1986

The VPRT: A Sequential Testing Procedure Dominating the SPRT

Noel Cressie

Peter B. Morgan

THE CENTRE FOR DECISION SCIENCES AND ECONOMETRICS

The VPRT: A Sequential Testing Procedure Dominating The SPRT

Noel Cressie
Peter B. Morgan

TECHNICAL REPORT NO. 11
MARCH 1986

Centre For Decision Sciences And Econometrics
Social Science Centre
The University of Western Ontario
London, Ontario N6A 5C2

ISSN: 0831-4454
ISBN: 0-7714-0731-9

1634

THE VPRT: A SEQUENTIAL TESTING PROCEDURE DOMINATING THE SPRT

By Noel Cressie and Peter B. Morgan

Iowa State University and University of Western Ontario

Abstract

Under more realistic assumptions than those usually imposed in the sequential analysis literature, a variable-sample-size sequential probability ratio test (VPRT) of two simple hypotheses is found which maximizes the expected net gain over all sequential decision procedures. The VPRT also minimizes the expected total sampling cost and, under slightly more general conditions than those imposed by Wald and Wolfowitz (1948), reduces to the one-observation-at-a-time sequential probability ratio test (SPRT). Finally, the ways in which the size and power of the VPRT depend upon the parameters of the decision procedure are examined.

AMS 1980 subject classifications. Primary 62C05; secondary 62F03.

Keywords and phrases. Cost function, discounting, expected net gain, finite truncation time, payoff function, sequential decision procedure, variable-sample-size sequential probability ratio test.

Running head: The VPRT.

1. Introduction

The sequential probability ratio test (SPRT) for testing two simple hypotheses is shown by Wald and Wolfowitz (1948) to be optimal in the following sense. Of all size α tests having the same power, the test which requires on average the fewest observations is the SPRT. How then can the claim in the title of our paper be made?

The assumptions underlying the Wald-Wolfowitz result need to be stated explicitly, to expose the rather specific nature of their result. More general (and more realistic) assumptions can then be entertained. First Wald and Wolfowitz assumed that the sequential experiment to determine the true hypothesis has no fixed truncation time T, within which a decision has to be made. The real world does not operate without deadlines, and indeed use of the SPRT in sequential clinical trials is always modified to ensure termination of the trial in a fixed time (Armitage, 1985). What optimality results are available for these truncated decision procedures?

Second Wald and Wolfowitz assumed that the cost of sampling each individual observation is a constant c and that there is no start-up cost c_0. More realistic cost functions show an often substantial start-up cost for an experiment (e.g. purchase of equipment, planning, pilot experiment) and a decreasing unit cost as time goes on and the experiment is made to run more efficiently. What optimality results are available for a more general cost structure?

Third Wald and Wolfowitz implicitly assume the payoff received for deciding upon one hypothesis or the other is independent of when that decision is taken. In some circumstances however, delay of the decision depreciates the payoff (e.g. indecisiveness on where to locate a rapid transit station affects all potential commuters detrimentally). Thus we have included the discount factor τ ($0 \leq \tau \leq 1$) as a control variable in the decision procedure; Wald and Wolfowitz implicitly assume $\tau = 1$. What happens to their result when the discount factor is less than one?

Finally, until recently it has always been assumed that a sequential procedure only takes one observation at a time. Group sequential sampling, usually with constant group size, has been suggested by Ghosh (1970, p. 224) and Pocock (1977). Gupta and Miescke (1984) present sequential selection procedures based on sample sizes that differ from one time point to another. In all these studies, the sample size is either assumed to be known in some way, or is chosen in a (perhaps stochastic) ad hoc manner. Is there an optimal choice of sample size? And is there an analogue of the Wald-Wolfowitz theorem for the variable-sample-size sequential procedure?

This article addresses all of the questions above, using dynamic programming to develop an optimal procedure we call the VPRT (variable-sample-size-sequential probability ratio test) that dominates the SPRT. It will be seen that under the assumptions made by Wald and Wolfowitz (1948), the VPRT and SPRT are identical. However it will be shown that under the more realistic model, the SPRT's disadvantage can be substantial.

At almost the same time as Wald and Wolfowitz published their optimality results, Arrow, Blackwell, and Girshick (1949) gave different proofs of many of the same results. Indeed their assumptions are more general (although they still do not take into account a discount factor and variable sample size), their proofs are more direct, and their approach foreshadows the use of dynamic programming in constructing optimal decision procedures. In addition, they establish optimality results for the SPRT in the multi-valued decision problem, something we will take up in a subsequent paper.

One of the assumptions we do not attempt to generalize here is the independent and identically distributed nature of the sampling from fixed distributions. Irle (1984) shows optimality of the SPRT under minimal assumptions on a stochastic process generating the observations (one-at-a-time); and Quang (1985) shows that robust (one-at-a-time) sequential testing of two contaminated distributions leads to an SPRT based on the least-favorable pair.

Section 2 sets out formally the notation and the model structure already alluded to above. Section 3 establishes the existence of a sequence of unique pairs of probabilities that completely describe the stopping rule of the optimal decision procedure. In Section 4 the optimal decision procedure is shown to be equivalent to a probability ratio test (the VPRT) which dominates the SPRT and which, under the Wald-Wolfowitz assumptions, reduces to the SPRT. Section 5 offers comments on how the size and power of the VPRT are constructed. Concluding remarks are given in Section 6. To enhance the readability of the paper we have collected the proofs of our results into an appendix, and left their statements and surrounding discussion to the main body of the paper.

2. MODEL STRUCTURE AND NOTATION

Consider the set $\Omega = \{\omega_0, \omega_1\}$ of the states of nature. Which of ω_0 or ω_1 is the true state is unknown, however observations drawn from an observation space X according to the cumulative distribution function (cdf) F provide information from which a decision can be made. Assume F has a density (or probability mass) function f.

Let T denote the number of decision points available to the decisionmaker and let $t \in \{1,\dots, T\}$ denote the t-th decision point, $1 \leq T \leq \infty$; when $T < \infty$ we have a truncated decision procedure. In other words a decision about the true state of nature can be made or put off at any of the $t = 1,\dots,T-1$ discrete time points, but a decision <u>has</u> to be made by time T.

The decisionmaker's intial information, at t=1, is: complete knowledge of $\{F(x;\omega): x \in X\}$ for given $\omega \in \Omega$, an initial information vector y_1, and a prior probability p_1 for the event $\{\omega=\omega_0\}$. The truncation time T is also known.

Sampling from X is costly. In particular, $c(n_t)$ is the cost incurred at time t if $n_t \geq 0$ observations are demanded at t. For $n_t \geq 1$ these observations are denoted by $x_{t1},\dots,x_{tn_t}$ and are received by the decisionmaker at time t+1. It is assumed that $c(0) = c0 \geq 0$, c(n) is monotonic strictly increasing with n, and that c(n) is unbounded above (c_0 is the "overhead" or "fixed cost of sampling").

Upon receiving $x_{t-1} \equiv (x_{t-1,1},\dots,x_{t-1,n_{t-1}})$ for $n_{t-1} \geq 1$ at t, the decisionmaker augments the information vector yt-1 by xt-1 to generate the current information vector y_t (for $t=2,\dots,T$); that is

$$y_t \in Y_t = \begin{cases} Y_1 \equiv \{y_1\} & , t=1 \\ Y_{t-1} & , t \geq 2 \text{ and } n_{t-1} = 0 \\ Y_{t-1} \times (\times_{i=1}^{n_{t-1}} X) & , t \geq 2 \text{ and } n_{t-1} \geq 1. \end{cases} \qquad (2.1)$$

It is not yet clear how the "right" decision about ω is to be made, but it is clear that a sequential rule will have to be considered.

DEFINITION 1. A terminal decision rule δ is a sequence $\{\delta_t\}_{t=1}^{T}$ where $\delta_t: Y_t \rightarrow \{0,1\}$; $\delta_t(\cdot) = 0$ if ω_0 is chosen and $\delta_t(\cdot) = 1$ if ω_1 is chosen, for $t = 1,\ldots,T$.

At any decision point t a choice is made between collecting more observations or terminating sampling. If sampling is terminated then δ is used to choose between ω_0 and ω_1. The payoffs to these choices are defined by the payoff function $U: \Omega \times \{0,1\} \rightarrow R$, where for given t and y_t,

$$U(\omega,\delta(y_t)) = \begin{cases} u_{00}, & \text{if } \delta(y_t) = 0 \text{ and } \omega = \omega_0 \\ u_{01}, & \text{if } \delta(y_t) = 0 \text{ and } \omega = \omega_1 \\ u_{10}, & \text{if } \delta(y_t) = 1 \text{ and } \omega = \omega_0 \\ u_{11}, & \text{if } \delta(y_t) = 1 \text{ and } \omega = \omega_1 . \end{cases} \qquad (2.2)$$

Typically $u_{00} > u_{10}$ and $u_{11} > u_{01}$, since correct decisions are usually rewarded more generously than incorrect decisions.

In the decision-theory literature, payoffs are often expressed as negative losses, by way of a prespecified loss function. Which one uses is really a matter of taste; those readers familiar with loss functions should have no difficulty thinking in terms of payoff functions.

Although most sequential procedures specify that one further observation is taken whenever sampling is continued, we wish to entertain the possibility of continuing by drawing none, one, or more observations at a time. We will see later that an optimal choice of these sample sizes may lead to a substantial improvement over the one-at-a-time sequential procedures.

DEFINITION 2. A sample-size rule ν is a sequence $\{\nu_t\}_{t=1}^{T}$ where $\nu_t: Y_t \rightarrow N_+ \cup \{0\}$, the set of non-negative integers, for $t = 1,\ldots,T$.

DEFINITION 3. A stopping rule S is a sequence $\{S_t\}_{t=1}^{T}$ where $S_t\colon Y_t \to \{0,1\}$; $S_t(\cdot) = 0$ if sampling continues at t and $S_t(\cdot) = 1$ if sampling is not continued at t, for t=1,...,T. Then by definition, $S_T(\cdot) \equiv 1$.

A decision rule is made up of an S, a ν, and a δ, and operates as follows. At time t the decisionmaker's observation vector is y_t. Two mutually exclusive alternatives are to be considered. One alternative is to terminate sampling at t and use δ_t to select a state of nature from Ω. The other is to continue to sample at t by drawing $\nu_t(y_t) = n_t$ additional observations from X. Which is chosen is determined by $S_t(y_t)$, and the payoff and cost functions offer a precise way of valuing the two alternatives.

