. f^b is called the *binomially transformed estimate* of f if f^b is binomial and $f^b(Y_{1,N}) = f(Y_{1,N}^b)$. Note that if f is deterministic then f^b is not deterministic unless f is constant on $\{0, 1\}^N$.

Loss W from making estimate $z \in \mathcal{G}$ given target g and mean vector μ is given by $W(g, \mu, z)$, later it will become meaningful to also allow loss to depend on the outcomes in the sample. *Risk* R from using estimate f when facing distribution P is defined as expected loss $E_P W(g(\mu), \mu, f(Y_{1,N}))$, so

$$R(f, P) := E_P W(g(\mu), \mu, f(Y_{1,N})) = \int W(g(\mu(P)), \mu, f(y_{1,N})) dP^N(y_{1,N}).$$

f^* attains *minimax risk* (Wald, 1950) if

$$f^* \in \arg \min_f \max_{P \in \mathcal{P}} R(f, P)$$

where $\min_f \max_{P \in \mathcal{P}} R(f, P)$ is called the *value of minimax risk*. An estimate is *exact* if the measurement of its performance does not rely on asymptotic properties. In

¹In the statistics literature “deterministic” is also sometimes called “non randomized”, however we use the term “random” to refer to all distributions including those that are degenerate.

this paper we show how to derive estimates that attain minimax risk, hence these estimates are exact.

We briefly recall some decision theoretic foundations of specific loss functions. Accordingly, a state of the world is described by P . In each state of the world a decision maker has to choose an action \tilde{g} which is some element of \mathcal{G} . Let $u(\tilde{g}, \mu(P))$ be the payoff of choosing \tilde{g} in state P .² If the target is to choose the action that maximizes payoffs in the true state of the world, so if $g(\mu) \in \arg \max_{\tilde{g} \in \mathcal{G}} u(\tilde{g}, \mu(P))$, and if loss is given by the difference between the target and the realized payoff, so if $W(g(\mu), \mu, z) = u(g(\mu), \mu) - u(z, \mu)$, then loss and risk are both referred to as *regret*. Savage (1951) introduced this original notion of regret based on his interpretations of Wald (1950). Milnor (1954) presented an axiomatic foundation of the objective to attain minimax regret. Wald (1950) originally proposed the alternative *maximin rule*, also axiomatized by Milnor (1954), under which the decision maker aims to find an estimate that maximizes the minimum payoff which means in our above notation that loss W is defined by $W(g(\mu), \mu, z) = -u(z, \mu)$.

2.1 Minimax Risk

We obtain the following result.

Proposition 1 (ia) $R(f^b, P) = R(f, P^b)$.

(ib) $\{R(f^b, P), P \in \mathcal{P}\} \subseteq \{R(f, P), P \in \mathcal{P}\}$, in particular $\max_{P \in \mathcal{P}} R(f^b, P) \leq \max_{P \in \mathcal{P}} R(f, P)$.

(ic) If $f^* \in \arg \min_f \max_{P \in \mathcal{P}} R(f, P)$ then $f^{*b} \in \arg \min_f \max_{P \in \mathcal{P}} R(f, P)$.

(ii) If $f^* \in \arg \min_f \max_{P^b \in \mathcal{P}^b} R(f, P^b)$ then $f^{*b} \in \arg \min_f \max_{P \in \mathcal{P}} R(f, P)$.

(iii) If $f^* \in \arg \min_f \max_{P \in \mathcal{P}} R(f, P)$ then $\max_{P \in \mathcal{P}} R(f^*, P) = \max_{P^b \in \mathcal{P}^b} R(f^*, P^b)$.

Part (ia) states that the risk attained by an estimate given a Bernoulli distribution with mean μ_0 is equal to the risk attained by the binomially transformed estimate given any distribution with mean μ_0 . The rest of the results are an immediate implication. (ib) The binomially transformed estimate attains weakly lower maximal risk (but also raises minimum risk). (ic) If an estimate attains minimax risk then so does its binomial transformation. (ii) To derive minimax risk it is sufficient to solve

²More specifically, u is the expected utility of choosing g in state P where this utility representation can be derived from the standard von Neumann and Morgenstern (1944) framework.

the simpler problem of finding a minimax risk estimate for Bernoulli distributions as the binomial transformation of this estimate will then attain minimax risk for all distributions with support in $[0, 1]$. (iii) Maximal risk under a minimax risk estimate can always be attained by some Bernoulli distribution. In other words, it is hardest to guarantee the lowest risk when facing Bernoulli distributions.

Proof. Consider the n -th element of the sample y_n . Under f^b , y_n is transformed into outcome 1 with probability y_n and transformed into 0 otherwise. As y_n itself was independently drawn from P we obtain that the ex-ante probability that the n -th element of the sample is transformed into 1 is equal to $\int_0^1 y dP(y) = \mu(P)$. Thus, the risk under f^b when facing P is equal to the risk under f when facing a Bernoulli distribution with the same mean which proves (ia). The rest of part (i) as well as part (ii) is a direct consequence of (ia). Part (iii) also follows immediately. Since f^* attains minimax risk,

$$\max_{P \in \mathcal{P}} R(f^*, P) \leq \max_{P \in \mathcal{P}} R(f^{*b}, P).$$

Using (ia) we obtain

$$\max_{P \in \mathcal{P}} R(f^{*b}, P) = \max_{P^b \in \mathcal{P}^b} R(f^*, P^b) \leq \max_{P \in \mathcal{P}} R(f^*, P).$$

Combining these two inequalities proves (iii). ■

Example 1 (Unbiased Estimate of a Mean) *It is easily shown that the binomial average \bar{y}^b attains minimax risk among the unbiased estimates of the mean of a Bernoulli distribution.³ Hence, the binomial average attains minimax risk among the unbiased estimates of the mean for any loss function W , formally $\bar{y}^b \in \arg \min_{f^u \in \mathcal{F}^u} \max_{P \in \mathcal{P}} R(f^u, P)$ where \mathcal{F}^u denotes the set of unbiased estimates of the mean.*

Combining Hodges and Lehmann (1950, Theorem 3.2) and Proposition 1(iii) we obtain a result on how to derive deterministic minimax risk estimates when loss is convex in the estimate.

Corollary 1 *Assume that \mathcal{G} is a Euclidean space and that loss W is convex in z . If $f_0 \in \arg \min_f \max_{P^b \in \mathcal{P}^b} R(f, P^b)$ and f^* is such that $f^*(y_{1,N}) = E(f_0^b | Y_{1,N} = y_{1,N})$ then f^* is a deterministic estimate that attains minimax risk.*

³In a first step one shows that there is a symmetric minimax risk estimate where symmetric means that all rounds are treated equally. In the second step, given that behavior of a symmetric estimate only depends on the total number of successes, one applies the identity theorem for polynomials to show that the sample mean is the unique symmetric unbiased estimate.

