Centrum voor Wiskunde en Informatica
Centre for Mathematics and Computer Science

T.A. Louis

Efficient monotone sequential design

Department of Mathematical Statistics Report MS-R8705 April

The Centre for Mathematics and Computer Science is a research institute of the Stichting Mathematisch Centrum, which was founded on February 11, 1946, as a nonprofit institution aiming at the promotion of mathematics, computer science, and their applications. It is sponsored by the Dutch Government through the Netherlands Organization for the Advancement of Pure Research (Z.W.O.).

Efficient Monotone Sequential Design

Thomas A. Louis

Department of Biostatistics, Harvard School of Public Health

1980 Mathematics Subject Classification: 62K05, 62L05, 62L15.

Key Words & Phrases: optimal design, sequential methods, monotone follower problem, regret, scale invariance.

Note: This work was supported by grants ES02709 from the National Institute of Environmental Health Sciences, and INT-8512148 from the National Science Foundation. Parts of this work were completed while visiting the Department of Mathematical Statistics, Centre for Mathematics and Computer Science, Amsterdam, The Netherlands. I thank Julia Bailey for helping with the computations, and John Orav for his advice.

Report MS-R8705

Centre for Mathematics and Computer Science

P.O. Box 4079, 1009 AB Amsterdam, The Netherlands

ABSTRACT

Optimal designs for non-linear problems depend on unknown parameters. With no constraints, sequential methods can pick the next design point using prior and accruing information and "home in" on the optimal design. In monotone designs, observations must be ordered in time (or other metameter), so design options decrease. For example, in rodent bioassay experiments where a group of n rodents are simultaneously put on test, sacrifices to discover the presence or absence of tumors can occur only at ages greater than or equal to the current age. So, if data are taken beyond the optimal age, it is not possible subsequently to go back to it.

For the class of problems studied, statistical information depends on the time data are taken and unknown parameters. In this report we consider a class of monotone designs based on a scale invariant design objective, and two structures for statistical information. One factors into a function of time and a function of the unknown parameter and results from data from a scale family for each t. The other depends on time divided by the unknown parameter, and results, for example, from destructive life tests. We develop and investigate two- and three-stage adaptive rules, compute the asymptotic order of the regret for these rules relative to the rule taking all data at the optimal time, and show that the order is close to the best possible. We find the rate at which the performance of these rules departs from scale invariance, and report on a simulation study comparing the current rule to others in the literature.

1. INTRODUCTION

In some sequential design problems the optimal design depends on unknown parameters and options constrain as time evolves. For example, the carcinogen bioassay in small rodents puts animals on test and follows them until natural or sacrificial death. For a specific design goal and a scalar parameter of interest, there is an optimal age of sacrifice. Adaptive interim sacrifice plans use accruing data to determine sacrifice ages, but options constrain since a rodent cannot be killed at an age younger than the current. Therefore, a balance must be struck between early sacrifices useful in determining the optimal sacrifice age, and sacrifices at the estimated optimal age. Bergman and Turnbull (1983) and Louis and Orav (1985) study this destructive life testing example, and we use it for motivation and a simulation study. Our general model applies to other irreversible processes such as timing measurements in followup studies, studying metal fatigue, and estimating the shelf-life of chemicals. In each of these applications observations can indicate exceedance of threshold level or a report a direct measurement.

Time-ordered designs generalize the monotone follower problem (Benes et al 1980, Karatzas 1981), where a Brownian motion is tracked by a non-decreasing function. Unlike the monotone follower problem, in the current setting adaptive rules use accruing information to update parameter estimates and to determine the next observation time. Each observation provides both (Fisher) information that improves the determination of the optimal time and produces an increment in the overall design objective function. But, the amount of information on parameters depends on the observation time. Under time-ordering, rules must strike a balance between the benefit of taking data at the current estimate of the optimal time and the risk of overshooting the true value. In the estimation context, Fisher information is also the design objective, but we analyze a more general model that includes hypothesis testing.

Though the free boundary differential equation for obtaining the optimal rule can be derived from a Brownian motion embedding, we have not obtained its solution. We do use the Brownian embedding to compute the asymptotic regret for two and three-stage rules. Though we do not prove that these rules are optimal, we show that their regret is close to the lowest possible for rules that are asymptotically equivariant. The order of regret depends on the rate at which the Fisher information for unknown parameters goes to zero as the design time goes to zero. The best order of regret for any monotone rule is $\sqrt{n}$, whereas without the monotonicity constraint the order is log(n).

We introduce the problem in section 2, section 3 develops single and multi-stage rules, section 4 discusses equivariance, section 5 presents simulation results for a destructive life testing model, section 6 contains discussion, and the appendices contain basic theorems.

2. THE MODEL AND EXAMPLES

2.1 General Setting

Let θ parameterize the design objective function $[i(t|\theta), \theta \in \Theta]$ for an observation taken at time t. A total of n observations are to be taken. If they are taken at time points $\mathbf{t} = (t_1, \cdots, t_n)$, $t_\nu \geq 0$, the experimenter gains:

$$\sum_{\nu=1}^{n} i(t_\nu|\theta) \tag{2.1}$$

In addition to contributing to (2.1), at time t we observe a random variable X_t with distribution $F_t(\cdot|\theta)$ and Fisher information $j(t|\theta)$. This observation augments the information on θ. Experimentation is constrained by the requirement that $t_1 \leq t_2 \leq \cdots \leq t_n$. The design goal is to maximize the expectation of (2.1) over samples in the frequentist setting and in addition over the parameter space Θ with respect to a prior in the Bayesian formulation. We assume throughout that:

$$i(t|\theta) = c(\theta)i(t/\theta), \tag{2.2}$$

and that either

$$j(t|\theta) = \theta^{-2}j(t), \tag{2.3a}$$

or

$$j(t|\theta) = \theta^{-2}j(t/\theta), \tag{2.3b}$$

with, as u→0

$$i(u) = O(u^q), j(u) = O(u^p); \text{ some } p,q \geq 0. \tag{2.4}$$

Further, we assume that "i" and "j" are non-negative, have unique maxima at "T" and "S" respectively, are locally quadratic at the maximum, and have a continuous, bounded second derivative. By rescaling time and normalizing we can assume that T = 1 and $i''(1) = 2$. In addition, we assume that the maximum likelihood estimate of θ ($\hat{\theta}_d$) satisfies the invariance Theorem A4.

This structure in encompasses many design goals including efficient hypothesis testing and parameter estimation. For estimation i=j. When X_t is scale invariant for all t, (2.3a) holds, while (2.3b) applies, for example, to destructive life testing models based on a scale invariant survival distribution (see section 2.3).