DEFINITION 4. A decision rule d is an ordered triple (S, ν, δ), of a stopping rule S, a sample-size rule ν, and a terminal decision rule δ.

We conclude this section by outlining the form of the optimal decision rule. Let p_t be the decisionmaker's posterior probability at t of the event $\{\omega=\omega_0\}$. By Bayes' Theorem, for t = 1,...,T-1 and $n_t > 0$,

$$p_{t+1} = \frac{f(x_t;\omega_0)p_t}{f(x_t;\omega_0)p_t + f(x_t;\omega_1)(1 - p_t)} = \frac{f(y_{t+1};\omega_0)p_1}{f(y_{t+1};\omega_0)p_1 + f(y_{t+1};\omega_1)(1-p_1)}. \tag{2.3}$$

If nt = 0 for any t ∈ {1,...,T-1} then pt+1 = pt.
Also the unconditional cdf over X at t is

$$F_t(x) := p_tF(x;\omega_0) + (1 - p_t)F(x;\omega_1)$$

where := means "is defined to be". At time t the expected net gain (taking into account both payoff and cost) of a terminal decision reached using

$$\delta_t(y_t) = \begin{cases} 0, \text{ if } u_{00}p_t + u_{01}(1 - p_t) \geq u_{10}p_t + u_{11}(1 - p_t) \\ 1, \text{ if } u_{00}p_t + u_{01}(1 - p_t) < u_{10}p_t + u_{11}(1 - p_t) \end{cases} \tag{2.4}$$

is

$$E_\omega[U(\omega,\delta_t(y_t))|y_t,p_t] = \max\{u_{00}p_t + u_{01}(1 - p_t),\ u_{10}p_t + u_{11}(1-p_t)\}, \tag{2.5}$$

where E_ω denotes the expectation with respect to the measure on Ω defined by the posterior probability p_t. Should the decisionmaker sample until $t=T$ is reached, then sampling must terminate and a terminal decision must be taken. Thus the expected net gain at $t=T$ of the decision procedure is

$$V_T^T(y_T,d,p_T) := E_\omega[U(\omega,\delta_T(y_T))|y_T,p_T]. \tag{2.6}$$

Backward recursion and the optimality principle of dynamic programming can now be used to construct a sequence of maximum expected net gain functions:

$$V_t^T(y_t,d,p_t) := \max\{E_\omega[U(\omega,\delta_t(y_t))|y_t,p_t], \tag{2.7}$$
$$\max_{n_t \geq 0}\{-c(n_t) + \tau E_{F_t}[V_{t+1}(y_{t+1},d,p_{t+1})|y_t,p_t,n_t]\}\};\ t=1,\ldots,T.$$

That is, $V_t^T(y_t,d,p_t)$ is the larger of the expected net gain of stopping at t and the expected net gain of continuing to sample by demanding n_t^* additional observations. The sample size n_t^* is itself optimal in that it maximizes the expected net gain of continuing. That a dollar tomorrow is worth less than a dollar today is expressed through τ, the discount factor between decision points; $0 \leq \tau \leq 1$. Spahn and Ehrenfeld (1974) introduce the idea of an optimal sample size but in the less realistic case of $T = \infty$, $\tau = 1$ and $u_{00} = u_{11}$.

Let D denote the set of decision rules. Formally, the problem we wish to solve is: Find a decision rule $d^* \in D$ such that for any given $y_t \in Y_t$, and any $p_t \in [0,1]$,

$$V_t^T(y_t, d^*, p_t) \geq V_t^T(y_t, d, p_t), \ \forall\, d \in D \ ; \ t = 1, \ldots, T. \tag{2.8}$$

Such a d^* is called an <u>optimal decision rule</u>.

Thus an <u>optimal</u> decision rule for the unknown parameter ω is a function of prior information and sample data that maximises expected net gain by optimising not only on stop/continue decisions but, in addition, on sample size choices. Morgan and Manning (1985) show that implicit in the construction of the maximum expected net gain functions described in (2.7) is the existence of an optimal decision rule $d^* = (S^*, \nu^*, \delta^*)$, where

(a) $\delta^* := \{\delta_t^*\}_{t=1}^T$ such that $\delta_t^*: Y_t \to \{0,1\}$ and

$$\delta_t^*(y_t) = \begin{cases} 0, \text{ if } u_{00}p_t + u_{01}(1-p_t) \geq u_{10}p_t + u_{11}(1-p_t) \\ \\ 1, \text{ if } u_{00}p_t + u_{01}(1-p_t) < u_{10}p_t + u_{11}(1-p_t) \end{cases} \tag{2.9a}$$

(b) $\nu^* := \{\nu_t^*\}_{t=1}^T$ such that $\nu_t^*: Y_t \to N_+ \cup \{0\}$ and

$$\nu_t^*(y_t) = \underset{n_t}{\text{argmax}} \{-c(n_t) + \tau E_{F_t}[V_{t+1}^T(y_{t+1}, d^*, p_{t+1}) | y_t, p_t, n_t]\}, \tag{2.9b}$$

for $t \in \{1, \ldots, T-1\}$, and $\nu_T^*(y_T) \equiv 0$

(c) $S^* := \{S_t^*\}_{t=1}^T$ such that $S_t^*: Y_t \to \{0,1\}$ and

$$S_t^*(y_t) = \begin{cases} 0, \text{ if } \max\{u_{00}p_t + u_{01}(1-p_t),\ u_{10}p_t + u_{11}(1-p_t)\} \\ \geq - c(\nu^*(y_t)) + \tau E_{F_t}[V_{t+1}^T(y_{t+1}, d^*, p_{t+1}) | y_t, p_t, \nu_t^*(y_t)] \\ 1, \text{ otherwise,} \end{cases} \quad (2.9c)$$

for $t \in \{1,\ldots,T-1\}$, and $S_T^*(y_T) \equiv 1$.

In effect, d^* is simply a statement that, to maximize the expected net gain from any stage of the decision problem, the decisionmaker should discover and choose the alternative which, at that stage, offers the highest expected net gain.

The reader may wish to note here that since $V_1^T(y_1, d^*, p_1)$ is a maximum expected-net-gain function constructed employing backward induction and the optimality principle of dynamic programming, it is equivalent to define an optimal decision rule as any d^* which satisfies:

$V_1^T(y_1, d^*, p_1) \geq V_1^T(y_1, d, p_1)$, $\forall\ d \in D$, any given y_1, and any $p_1 \in [0,1]$. (2.10)

For an example of this construction applied to the usual sequential decision procedure, see Ferguson (1967, pp. 315-317).

3. THE OPTIMAL DECISION PROCEDURE

The problem described in Section 2 is formally a special case of a decision problem analyzed by Morgan and Manning (1985). There they provide a proof (1985, Theorem 2, p. 937) of the existence and a description of a maximum-expected-net-gain decision procedure under conditions which include the decision problem described in Section 2. Accordingly, we refer the reader to this result and do not consider the existence question further.

The goal of this section is to establish properties of the maximum-expected-net-gain functions $\{V_t^T(\cdot): t=1,\ldots,T\}$ from which it follows that the optimal decision procedure has the same very simple and intuitively appealing form that the usual sequential decision procedures have. The results presented are linked by informal and intuitive explanations of their validity and importance; the formal proofs are collected into the Appendix. Briefly, we show that each of the expected-net-gain functions $V_t^T(y_t,d^*,p_t)$ is the maximum of functions which are both continuous and convex with respect to p_t. These results allow us to establish the existence of two unique critical probabilities p_{tL} and p_{tU} that partition the unit interval into three intervals: an interval where ω_0 is chosen, an interval where ω_1 is chosen, and an interval where sampling is continued.

We begin by reminding the reader that for finite T, the last member of the sequence of maximum-expected-net-gain functions is

$$V_T^T(y_T,d^*,p_T) = \max\{u_{00}p_T + u_{01}(1 - p_T),\ u_{01}p_T + u_{11}(1 - p_T)\}. \quad (3.1)$$

This is clearly convex and continuous with respect to p_T, and only depends on the data y_T through p_T. Now notice that

$$V_{T-1}^T(y_{T-1},d^*,p_{T-1}) = \max\{u_{00}p_{T-1} + u_{01}(1 - p_{T-1}),\ u_{10}p_{T-1} + u_{11}(1 - p_{T-1}),$$

$$\max_{n_{T-1}\geq 0}\{-c(n_{T-1}) + \tau E_{F_{T-1}}[V_T^T(y_T,d^*,p_T)|y_{T-1},p_{T-1},n_{T-1}]\}\} \quad (3.2)$$

is the maximum of three functions, the first two of which are linear with respect to p_{T-1} and the third is the maximum of functions which are each linear with respect to p_{T-1} since they depend upon p_{T-1} linearly through $F_{T-1}(\cdot)$. Because $V_{T-1}^T(y_{T-1},d^*,p_{T-1})$ is the maximum of functions linear in p_{T-1}, it is a function which is convex and thus continuous with respect to p_{T-1}. Once again dependence on the data y_{T-1} is only through p_{T-1}. This argument may be extended to establish the general result for V_t^T.

LEMMA 3.1. Given $T \leq \infty$, $t \in \{1,\ldots,T\}$ and $y_t \in Y_t$, then $V_t^T(y_t,d^*,p_t)$ is a convex and continuous function of $p_t \in [0,1]$.

Proof: See the Appendix.