Example 2 (Another Unbiased Estimate of a Mean) *Return to Example 1 and assume additionally that loss is convex in z , e.g. loss is quadratic so $W(g(\mu), \mu, z) = (\mu - z)^2$. Since $E(\bar{y}^b | Y_{1,N} = y_{1,N}) = \bar{y}$ we obtain that the sample mean also attains minimax risk among the unbiased estimates of the mean. Note that this does not contradict the fact that the binomial average is a minimax risk unbiased estimate as both of these estimates attain the same maximal risk on the set of Bernoulli distributions, specifically when the underlying mean is equal to $1/2$.*

Example 3 (Biased Estimate of a Mean) *Consider quadratic loss. Hodges and Lehmann (1950, Problem 1) establish $\frac{1}{1+\sqrt{N}} \left(\sqrt{N}\bar{y} + \frac{1}{2} \right)$ as a minimax risk estimate for Bernoulli distributions. They then continue to use special properties of quadratic loss to show that this estimate also attains minimax risk among all distributions $P \in \mathcal{P}$ (Hodges and Lehmann, 1950, Theorem 6.1). Following Proposition 1(ic), the binomial transformation of this estimate, namely $\frac{1}{1+\sqrt{N}} \left(\sqrt{N}\bar{y}^b + \frac{1}{2} \right)$, is also a minimax regret estimate. The expected estimate given the sample equals $\frac{1}{1+\sqrt{N}} \left(\sqrt{N}\bar{y} + \frac{1}{2} \right)$ which demonstrates by Corollary 1 an alternative proof of its minimax risk properties by only using convexity of the quadratic loss function. Notice that these two estimates attain the same risk for all $P \in \mathcal{P}$, unlike above where the risk under binomial average only coincides with the risk of the sample mean on the set of Bernoulli distributions.*

It is now feasible (albeit possibly tedious) to derive estimates for alternative convex loss functions, e.g. $W(g(\mu), \mu, z) = |\mu - z|$, by first solving the case of Bernoulli distributions which is in turn solved by setting up a zero-sum game between the decision maker and nature (Savage, 1954).

Example 4 (Minimax Regret Estimate of Single Unknown Mean) *Consider the choice between two random variables Y and Y_0 with the objective to choose the random variable with the higher mean where Y has unknown mean μ while Y_0 has known mean μ_0 . Set $\mathcal{G} = \{0, 1\}$ and identify 1 with the selection of Y and 0 with the selection of Y_0 . Consider loss in terms of regret so $W(g(\mu), \mu, z) = \max\{\mu_0, \mu\} - z\mu - (1 - z)\mu_0$. Manski (2004) provides the minimax regret estimate among Bernoulli distributions (see Stoye, 2005 for a simpler formula). Proposition 1 shows that the binomial transformation of this estimate attains minimax regret for all distributions $P \in \mathcal{P}$.⁴*

⁴This result was first stated in (Schlag, 2006).

2.2 Existence

Proposition 1 can also be used to ensure existence of a minimax risk estimate under very weak conditions. The proof follows the idea Savage (1954) to calculate minimax risk by investigating a specific zero-sum game. This also allows us to interpret any minimax risk estimate as a particular Bayes solution.

Consider the more classic decision theoretic setting (von Neumann and Morgenstern, 1944) where the true distribution P is drawn from a known probability distribution (or prior) Q , so $Q \in \Delta\mathcal{P}$. The estimate f that minimizes expected risk $R(f, Q) = \int_{\mathcal{P}} R(f, P) dQ(P)$ is called a *Bayes solution*, the value of the minimum expected risk is called the *Bayes risk*. Q^* is called a *least favorable prior* if $Q^* \in \arg \max_{Q \in \Delta\mathcal{P}} \min_f R(f, Q)$.

Proposition 2 *Assume that W as a function of μ and z is continuous and that \mathcal{G} is a metric space.*

(i) *There exists a minimax risk estimate and a least favorable prior containing only Bernoulli distributions in its support.*

(ii) *Any minimax risk estimate minimizes Bayes risk under any least favorable prior.*

Proof. Following Savage (1954), we solve minimax risk by considering the following simultaneous move zero-sum game between the decision maker and nature. The decision maker chooses an estimate f and nature chooses a distribution $Q^b \in \Delta\mathcal{P}^b$ over Bernoulli distributions $P^b \in \mathcal{P}^b$. The payoff to the decision maker is given by $-R(f, Q^b)$ while that of nature by $R(f, Q^b)$. Under the above assumptions there exists a Nash equilibrium (or saddle point) (f^*, Q^{b*}) of this game (Glicksberg, 1952) which means that $R(f^*, Q^b) \leq R(f^*, Q^{b*}) \leq R(f, Q^{b*})$ holds for all f and all Q^b . Now we use the well known minimax theorem of zero-sum games, to conclude that $R(f^*, Q^{b*}) = \min_f \max_{Q^b \in \Delta\mathcal{P}^b} R(f, Q^b) = \min_f \max_{P^b \in \mathcal{P}^b} R(f, P^b)$. Following Proposition 1(ic), the binomial transformation f^{*b} of f^* attains minimax risk.

Concerning (ii), note that (f^*, Q^{b*}) is also a saddle point of the game where nature is allowed to choose any prior $Q \in \Delta\mathcal{P}$. Applying the minimax theorem for zero-sum games we find that $R(f^*, Q^{b*}) = \max_{Q \in \Delta\mathcal{P}} \min_f R(f, Q)$ which means that Q^{b*} is a least favorable prior. ■

Combining the above with Corollary 1 we obtain:

Corollary 2 *Assume that \mathcal{G} is a Euclidean space, W is continuous in μ and z and convex in z . Then a deterministic minimax risk estimate exists.*

2.3 Hypothesis Testing

Consider the objective of testing a null hypothesis $H_0 : \mu \in \Omega_0$ against an alternative hypothesis $H_1 : \mu \in \Omega_1$ where $\Omega_0, \Omega_1 \subset [0, 1]$, $\Omega_0, \Omega_1 \neq \emptyset$ while $\Omega_0 \cap \Omega_1 = \emptyset$. Selecting hypothesis H_i will be identified with i and hence $\mathcal{G} = \{0, 1\}$. The *power function* β of the estimate or *test* f is then given by $\beta_f(P) = \Pr_P(f = 1)$. f is a *level α test* if $\beta_f(P) \leq \alpha$ whenever $\mu(P) \in \Omega_0$ where f is *unbiased* if additionally $\beta_f(P) \geq \alpha$ whenever $\mu(P) \in \Omega_1$. f is *uniformly most powerful* (UMP) if for any level α test \tilde{f} and any P such that $\mu(P) \in \Omega_1$ we find that $\beta_f(P) \geq \beta_{\tilde{f}}(P)$. We call f *parameter most powerful* (PMP) if for any level α test \tilde{f} and for any $\tilde{\mu} \in \Omega_1$ we have $\min_{P:\mu(P)=\tilde{\mu}} \beta_f(P) \geq \min_{P:\mu(P)=\tilde{\mu}} \beta_{\tilde{f}}(P)$. Thus, f is parameter most powerful if it is a maximin test in the sense of Lehmann and Romano (2005, chapter 8) for any set of alternatives that only depends on the mean of the underlying distributions.

While the following results can also be phrased in terms of risk we choose to present them directly using the terminology above. Remember that f^b denotes the binomial transformation of the estimate f and that given P we have defined P^b as the Bernoulli distribution that has the same mean as P .

Proposition 3 (i) $\beta_f(P^b) = \beta_{f^b}(P)$ so $\{\beta_{f^b}(P), P \in \mathcal{P}\} \subseteq \{\beta_f(P), P \in \mathcal{P}\}$.

(ii) If f is a level α test for all $P \in \mathcal{P}^b$ then f^b is a level α test for all $P \in \mathcal{P}$.

(iii) If f is unbiased for all $P \in \mathcal{P}^b$ then f^b is unbiased for all $P \in \mathcal{P}$.

(iv) If f is a uniformly most powerful test for all $P \in \mathcal{P}^b$ then f^b is a parameter most powerful test for all $P \in \mathcal{P}$.

(v) If f is a uniformly most powerful unbiased test for all $P \in \mathcal{P}^b$ then f^b is a parameter most powerful unbiased test for all $P \in \mathcal{P}$.

More generally, even if there is no UMP test for the Bernoulli case, we find that any test for the Bernoulli case can be extended to a randomized test for the nonparametric setting of this paper where power can be derived from the Bernoulli setting.