For any rule **R** we define the scale adjusted regret (Δ) and the relative efficiency (RE) by:

$$\Delta_n(\mathbf{R},\theta) = n\cdot i(1) - \sum_{\nu} E_\theta\left\{i(t_\nu/\theta)\right\}$$

$$RE_n(\mathbf{R},\theta) = 100\frac{\sum_{\nu} E_\theta\left\{i(t_\nu/\theta)\right\}}{n\cdot i(1)}.$$

Chernoff (1972) and Abdelbasit and Plackett (1983) consider a similar problem, but without the time-order constraint. They show that a rule taking the next observation at $\hat{\theta}$, where $\hat{\theta}$ is the MLE is asymptotically efficient in that RE_n goes to 100 as $n\to\infty$. The rule has a regret of order log(n).

In the sequel we use the function best[d,A_d] to represent an unconstrained estimated optimal time. It depends on the number of observations taken (d) and the accumulated information on θ,(A_d). For example, using the MLE "best"$=\hat{\theta}_d$. A Bayesian computation would find the time that maximized the conditional expected increment in (2.1). Our large-sample results do not depend on the form of "best" so long as it is asymptotically $\sqrt{n}$ equivalent to the MLE. Information A_0 can be used to stabilize the maximum likelihood estimate or represent a prior distribution.

2.2 A basic example

To fix ideas and gain insight into the asymptotic performance of time-ordered rules consider the case with a quadratic gain function "i" and with F_t being exponential independent of time. Specifically, let:

$$f_t(x|\theta)=\theta^{-1}e^{-x/\theta}$$

$$i(u)=1-(1-\frac{u}{\theta})^2, 0\leq\theta, u\leq\infty,$$

so $c(\theta)\equiv 1, j(u)\equiv 1$, p=0, we are in situation (2.3a). Consider the adaptive rule that first takes m_n observations at 0 and then uses:

$$t_{v+1}=max[t_v,\bar{X}_v],\ m_n\leq v\leq n,$$

where

$$\bar{X}_v=\frac{1}{v}\sum_{j=1}^{v}X_{t_j}.$$

Assume $\theta=1$. The first m observations contribute m units to the regret, and Theorem A2 proves that the remaining (n-m) observations produce an expected regret: (n-m)/m. Therefore, the optimal m_n is $\sqrt{n}$, and produces a regret of $2\sqrt{n}$. Table 1 presents a comparison of this computation with values obtained by simulation. The agreement is excellent over a wide range of sample sizes. For comparison, the expected regret for an unconstrained rule is easily shown to be $\sum d^{-1}\approx\log(n)$. So, the penalty produced by the order constraint is the difference between $2\sqrt{n}$ and log(n).

Rules with $t_1 \equiv 0$, and

$$t_{d+1} = max[t_d, s_{d+1}X_d] = \max_{1 \le j \le d} \left[s_{j+1}X_j\right],$$

where the s's are non-negative are scale equivariant. Theorem A3 develops conditions on the s's for the regret to be as small an order as possible, and proves that the minimum is $O(\sqrt{n})$. So, the two-stage rule described above produces the best order of regret, though it may not be optimal.

In this example we can achieve equivariance, since initial data can be taken at 0 (safely below the optimal time for any θ), and then the scale family X's coupled with the scale invariant design objective function produce invariant relative efficiencies and scale adjusted regrets. In section 3 we generalize these findings to situations where the information $j(t)$ goes to zero as t goes to zero, so observations at 0 are uninformative.

2.3 Destructive Life Testing

To see the complications when p>0, consider the following destructive life testing model studied by Chernoff (1972) without the order constraint and by Bergmen and Turnbull (1983) and Louis and Orav (1985) with the constraint. Let θ be the mean of an exponential holding time in a state. At time t a destructive test can be performed to determine if a study unit is still in the state. For example, in carcinogenicity testing a rodent can be sacrificed at time t to see if it has left a disease-free state and entered a tumor state. Our goal is to estimate θ maximizing Fisher information as the design objective.

For this model we have at time t that information comes from a Bernoulli random variable with probability mass function depending on the exponential cdf:

$$f_t(x|\theta)=\left[1-e^{-t/\theta}\right]^x e^{\frac{-(1-x)t}{\theta}},\quad x=0\ or\ 1,$$

where x=0 indicates the absence and x=1 indicates the presence of a tumor. Here we are in situation (2.3b), $i=j$ and:

$$j(u)=\frac{u^2e^{-u}}{1-e^{-u}}.$$

In this example $T\approx$ 1.5936 (and can be made 1 by changing the time scale) and p=1. The example can be generalized by letting $f_t(x|\theta)$ equal $H(t/\theta)^x[1-H(t/\theta)]^{1-x}$, where H is a distribution function. In section 5 we compare the performance of the rules proposed by Bergman and Turnbull and Louis and Orav to the three- stage rule of the next section.

3. TWO AND THREE STAGE RULES

For model (2.3a) the two-stage rule (a special case of the three-stage) puts a fixed number of observations at a fixed time and then takes all remaining observations at the minimum of the current time and the estimated optimal time (a terminal sacrifice in the bioassay application). After an initial group at a fixed time, the three-stage rule assigns at the sequentially updated maximum of the current time and a fraction of the current "best" for a fixed number of observations, and then performs a terminal sacrifice. The three-stage rule produces better invariance properties (see section 4). For each rule, the terminal sacrifice can be replaced by assignments at the sequentially updated maximum of the current time and the current "best" without changing asymptotic performance. This replacement does improve finite sample performance. We first assume (2.3a) and then discuss (2.3b).

3.1 A Two-stage Rule

Defn: R(a,m): Take the first "m" observations at "a" and the remaining (n-m) at $t_{m+1}=\max[t_m,\text{best}(m,A_m)]$. Rule R^*(m,n) takes the remaining (n-m) at the sequentially updated $t_{v+1}=max[t_v,\text{best}(v,A_v)]$.

Theorem 3.1: If $m_n a_n^p \to \infty$, for each $\theta > a_\infty$, then for R and R* the asymptotic regret is of order: $max[m_n, \frac{n-m_n}{m_n a_n^p}]$.

Proof: By definition $t_m \equiv a_n$ so for $\theta > a_n$ the first m_n observations contribute m_n to the order of regret. We wish to apply Theorem A2, and have to consider:

i) Holding at a_n because $\hat{\theta}_v$ remains too small.

ii) Overshooting θ and generating too large a regret for the remaining $(n-m_n)$ observations.

Holding at a_n : Let $J_v = \sum_{k=1}^{v} j(t_v)$. If we hold at a_n, we continue to accrue information on θ at a rate at least as large as a_n^p per observation. Now, since $\hat{\theta}_v$ converges to θ, we have for $\epsilon > 0$, and $\tau = inf[v > m_n : \hat{\theta}_v > \theta \cdot (1-\epsilon)]$ that $E(\tau) = m_n + B/a_n^p$, $B < \infty$. So, we hold for at most a finite time, with a finite expectation.