Figure 1 shows how these results translate into plots, given here for t=1. The figure shows $V_1^T(y_1,d^*,p_1)$ as the solid line. This is

$$V_1^T(y_1,d^*,p_1) = \begin{cases} u_{10}p_1 + u_{11}(1-p_1) & ;\ p_1 \in [0,p_{1L}] \\ \max_{n_1 \geq 0} \{-c(n_1) + \tau E_{F_1}[V_2^T(y_2,d^*,p_2)|y_1,p_1,n_1]\}; & p_1 \in (p_{1L},p_{1U}) \\ u_{00}p_1 + u_{01}(1-p_1) & ;\ p_1 \in [p_{1u},1]. \end{cases} \quad (3.3)$$

Figure 1 here

So we see that, for the case illustrated, the optimal decision procedure says "choose ω_1 if $p_1 \equiv \Pr(\omega=\omega_0)$ is smaller than or equal to p_{1L}," "choose ω_0 if p_1 is larger than or equal to p_{1U}," and "collect n_1^* extra data if p_1 is between

p_{1L} and p_{1U}". The quantity $\hat{p}$ is the value of the prior probability p_1 for which the decisionmaker values equally the terminal alternatives of choosing ω_0 or of choosing ω_1. The figure suggests that there will exist only one root to each of the equations

$$u_{00}p_1 + u_{01}(1-p_1) = \max_{n_1 \geq 0} \{-c(n_1) + \tau E_{F_1}[V_2^T(y_2,d^*,p_2)|y_1,p_1,n_1]\}, \quad (3.4)$$

and

$$u_{01}p_1 + u_{11}(1-p_1) = \max_{n_1 \geq 0} \{-c(n_1) + \tau E_{F_1}[V_2^T(y_2,d^*,p_2)|y_1,p_1,n_1]\}. \quad (3.5)$$

Let p1L be the root of (3.4) and p1U the root of (3.5). These solutions exist provided

Figure 1

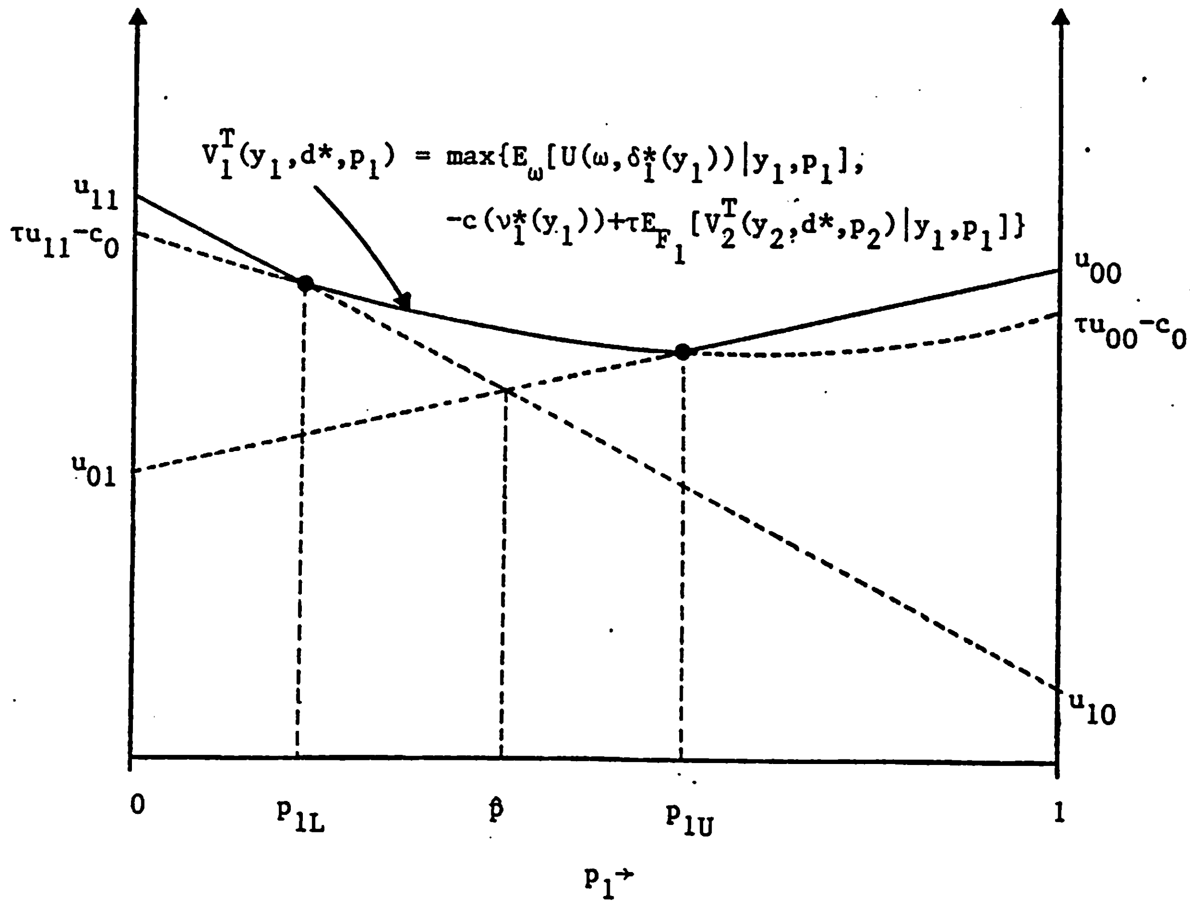

The maximum-expected-net-gain function for t=1. Intervals $[0,p_{1L}]$, (p_{1L},p_{1U}), $[p_{1U},1]$ are respectively where ω_1 is chosen, where sampling is continued, and where ω_0 is chosen.

$$u_{00}\hat{p} + u_{01}(1 - \hat{p}) = u_{01}\hat{p} + u_{11}(1 - \hat{p})$$

$$< \max_{n_1 \geq 0} \{-c(n_1) + \tau E_{F_1}[V_2^T(y_2, d^*, p_2) | y_1, \hat{p}, n_1]\}. \qquad (3.6)$$

The essence of the proof of the existence and uniqueness of p_{1L} and p_{1U} subject to (3.6) is simply that the endpoints of the convex function of p_1,

$$\max_{n_1 \geq 0} \{-c(n_1) + \tau E_{F_1}[V_2^T(y_2, d^*, p_2) | y_1, p_1, n_1]\},$$

are $\tau u_{11} - c_0$ (for $p_1 = 0$) and $\tau u_{00} - c_0$ (for $p_1 = 1$); since the discount factor $\tau \leq 1$, and the overhead $c_0 \geq 0$, these endpoints are smaller than or equal to u_{11} and u_{00} respectively. This, (3.6), and the convexity of the curve make it obvious that p_{1L} and p_{1U} must exist and be unique. The argument may be repeated for any $t = 2,\ldots,T-1$, giving the following result.

<u>THEOREM 3.1.</u> Given $2 \leq T \leq \infty$, $t \in \{1,\ldots,T-1\}$ and $y_t \in Y_t$, if

$$-c(\nu_t^*(y_t)) + \tau E_{F_t}[V_{t+1}^T(y_{t+1}, d^*, p_{t+1}) | y_t, p_t = \hat{p}]$$

$$> u_{00}\hat{p} + u_{01}(1 - \hat{p}) = u_{10}\hat{p} + u_{11}(1 - \hat{p}),$$

then there exist p_{tL}, $p_{tU} \in [0, 1]$, $p_{tL} < \hat{p} < p_{tU}$, such that the maximum-expected-net-gain (optimal) decision procedure d^* is:

stop sampling and choose ω_1, if $p_t \in [0, p_{tL}]$

defer a terminal decision and

collect $\nu_t^*(y_t)$ further observations, if $p_t \in (p_{tL}, p_{tU})$

stop sampling and choose ω_0, if $p_t \in [p_{tU}, 1]$.

Proof: See the Appendix.

The reader familiar with the literature on the development of SPRT's will now have recognised that the above result opens the way to re-expressing the optimal decision procedure d^* as a likelihood-ratio test akin to the SPRT; in this form we call d^* the VPRT.

We now establish a monotonicity property of the $V_t^T(\cdot)$ functions and show that this implies the sequence $\{(p_{tL}, p_{tU})\}_{t=1}^{T-1}$ of continue-sampling intervals is nested. It is intuitively clear that the maximum value functions $V_t^T(p, d^*)$ are monotonically increasing in T; that is, having the opportunity to draw more samples of data cannot reduce the expected-net-gain of the optimal decision procedure.

LEMMA 3.2 Given $t \in \{1, \ldots, T\}$ and $p_t \in [0,1]$, $V_t^T(y_t, d^*, p_t)$ is a monotonic increasing function of T.

Proof: See the Appendix.

LEMMA 3.3 $\{V_t^T\}_{T=t}^{\infty}$ is uniformly convergent to V_t^{∞}, $\forall\, t \geq 1$.

Proof: See the Appendix

Figure 2 illustrates the content of the next two results. With an arbitrary choice of t=1 the figure shows some members of the monotonically increasing sequence of V_t^T functions. The continue-sampling regions $(p_{tL}(T), p_{tU}(T))$ increase with T, approaching a limiting interval

Figure 2

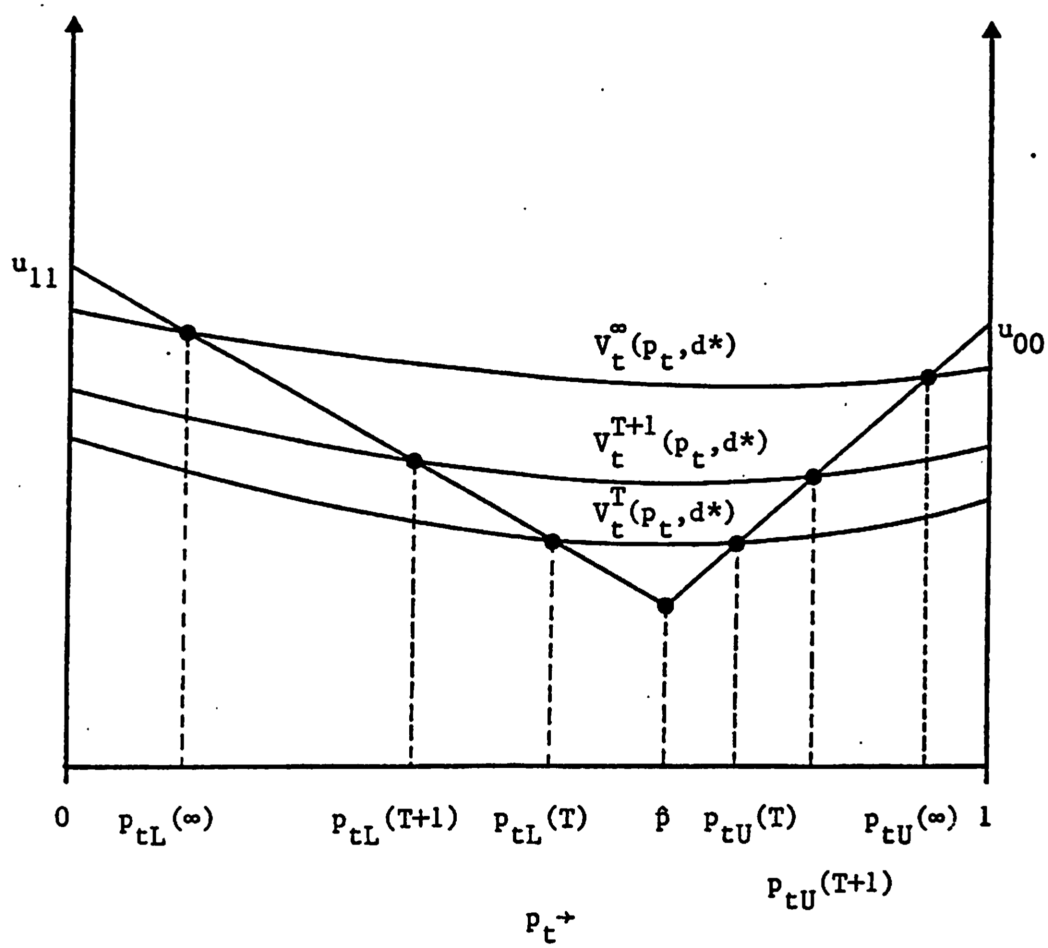

The continue-sampling intervals at t for termination times T, T+1 and ∞ are respectively $(p_{tL}(T),p_{tU}(T))$, $(p_{tL}(T+1),\ p_{tU}(T+1))$ and $(p_{tL}(\infty),p_{tU}(\infty)) \equiv (p_L^*,p_U^*)$. The continue-sampling intervals increase in a nested manner with T due to $V_t^T(p_t,d^*)$ being an increasing function of T.