Proof. Part (i) follows directly from the definitions as $\beta_f(P^b) = \Pr_{P^b}(f = 1) = \Pr_P(f^b = 1) = \beta_{f^b}(P)$. Concerning part (ii) let f be a level α test for Bernoulli distributions and consider P such that $\mu(P) \in \Omega_0$ so $\beta_f(P^b) \leq \alpha$. Part (i) implies that $\beta_{f^b}(P) \leq \alpha$ which implies that f^b is a level α test for all $P \in \mathcal{P}$. For part

(iii) let f be unbiased for all $P \in \mathcal{P}^b$ and let $\mu(P) \in \Omega_1$ so $\beta_f(P^b) \geq \alpha$. Then part (i) shows that $\beta_{f^b}(P) \geq \alpha$ so together with part (ii) it follows that f^b is unbiased for all $P \in \mathcal{P}$. For proof of part (iv) let f be a uniformly most powerful test for Bernoulli distributions and let $\tilde{\mu} \in \Omega_1$. Let \tilde{f} be a level α test for all $P \in \mathcal{P}$. Then $\beta_{f^b}(P) = \beta_f(P^b) \geq \beta_{\tilde{f}}(P^b) \geq \min_{P:\mu(P)=\tilde{\mu}} \beta_{\tilde{f}}(P)$ and hence f is a parameter most powerful test for all $P \in \mathcal{P}$. Part (v) follows similarly using (iii) and the inequality used in the proof of part (iv). ■

One may choose to comment on the efficiency of a test in terms of the number of samples it needs. Let $N(\alpha, \beta_0, \mu_0, f)$ be the smallest number of observations needed by the level α test f to achieve power of at least β_0 for all distributions P that have mean μ_0 where μ_0 belongs to the set of alternatives.⁵ Holding α , β_0 and μ_0 fixed we call $N(\alpha, \beta_0, \mu_0, f) / N(\alpha, \beta_0, \mu_0, \tilde{f})$ the *relative parameter efficiency* of f relative to \tilde{f} . Here we adapt the standard definition of relative efficiency (cf. Hodges and Romano, 2005, chapter 13.2, p. 534) to the set of alternatives that all have the same mean. We call f^* *parameter efficient* if there is no alternative test f such that $N(\alpha, \beta_0, \mu_0, f) < N(\alpha, \beta_0, \mu_0, f^*)$. This leads to the following observation.

Corollary 3 *The binomial transformation of any parameter efficient test is also parameter efficient. Any parameter most powerful test is a parameter efficient test. Any parameter most powerful unbiased test is a parameter efficient unbiased test.*

For the nonparametric setting we obtain the first unbiased one-sided test and the first unbiased two-sided test for a mean.

Example 5 (Test for a Mean) (i) *Assume that $H_0 : \mu = \mu_0$ (or equivalently $H'_0 : \mu \leq \mu_0$) should be tested against $H_1 : \mu > \mu_0$ for some given $\mu_0 \in [0, 1]$. Clearly a uniformly most powerful test does not exist for all $P \in \mathcal{P}$. The most powerful test f^* for $P \in \mathcal{P}^b$ is given by $f^* = 1$ if $\bar{y} > \frac{k_\alpha}{N}$, $f^* = 1$ with probability γ_α if $\bar{y} = \frac{k_\alpha}{N}$ and $f^* = 0$ otherwise where $k_\alpha \in \mathbb{N}_0$ and $\gamma_\alpha \in [0, 1]$ are such that $\Pr_P(f^* = 1) = \alpha$ if $\mu(P) = \mu_0$ (e.g. see Rohtagi, 1976, Example 2, p. 415). Following Proposition 3(iv) and Corollary 3, a parameter most powerful test which is parameter efficient is given by the binomial transformation f^{*b} of f^* which means that $f^{*b} = 1$ if $\bar{y}^b > \frac{k_\alpha}{N}$, $f^{*b} = \gamma_\alpha$ if $\bar{y}^b = \frac{k_\alpha}{N}$ and $f^{*b} = 0$ otherwise where k_α and γ_α are given above. Notice that f^{*b} is*

⁵We consider only deterministic sample sizes to simplify exposition. Performance can sometimes be (marginally) improved by selecting sample size randomly.

unbiased. In the appendix we illustrate how to accommodate with the random nature of this test and compare its performance to an alternative exact deterministic test.

(i) For $H_0 : \mu = \mu_0$ against $H_1 : \mu \neq \mu_0$ a uniformly most powerful test for $P \in \mathcal{P}^b$ does not exist. However, a uniformly most powerful unbiased test exists for $P \in \mathcal{P}^b$ (Lehmann and Romano, 2005, Example 4.2.1) and we can apply Proposition 3(v) to obtain a parameter most powerful unbiased test and hence a parameter efficient unbiased test.

3 Multiple Sample Problem

Consider now the case of $K \geq 2$ random variables Y_1, \dots, Y_K where $Y_k \in [0, 1]$ for all $k = 1, \dots, K$.⁶ Let P denote the (joint) distribution of $(Y_k)_{k=1}^K$, let μ_k denote the mean of Y_k and let μ denote the vector of means. All results below hold if one makes no assumptions on the relationship between these random variables or instead if one assumes that Y_1, \dots, Y_K are known to be independent.

Let P^b be the distribution with support in $\{0, 1\}^K$ that has the same marginal distributions as P and where the associated random variables are independent.

3.1 Exogenous Sampling

Consider a random sample that consists of N_k independent observations of random variable Y_k where $(N_k)_{k=1}^K \in \mathbb{N}^K$ is given. Let $N = \sum_{k=1}^K N_k$ be the total sample size. The sample is *balanced* if $N \bmod K = 0$ and $N_k = N/k$ for all $k = 1, \dots, K$. To keep notation simple, assume that the observed outcomes are ordered in terms of increasing index of the underlying random variable. We maintain the notation that $y_{1,N}$ denotes a typical realization so the first N_1 elements are realizations of random variable Y_1 . The other definitions from the earlier sections extend immediately. Parameter efficiency and relative parameter efficiency refer to the total sample size N . Since none of the proofs of the previous results relied on the single dimensionality of data they all extend immediately.

⁶A more general fixed compact range of outcomes $\mathcal{Y} \subset \mathbb{R}^K$ can be dealt with as in the case of $K = 1$ by first affinely transforming all components of outcome $y = (y_1, \dots, y_K)$ into $[0, 1]$. So y_t needs to be replaced by $\frac{y_k - d_k^0}{d_k^1 - d_k^0}$ where $d_k^0 = \inf \{z : y_k = z \text{ for some } y \in \mathcal{Y}\}$ and $d_k^1 = \sup \{z : y_k = z \text{ for some } y \in \mathcal{Y}\}$, $k = 1, \dots, T$.