Overshooting: After τ theorem A2 takes over, giving that the rule has a regret of order $\max[m_n + \frac{B}{a_n^p}, \frac{n-m_n}{m_n a_n^p + B}]$, whether we use R or R*.

Corollary 3.1.1: For a given a_n the best order is achieved by picking $m_n = (n/a_n^p)^{1/2}$, which is also the best order.

Proof: To control the order both terms in the "max" above must be of the same order. Solving for d gives $(n/a_n^p)^{1/2} - B/a_n^p$. But the second term is negligible, since $na_n^p \to \infty$.

Corollary 3.1.2: If we require asymptotic efficiency for each $\theta > 0$, then for p>0 the order of regret can be arbitrarily close to but not equal $\sqrt{n}$.

Proof: For asymptotic efficiency it is necessary that $a_n \to 0$. The regret is $(n/a_n^p)^{1/2}$, and we can let $a_n \to 0$ arbitrarily slowly.

Corollary 3.1.3: If we require asymptotic efficiency only for $\theta \geq \theta_0 > 0$, we can achieve regret of order $\sqrt{n}$.

Proof: Set $a_n \equiv \theta_0$, and $m_n = \sqrt{n}$.

3.2 Three-stage rules

The two-stage rule stays at a fixed time for the first m_n observations, no matter the value of θ. The following rule allows adaptation to the parameter estimate, and produces regrets closer to scale invariant.

Defn: R(s,d;a,m), m $\leq$ d: Take "m" observations at "a" and then successively at:

$$t_{\nu+1} = max[t_\nu, s \cdot best(\nu, A_\nu)], \quad m < \nu \leq d,$$

and the remaining (n-d) at $max[t_d, best(d, A_d)]$.

Rule R^* replaces the last stage by the sequentially updated $t_{\nu+1} = max[t_\nu, best(\nu, A_\nu)], \ d < \nu \leq n.$

The initial m observations are used to stabilize the second stage, and can be replaced by a prior distribution. Note that the two-stage rule is equivalent to setting s=a and m=d.

Let $\xi =$ -log(s), $\lambda = -\log(\theta)$, $\alpha =$-log(a), and $\delta = \log(t) + \lambda$.

Theorem 3.2: If $d_n s_n^p \to \infty$, $O(a_n) < O(s_n) \to 0$, and if there exists a $J < \infty$ such that for $J_\nu \geq J$, $E[(\hat{\lambda}_\nu - \lambda)^2] < \infty$, then with $m_n \geq J a_n^{-p}$, for all $\theta > 0$ R and R* have asymptotic log regret of order:

$$\log(\Delta_n) = .5 log(n) + \frac{p \cdot log(n)}{p + [p^2 + 4 log(n)]^{½}}.$$

Proof: We prove the theorem for rule R^*. Modifications are straightforward for rule R. Since the X_t are scale invariant, the distribution of $\hat{\lambda}_\nu - \lambda$ depends on λ only through the influence of sampling times on accrued information. Further, the variance of $\hat{\lambda}_\nu$ conditional on J_ν is proportional to J_ν^{-1}.

Let $M_{l,h} = \max_{l \le k \le h}(\hat{\lambda}_k)$, and define:

$$\delta_\nu(\lambda) = \begin{cases} \lambda - \alpha_n & ,nu \le m_n \\ \max(\lambda - \alpha_n, M_{m_n, d_n - 1} - \xi_n) & ,m_n < \nu \le d_n \\ \max(\delta_{m_n}, \delta_{d_n}, M_{d_n, n-1}) & ,d_n < \nu \le n \end{cases} \quad . \tag{3.1}$$

Except for the $(\lambda - \alpha_n)$ term, the sequence depends on λ only through the sampling times.

Using the Brownian motion embedding, we have that for $(h-l) \to \infty$:

$$M_{l,h} \overset{L}{=} W^*(J_l^{-1}),$$

where $\overset{L}{=}$ denotes convergence in distribution, and W^* is the positive part of a Brownian motion. Also,

Lemma: If $J_{m_n} = o(J_{d_n})$, and $J_{d_n} \to \infty$, then $M_{m_n, d_n - 1}$ and $M_{d_n, n}$ are asymptotically independent.

Proof: We have, for example, that $M_{m_n, d_n - 1} \overset{L}{=} M_{m_n, m_n + b_n}$, where $b \to \infty$ but is $o(d_n)$. The Brownian trajectories contributing to $M_{m_n, m_n + b_n}$ and $M_{d_n, n}$ are asymptotically uncorrelated, and therefore independent.

Now, back to the theorem. To simplify notation, let J_m = J. For fixed θ and $\alpha_n \to \infty$, eventually the $\lambda - \alpha_n$ term is negligible. Consider first the contribution to the regret of the last $(n - d_n)$ observations. This contribution is of order $(n - d_n)E[D(\delta_{d_n + 1})]$, where

$D(\delta) = i(1) - i(e^{\delta})$. And, from (3.1) coupled with $W^* \geq 0$, this is:

$$(n - d_n) \cdot E[D(\max\{[W^*(J^{-1}) - \xi_n]^+, W^*(J_{d_n}^{-1})\})], \tag{3.2}$$

where "+" denotes the positive part.

Now, (3.2) is greater than or equal to the maximum of the expectation of each component, so to control the order it is necessary that each of these be controlled. Since we have the maximum of two asymptotically independent terms, the order of (3.2) will equal the maximum of the component- wise expectations. So, we want:

$$E[D([W^*(J^{-1}) - \xi_n]^+)] = E[D(W^*(J_{d_n}^{-1}))]. \tag{3.3}$$

Since $\hat{\theta}$ converges to θ, $J_{d_n} = O(s_n^p d_n)$. So, as $n \to \infty$ we can use a Taylor series expansion to rewrite (3.3) approximately (recall that $i'' = 2$, and is bounded) as:

$$E[\{W^*(J^{-1}) - \xi_n)^+\}^2] = E[W^*(J_{d_n}^{-1})^2]. \tag{3.4}$$

Equation (3.4) gives one relation for determining (a_n, m_n, s_n, d_n). For a second relation we consider the first d_n observations. Since $\xi_n \to \infty$, the probability that the first component of (3.2) is postive goes to zero. This implies that the contribution of the first d_n observations is O(d_n). For fixed ξ_n the right-hand side of (3.4) is decreasing in d_n, implying that both terms in (3.4) should be O(d_n). We now set out to solve these relations.

Using approximations to the expectation and variance of a folded Gaussian variable conditional on its being greater than a constant, (see Theorem A2 and corollaries) we find that:

$$E[\{(W^*(J^{-1}) - \xi_n)^+\}^2] \overset{O}{=} 2\Phi(-J^{1/2}\xi_n)\frac{1}{J}\min[1, \frac{3}{\xi_n^2 J - 1}]. \tag{3.5}$$

and

$$E[W^*(J_{d_n}^{-1})^2] \overset{O}{=} \frac{1}{J_{d_n}} = \frac{e^{p\xi_n}}{d_n}, \tag{3.6}$$

where $\overset{O}{=}$ denotes equality in order.