$(p_{tL}(\infty), p_{tU}(\infty))$ which is the continue-sampling region when the decision problem has no finite truncation time ($T=\infty$).

<u>THEOREM 3.2</u> For given $t \geq 1$ the sequence $\{(p_{tL}(T), p_{tU}(T))\}_{T=t}^{\infty}$ is nested, increasing and convergent to $(p_{tL}(\infty), p_{tU}(\infty))$.

<u>Proof</u>: See the Appendix.

Figure 2 here

That the continue-sampling intervals $(p_{tL}(T), p_{tU}(T))$ widen in a nested manner as T increases follows from the expected net gain of continuing to sample being an increasing function of the number of intervals remaining to the truncation time T. Of course when $T = \infty$ then $p_{tL}(\infty) \equiv p_L^*$ and $p_{tU}(\infty) \equiv p_U^*$ are independent of T because the passage of time does not reduce the number of periods remaining for the decisionmaker. This limiting case of time-invariant boundaries for the continue-sampling interval is not unfamiliar to those studying the limiting case of the one-at-a-time sequential decision procedures (see Ferguson, 1967, pp. 355–356) with the difference that the one-at-a-time procedure's continue-sampling region $(p_L', p_U') \subseteq (p_L^*, p_U^*)$. The containment will be strict in many circumstances since the optimal sample size to draw when continuing sampling will often not be unity. A class of problems for which $(p_L', p_U') = (p_L^*, p_U^*)$ is described in Theorem 4.2.

In the next section we write out the VPRT explicitly in terms of the likelihood ratio, develop its optimality properties, compare these to the properties of the SPRT, and demonstrate that the VPRT always dominates the SPRT.

4. THE VARIABLE-SAMPLE-SIZE-SEQUENTIAL PROBABILITY RATIO TEST

The goal of this section is to extend the pioneering works of Wald and Wolfowitz (1948) and Arrow, Blackwell and Girshick (1949) from the class of pure sequential (one-at-a-time-sampling strategy) decision procedures to the more general class of decision procedures described in Section 3. This is done by first introducing the variable-sample-size-sequential probability ratio test (VPRT), second explaining the optimality properties of the VPRT, and third demonstrating the dominance of the VPRT over the SPRT in both the Arrow-Blackwell-Girshick sense of maximizing expected-net-gain and in the Wald-Wolfowitz sense of minimizing the expected total number of observations needed to reach a choice for ω, under either state of nature and amongst all decision procedures of a given size.

Given $T \leq \infty$, any $t \in \{1,\ldots,T\}$ and $y_t \in Y_t$, define the likelihood ratio

$$\lambda_t := \frac{f(y_t;\omega_0)}{f(y_t;\omega_1)} .$$

We define a wide-sense sequential probability ratio test procedure (WSPRT) for given $T \leq \infty$, as a sequence $\{(S_t,\nu_t,\delta_t,A_t,B_t)\}_{t=1}^{T}$ where $S_t\colon Y_t \to \{0,1\}$; $\nu_t\colon Y_t \to N_+ \cup \{0\}$; $\delta_t\colon Y_t \to \{0,1\}$; $A_t,B_t \in R_+ \cup \{0\}$ with $0 \leq A_t \leq B_t \leq \infty$, and where the decision rule is, for any given $t \in \{1,\ldots,T-1\}$ and $y_t \in Y_t$,

$$
\begin{aligned}
&S_t(y_t) = 1 \text{ and } \delta_t(y_t) = 1 && \text{, if } 0 \leq \lambda_t \leq A_t \\
&S_t(y_t) = 0 \text{ and } \nu_t(y_t) \text{ extra observations collected} && \text{, if } A_t < \lambda_t < B_t \quad (4.1) \\
&S_t(y_t) = 1 \text{ and } \delta_t(y_t) = 0 && \text{, if } B_t \leq \lambda_t .
\end{aligned}
$$

In words, (4.1) says that if $\lambda_t \in [0,A_t]$ then sampling should stop and ω_1 be chosen; if $\lambda_t \in (A_t,B_t)$ then sampling should continue by collecting $n_t = \nu_t(y_t)$ additional observations; and if $\lambda_t \in [B_t,\infty)$ then sampling should stop and ω_0 be chosen.

Let W denote the class of wide-sense sequential probability ratio tests. We wish to demonstrate the breadth of this class. Obviously the usual one-at-a-time sequential probability ratio tests are in W, with $\nu_t \equiv 1, \forall\ t \geq 1$. Less obviously, the Neyman-Pearson test (Neyman and Pearson, 1933) based on a fixed sample size n is also in W, with $\nu_1 \equiv n$ and $S_2 \equiv 1$. Moreover, truncated sequential procedures such as a 2-sample test based on the likelihood ratio are in W, with $S_3 \equiv 1$.

Due to Wald and Wolfowitz's (1948) optimality result, we use the SPRT as the "yardstick" against which to measure the performance of any member of W. By relaxing the sampling restriction $\nu_t \equiv 1, \forall\ t \geq 1$, we create the possibility of discovering a more flexible form of WSPRT which dominates the SPRT. Below we demonstrate that the maximum-expected-net-gain decision procedure d* described in Section 3 is <u>equivalent</u> to such a test; this test is the VPRT, which we now describe.

From Theorem 3.1, the optimal decision at t is as follows:

Choose ω_0 iff $p_t \in [p_{tU},1]$ which, rewritten, is

$$\frac{f(y_t;\omega_0)p_1}{f(y_t;\omega_0)p_1 + f(y_t;\omega_1)(1-p_1)} \in [p_{tU},1]. \qquad (4.3)$$

That is,

$$\text{choose } \omega_0 \text{ iff } 0 \leq \lambda_t \equiv \frac{f(y_t;\omega_0)}{f(y_t;\omega_1)} \leq \frac{(1-p_{tU})p_1}{p_{tU}(1-p_1)}. \qquad (4.4)$$

Similarly,

$$\text{choose } \omega_1 \text{ iff } \frac{(1-p_{tL})p_1}{p_{tL}(1-p_1)} \leq \lambda_t \tag{4.5}$$

and

continue to sample by collecting $n_t^* = \nu_t^*(y_t)$ further observations iff

$$\frac{(1-p_{tU})p_1}{p_{tU}(1-p_1)} < \lambda_t < \frac{(1-p_{tL})p_1}{p_{tL}(1-p_1)} . \tag{4.6}$$

These simple rearrangements show that the optimal decision procedure d* can be rewritten as the following WSPRT: Given $T \leq \infty$, define the sequence $\{(S_t^*, \nu_t^*, \delta_t^*, (1-p_{tU})p_1/p_{tU}(1-p_1), (1-p_{tL})p_1/p_{tL}(1-p_1))\}_{t=1}^{T}$ where, given $t \in \{1,\ldots,T-1\}$ and $y_t \in Y_t$, the functions $S_t^*: Y_t \to \{0,1\}$, $\nu_t^*: Y_t \to N_+ \cup \{0\}$, $\delta_t^*: Y_t \to \{0,1\}$ are defined by (2.9), where p_{tL} and p_{tU} are as defined in Theorem 3.1 and where

$$S_t^*(y_t) = 1 \text{ and } \delta_t^*(y_t) = 1 \quad , \text{ if } 0 \leq \lambda_t \leq \frac{(1-p_{tU})p_1}{p_{tU}(1-p_1)}$$

$$S_t^*(y_t) = 0 \text{ and } \nu_t^*(y_t)$$
$$\text{extra observations collected, if } \frac{(1-p_{tU})p_1}{p_{tU}(1-p_1)} < \lambda_t < \frac{(1-p_{tL})p_1}{p_{tL}(1-p_1)} \tag{4.7}$$

$$S_t^*(y_t) = 1 \text{ and } \delta_t^*(y_t) = 0 \quad , \text{ if } \lambda_t \leq \frac{(1-p_{tL})p_1}{p_{tL}(1-p_1)} .$$

Define (4.7) to be the <u>VPRT</u>.

Since the set of all WSPRT's $W \subset D$, the set of all decision procedures, and since the VPRT is an element of W, it follows that the VPRT maximizes the expected net gain over all WSPRT's.

The model's assumptions are very general; virtually the only restrictions put on the decision problem so far are that $c_0 \geq 0$ and c(n) is strictly increasing (even this can be relaxed, provided expected values remain well-defined). The optimality theorems of Morgan and Manning (1985, Theorems 2 and 3) show that the VPRT dominates the SPRT. This should not be surprising. The SPRT is a very special type of WSPRT where $\nu_t(\cdot) \equiv 1$ at every decision point t, and this restriction will be inconsistent with achieving highest expected net gain in many circumstances. What then of the famous Wald-Wolfowitz result that the SPRT dominates "all" tests of the same or smaller size whatever the state of nature? In the remainder of this section we show how the Wald-Wolfowitz theorem is a special case of the results of our more general model.