Proposition 4 *The statements in Propositions 1, 2,3 and in Corollaries 1 and 3 also hold for the setting of K random variables with N_k observations of Y_k , $k \in \{1, \dots, K\}$.*

Example 6 *Consider the objective to find the minimax risk estimate of the maximal mean under absolute or quadratic loss, so consider $W(g(\mu), \mu, z) = |\max\{\mu_1, \dots, \mu_K\} - z|^\gamma$ where $\gamma \in \{1, 2\}$. While the solution to this problem is unknown, its solution is feasible given the results in this paper. The extension of Corollary 1 to multiple random variables shows that a deterministic minimax risk estimate exists and that it can be derived by first finding the minimax risk estimate f_0 for Bernoulli distributions and then deriving the expected estimate of the binomial transformed estimate f_0^b .*

Example 7 (Regret Estimate under Costly Testing) *(i) Canner (1970) considers a model with $K = 2$ in which each random variable is associated to the outcome of a treatment on some subject. Each treatment is tested on n subjects, so $N = 2n$, after which the decision maker has to recommend a treatment for the remaining $M - 2n$ subjects. Tests cost $c \geq 0$ per subject. The aim is to minimize maximum regret where performance is measured using the sum of the outcomes of each treatment. So $K = 2$, $\mathcal{G} = \{1, 2\}$, $g(\mu) = \min\{\arg \max\{\mu_1, \mu_2\}\}$ and $W(g(\mu), \mu, z, x) = M\mu_{g(\mu)} - \left(\sum_{n=1}^N y_{n,k_n} + (M - 2n)\mu_z - 2nc\right)$ where k_n is the treatment chosen in round k . For Bernoulli distributions, Canner (1970) derives a minimax regret estimate that is based on recommending to the $M - 2n$ remaining subjects the treatment that yielded more successes during the $2n$ tests. Following Proposition 4, the binomial transformation of this rule, called the binomial average rule (Schlag, 2006), attains minimax regret for all distributions with support in $[0, 1]$.*

(ii) The model of Canner (1970) is easily extended to allow for samples that are not balanced and for more than two treatments. Proposition 4 shows that there exists a binomial estimate that attains minimax regret for this more general setup.

(iii) The case of M infinitely large is considered by Manski (2004, cf Stoye, 2005, Schlag, 2006) in which case average loss is given by $W(g(\mu), \mu, z, x) = \mu_{g(\mu)} - \mu_z$. Schlag (2006) shows using (Canner, 1970) that the binomial average rule attains minimax regret.

Example 8 (Test for Equality of Two Means) *Consider $K = 2$ and the one sided test of the null hypothesis $H_0 : \mu_1 \geq \mu_2$ (or $\mu_1 = \mu_2$) against the alternative $H_1 : \mu_1 < \mu_2$ where μ_1 and μ_2 are unknown. The randomized version of the Fisher's exact test (Tocher, 1950, Fisher, 1935) is a uniformly most powerful unbiased test*

for $P^b \in \mathcal{P}^b$.⁷ Proposition 4 shows that applying the Fisher test to the Bernoulli transformed payoffs yields a randomized test that is parameter most powerful among the unbiased tests for $P \in \mathcal{P}$. Two sided tests can similarly be adopted to tests for all $P \in \mathcal{P}$ albeit of course with weaker properties in terms of power than the alternative is true due to the non existence of a uniformly most powerful unbiased test for the Bernoulli case.

3.2 Endogenous Simultaneous Sampling

In the following we consider the setting in which data is gathered by a so-called decision maker who is allowed to determine ex-ante the number of independent observations of each random variable and who is confined by a given maximal number of observations M .

Let n_k denote the number of independent observations gathered of random variable Y_k for $k = 1, \dots, K$ and let $N = \sum_{k=1}^K n_k$ be the total number of observations so $n_k \geq 0$ and $N \leq M$. Note that we explicitly allow for collecting strictly less than M observations, relevant when data collection is costly, thereby also generalizing the setting of $K = 1$ in Section 2. So the decision maker has to first determine $n = (n_k)_k$ and then has to make an estimate f . We allow for mixed strategies in the initial assignment and let $h(n)$ be the probability of selecting n . So a strategy of the decision maker is a tuple (h, f) . We refer to this scenario as *endogenous simultaneous sampling*. Our previous definitions are immediately extended by replacing f by (h, f) .

As the decision maker influences what is observed we allow for loss to depend on the outcomes during sampling so $W = W(g(\mu), \mu, z, y_{1,N})$.

Proposition 5 *If $W(g(\mu), \mu, z, y_{1,N})$ is linear in y_n for each $n = 1, \dots, N$ then Proposition 4 also holds for endogenous simultaneous sampling. If $W(g(\mu), \mu, z, y_{1,N})$ is convex in y_n for each $n = 1, \dots, N$ then all statements extend except for the extension of Proposition 1 (ia) that has to be replaced by “ $R(f^b, P) \leq R(f, P^b)$ (with equality holding when $P \in \Delta\mathcal{P}^b$)”.*

Proof. The only real thing to check is that Proposition 1 (ia) extends. If W is linear in the observed outcomes then clearly the risk of a binomial transformation

⁷As samples are independent, the random variables need not be independent. In particular, any additional assumptions can be made on the correlation of the two random variables without affecting the validity of the test.

f^b under some distribution P is equal to the risk of f under the distribution P^b . If instead we only assume that W is convex in each observed outcome then risk weakly increases if the observation is replaced by the Bernoulli transformation where risk in fact does not change if P is Bernoulli. ■

Example 9 *To continue Example 7, assume now that the decision maker may decide on how many tests are made of each treatment. Canner (1970) is in fact interested in this scenario but has no analytical results for this case. Only few analytic results exist even for the Bernoulli case. In analogy to the proof of Proposition 6 in Schlag (2006) it is easy to show that the minimax risk estimate will consist of choosing the sample size N , possibly at random, and then sampling according to the binomial average rule. Schlag (2006, Proposition 2) explicitly provides the answer and proof for the setting of Example 7(iii).*

3.3 Sequential Sampling

Consider again the case where the decision maker decides which variables are sampled but assume now that data is gathered sequentially in N rounds. After each observation the decision maker decides which variable to observe next.

Let k_n be the index of the random variable to be observed in round n of the sampling where k_n may depend on all previous observations. Let $k = (k_1, \dots, k_N)$ be the sampling strategy of the decision maker. So now the decision maker chooses the tuple (k, f) , a setting we call sequential sampling. While strategies are more intricate, there are no substantial differences in formalities as compared to the endogenous simultaneous sampling scenario and hence we can state results without proofs.

Proposition 6 *Proposition 5 also holds for sequential sampling.*

Example 10 *Assume that the decision maker aims at maximizing discounted payoffs and measures loss using regret so $W(g(\mu), \mu, z, y_{1,N}) = \frac{1-\delta^N}{1-\delta} \left(\max_k \mu_k - \sum_{n=1}^N \delta^{n-1} y_{n,k_n} \right)$ for some $\delta \in (0, 1]$ where N can also be replaced by ∞ when $\delta < 1$. This setting is called the multi-armed bandit and originates in Robbins (1952). Using the same basic approach as in Proposition 2 Schlag (2003) proves existence of a minimax regret sequential sampling strategy for the case of $N = \infty$ and derives an explicit solution for the case where $\delta \leq \frac{1}{2}(\sqrt{5} - 1) \approx 0.618$. Schlag (2003) was the first paper that used the Bernoulli transformation technique of this paper.*

Example 11 For the setting of Example 7(iii) Schlag (2006) shows that sequential sampling cannot be used to lower minimax regret when $K = 2$ but that it can when $K = 3$.

Example 12 The sequential probability ratio test for Bernoulli distributions (Wald, 1947) can be extended to a sequential nonparametric test by considering the binomial transformation, as suggested by Cucconi (1968).

4 Extensions

In the following we briefly present some extensions.

4.1 Adversarial Nature

One may choose to consider a more general form of payoff realizations that are not iid. Assume that nature chooses the entire sequence of payoffs before the decision maker takes any action. Let P^N be the unknown distribution of payoff sequences so $P^N \in \Delta\left([0, 1]^{K \times N}\right)$. Let $\mu_{n,k}$ denote the mean of random variable k in round n . While one may also want to consider estimation or hypothesis testing in this environment, we only discuss the extension of the multi-armed bandit (see Example 10) called the *adversarial multi-armed bandit* (Auer et al., 1995). Sampling is sequential and the decision maker is interested in maximizing the present value of future payoffs.