To simplify the notation we assume that J=1, that $\xi_n > 1$, and in (3.2) replace $(n-d_n)$ by n. This substitution has negligible effect when $d_n = o(n)$. Using this modified (3.2), and that the first d_n observations contribute $O(d_n)$ to the regret, we require that $d_n = n^{½}\exp[.5p\xi_n]$. Then formulae (3.4-6) with $\Phi(-\xi) \approx \xi^{-1}\exp(-\xi^2/2)$, imply that

$$\xi_n^2 + p\xi_n - \log(n) + 2\log(\xi_n) + 2\log(\xi_n^2 - 1) + \text{const} = 0.$$

Ignoring the last three terms (easily shown to be negligible) gives:

$$\xi_n = \frac{2\log(n)}{p + (p^2 + 4\log(n))^{½}}. \tag{3.7}$$

From this we obtain for d_n, and therefore for the regret:

$$\log(d_n) \overset{O}{=} \left[\frac{1}{2} + \frac{p}{p + (p^2 + 4\log(n))^{½}}\right] \cdot \log(n). \tag{3.8}$$

Notice that ξ_n is decreasing in p, that it is of asymptotic order $\sqrt{log(n)}$, and that the order of regret converges to $n^{½}exp[p(log(n))^{½}]$ for all p. As $p \to \infty$ the order of regret goes to n, and as $p \to 0$, the order of regret goes to $\sqrt{n}$.

Since we know that we can start at 0 when p=0, we would like $\lim_{p \to 0} \xi_n = \infty$. We can accomplish this by premultiplying the "log(n)" in the square root by p to a suitable power, producing:

$$\xi_n = \frac{2\log(n)}{p + (p^2 + 4p^{\epsilon_n}\log(n))^{½}}, \tag{3.9}$$

where $\epsilon_n < 2$, and is $o(1/\sqrt{\log(n)})$. These conditions on ϵ_n ensure that the modification will be negligible as either n or p goes to infinity.

Notice that this three stage rule produces a regret close to the absolute best possible order $(\sqrt{n})$. So, even if multi-stage rules improve on (3.8), the improvement should be negligible for most situations. We can modify the rule to improve small sample performance by adding a constant to (3.9) and multiplying "a" and "m" by constants. Equation (3.4) can be derived under weaker conditions on i''.

For $p>0$ we have no hope of attaining a regret of uniform order over θ, for as $\theta \to 0$ ($\lambda \to \infty$), the $(\lambda - \alpha_n)$ in (3.1) dominates and the sup regret equals n. In a Bayes setting , we can use a_n to control the order of the a priori expected regret. Let G be the prior. For $\theta > a_n$, the order of regret is given by (3.8). For $\theta < a_n$, the regret is of order n. So, we want $nG(a_n)$ to equal the regret, or:

$$a_n = G^{-1}\left(n^{-\frac{1}{2}} e^{.5p\xi_n}\right).$$

For example, if $\lim_{a \to 0} G(a)/a^{\beta} = 1$ some $\beta > 0$, then $\log(a_n) = [p\xi_n - \log(n)]/2\beta$. We can easily show that this a_n is of order less than or equal to s_n. Recall that $m_n \geq J a_n^{-p}$, so $m_n = o(d_n)$ if and only if $p \leq \beta$. If the prior puts too much mass near zero, the a priori Bayes expected order must be larger than (3.8).

3.3 Modifications for 2.3b

We consider (2.3b) for the situation where $f_t(x|\theta) = H(t/\theta)^x[1-H(t/\theta)]^{1-x}$, with H a cumulative distribution functon and x = 0 or 1. In this discrete data setting we have to modify the theorems in the previous section. First, it is straightforward to show that if H is $O(u^{\epsilon})$ as $u \to 0$, then $p = \epsilon$. The initial stage of the three- stage rule must be modified, since having $m_n a_n = J$ does not ensure that $E[(\hat{\lambda}_\nu - \lambda)^2] < \infty$. We have two choices: either let $a_n m_n \to \infty$, or require that both the number of "successes" and "failures" each exceed a sufficiently large, finite constant. For the latter the expected waiting time is $O(a_n^{-p})$, producing the same asymptotic order as in the invariant sampling case.

The remaining proof of asymptotic order is basically identical to the foregoing theorems. Again, we do not have uniform convergence as $\theta \to 0$,($\lambda \to \infty$). But, since the cumulative information for λ now depends on λ and goes to zero as $\lambda \to -\infty$ (see 2.3b), we lose uniformity for $\theta \to \infty$. Further, due to the extra dependency of information on θ (refer to formula 2.3b), the departure from invariance will be greater than for (2.3a), but of the same asymptotic order.

4. ASYMPTOTIC INVARIANCE

For simplicity of notation we consider situation (2.3a) and study the departure from invariance:

$$|\Delta_n(\cdot,\theta) - \Delta_n(\cdot,\psi)|. \tag{4.1}$$

Results for (2.3b) are similar. For either the two or three stage rule, and for sufficiently large n, (4.1) has order:

$$m_n \cdot \left[i(\frac{a_n}{\theta}) - i(\frac{a_n}{\psi}) \right], \tag{4.2}$$

since the difference in the regrets for the remaining ($n-a_n$) observations is of a smaller order. Under (2.4), for small a_n, (4.2) is of order $m_n a_n^q$, and so this quantity controls the departure from asymptotic invariance. Since $a_n \to 0$ and m_n is of order less than the regret, (4.1) is of order less than the regret.

For the two-stage rule $m_n = (n/a_n^p)^{1/4}$, so $n^{1/4}\exp[\alpha_n(p-2q)/2]$. This term will go to infinity, if $p \geq 2q$, and for $p < 2q$ it remains finite if $\alpha_n \geq \log(n)/(2q-p)$. Thus, for (4.1) to be finite, $\log(m_n) \geq q\log(n)/(2q-p)$, $p < 2q$. But, if $p \geq 2q$ the order is greater than or equal to n, and an asymptotically non- negligible fraction of the data must be taken at a_n, producing an efficiency less than 100. More generally, (4.1) can be made finite only by producing a large regret, and as a practical matter $p \geq q$ is a problem.

For the three-stage rule (with J = 1), $m_n = a_n^{-p}$, and (4.1) is controlled by a_n^{q-p}. For $p<q$ (4.1) goes to 0; for $p=q$ (4.1) remains finite; and for $p>q$ it goes to infinity at a rate controlled by a_n. For example, setting $a_n = s_n$ (the largest order we allow), gives:

$$a_n^{(q-p)} \approx e^{(p-q)\sqrt{\log(n)}}, \tag{4.3}$$

a slow departure from invariance.