The Wald-Wolfowitz theorem finds a test procedure of a given size and power that minimizes the expected total number of observations needed to choose from $\Omega = \{\omega_0, \omega_1\}$, when either one of ω_0 or ω_1 is the true state of nature. Their result that the SPRT is optimal in this sense relies on some special structure in addition to that of just one-at-a-time sampling; that is, $c(n) \equiv cn$, $\tau = 1$ and $T = \infty$. Our explanation of their result is to first show that if $\tau = 1$ then, for either state of nature, the VPRT (given by (4.7)) minimizes over W the expected total sampling cost incurred before a choice for ω is made. When $c(n) \equiv cn$ is imposed, this is equivalent to minimizing the expected total number of observations. When the extra restriction $T = \infty$ is added as well, we are able to show the VPRT coincides with the SPRT; thus the Wald-Wolfowitz theorem emerges as a special case of our more general optimality results.

We have seen that the optimal (expected-net-gain maximizing) decision procedure d* is equivalent to the VPRT given by (4.7). The following result shows the VPRT is also optimal in the sense that it minimizes the expected total cost of sampling.

<u>THEOREM 4.1</u>: Let α^* and β^* respectively denote the type I and type II error probabilities for the VPRT with given values of $u_{00}, u_{01}, u_{10}, u_{11}$, and $1 \leq T \leq \infty$, $\tau = 1$, and a given non-negative monotonic strictly increasing sampling cost function c(n). Let $D(\alpha^*,\beta^*)$ denote the set of all the WSPRT's with type I error probability no larger than α^* and type II error probability no larger than β^*. Let T_s denote the stopping time variable of the WSPRT. The VPRT minimizes the expected total sampling cost $E[\sum_{t=1}^{T_s-1} c(n_t)|\omega]$ over $D(\alpha^*,\beta^*)$ under either state of nature, $\omega = \omega_0$ or $\omega = \omega_1$.

<u>Proof</u>: See the Appendix.

If the additional restriction $c(n) \equiv cn$ is imposed, so that minimizing $E[c(n)]$ is equivalent to minimizing $cE[n]$, then a corollary reminiscent of, but still more general than, the Wald-Wolfowitz theorem is obtained.

<u>COROLLARY 4.1:</u> Under the conditions of Theorem 4.1 with $c(n) \equiv cn$ and $c > 0$, the VPRT minimizes the expected total sample size $E[\sum_{t=1}^{T_s-1} n_t|\omega]$ over $D(\alpha^*,\beta^*)$ under either state of nature, $\omega = \omega_0$ or $\omega = \omega_1$.

<u>Proof</u>: See the Appendix.

Corollary 4.1 achieves a generality beyond that of the Wald-Wolfowitz theorem by establishing the minimum-expected-total-sample-size result for either a finite or an infinite truncation time T, and by allowing possibly more than one observation to be collected at each decision point.

Now notice that when $c(n) \equiv cn$ it is never optimal for the decision-maker to defer making a terminal decision by doing nothing; that is, at <u>least</u> one extra datum is always gathered whenever the decision is to not stop. When the SPRT is used $\nu_t(\cdot) \equiv 1$, $\forall$ t = 1,...,T-1, so that $E[\sum_{t=1}^{T_s-1} n_t|\omega]$ is also the expected number of periods before stopping. However, when the VPRT is used and $c(n) \equiv cn$, $\nu_t^*(y_t) \geq 1$, $\forall$ t = 1,...,T-1 and the expected total sample size is minimized so, on average, the VPRT user not only needs fewer observations but waits fewer periods before stopping than does the SPRT user.

<u>COROLLARY 4.2</u>: Under the conditions of Theorem 4.1 with $c(n) \equiv cn$, $c > 0$, the expected time taken by the VPRT to reach a terminal decision is smaller than or equal to the expected time taken by the SPRT to reach a terminal decision, under either state of nature, $\omega = \omega_0$ or $\omega = \omega_1$.

<u>Proof</u>: See the Appendix.

To the above results we could now add the remaining special structure of the Wald-Wolfowitz theorem and demonstrate their result to be a special case of the above. However, a more informative route is to proceed by establishing a set of conditions more general than those of Wald-Wolfowitz and sufficient for the VPRT to reduce to the SPRT. The Wald-Wolfowitz theorem then follows as a special case.

THEOREM 4.2: The VPRT and the SPRT coincide under the conditions of Theorem 4.1 with $T = \infty$, $c_0 = 0$, $c(n) > 0$, and $c(n) - c(n-1)$ a non-decreasing function for $n \geq 1$.

Proof: See the Appendix.

Theorems 4.1 and 4.2 show the Wald-Wolfowitz result is unnecessarily restrictive when it assumes $c(n) \equiv cn$ with $c > 0$. The SPRT minimizes expected total sample size under the above conditions for quite general cost structures.

The intuition behind Theorem 4.2 is rather simple. First consider the effect of finite T. Suppose $T = 2$ only. Then an SPRT will allow the decisionmaker just one datum before a terminal decision has to be made. In contrast, the VPRT would typically dictate a single sample of several observations be collected. In cases when $T < \infty$ (and especially when T is small) we should therefore anticipate the VPRT's expected net gain to exceed that of the SPRT. Second, the effect of a discount rate $\tau < 1$ is to penalize delay in reaching a terminal choice for ω by reducing the present value of terminal payoffs. The way to reach a terminal decision more quickly is to gather observations more rapidly than one-at-a-time. Thus when $\tau < 1$ we should see the VPRT's expected net gain exceed that of the SPRT. Third, the effect of a positive fixed cost of sampling $c_0 > 0$ is also to penalize delays in reaching a terminal choice for ω since delay means at least one more fixed cost c_0 must be suffered. Thus, again the SPRT's one-at-a-time-sampling strategy will typically have a lower expected net gain than the VPRT's flexible sampling strategy. Finally, when the incremental sampling cost function $c(n) - c(n-1)$ is not monotonic increasing it is possible for the average cost per observation to be lower for some $n \geq 2$ than it is for $n = 1$.

These cost savings can be exploited by the VPRT since its more flexible sampling strategy allows the drawing of samples of more than one datum, thereby lifting the expected net gain of the VPRT above that of the SPRT.

5. THE ENDOGENEITY OF SIZE AND POWER

The size and power of the VPRT are respectively

$$\phi(\omega_0) \equiv \alpha := \sum_{t=1}^{T} \Pr(\text{choose } \omega_1 \text{ at } t|\omega_0) \tag{5.1}$$

and

$$\phi(\omega_1) \equiv \pi := \sum_{t=1}^{T} \Pr(\text{choose } \omega_1 \text{ at } t|\omega_1). \tag{5.2}$$

Using (4.7) we can rewrite (5.1) and (5.2) as

$$\phi(\omega) = \Pr(\lambda_1 \leq \frac{(1-p_{1U})p_1}{p_{1U}(1-p_1)} |\omega)$$

$$+ \sum_{t=2}^{T-1} \Pr(\bigcap_{i=1}^{t-1} \{\frac{(1-p_{iU})p_1}{p_{iU}(1-p_1)} < \lambda_i < \frac{(1-p_{iL})p_1}{p_{iL}(1-p_1)}\} \cap \{\lambda_t \leq \frac{(1-p_{tU})p_1}{p_{tU}(1-p_1)}\}|\omega)$$

$$+ \Pr(\bigcap_{i=1}^{T-1} \{\frac{(1-p_{iU})p_1}{p_{iU}(1-p_1)} < \lambda_i < \frac{(1-p_{iL})p_1}{p_{iL}(1-p_1)}\} \cap \{\lambda_T < \frac{(1-\hat{p})p_1}{\hat{p}(1-p_1)}\}|\omega) \tag{5.3}$$

for $\omega \in \{\omega_0, \omega_1\}$.

<u>REMARK 5.1</u>: $\alpha = \phi(\omega_0)$ and $\pi = \phi(\omega_1)$ are entirely determined by the values of the payoffs $u_{00}, u_{01}, u_{10}, u_{11}$, the discount rate τ, the truncation time T, the decisionmaker's prior probability p_1, and the cost function $c(n)$.

The remark is obvious from (5.3) once it is realized that each of (p_{1L},p_{1U}), $(p_{2L},p_{2U}),\ldots,$ $(p_{T-1,L},p_{T-1,U})$ are functions only of these parameters of the decision problem and not of the data. The implications of the remark are rather sweeping; the usual hypothesis testing procedure of choosing a test of a certain size α which achieves a power π for testing $\omega = \omega_0$ against $\omega = \omega_1$ is implicitly a selection of a subset of values for $u_{00},u_{01},u_{10},u_{11}$, τ, T, p_1 and c(n) that satisfy the restrictions (5.1) and (5.2). This subset has uncountably many members since (5.3) is homogeneous of degree zero with respect to $u_{00},u_{01},u_{10},u_{11}$ and c(n); for example, doubling each of these leaves $\{(p_{tL},p_{tU})\}_{t=1}^{T-1}$ unaffected and so has no effect upon α or π. The essential point we wish to make here is that using the VPRT _and_ choosing α and π is _equivalent_ to the imposition of a restriction on the payoffs, discount rate, decision horizon, prior probability and sampling cost function in a _Bayesian_ decision problem where the goal is to choose a decision procedure with highest possible expected-net-gain. It is possible then that the classical testing procedure of (implicitly) choosing the uninformative prior p_1 = ½ as well as a value for the type I error probability α, is inconsistent with optimal decisionmaking whenever (as is often the case) the payoffs, discount rate, truncation time and sampling cost function are exogenously determined for the decisionmaker; Cressie and Morgan (1986) use the results developed here to investigate the classical fixed-sample-size decision procedure and are able to improve upon the Neyman-Pearson (Neyman and Pearson, 1933) approach to testing hypotheses.

Equations (5.1), (5.2) and (5.3) yield an exact description of how α and π are allocated over the decision points t = 1,...,T. Notice that the sequence of probabilities summed in (5.3) is a monotonic decreasing sequence since each subsequent member of the sequence is the probability of the preceeding event intersected with another event. This reveals that α is "spent" and π is "generated" most rapidly at early decision points; stated informally, first impressions matter most.

6. CONCLUDING REMARKS

Under more realistic and general assumptions than those usually imposed in the sequential analysis literature, we have found a sequential decision procedure that is optimal amongst all procedures that choose between two possible states of nature using inferences based upon sequentially collected data. "Optimal" here means maximizing the expected net gain (payoff net of sampling costs) from choosing a state of nature (formulating the problem in terms of losses leads to identical results). We call the optimal decision procedure the "variable-sample-size-sequential probability ratio test" (VPRT).