Consider first regret in its original sense (Savage, 1951, Milnor, 1954), defined as the difference between payoffs achieved when the true state of nature is known and the actual payoffs. Identifying the P^N with the state of nature, regret is given by

$$\frac{1 - \delta^N}{1 - \delta} \sum_{n=1}^N \delta^{n-1} \left(\max_k \{\mu_{n,k}\} - y_{n,k_n} \right). \quad (1)$$

Notice that if P^N is a degenerate distribution that selects $y_{1,N}$ almost surely then regret is given by

$$W = \frac{1 - \delta^N}{1 - \delta} \sum_{n=1}^N \delta^{n-1} \left(\max_k \{y_{n,k}\} - y_{n,k_n} \right). \quad (2)$$

Our results in the previous sections extend immediately to show that a minimax regret estimate (for each N) exists. It is easily argued that regret can always be maximized by a deterministic payoff sequence (cf. Auer et al, 1995, footnote 1). Thus, one can

also derive minimax regret by solving for minimax risk with loss defined by (2) and hence being independent of the true state of nature. This observation immediately provides a simple way to see why the value of minimax regret is equal to $\frac{K-1}{K}$. The idea is to verify that the following pair of strategies constitutes an equilibrium of the associated zero sum game between the decision maker and nature. Nature chooses P^N generated by an iid process in which in each round each random variable is selected equally likely, the selected random variable is assigned payoff 1 and the others not selected are assigned payoff 0. The decision maker chooses independently in each round each treatment equally likely.

Possibly in light of the constant value of minimax regret, the literature starting with Hannan (1957) has considered an alternative loss function we refer to as *Hannan regret*, which is given by⁸

$$H(f, P^N) := \frac{1 - \delta^N}{1 - \delta} \left(\max_j \left\{ \sum_{n=1}^N \delta^{n-1} \mu_{n,j} \right\} - \sum_{n=1}^N \delta^{n-1} y_{n,k_n} \right). \quad (3)$$

Own performance is now only compared to the truly best constant choice over time. Analogous to our above arguments for the case of regret, any minimax Hannan regret estimate also attains minimax risk (and vice versa) if loss is instead defined as

$$\frac{1 - \delta^N}{1 - \delta} \left(\max_j \left\{ \sum_{n=1}^N \delta^{n-1} y_{n,j} \right\} - \sum_{n=1}^N \delta^{n-1} y_{n,k_n} \right). \quad (4)$$

The equivalent definition of loss in (4) is actually the way Hannan regret is used in the literature. Looking however at the expression in (3), we see that our results obtained for the multi-armed bandit extend directly. The maximal Hannan regret of any sequential decision rule is weakly lowered by considering the binomial transformation of the sequential decision rule.

Up to now we considered the setting of the so-called *oblivious adversary* (Auer et al., 1995) where the payoff sequence is determined ex-ante by nature. In the following we consider the *non oblivious adversary* where payoffs are chosen by nature in each round conditional on previous choices of the decision maker. In round n the choice of a payoff vector in $[0, 1]^K$ is a function of $\{1, \dots, K\}^{n-1}$ representing the $n - 1$ previous choices so the distribution we denote by P^{N*} is formally an element of $\Delta \left(\bigcup_{n=1}^N \left([0, 1]^K \right)^{\{1, \dots, K\}^{n-1}} \right)$. As $\mu_{n,j}$ is now a function of k_1, \dots, k_{n-1} we formally

⁸A decision theoretic justification of this type of loss function is not known.

define Hannan regret by

$$H(f, P^{N*}) := \frac{1 - \delta^N}{1 - \delta} \left(\max_j \left\{ \sum_{n=1}^N \delta^{n-1} \mu_{n,j}(k_1, \dots, k_{n-1}) \right\} - \sum_{n=1}^N \delta^{n-1} y_{n,k_n} \right). \quad (5)$$

A deterministic behavior of nature (and an equivalent alternative definition of loss) is again given by (4). The interpretation is however now a “bit strange” as the benchmark is no longer the truly best constant choice over time. Different choice possibly changes later payoffs (cf. Auer et al., 1995, page 4). The importance of the representation in (5) is immediate. As under the oblivious adversary, maximal Hannan regret is weakly reduced by considering the binomial transformation of a sequential rule. In particular, there is no need to specify how the decision maker reacts to interior payoffs.

Summary 1 *Whether nature is an oblivious or a non oblivious adversary, the Hannan regret of any sequential rule is weakly reduced by considering the binomial transformation of this rule.*

4.2 Perfect Information

In the above settings involving multiple random variables we assumed up to now that only a single observation is made in each test. Assume now instead that the outcome of each random variable is observed in each test. So in round n we now assume that $y_n \in [0, 1]^K$ is observed where y_n is drawn according to the joint distribution $P \in \mathcal{P}$. In the terminology of decision theory, *foregone payoffs* are observable. In the terminology of hypothesis testing, for $K = 2$, this is called a *paired experiment* or the setting of *matched pairs*. For the sequential sampling setting we let $c_n \in \{1, \dots, K\}$ denote the choice in round n that may now depend on the previously observed outcomes and choices. The definitions of regret and Hannan regret carry over.

Again we derive results by considering a random transformation of outcomes. Assume that each component of an observation $y \in [0, 1]^K$ is independently transformed into a binary outcome as in the previous sections. So if $z \in \{0, 1\}^K$ then y is transformed into z with probability $\prod_{k=1}^K [z_k y_k + (1 - z_k)(1 - y_k)]$. All propositions are easily adapted to this setting. The previous examples of loss functions are still valid as the only addition is that the decision maker now has more information.

Example 13 *Schlag (2006) shows for the setting of Example 7 (iii) for $K = 2$ that the ability to observe foregone payoffs will not reduce minimax regret.*

Example 14 Consider the adversarial multi-armed bandit with a non oblivious adversary when foregone payoffs are observable. A sequential decision strategy f is called Hannan consistent (following Hannan, 1957, termed by Hart and Mas-Colell, 2001) if $\lim_{N \rightarrow \infty} \sup_{P^{N*}} H(f, P^{N*}) \leq 0$.⁹ It follows from our analysis that the binomial transformation of a Hannan consistent strategy is also Hannan consistent with maximal Hannan regret weakly reduced for any given horizon N . Similarly, although outside the scope of the present paper, it can be shown that binomial transformation can be used to reduce conditional Hannan regret as considered in Hart and Mas-Colell (2000). Future research will investigate whether estimates are better calibrated (sensu Foster and Vohra, 1998) when considering their binomial transformation.

Example 15 For the setting of Example 8 as a paired experiment, the randomized Fisher test is no longer UMP unbiased. Instead McNemar's test is a UMP unbiased test (McNemar, 1947, see Lehmann and Romano, 2005, p. 138). Given our results this test can be extended to a parameter most powerful unbiased test.

4.3 Dependent Random Variables

In the following we note how binomial transformations can also be used when a bound on the degree of independence between the random variables is imposed. Assume that there is some $c > 0$ such that the decision maker knows that $Cov(Y_k, Y_r) \geq c$ for all $k, r \in \{1, \dots, K\}$. Since covariance remains unchanged under the binomial transformation it follows that the decision maker can solve the problem for the Bernoulli case to then obtain via the binomial transformation a solution for general distributions with support in $[0, 1]$.