Notice that for estimation $i=j$, so $p=q$, and we have a finite departure from invariance. Multi-stage rules [with a series of stages like (s,d), with the s's increasing to 1] may allow (4.1) to remain finite for $p \leq P(q,K)$, where K is the number of stages, $P(q,2)=q$, and P is increasing in K. But, in the light of (4.3), the improvement over the three-stage rule should not be large.

5. SIMULATION STUDY

We now return to the destructive life testing example of section 2.3. Bergman and Turnbull (1983) propose a rule where sacrifices are made at predetermined times with the decision to move to the next age based on a sequential test (their "ratio" rule, formula 3.1). If the current time is less than the optimal, the test tends to terminate quickly; if the current time is greater than the optimal, the stopping time is defective. Accruing information is not used to construct the sequential test, so each time-point is considered separately. They give conditions for asymptotic efficiency of the rule and present simulation results. Louis and Orav (1985) study a rule that adjusts the damping factor after each observation, by comparing the expected increment in (2.1) produced by taking all remaining observations at the maximum of the current time and the estimated optimal time to the increment produced by the rule that takes one more observation at some (best) future time and then takes the remaining observations (one fewer than before) at the updated estimated optimal time.

We compared our three-stage rule R^s to these by simulating the cases n=50, 100 (100), 500 for values of θ used in the previous publications. The simulator was run with 15 replications, producing standard errors no greater than 1.2. The rule used a Gamma prior on θ^{-1} with a shape parameter equal to 3 and a mean of 1, set $a=1.5s$ (s divided by the mode of the prior), and J (of Theorem 3.2) equal to 2. Table 2 and Figure 1 compare performance of the three-stage rule with values taken from Table 2 in Bergman and Turnbull ($b=4$, $Z^*=4$), and from Orav and Louis. For clarity $\theta = .125$ is omitted from the plot. The table shows that the three-stage rule effectively dominates the Orav and Louis, but that it does not dominate the Bergman and Turnbull. The three-stage rule comes close, and does produce a reasonably flat plot over a broad range of θ values. Depending on the initial testing time and the rule, there will be a set of θ's with high efficiency. Simulations of the three-stage rule R (with a terminal sacrifice) indicate that it does not perform as well as R^* for the sample sizes considered, though the theory shows it has the same asymptotic order of regret. It has the practical advantage of shortening the experiment.

The Bergman and Turnbull rule starts at t=1, so for small θ its performance will be poor. By assuming that for small θ all observations are taken at 1, we obtain a coarse upper bound for efficiency. If a significant fraction of the observations are taken at 2, the efficiency is greatly reduced. Since the efficiency plots are concave, we have tabluated and plotted (with a dotted line) the minimun of the computed upper bound and the performance at $\theta = 1$. As can be seen, the new rule performs extremely well. Of course, the initial time for the Bergman and Turnbull rule could be reduced, improving performance for small θ, but only at the cost of reduced efficiency for larger θ's.

Notice that for all rules, performance degrades more rapidly as θ decreases from 1 than as it increases (compare equal distances from 1 on the log scale). For large θ, rules can "catch up" with the optimal value, but for small θ testing times are likely to be too large. As can be seen in Figures 1 and 2, increasing the number of rodents on test does not uniformly improve efficiency, though monotonicity does hold for extreme θ's. This lack of monotonicity will occur for all rules, as can be seen by considering increasing from n=1 to n=2. Under either a Bayesian or frequentist approach, one θ -value (θ_1) will have 100% efficiency for n=1. Generally, the observation time for n=1 will fall between the pair for n=2, degrading the performance at θ_1 . This phenomenon occurs for all sample sizes, but for the rules studied,

$RE_n \to 100$ for all θ's, and the curves are nearly monotone in n.

6. DISCUSSION

Although the proposed rules produce close to the best order of asymptotic regret, they fully exploit neither the detailed relation between "i" and "j" (eg. the relation between S and T), nor a prior distribution. The damping factors (the s's) should depend on accrued (monotonized) observed information, and not simply on sample size. More attention to these details would improve performance, but our simulations show that rules R and R^* perform very well, even for situation (2.3b). The rules exhibit flatter curves than might be expected in the discrete data setting. With discrete data, if testing time is too small, few or no events will be observed ($\hat{\theta}$ is large) and the sequentially updated time will increase rapidly. This rapid adaptation provides a partial explanation.

The class of rules considered can be extended to multi- stage, but we expect little improvement in performance. Theorem A3 shows we can do no better than order $\sqrt{n}$ (attained when $p=0$). Our rules attain close to this order for all p. We conjecture that the optimal order is increasing in p, but haven't obtained a formal proof.

In situation (2.3a), if $\theta_0 > 0$, we have uniform asymptotic order, and one could try to characterize rules with minimax regret. Then, the minimax rule that is "closest" to invariant becomes an attractive candidate. In situation (2.3b) θ's of interest must also be bounded from above before the minimax approach can be applied.

Our rules have been developed assuming that data are instantaneously available and that the next observation can be taken at any continuous time point. For practical application both of these assumptions must be relaxed. Also, robustness with respect to specification of "i" and "j" needs to be considered. See Louis and Orav (1985) for a discussion of these issues for the bioassay.

As in other sequential settings, we are left with the problem of producing valid frequentist inferences (hypothesis tests and confidence intervals). The most straightforward approach uses reciprocal observed information for the variance, and computes Gaussian intervals. The validity of this and likelihood- based methods requires investigation, though the near- ancillarity of the observed information and the Brownian embedding for the rules considered, suggest that these approaches will be asymptotically valid.

Finally, unlike the scalar parameter case, with vector parameters optimal designs require at least two time points. While unconstrained adaptive rules are easily generalized to this case, solving the monotone problem will be difficult, and the monotonicity constraint will produce regrets with larger orders than for the scalar case.

REFERENCES

Abdelbasit K., Plackett RL (1983). Experimental Design for binary data. *J. Am. Statist. Assoc 78:* 90-98.

Benes VE, Shepp LA, Witsenhausen HS (1980). Some solvable stochastic control problems. *Stochastics 4:* 39-83.

Bergman SW, Turnbull BW (1983). Efficient sequential designs for destructive life testing with application to animal serial sacrifice experiments. *Biometrika 70:* 305-314.

Berry DA, Fristedt B (1985). *Bandit Problems, Sequential Allocation of Experiments.* Chapman & Hall: London.

Billingsley P (1968). *Convergence of Probability Measures.* Wiley, New York.

Breiman L (1968). *Probability.* Addison Wesley, New York.

Chernoff H (1972). *Sequential Analysis and Optimal Design.* Soc. Indust. Appl. Math, Philadelphia.

Karatzas I (1981). The monotone follower problem in stochastic decision theory. *Appl. Math. Optim. 7:* 175-189.