The VPRT is a "wide-sense sequential probability ratio test" (WSPRT); that is, a sequential probability ratio test where the number of additional data that may be collected at any given decision point need not be 1, as is required by the well-known SPRT initially developed by Wald and Wolfowitz (1948). Let α^* and β^* respectively denote the VPRT's type I and type II error probabilities and $D(\alpha^*, \beta^*)$ denote the class of WSPRT's with type I and type II error probabilities bounded above by α^* and β^*. Then the VPRT is also optimal in the sense that it minimizes the expected total sampling cost over $D(\alpha^*, \beta^*)$ under either state of nature. In particular, the SPRT's expected total sampling cost is at least as large as, and often strictly larger than the VPRT's. When $T = \infty$, $\tau = 1$, $c(0) = 0$, $c(n) > 0$ and $c(n) - c(n-1)$ is monotonic increasing for $n \geq 1$, the VPRT coincides with the SPRT. Thus Wald and Wolfowitz's (1948) optimality result can be extended to a class of cost functions more general than $c(n) \equiv cn$ with $c > 0$.

Size α^* and power $1-\beta^*$ are endogenous, being functions of the parameters of the decision problem (e.g. payoffs, priors etc.). Thus fixing either the size or the power of a test puts restrictions on the values for payoffs, priors etc. which are consistent with the test being an optimal decision procedure (see Cressie and Morgan, 1986).

The determination of stop-continue boundaries for any particular problem is a computationally intensive task which is currently being considered by the authors.

Acknowledgements

In this research Cressie was partially supported by the National Science Foundation under grant SES 84-01460, "On-Site Research to Improve Government-Generated Social Science Data Base." The research was partially conducted at the U.S. Bureau of the Census while the author was a participant in the American Statistical Association/Census Bureau Research Program, which is supported by the Census Bureau and through the NSF grant.

Morgan was partially supported by the Social Sciences and Humanities Research Council of Canada under grant 410-84-1211, "Learning and the Search for Information" and by financial assistance from the Department of Economics at the University of Western Ontario.

APPENDIX

This appendix contains the proofs of Lemmas 3.1, 3.2 and 3.3, Theorems 3.1, 3.2, 4.1 and 4.2, and Corollaries 4.1 and 4.2. Since the maximum expected-net-gain functions $V_t^T(y_t,d^*,p_t)$ depend on the data y_t only through p_t for $d = d^*$ (see Section 3), we streamline the notation here and write $V_t^T(y_t,d^*,p_t)$ as $V_t^T(p_t,d^*)$.

<u>LEMMA 3.1</u>: Given $T \leq \infty$, $t \in \{1,\ldots,T\}$ and $y_t \in Y_t$, $V_t^T(p_t,d^*)$ is a convex and continuous function of $p_t \in [0,1]$.

<u>Proof</u>: Since a convex function is necessarily continuous the proof need establish only that $V_t^T(p_t,d^*)$ is convex with respect to p_t.

Choose values $p_t', p_t'' \in [0,1]$, $p_t' \neq p_t''$, and choose $\theta \in [0,1]$. Define $p_\theta := \theta p_t' + (1-\theta)p_t''$. Then $V_t^T(p_\theta,d^*)$

$$= \max \{E_\omega [U(\omega,\delta_t^*(p_\theta))|p_\theta], -c(\nu_t^*(p_\theta)) + \tau E_{F_t} [V_{t+1}^T(p_{t+1},d^*)|p_\theta]\}. \quad \text{(A.1)}$$

Now, the first component of (A.1) is convex with respect to p_t since

$$\begin{aligned}
&E_\omega[U(\omega,\delta_t^*(p_\theta))|p_\theta] \\
&= \max\{u_{00}(\theta p_t' + (1-\theta)p_t'') + u_{01}(1-\theta p_t' - (1-\theta)p_t''), \\
&\qquad u_{10}(\theta p_t' + (1-\theta)p_t'') + u_{11}(1-\theta p_t' - (1-\theta)p_t'')\} \\
&= \max\{\theta[u_{00}p_t' + u_{01}(1-p_t')] + (1-\theta)[u_{00}p_t'' + u_{01}(1-p_t'')], \\
&\qquad \theta[u_{10}p_t' + u_{11}(1-p_t')] + (1-\theta)[u_{10}p_t'' + u_{11}(1-p_t'')]\} \\
&\leq \theta\max\{u_{00}p_t' + u_{01}(1-p_t'), u_{10}p_t' + u_{11}(1-p_t')\} \\
&\qquad + (1-\theta)\max\{u_{00}p_t'' + u_{01}(1-p_t''), u_{10}p_t'' + u_{11}(1-p_t'')\} \\
&= \theta E_\omega[U(\omega,\delta_t^*(p_t'))|p_t'] + (1-\theta)E_\omega[U(\omega,\delta_t^*(p_t''))|p_t'']. \qquad \text{(A.2)}
\end{aligned}$$

Similarly we can use the fact that $E_{F_t}[V^T_{t+1}(p_{t+1},d^*)|p_t]$ is a linear function of p_t to prove

$$E_{F_t}[V^T_{t+1}(p_{t+1},d^*)|p_\theta]$$

$$= \theta E_{F_t}[V^T_{t+1}(p_{t+1},d^*)|p_t^{'}] + (1-\theta)E_{F_t}[V^T_{t+1}(p_{t+1},d^*)|p_t^{''}], \qquad \text{(A.3)}$$

from which it follows that

$$-c(\nu^*_t(p_\theta)) + \tau E_{F_t}[V^T_{t+1}(p_{t+1},d^*)|p_\theta]$$

$$= \theta\{-c(\nu^*_t(p_\theta)) + \tau E_{F_t}[V^T_{t+1}(p_{t+1},d^*)|p_t^{'}]\}$$

$$+ (1-\theta)\{-c(\nu^*_t(p_\theta)) + \tau E_{F_t}[V^T_{t+1}(p_{t+1},d^*))|p_t^{''}]\} \qquad \text{(A.4)}$$

$$\le \theta\{-c(\nu^*_t(p_t^{'})) + \tau E_{F_t}[V^T_{t+1}(p_{t+1},d^*)|p_t^{'}]\}$$

$$+ (1-\theta)\{-c(\nu^*_t(p_t^{''})) + \tau E_{F_t}[V^T_{t+1}(p_{t+1},d^*)|p_t^{''}]\}. \qquad \text{(A.5)}$$

The inequality between (A.5) and (A.4) follows because $\nu^*(p_t^{'})$ and $\nu^*(p_t^{''})$ maximize the expected value of continuing, for $p_t = p_t^{'}, p_t^{''}$ respectively. Combining (A.1), (A.2) and (A.5) shows

$$V^T_t(p_\theta,d^*) \le \theta V^T_t(p_t^{'},d^*) + (1-\theta)V^T_t(p_t^{''},d^*). \qquad \text{(A.6)}$$

Q.E.D.

<u>**THEOREM 3.1**</u>: Given $2 \leq T \leq \infty$, $t \in \{1,\ldots,T-1\}$ and $y_t \in Y_t$, if

$$-c(\nu_t^*(\hat{p})) + \tau E_{F_t}[V_{t+1}^T(p_{t+1},d^*)|p_t = \hat{p}]$$

$$> u_{00}\hat{p} + u_{01}(1-\hat{p}) = u_{10}\hat{p} + u_{11}(1-\hat{p}), \qquad \text{(A.7)}$$

then there exist $p_{tL}, p_{tU} \in [0,1]$, $p_{tL} < \hat{p} < p_{tU}$ such that the maximum-expected-net-gain (optimal) decision procedure d* is:

stop sampling and choose ω_1	, if $p_t \in [0,p_{tL}]$
defer a terminal decision and collect $\nu_t^*(p_t)$ further observations	, if $p_t \in (p_{tL},p_{tU})$
stop sampling and choose ω_0	, if $p_t \in [p_{tU},1]$.

<u>Proof</u>: For $p_t \in [0,\hat{p}]$ define the function

$$g(p_t) := u_{10}p_t + u_{11}(1-p_t) + c(\nu_t^*(p_t)) - \tau E_{F_t}[V_{t+1}^T(p_{t+1},d^*)|p_t]; \qquad \text{(A.8)}$$

g is the sum of two functions which are convex (and continuous) with respect to p_t, so g is convex (and continuous) with respect to p_t.

Comparing (A.7) to (A.8) shows $g(\hat{p}) < 0$. Also,

$$g(0) = u_{11} + c(\nu_t^*(0)) - \tau u_{11} = (1-\tau)u_{11} + c_0 \geq 0. \qquad \text{(A.9)}$$

Therefore there exists exactly one root solving $g(p_t) = 0$ for $p_t \in [0,\hat{p})$. Denote this root by p_{tL}. Then

$$u_{10}p_t + u_{11}(1-p_t) \; \begin{matrix} > \\ = \\ < \end{matrix} \; - c(\nu_t^*(p_t)) + \tau E_{F_t}[V_{t+1}^T(p_{t+1},d^*)|p_t]$$

$$\text{as} \quad p_t \; \begin{matrix} \in [0,p_{tL}) \\ = p_{tL} \\ \in (p_{tL},\hat{p}] \end{matrix} \; . \qquad \text{(A.10)}$$

An analogous argument establishes that (A.7) implies the existence of a unique value $p_{tU} \in (\hat{p},1]$ such that

$$u_{00}p_t + u_{01}(1-p_t) \begin{matrix} > \\ = \\ < \end{matrix} - c(\nu_t^*(p_t)) + \tau E_{F_t}[V_{t+1}^T(p_{t+1},d^*)|p_t] \quad \text{as} \quad p_t \begin{matrix} \in (p_{tU},1] \\ = p_{tu} \\ \in [\hat{p},p_{tU}) \end{matrix} . \tag{A.11}$$

Combining (A.10) and (A.11) establishes the result.

Q.E.D.

<u>LEMMA 3.2</u>. Given $t \in \{1,\ldots,T\}$ and $p_t \in [0,1]$, $V_t^T(p_t,d^*)$ is a monotonic increasing function of T.