4.4 Covariates

For completeness we also mention how to deal with covariates. To simplify exposition we associate random variable Y_k to a treatment indexed by k that is given to a random subject in an infinite population. Let X be a finite set of covariates that possibly influence the outcome of the random variable. Let q_ξ be the proportion of subjects that have covariate ξ . So Y_k is now $|X|$ dimensional random vector where $Y_{k,\xi}$ is the outcome when treatment k is given to a subject that has covariate ξ . It is important

⁹lim sup is used as traditionally only the case $\delta = 1$ is considered where standard limits need not exist.

to note that no assumptions are made on how treatments possibly affect covariates differently. In particular, the decision maker cannot rule out that $Y_{k,\xi}$ are independent across k and ξ . The decision maker now has to make an estimate for each covariate and similarly has to accept or reject the null hypothesis for each covariate.

The insight of this paper is that one only needs to solve for Bernoulli distributions and then binomially transform the solution. See Schlag (2006) for a specific solution in the context of minimax regret.

5 Conclusion

In this paper we expand on a specific randomization technique used in Schlag (2003, 2006) and show how it can be used to derive minimax risk estimates and hypothesis tests.

The setting is nonparametric except that we assume exogenous known bounds on the underlying random variables. The existence of such bounds means decision-theoretically that there is a most preferred and a least preferred outcome. Whether such bounds exist depends on the specific application. Natural bounds are given when there is some fixed measurement scale such as grades. Even if there is a scale that is unbounded above then one can assess a maximal value, e.g. 20% above the maximal temperature given historical data. However, the bounds may not be chosen as the maximal and minimal outcome observed in the sample.

In this nonparametric setting one needs to count on analytic results as trustworthy simulations of all distributions with support in $[0, 1]$ are hardly feasible. There are n^n possible distributions of a single random variable for grid size $1/n$ (in contrast to only n for the case of Bernoulli distributions). In this context, notice that the advantage of binomial estimates or tests is that one can simulate nonparametric performance by exploring the much smaller set of Bernoulli distributions when analytic results are lacking.

The binomial transformation is powerful as it allows to solve for minimax risk. The method involves strategic use of randomization. We recall some roles of randomization. Randomization can be used to convexify the state space, e.g. used in the fundamental lemma of Neymann Pearson (e.g. see Lehmann and Romano, 2005, Theorem 3.2.1), used by Samuels (1968) in a multi-armed bandit setting and is fundamental to the proof of the existence of Nash equilibria via incorporating mixed

strategies (e.g. see Glicksberg, 1952). Randomization helps to reduce expected risk by convexifying risk as in the case of the Hodges and Lehmann (1950) estimation of mean under concave risk. It can be used to stay unpredictable as in Foster and Vohra (1998). In this paper randomization plays the novel role of fictively reducing the set of possible environments faced by the decision maker which consequently facilitates learning. The decision maker transforms observations and then acts as if the transformed environment is the true one. This only works if the original measure of risk is invariant to this transformation, for the binomial transformation this means that risk may only depend on the state of nature via the underlying means.

By its nature the binomial transformation creates a randomized estimate or a randomized test. Whether or not estimates or tests contained in this paper are useful for applications depends on whether or not one likes to randomize per se. Notice that randomization lies at the heart of decision theory (e.g. see convexity axiom in Milnor, 1954) and at the heart of estimation and inference based on a sample as the sample itself is random. As none of the results in this paper are stated in terms of uniqueness, future research should investigate the degree of randomness needed to obtain such tight results. For alternative estimates or tests the results herein maintain their usefulness as benchmarks. For example, the bounds used to evaluate the deterministic empirical success rule in the context of minimax regret (see Example 7(iii), Schlag, 2006) require more than 12 times the data as compared to the minimax regret rule to guarantee the same maximal regret.

This paper shows for the first time how to derive exact nonparametric tests relating to means. Specifically, given a UMP (unbiased) test one can derive a parameter most powerful (unbiased) test. This test is randomized so for its application one should specify the parameters of the test so that the designer is indifferent between type I and type II errors. Minimal sample sizes can be derived in order to secure given bounds on size and power. Performance of alternative exact tests can be evaluated using the concept of relative parameter efficiency.

Similarly novel are the results herein that demonstrate how to improve performance of Hannan consistent rules. Future research should investigate whether binomial transformation will also improve the calibration of forecasting methods.

References

- [1] Anderson, T.W. (1967), “Confidence Limits for the Expected Value of an Arbitrary Bounded Random Variable with a Continuous Distribution Function,” *Bull. ISI* **43**, 249–251.
- [2] Auer, P., Cesa-Bianchi, N., Freund, Y., and R. Schapire (1995), “Gambling in a Rigged Casino: The Adversarial Multi-Armed Bandit Problem,” in 36th Annual IEEE Symposium on Foundations of Computer Science, November.
- [3] Bahadur, R.R. and L.J. Savage (1956), “The Nonexistence of Certain Statistical Procedures in Nonparametric Problems,” *Ann. Math. Stat.* **27(4)**, 1115-1122.
- [4] Bickel, P. (1992), “Inference and auditing: the Stringer Bound,” *Int. Stat. Rev.* **60**, 197–209.
- [5] Bickel, P., J. Godfrey, J. Neter and H. Clayton (1989), “Hoeffding Bounds for Monetary Unit Sampling in Auditing,” *Contr. Pap. I.S.I. Meeting*, Paris.
- [6] Canner, P.L. (1970), “Selecting one of Two Treatments when the Responses are Dichotomous,” *J. Amer. Stat. Assoc.* **65(329)**, 293-306.
- [7] Cucconi, O. (1968), “Contributi all’Analisi Sequenziale nel Controllo di Accettazione per Variabili,” *Atti dell’ Ass. Italiana per il Controllo della Qualità* **6**, 171-186 (in Italian).
- [8] Fisher, R.A. (1935), “The Logic of Inductive Inference,” *J. Roy. Stat. Soc.* **98(1)**, Ser. A, 39-54.
- [9] Foster, D. and R. Vohra (1998), “Asymptotic Callibration,” *Biometrika* **85(2)**, 379-390.
- [10] Glicksberg, I.L. (1952), “A further generalization of the Kakutani fixed point theorem with application to Nash equilibrium points,” *Proc. Nat. Acad. Sci.* **38**, 170-174.
- [11] Hannan, J. (1957), “Approximation to Bayes Risk in Repeated Plays,” in *Contributions to the Theory of Games of Games*, (M. Dresher, A. W Tucker, and P. Wolfe, Eds.), Vol. 3, pp. 97–139. Princeton, NJ: Princeton Univ. Press.

- [12] Hart, S. and A. Mas-Colell (2000), "A Simple Adaptive Procedure Leading to Correlated Equilibrium," *Econometrica* **68**, 1127-1150.
- [13] Hart, S. and A. Mas-Colell (2001), "A General Class of Adaptive Strategies," *J. Econ. Theory* **98**, 26-54.
- [14] Hodges, J.L.Jr. and E.L. Lehmann (1950), "Some Problems in Minimax Point Estimation," *Ann. Math. Stat.* **21(2)**, 182-197.
- [15] Hoeffding, W. (1963), "Probability Inequalities for Sums of Bounded Random Variables," *J. Amer. Stat. Assoc.* **38(301)**, 13-30.
- [16] Manski, C. (2004), "Statistical Treatment Rules for Heterogeneous Populations," *Econometrica* **72(4)**, 1221-1246.
- [17] McNemar, Q. (1947), "Note on the Sampling Error of the Difference between Correlation Proportions or Percentages," *Psychometrika* **12**, 153-157.
- [18] Milnor, J. (1954), Games Against Nature. In Decision Processes, ed. R.M. Thrall, C.H. Coombs & R.L. Davis. New York: John Wiley & Sons.
- [19] von Neumann, J. and O. Morgenstern (1944), *Theory of Games and Economic Behavior*, Princeton, NJ: Princeton University Press.
- [20] Pesarin, F. (1984), "On Randomized Statistical Procedures," *Proc. Seventh Conf. Prob. Th.*, Editura Academiei Republicii Romania, M. Iosifescu Ed., 295-306.
- [21] Robbins, H. (1952), "Some Aspects of the Sequential Design of Experiments," *Bull. Amer. Math. Soc.* **55**, 527-535.
- [22] Rohtagi, V.K. (1976), *An Introduction to Probability Theory and Mathematical Statistics*, New York: John Wiley & Sons..
- [23] Romano, J.P. and M. Wolf (2000), "Finite Sample Non-Parametric Inference and Large Sample Efficiency," *Ann. Stat.* **28(3)**, 756-778.
- [24] Samuels, S.M. (1968), "Randomized Rules for the Two Armed Bandit Setting with Finite Memory," *Ann. Math Stat.* **39(6)**, 2103-2107.
- [25] Savage, L. J. (1951), "The Theory of Statistical Decision," *J. Amer. Stat. Assoc.* **46(253)**, 55-67.