Louis TA, Orav J (1985). Sacrifice plans for the carcinogen bioassay. *Proc. Long term animal carcinogenicity studies: A Statistical Perspective.* Am. Statistical Assoc, Washington DC: 36-42.

Siegmund D (1985). *Sequential Analysis.* Springer Verlag: New York

APPENDIX

For model (2.3a) we analyze the class of adaptive designs characterized by a fixed number m_n of observations at an initial time a_n, followed by allocation according to a general adaptive rule with damping factors $s_n(d,A_d)$, where A_d is the accrued information. A similar development applies to (2.3b). Assuming the invariance result for the maximum likelihood estimate of θ (Theorem A4), we write $\hat{\theta}_{J_n(d)}$ for a random process with the same distribution as the sequence of maximum likelihood estimates, where $J_n(d)$ is the accrued information (denoted by J_d in section 3). Then, we can define for $0 \leq u \leq 1$, with "[ ]" the greatest integer function:

$$\tau_n(u) = n^{-1} J_n^*([nu])$$

$$\psi_n[u,\tau_n(u)] = s_n([nu], n\tau_n(u))$$

and

$$y_n(\tau_n(u)) = \hat{\theta}_{n\tau_n(u)} .$$

Theorem A1: The order of regret is asymptotically equal to the expectation of:

$$m_n + n \int_{n^{-1}m_n}^{1} W_n^2(t)\, dt$$

where,

$$W_n(t) = n^{1/2} \sup_{n^{-1} \leq u \leq t} \left[\psi_n[u,\tau_n(u)]\, b\{\frac{1}{\tau_n(u)}\} - \sqrt{n}(1-\psi_n[u,\tau_n(u)]) \right]$$

and $b(\cdot)$ is a Brownian motion.

Proof: Assume that $\theta = 1$, and let:

$$W_n(u) = t_n([nu],\tau_n(u)) - 1 ,$$

where $t_n([nu],\tau_n(u))$ is the time when the [nu]th observation is taken. Then, since $\sqrt{n}\,\tau_n(u) y_n(\tau_n(u)) \to b(\tau_n(u))$ and the probability law of b(t)/t is the same as that of b(t^{-1}), we have the stated asymptotic representation for $W_n(\cdot)$.

For $s_n \neq 1$ the first d_n observations contribute order d_n to the regret. Now, the remaining order of regret depends on $i(1) - i[1 + W_n(u)]$. But, $i'(T) = 0$, and is the global maximum. Thus, by the bounded second derivative and boundedness of i, we have that the above difference can be bounded by a constant times $W_n^2(u)$; proving the theorem.

Theorem A2: Let X be N(0,σ^2). Then,

$$E[X|X>c] = \sigma \frac{\phi(c/\sigma)}{1-\Phi(c/\sigma)},$$

and

$$E[X^2|X>c] = \sigma^2 \left[1 + \frac{\frac{c}{\sigma}\phi(c/\sigma)}{1-\Phi(c/\sigma)} \right],$$

so

$$E[(X-c)^2|X>c] = \sigma^2 + c^2 - \frac{c\sigma\phi(c/\sigma)}{1-\Phi(c/\sigma)}. \tag{A.1}$$

proof: The proof is a straightforward application of evaluations of Gaussian integrals.

Proofs of the three following corollaries are straightforward.

Corr A2.1: Formula A.1 is monotone decreasing in c.

Corr A2.2: If c $\geq$ 0, theorem A2 applies to a folded Gaussian variable.

Corr A2.3: Theorem A2 applies to the supremum of a Browian Motion ($M_{l,h}$).

Corr A2.4: For all c$\geq$ 0 and σ^2:

$$E[(X-c)^2|X>c] \leq \sigma^2 \min[1, \frac{1}{(c^2-\sigma^2)^+}].$$

proof: Since (A.1) is monotone decreasing in c, for c ≥ 0 it is less than or equal to σ^2. For large c, apply l'Hopital's rule three times to obtain the bound.

Theorem A3: For p=0, the smallest order of regret for rules asymptotically efficient for all $\theta > 0$ is $\sqrt{n}$.

proof: Without loss of generality assume that $\theta = 1$. By theorem 3.1 the rules R(0,$\sqrt{n}$) and R*(0,$\sqrt{n}$) have regret O($\sqrt{n}$), so the order for the optimal rule is no greater than $\sqrt{n}$. Now, write (A.1) as:

$$n \int_{n^{-1}}^{n^{-1/4}} W_n^2(t)dt + n \int_{n^{-1/4}}^{1} W_n^2(t)dt = A_n + B_n. \tag{A.2}$$

A_n contributes at least order $\sqrt{n}$ to the regret, for if in the range of integration ψ ever gets greater than order $n^{-.25}$, the overshoot of the optimal time is of order at least $\sqrt{n}$. In this case B_n is of order at least $\sqrt{n}$.

Corr A3.1: For all p≥0, if $J_n(d) \geq O(\sqrt{n})$, then the order of regret obtained by setting $s_n(d, A_d) \equiv 1$ (or a terminal kill) is as small as possible.

proof: It takes at least O($\sqrt{n}$) to reach the specified level of information, and the regret contributed by the s ≡ 1 rule is no greater than order $\sqrt{n}$.

Note: Theorem 3.1 and Corollary 3.1.1 show that if p>0, then the optimal rule sets s ≡ 1 before $J_n(d)$ reaches order $\sqrt{n}$.

Theorem A4: For all t and θ let $f_t(\cdot|\theta)$ satisfy the usual regularity conditions for asymptotic normanlity of the mle. Then, as $n \to \infty$:

$$Y_n(\{n\tau_n(u)\}^{-1}) \to b(\tau_n(u)^{-1}), 0 \leq u \leq 1,$$

where τ_n is defined in Theorem A.1, $b(\cdot)$ is a Brownian motion, and:

$$Y_n(\zeta^{-1}) = \{\hat{\theta}_\zeta - \theta\}.$$

proof: The proof uses the results in Siegmund (1985, III.9 p. 63), and Billingsley (1968, ch 17, p. 143). The score function obeys an invariance theorem with the random time change. Then, the implicit function theorem along with tb(1/t) also being a Brownian motion, proves the theorem.

Table 1: A comparison of the simulated and computed relative efficiencies for the case p=0. See scetion for details.

n->	10	20	30	40	50	75	100	150	200	300	400	500	1000
SIMU	53	64	70	73	76	80	82	85	87	89	90	91	94
COMP	47	60	67	71	74	78	81	84	86	88	90	91	94

FIGURE CAPTIONS

Figure 1: Relative efficiencies for R^* and the Bergman and Turnbull rule. The dotted line denotes Asterisks denote upper bounds. See section 5 for details.

Figure 2: Relative efficiencies for R^*.