<u>Proof</u>: $V_t^T(p_t,d^*)$

$$= \max\{E_\omega[U(\omega,\delta_t^*(p_t))|p_t],$$

$$-c(\nu_t^*(p_t)) + \tau E_{F_t}[\max\{E_\omega[U(\omega,\delta_{t+1}^*(p_{t+1}))|p_{t+1}],$$

$$-c(\nu_{t+1}^*(p_{t+1})) +\ldots+ \tau E_{F_{T-1}}[E_\omega[U(\omega,\delta_T^*(p_T))|p_T]|p_{T-1}]\ldots\}|p_t]\}$$

$$\leq \max\{E_\omega[U(\omega,\delta_t^*(p_t))|p_t],$$

$$-c(\nu_t^*(p_t)) + \tau E_{F_t}[\max\{E_\omega[U(\omega,\delta_{t+1}^*(p_{t+1}))|p_{t+1}],$$

$$-c(\nu^*_{t+1}(p_{t+1})) + \ldots + \tau E_{F_{T-1}} [\max\{E_\omega [U(\omega,\delta^*_T(p_T))|p_T],$$

$$-c(\nu^*_T(p_T)) + \tau E_{F_T} [E_\omega [U(\omega,\delta^*_{T+1}(p_{T+1}))|p_{T+1}]|p_T]\}|p_{T-1}]\ldots\}|p_t]\}$$

$$= V_t^{T+1}(p_t,d^*), \ \forall \ p_t \in [0,1]. \qquad \text{Q.E.D.}$$

<u>LEMMA 3.3.</u> $\{V_t^T\}_{T=t}^\infty$ is uniformly convergent to V_t^∞, $\forall$ $t \geq 1$.

<u>Proof</u>: The result follows immediately from Dini's Theorem: monotonicity is established in Lemma 3.2, and continuity with respect to $p \in [0,1]$ of $V_t^T(p,d^*)$, $\forall$ $T \geq t \geq 1$, is established in Lemma 3.1.

Q.E.D.

<u>THEOREM 3.2.</u> For given $t \geq 1$ the sequence $\{(p_{tL}(T),\ p_{tU}(T))\}_{T=t}^\infty$ is nested, increasing and convergent to $(p_{tL}(\infty),\ p_{tU}(\infty))$.

<u>Proof</u>: For any $T \geq t+1$, $p_{tL}(T)$ uniquely solves

$$g(p_t) = u_{10}p_t + u_{11}(1-p_t) + c(\nu^*_t(p_t)) - \tau E_{F_t} [V^T_{t+1}(p_{t+1},d^*)|p_t] = 0.$$

By Lemma 3.2, $V^{T+1}_{t+1}(p_{t+1},d^*) \geq V^T_{t+1}(p_{t+1},d^*)$, $\forall$ $p_{t+1} \in [0,1]$.

Define

$$h(p_t) = u_{10}p_t + u_{11}(1-p_t) + c(\nu^*_t(p_t)) - \tau E_{F_t} [V^{T+1}_{t+1}(p_{t+1},d^*)|p_t]. \qquad (A.12)$$

Then $g(p_t) \geq h(p_t)$, $\forall$ $p_t \in [0,1]$. In particular,

$$h(p_{tL}(T+1)) = 0 = g(p_{tL}(T)) \geq h(p_{tL}(T)). \qquad (A.13)$$

In the proof of Theorem 3.1, g (and thus h) is shown to be convex and monotonic decreasing. Therefore (A.13) implies

$$p_{tL}(T+1) \leq p_{tL}(T), \ \forall \ T \geq t+1. \qquad (A.14)$$

An analogous argument establishes that

$$p_{tU}(T+1) \geq p_{tU}(T), \ \forall\ T \geq t+1 \qquad \text{(A.15)}$$

which, with (A.14), establishes the nesting of the sequence of the continue-sampling regions.

By definition $p_L^* \equiv p_{tL}(\infty)$ solves

$$g(p_L^*) = u_{10}p_L^* + u_{11}(1-p_L^*) + c(\nu_t^*(p_L^*)) - \tau E_{F_t}[V_t^{\infty}(p_{t+1},d^*)|p_L^*] = 0 \qquad \text{(A.16)}$$

and is unique. Since g is continuous in p_t it follows that $\lim_{T\to\infty} g(p_{tL}(T)) = g(\lim_{T\to\infty} p_{tL}(T)) = 0$; that is

$$\lim_{T\to\infty}\{u_{10}p_{tL}(T)+u_{11}(1-p_{tL}(T))+c(\nu_t^*(p_{tL}(T)))-\tau E_{F_t}[V_{t+1}^{T}(p_{t+1},d^*)|p_{tL}(T)]\}$$

$$= u_{10}\lim_{T\to\infty} p_{tL}(T) + u_{11}(1-\lim_{T\to\infty} p_{tL}(T)) + c(\nu_t^*(\lim_{T\to\infty} p_{tL}(T)))$$

$$+ \tau E_{F_t}[V_{t+1}^{\infty}(p_{t+1},d^*)|\lim_{T\to\infty} p_{tL}(T)] = 0. \qquad \text{(A.17)}$$

Comparing (A.16) to (A.17) shows

$$p_L^* = \lim_{T\to\infty} p_{tL}(T). \qquad \text{(A.18)}$$

An analogous argument proves $p_U^* = \lim_{T\to\infty} p_{tU}(T)$.

Q.E.D.

<u>THEOREM 4.1</u>: Let α^* and β^* respectively denote the type I and type II error probabilities for the VPRT with given values of $u_{00}, u_{01}, u_{10}, u_{11}$, and $1 \leq T \leq \infty$, $\tau = 1$, and a given non-negative monotonic strictly increasing sampling cost function c(n). Let $D(\alpha^*,\beta^*)$ denote the set of all the WSPRT's with type I error probability no larger than α^* and type II error probability no larger than β^*. Let T_s denote the stopping time variable of the WSPRT.

The VPRT minimizes the expected total sampling cost

$E[\sum_{t=1}^{T_s -1} c(n_t)|\omega]$ over $D(\alpha^*, \beta^*)$ under either state of nature, $\omega = \omega_0$ or $\omega = \omega_1$.

<u>Proof</u>: Define a sequence of functions $\{\psi_t^*\}_{t=1}^{T}$ where $\psi_t^*: [0,1] \rightarrow \{0,1\}$ is such that

$$\psi_t^*(p_t) = \begin{cases} 1, & \text{if the VPRT makes a terminal decision at } t \\ 0, & \text{otherwise,} \end{cases} \qquad \text{(A.19)}$$

Then the stopping time variable T_s^* of the VPRT is $T_s^* \equiv \sum_{t=1}^{T} t\psi$.

Furthermore the expected net gain of the VPRT, given $\omega = \omega_0$, is

$$V_1^T(p_1, d^*; \omega_0)$$

$$= \psi_1^*(p_1)[(1-\delta_1^*(p_1))u_{00} + \delta_1^*(p_1)u_{10}]$$

$$+ (1-\psi_1^*(p_1))\{-c(\nu_1^*(p_1)) + \int_{X_2} \psi_2^*(p_2)[(1-\delta_2^*(p_2))u_{00} + \delta_2^*(p_2)u_{10}]$$

$$+ (1-\psi_2^*(p_2))\{-c(\nu_2^*(p_2)) + \int_{X_3} \psi_3^*(p_3)[(1-\delta_3^*(p_3))u_{00} + \delta_3^*(p_3)u_{10}]$$

$$+ \ldots$$

$$+ (1-\psi_{T-1}^*(p_{T-1}))\{-c(\nu_{T-1}^*)) + \int_{X_T} \psi_T^*(p_T)[(1-\delta_T^*(p_T))u_{00} + \delta_T^*(p_T)u_{10}].$$

$$f(x_T;\omega_0)dx_T\}\ldots f(x_3;\omega_0)dx_3\}f(x_2;\omega_0)dx_2\}, \qquad \text{(A.20)}$$

where X_t is the support of the (vector-valued) variable x_t sampled at the t-th stage. Then (A.20) equals

$$u_{00}\Big[\psi_1^*(p_1)(1-\delta_1^*(p_1)) + \sum_{t=2}^{T-1} \int_{Y_t} \psi_t^*(p_t)(1-\delta_t^*(p_t)) \prod_{i=1}^{t-1}(1-\psi_i^*(p_i))f(x_i;w_0)dx_i$$

$$+ \int_{Y_T} (1-\delta_T^*(p_T)) \prod_{i=1}^{T-1}(1-\psi_i^*(p_i))f(x_i;\omega_0)dx_i\Big]$$

$$+ u_{10}\Big[\psi_1^*(p_1)\delta_1^*(p_1) + \sum_{t=2}^{T-1} \int_{Y_t} \psi_t^*(p_t)\delta_t^*(p_t) \prod_{i=1}^{t-1}(1-\psi_i^*(p_i))f(x_i;\omega_0)dx_i$$

$$+ \int_{Y_T} S_T^*(p_T) \prod_{i=1}^{T-1}(1-\psi_i^*(p_i))f(x_i;\omega_0)dx_i\Big]$$

$$- \sum_{t=1}^{T-1} \int_{Y_{t+1}} \sum_{k=1}^{t} c(\nu_k^*(p_k))\psi_{t+1}^*(p_{t+1}) \prod_{i=1}^{t}(1-\psi_i^*(p_i))f(x_i;\omega_0)dx_i\, f(x_{t+1};\omega_0)dx_{t+1} \qquad \text{(A.21)}$$

$$= u_{00} \sum_{t=1}^{T} \text{Pr(stop at t and choose } \omega_0|\omega_0)$$

$$+ u_{10} \sum_{t=1}^{T} \text{Pr(stop at t and choose } \omega_1|\omega_0)$$

$$- \sum_{t=1}^{T-1} E\Big[\sum_{k=1}^{t} c(\nu_k^*(p_k))|\text{stop at t+1}\Big] \text{Pr(stop at t+1}|\omega_0) \qquad \text{(A.22)}$$

$$= u_{00}(1-\alpha^*) + u_{10}\alpha^* - E\Big[\sum_{k=1}^{T_s^*-1} c(\nu_k^*(p_k))|\omega_0\Big]. \qquad \text{(A.23)}$$

Similarly, the expected net gain of the VPRT, given $\omega=\omega_1$, is

$$V_1^T(p_1,d^*;\omega_1) = u_{01}\beta^* + u_{11}(1-\beta^*) - E\Big[\sum_{k=1}^{T_s^*-1} c(\nu_k^*(p_k))|\omega_1\Big]. \qquad \text{(A.24)}$$