- [26] Savage, L.J. (1954), *The Foundations of Statistics*, New York: John Wiley & Sons..
- [27] Schlag, K. H. (2003), *How to Minimize Maximum Regret under Repeated Decision-Making*, Mimeo, European University Institute, <http://www.iue.it/Personal/Schlag/papers/regret7.pdf>.
- [28] Schlag, K.H. (2006a), *Eleven - Tests needed for a Recommendation*, European University Institute Working Paper ECO **2006-2**, January 17.
- [29] Shapiro, S.S., M.B. Wilk and H.J. Chen (1968), "A Comparative Study of Various Tests for Normality," *J. Amer. Stat. Ass.* **63(324)**, 1343-1372.
- [30] Spiegelhalter, D.J. (1980), "An Omnibus Test for Normality for Small Samples," *Biometrika* **67(2)**, 493-496.
- [31] Statistical Models and Analysis in Auditing (1989), *Stat. Sci.* **4(1)**, 2-33.
- [32] Stoye, J. (2005), *Minimax Regret Treatment Choice with Finite Samples*, Mimeo, November 15.
- [33] Stringer, K.W. (1963), "Practical Aspects of Statistical Sampling in Auditing," *Proc. Bus. Econ. Stat. Sec.*, Amer. Stat. Ass..
- [34] Tocher, K.D. (1950), "Extension of the Neyman-Pearson Theory of Tests to Discontinuous Variates," *Biometrika* **37(1/2)**, 130-144.
- [35] Wald, A. (1947), *Sequential Analysis*, New York: John Wiley.
- [36] Wald, A. (1950), *Statistical decision functions*, New York: John Wiley & Sons.

6 Appendix: A PMP Test for a Mean

Consider a single random variable Y generating outcomes in $[0, 1]$ with unknown distribution where Y can attain more than two different outcomes, possibly is even continuous. Let μ be the mean of Y and let $\mu_0 \in (0, 1)$ be some given value of interest. Given a sample of N independent observations of outcomes of Y we wish to test the null hypothesis $H_0 : \mu \leq \mu_0$ against the alternative hypothesis $H_1 : \mu > \mu_0$. Before illustrating how to apply the parameter most powerful (PMP) test of Example 5(i) for this setting we first further motivate the nonparametric approach.

6.1 Nonparametric Statistics

The by far most common test for the above hypotheses is the t test. The t test is of course arguable the best test if Y is known to be normally distributed as it is a uniformly most powerful unbiased test of these hypotheses.¹⁰ However, what if Y is not known to normally distributed? In fact, it is difficult to imagine a situation in which it is actually known that Y is normally distributed. A common argument for applying the t test the claim that the random variable Y is approximately normally distributed, referring either to the central limit theorem given a large sample or to the outcome of an initial test for normality.

Consider first the argument using the central limit theorem. Accordingly, the t test statistic is approximately normally distributed if the sample is sufficiently large. However this approximation is pointwise and not uniform. For any sample size N there is a distribution P with support contained in $[0, 1]$ such that the hypothesis will be rejected with probability close to 1. The size of the t test is thus equal to 1 for any given N (Lehmann and Loh, 1990). This worst case distribution has mean μ_0 , puts a very large mass at an outcome very close to but larger than μ_0 and puts the remaining mass at the lower end of the possible support 0. With high probability all outcomes in the sample are identical and slightly larger than μ_0 and since there is no variation the t test rejects H_0 . This shows that without further knowledge of the distribution and citation of appropriate results, significance results based on the t test are not viable. Notice that assuming continuity will not mend this problem. There do exist results showing how the class of distributions can be limited to obtain uniform convergence of the t statistic to the normal distribution (see Lehmann and Loh, 1990), however it is left to be seen whether these particular conditions can be justified in applications in terms of a priori knowledge.

Now consider the alternative response to only apply the t test if the data first passes a test for normality (e.g. based on the Shapiro-Wilk test). These tests for normality are actually tests for non normality as the null hypothesis is that the data is normally distributed. Simulations show that their power can be below 50% even when the sample size is 50 (Shapiro et al. 1968, Spiegelhalter, 1980). In other words,

¹⁰The t test statistic

$$\frac{\sqrt{n}(\bar{y} - \mu)}{\sqrt{\frac{1}{n-1} \sum_{k=1}^n (y_k - \bar{y})^2}}$$

is distributed according to the Student t distribution with $n - 1$ degrees of freedom.

there are non normal distributions (such as the double exponential) under which these tests wrongfully accept the hypothesis of normality and thus wrongfully justify that the analysis of the data with the t test for the majority of the samples gathered.

Notice that sometimes the data generating process is known to be very skewed so that normality is not even postulated. This is for instance the case in auditing where alternative methods have been developed early (Stringer, 1963, for an overview see *Statistical Models and Analysis in Auditing*, 1989).

The above can be seen as a plea for exact tests. There are only few exact tests in terms of size, for our one sided hypothesis test applied to auditing scenarios see Stringer (1963) and Bickel (1992), for two sided tests see Anderson (1967) and Romano and Wolf (2000). However, no previous tests are known that are unbiased. No results on the power of tests for finite samples have been provided. In particular, tests have only been compared based on their asymptotic power properties. However, asymptotically the correct hypothesis can be learned with probability one so these asymptotic comparisons are only concerned with comparing infinitely small intervals.

6.2 Our PMP Test

In the following we briefly illustrate how one can apply the parameter most powerful test for the case of $\mu_0 = \frac{1}{2}$ so we are testing $H_0 : \mu \leq 0.5$ against $H_1 : \mu > 0.5$.

As the parameter most powerful test is non deterministic the decision maker should choose the parameters of the test to incorporate the fact that typically the recommendation will put strictly positive probability on accepting the null hypothesis as well as on rejected the null hypothesis. We set up the hypothesis test as a decision problem which will result in designing a test that up to integer constraints makes the decision maker indifferent between accepting and rejecting the null hypothesis.

First the decision maker assesses an *invariance zone* (as in Hodges and Romano, 2005, ch. 8.1) which is a subset of the parameters for which the decision maker is indifferent between accepting and rejecting the null. In this example we assume that the invariance zone contains all distributions such that $0.5 < \mu < 0.7$. So for any distribution with mean $\mu \in (0.5, 0.7)$ the decision maker does not care whether the null hypothesis is rejected or accepted.

Next the decision maker determines $\alpha_1 \in (0, 1]$ such that the decision maker is indifferent between wrongly rejecting the null with probability α_1 and wrongly accepting the null with certainty. Such a threshold of indifference exists by the continuity

axiom underlying decision theory based on the von Neumann Morgenstern axioms. Invoking the independence axiom, it follows that the decision maker is indifferent between wrongly rejecting the null with probability α and wrongly accepting the null with probability α/α_1 . In our example we assume that $\alpha_1 = \frac{1}{4}$.