N	theta 10	8	6	4	2	1.33	1	.8	.67	.5	.25	.125
50	78		81	86	89	91	92	97	97	90	44	3
	73	78	83	89	93		91	91*	91*	91*	46*	3*
		81	84	91	98		88			35		
100	83		85	87	91	92	93	94	94	96	59	5
200	87		90	88	92	92	94	96	95	96	73	14
	92	94	95	95	94		89	89*	89*	89*	46*	3*
		81	85	88	95		96			86		
300	88		90	91	94	92	95	96	96	98	80	19
		83	85	87	94		98			88		
400	92		92	92	93	93	94	94	97	98	84	23
		84	86	89	94		98			95		
500	91		91	92	93	94	95	97	96	98	87	27
		84	87	90	94		97			94		

Table 2: Relative Efficiency (with respect to the optimal design for known theta) of the Three-stage rule (see the text for details), the B&T rule (with b=4, Z*=4), and the Orav& Louis rule. The B&T values result from numerical computations and are exact. The three-stage and Orav & Louis values result from simulations and have standard errors no greater than 1.2 .

* Upper bounds for relative efficiency (see section 5).

Figure 1

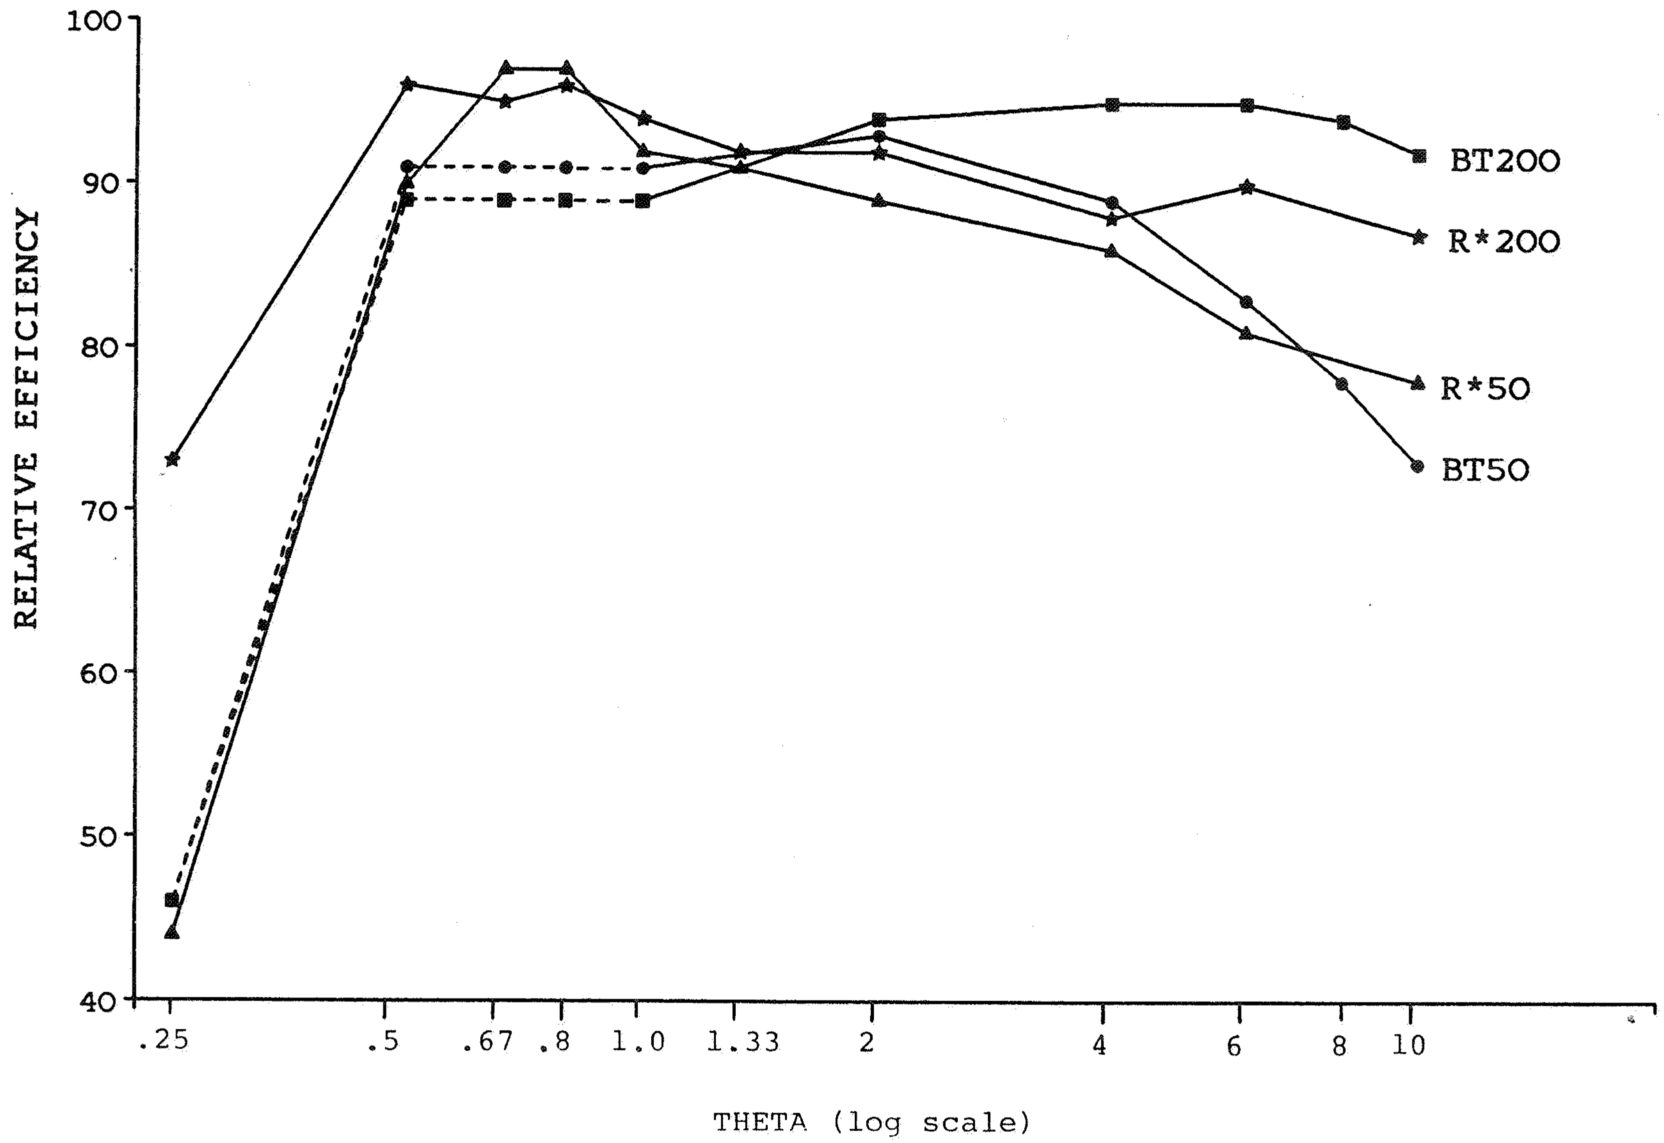
RELATIVE EFFICIENCY
100
90
80
70
60
50
40
.25
.5
.67
.8
1.0
1.33
2
4
6
8
10
THETA (log scale)
BT200
R*200
R*50
BT50