Finally the decision maker chooses the parameters of the test. If the number of tests N is given then k and γ are chosen to minimize $\max\{\alpha, (1 - \bar{\beta})/\alpha_1\}$ where $\alpha = \Pr(f^* = 1 | \mu(P) = \frac{1}{2})$ is the maximal type I error and $1 - \bar{\beta}$ is the maximal type II error so $\bar{\beta} = \Pr(f^* = 1 | \mu(P) = 0.7)$. So α is then the lowest size of the test under which the decision maker will be indifferent between the two types of errors. To illustrate, assume $N = 20$. Then it turns out that $k = 13$, $\gamma \approx 0.35$ and $\alpha = (1 - \bar{\beta})/4 = 0.0836$. Given the preferences of the decision maker assumed above who counts type I errors four times as important as type II errors we find for sample size 20 that the size of the test will be slightly above 8%.

If instead the number of tests N can be chosen then the decision maker can choose the minimal number of tests N to guarantee some value of α and of $(1 - \bar{\beta})$. For instance, if this maximum should be at most 0.05 to guarantee size at most 0.05 and power at least 0.8 then set $N = 37$, $k = 23$, $\gamma \approx 0.013$ and $\alpha = 0.05$. It follows that $1 - \bar{\beta} \approx 0.191$. So 37 observations will suffice to ensure size 0.05 and at the same time make the decision maker indifferent between accepting and rejecting the null hypothesis.

Two comments are made. (i) Since $1 - \bar{\beta} < 0.2$ holds when $N = 37$ the decision maker can use slightly less tests in expectation by mixing between choosing $N = 36$ and $N = 37$. (ii) Since our test is parameter most powerful, the above shows that no level 0.05 test can achieve a type II error below 0.2 with 36 or less tests.

It is beyond the scope of this paper to present all properties of the above test. Here we illustrate two other values. If the aim is to achieve $\alpha \leq 0.05$ and $1 - \bar{\beta} \leq 0.1$ then $N = 153$ is the smallest number of tests needed, for $\alpha \leq 0.02$ and $1 - \bar{\beta} \leq 0.2$ it is easily shown that $N = 51$ data points have to be gathered.

6.3 An Alternative Deterministic Test

Next we construct and analyze a deterministic test and compare its performance to the PMP test above.

6.3.1 Construction

In the above PMP test the null hypothesis was rejected if the binomial average was above some cutoff $\frac{k}{N}$. In the following we investigate the analogous but deterministic test that is based on the sample mean. So there is some m with $m > \mu_0$ such that the null hypothesis is rejected if and only if $\bar{y} \geq m$.

We first show when α is sufficiently small that such a test is necessarily extremely conservative as the null hypothesis will be never be rejected. Assume $m \leq 1$ and consider the distribution P with $\Pr_P(Y = m) = \frac{\mu_0}{m} = 1 - \Pr_P(Y = 0)$. Then $\Pr_P(\bar{Y} = m) = \left(\frac{\mu_0}{m}\right)^N$ so if this is a level α test then $m = m_\alpha$ has to satisfy $\left(\frac{\mu_0}{m_\alpha}\right)^N \leq \alpha$. Consequently, the only such deterministic test that has size below α when $\alpha < (\mu_0)^N$ requires to never reject the null (so set $m_\alpha > 1$).

Analogous to Bickel et al. (1989) we set the parameter m of this test by calculating size using using the Hoeffding (1963) bounds.¹¹ Assume $\alpha > (\mu_0)^N$. Hoeffding (1963) shows (for random variables Y with range $[0, 1]$) that $\Pr_P(\bar{Y} \geq m) \leq g(m, \mu)$ for $\mu < m < 1$ where $g(m, \mu) = \left(\left(\frac{\mu}{m}\right)^m \left(\frac{1-\mu}{1-m}\right)^{1-m}\right)^N$. Note that g is continuous, strictly increasing in μ for given m , strictly decreasing in m for given μ , $g(\mu, \mu) = 1$ and $\lim_{m \rightarrow 1} g(1, \mu) = \mu^N$. Consequently, given $\alpha \geq (\mu_0)^N$ there exists a unique $m_\alpha \in (\mu_0, 1)$ such that $g(m_\alpha, \mu_0) = \alpha$. Now consider Y with mean $\mu \leq \mu_0$. Then $\Pr_P(\bar{Y} \geq m_\alpha) \leq g(m_\alpha, \mu) \leq g(m_\alpha, \mu_0) = \alpha$. So if we reject the null hypothesis if and only if $\bar{y} \geq m_\alpha$ then we are wrongly rejecting it with probability at most α . With this value m_α as cutoff for the test we obtain the smallest cutoff value under the Hoeffding bound that ensures that our test has level α .

Next we investigate the type II error $1 - \beta_f(P)$. Consider Y with mean $\mu > \mu_0$. If $\mu \leq m_\alpha$ then the minimal power is clearly equal to 0. Assume $\mu > m_\alpha$. Then using the Hoeffding bound we obtain

$$\Pr_P(\bar{Y} \leq m_\alpha) = \Pr_P(1 - \bar{Y} \geq 1 - m_\alpha) \leq g(1 - m_\alpha, 1 - \mu) = g(m_\alpha, \mu)$$

so $1 - \beta_f(P) \leq g(m_\alpha, \mu)$ when $\mu(P) > m_\alpha$.

Unlike the parameter most powerful test, this deterministic test is not unbiased. Unbiased requires $\left(\frac{\mu_0}{m_\alpha}\right)^N = \alpha$. On the other hand, the probability that \bar{X} is at least m_α under a binomial distribution with mean μ_0 does not vary locally in m_α provided

¹¹While Bickel et al. (1989) (see also Bickel, 1992) are interested in upper confidence bounds for the mean we focus on finding a test that yields a specific size and power.

$Nm_\alpha \notin \mathbb{N}_0$. Thus this proves analytically that this test is not unbiased for almost all values of α .

6.3.2 Comparison in an Example

Return now to our example where $\mu_0 = \frac{1}{2}$, $\alpha = 0.05$ and the invariance zone is equal to $(0.5, 0.7)$. To compare to the PMP test consider first $N = 37$. When deriving the cutoff m_α by numerically solving $g(m_\alpha, \frac{1}{2}) = 0.05$ we find $m_\alpha \approx 0.698$ and hence $Nm_\alpha \approx 25.8$. Facing a binomial distribution the deterministic rule is more conservative. If the sample yields $k \in \{24, 25\}$ successes our parameter most powerful rule rejects the null while the alternative deterministic rule accepts the null. Notice that the choice of cutoff $k = 25$ would imply under our PMP rule that $\alpha \approx 0.01$. As $m_\alpha \approx 0.7$ the lower bound on power derived above is not very useful (it yields $2 * 10^{-4}$). It turns out that a sample size of $N = 106$ (where $m_\alpha \approx 0.618$) is needed to guarantee $\alpha \leq 0.05$ and type II error $1 - \bar{\beta} \leq 0.2$. Thus, over 3 times the sample is needed to ensure the same performance as the PMP test. In other words, the relative parameter efficiency of this deterministic test at $\alpha = 0.05$ and $\mu = 0.7$ is equal to $\frac{37}{106} \approx 0.35$.

This divergence between the performance of the non deterministic rule based on binomial transformations and the deterministic rule based on the sample mean with performance measured using the Hoeffding (1963) bound is reminiscent of the results in Schlag (2006). Schlag (2006) shows that in order to ensure maximal regret below 0.05 the binomial average rule requires $N = 11$ while the rule based on the sample mean requires $N = 148$.