Figure 2

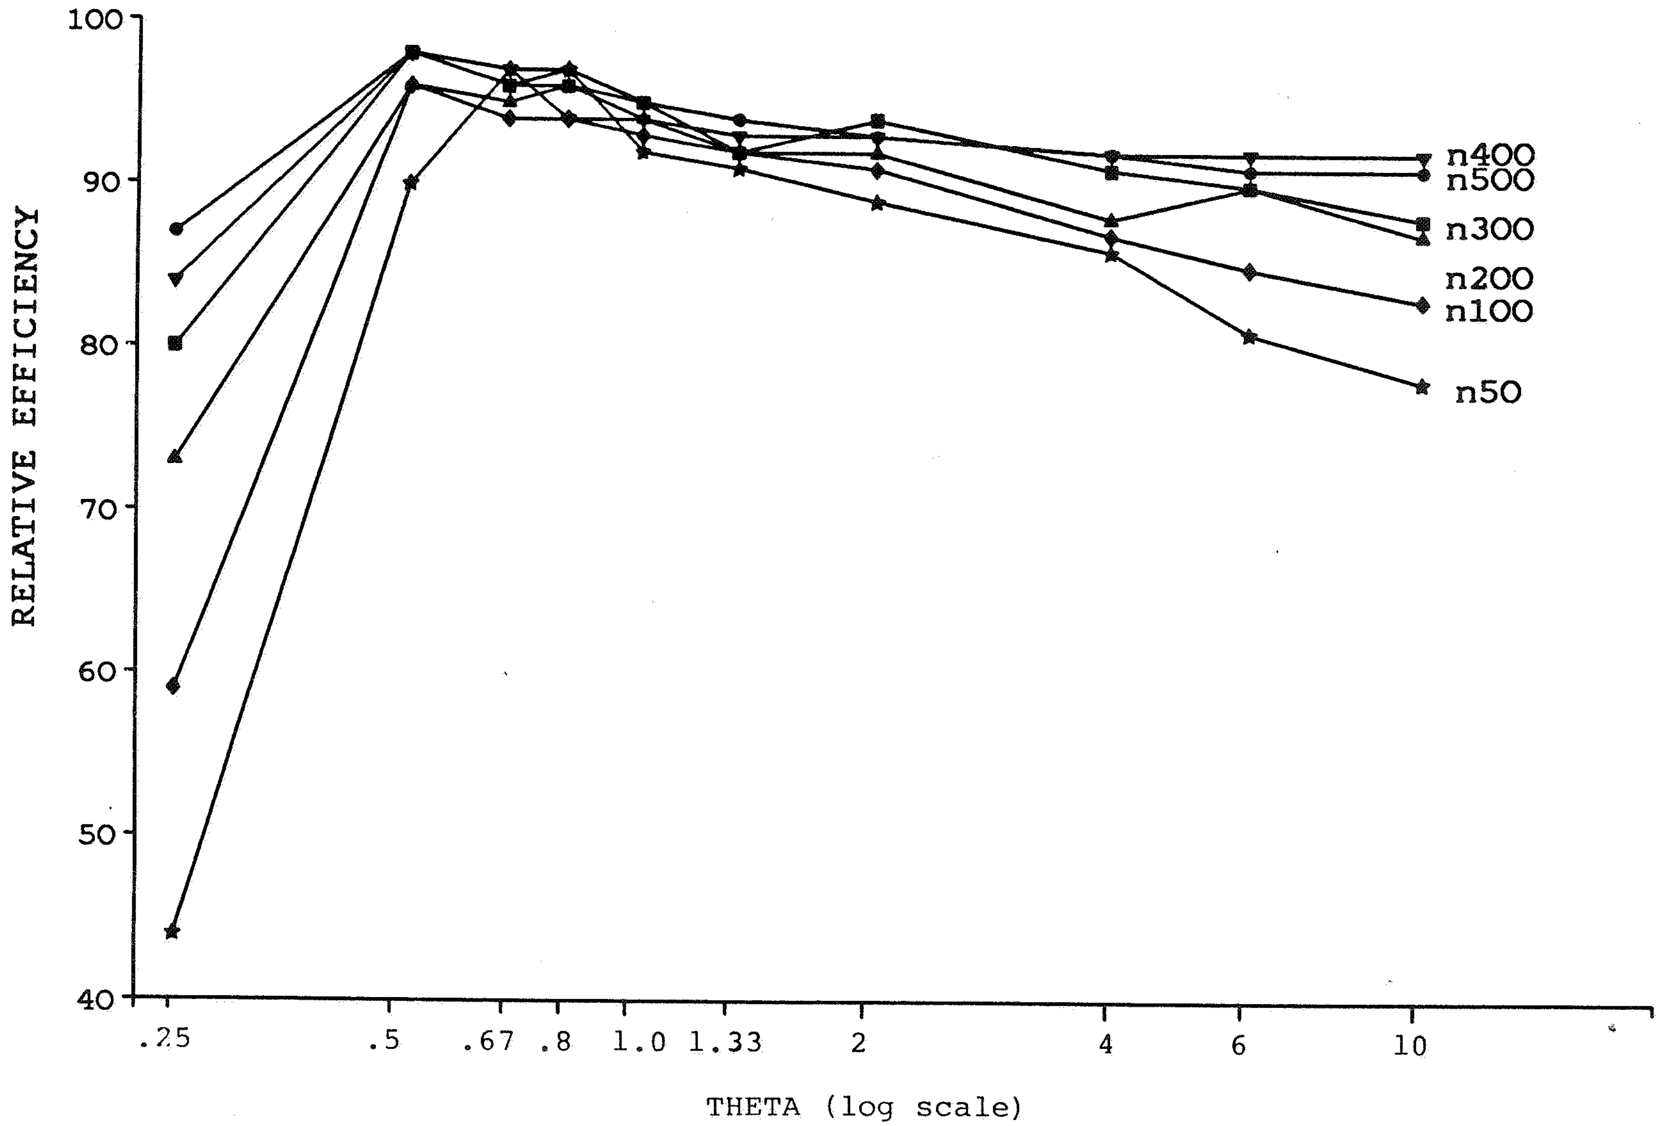
100
90
80
70
60
50
40
RELATIVE EFFICIENCY
.25
.5
.67
.8
1.0
1.33
2
4
6
10
THETA (log scale)
n400
n500
n300
n200
n100
n50