Repeating the above arguments for the expected net gain of any other WSPRT $d \in D(\alpha^*, \beta^*)$ shows this expected net gain is either

$$V_1^T(p_1, d; \omega_0) = u_{00}(1-\alpha) + u_{10}\alpha - E[\sum_{k=1}^{T_s-1} c(\nu_k(p_k))|\omega_0]; \ \alpha \le \alpha^*, \quad \text{(A.25)}$$

or

$$V_1^T(p_1, d; \omega_1) = u_{01}\beta + u_{11}(1-\beta) - E[\sum_{k=1}^{T_s-1} c(\nu_k(p_k))|\omega_1]; \ \beta \le \beta^*. \quad \text{(A.26)}$$

The VPRT d^* maximizes expected net gain over $D(\alpha^*, \beta^*)$ so, from (A.23), (A.24), (A.25) and (A.26),

$$(u_{00}(1-\alpha^*) + u_{10}\alpha^* - E[\sum_{k=1}^{T_s^*-1} c(\nu_k^*(p_k))|\omega_0])p_1$$

$$+ (u_{01}\beta^* + u_{11}(1-\beta^*) - E[\sum_{k=1}^{T_s^*-1} c(\nu_k^*(p_k))|\omega_1])(1-p_1)$$

$$\ge (u_{00}(1-\alpha) + u_{10}\alpha - E[\sum_{k=1}^{T_s-1} c(\nu_k(p_k))|\omega_0])p_1$$

$$+ (u_{01}\beta + u_{11}(1-\beta) - E[\sum_{k=1}^{T_s-1} c(\nu_k(p_k))|\omega_1])(1-p_1). \quad \text{(A.27)}$$

Since $\alpha \le \alpha^*$ and $\beta \le \beta^*$, $u_{00} \ge u_{10}$ and $u_{11} \ge u_{01}$,

$$u_{00}(1-\alpha^*) + u_{10}\alpha^* \le u_{00}(1-\alpha) + u_{10}\alpha \quad \text{(A.28)}$$

and

$$u_{01}\beta^* + u_{11}(1-\beta^*) \le u_{01}\beta + u_{11}(1-\beta). \quad \text{(A.29)}$$

Using (A.28) and (A.29) in (A.27) reduces (A.27) to

$$0 \geq \left(E\left[\sum_{k=1}^{T_s^*-1} c(\nu_k^*(p_k))|\omega_0\right] - E\left[\sum_{k=1}^{T_s-1} c(\nu_k(p_k))|\omega_0\right]\right)p_1$$

$$+ \left(E\left[\sum_{k=1}^{T_s^*-1} c(\nu_k^*(p_k))|\omega_1\right] - E\left[\sum_{k=1}^{T_s-1} c(\nu_k(p_k))|\omega_1\right]\right)(1-p_1). \qquad \text{(A.30)}$$

But (A.30) holds $\forall\ p_1 \in [0,1]$ so allowing $p_1 \to 0$ shows

$$E\left[\sum_{k=1}^{T_s-1} c(\nu_k(p_k))|\omega_1\right] \geq E\left[\sum_{k=1}^{T_s^*-1} c(\nu_k^*(p_k))|\omega_1\right] \qquad \text{(A.31)}$$

and allowing $p_1 \to 1$ shows

$$E\left[\sum_{k=1}^{T_s-1} c(\nu_k(p_k))|\omega_0\right] \geq E\left[\sum_{k=1}^{T_s^*-1} c(\nu_k^*(p_k))|\omega_0\right]. \qquad \text{(A.32)}$$

Together (A.31) and (A.32) are a statement that the VPRT d* minimizes the expected total sampling cost under either state of nature, ω_0 or ω_1.

Q.E.D.

<u>**COROLLARY 4.1**</u>: Under the conditions of Theorem 4.1 with $c(n) \equiv cn$ and $c > 0$, the VPRT minimizes the expected total sample size

$E[\sum_{t=1}^{T_s-1} n_t|\omega]$ over $D(\alpha^*, \beta^*)$ under either state of nature, $\omega = \omega_0$ or $\omega = \omega_1$.

<u>Proof</u>: With $c(n) \equiv cn$ the statement of Theorem 4.1 is that the VPRT d^* minimizes $cE[\sum_{k=1}^{T_s-1} \nu_k(p_k))|\omega]$, for $\omega = \omega_0, \omega_1$. The result follows immediately.

Q.E.D.

<u>COROLLARY 4.2</u>: Under the conditions of Theorem 4.1 with $c(n) \equiv cn$, $c > 0$, the expected time taken by the VPRT to reach a terminal decision is smaller than or equal to the expected time taken by the SPRT to reach a terminal decision, under either state of nature, $\omega = \omega_0$ or $\omega = \omega_1$.

<u>Proof</u>: Let $E[T_s^*|\omega]$ and $E[T_s'|\omega]$ respectively denote the expected times taken by the VPRT and the SPRT to reach terminal decisions. The VPRT's sample size function ν^* has the property that if sampling is continued at t then, necessarily, $\nu_t^*(p_t) \geq 1$. Therefore

$$E[T_s^*-1|\omega] \leq E[\sum_{k=1}^{T_s^*-1} \nu_k^*(p_k)|\omega], \text{ for } \omega = \omega_0,\omega_1. \tag{A.33}$$

The SPRT's sample size function ν' is restricted so that $\nu_t' \equiv 1$, $\forall\ t = 1,\ldots,T_s'$. Therefore

$$E[T_s'-1|\omega] \equiv E[\sum_{k=1}^{T_s'-1} \nu_t'(p_t)|\omega], \text{ for } \omega = \omega_0,\omega_1. \tag{A.34}$$

However, Corollary 4.1 establishes that

$$E[\sum_{k=1}^{T_s^*-1} \nu_k^*(p_k)|\omega] \leq E[\sum_{k=1}^{T_s'-1} \nu_t'(p_t)|\omega], \text{ for } \omega = \omega_0,\omega_1. \tag{A.35}$$

Thus (A.33), (A.34) and (A.35) together establish

$$E[T_s^*|\omega] \leq E[T_s'|\omega], \text{ for } \omega = \omega_0,\omega_1. \tag{A.36}$$

Q.E.D.

<u>THEOREM 4.2</u>: The VPRT and the SPRT coincide under the conditions of Theorem 4.1 with $T = \infty$, $c_0 = 0$, $c(n) > 0$ and $c(n) - c(n-1)$ a non-decreasing function for $n \geq 1$.

Proof: At any decision point t for which the VPRT's action of highest expected net gain is to continue sampling, it is necessary that the VPRT's sampling rule satisfies $\nu_t^* \geq 1$. The essence of the following proof is to demonstrate that if $\nu_t^*(p_t) \geq 2$ then there exists another WSPRT with a strictly larger expected net gain than that of the VPRT d^*. This contradicts the fact that d^* is a maximum-expected-net-gain decision procedure and so establishes $\nu_t^* = 1$ whenever it is optimal to continue.

For any $t \geq 1$ and any $p_t \in [0,1]$,

$$V_t^{\infty}(p_t,d^*) = \max\{E_\omega[U(\omega,\delta_t^*(p_t))|p_t],$$

$$- c(\nu_t^*(p_t)) + E_{F_t}[\max\{E_\omega[U(\omega,\delta_{t+1}^*(p_{t+1}))|y_{t+1}=(y_t:x_t),p_t],$$

$$- c(\nu_{t+1}^*(p_{t+1})) + E_{F_{t+1}}[\max\{E_\omega[U(\omega,\delta_{t+2}^*(p_{t+2}))|y_{t+2},p_t],$$

$$\ldots\}|y_{t+1}=(y_t:x_t),p_t]\}|p_t]\}. \qquad (A.37)$$

Also c(n) - c(n-1) is a nondecreasing function so

$$c(1) + c(\nu_t^*(p_t)-1) \leq c(\nu_t^*(p_t)). \qquad (A.38)$$

Substituting (A.38) into (A.37) shows

$$V_t^{\infty}(p_t,d^*)$$

$$\leq \max\{E_\omega[U(\omega,\delta_t^*(p_t))|p_t],$$

$$-c(1)-c(\nu_t^*(p_t)-1) + E_{F_t}[\max\{E_\omega[U(\omega,\delta_{t+1}^*(p_{t+1}))|y_{t+1}=(y_t:x_t),p_t],$$

$$- c(\nu_{t+1}^*(p_{t+1})) + E_{F_{t+1}}[\max\{E_\omega[U(\omega,\delta_{t+2}^*(p_{t+2}))|y_{t+2},p_t],$$

$$\ldots\}|y_{t+1}=(y_t:x_t),p_t]\}|p_t]\}$$

$$\leq \max\{E_{\omega}[U(\omega,\delta^*_t(p_t))|p_t],$$

$$- c(1)-c(\nu^*_t(p_t)-1) + E_{F_t}[\max\{E_{\omega}[U(\omega,\delta^*_{t+1}(p_{t+1}))|y_{t+1}=(y_t:x_{t1}),p_t],$$

$$E_{\omega}[U(\omega,\delta^*_{t+1}(p_{t+1}))|y_{t+1}=(y_t:x_t),p_t],$$

$$- c(\nu^*_{t+1}(p_{t+1})) + E_{F_{t+1}}[\max\{E_{\omega}[U(\omega,\delta^*_{t+2}(p_{t+2}))|y_{t+2},p_t],$$

$$\ldots\}|y_{t+1}=(y_t:x_t),p_t]\}|p_t]\}$$

$$= \max\{E_{\omega}[U(\omega,\delta^*_t(p_t))|p_t], - c(1) + E_{F_t}[\max\{-c(\nu^*_t(p_t)-1)$$

$$+ E_{\omega}[U(\omega,\delta^*_{t+1}(p_{t+1}))|y_{t+1}=(y_t:x_{t1}),p_t],$$

$$- c(\nu^*_t(p_t)-1) + E_{\omega}[U(\omega,\delta^*_{t+1}(p_{t+1}))|y_{t+1}=(y_t:x_t),p_t],$$

$$-c(\nu^*_t(p_t)-1)-c(\nu^*_{t+1}(p_{t+1}))+E_{F_{t+1}}[\max\{\ldots\}|y_{t+1}=(y_t:x_t),p_t]\}|p_t]\}$$

$$\leq \max\{E_{\omega}[U(\omega,\delta^*_t(p_t))|p_t],$$

$$-c(1) + E_{F_t}[\max\{-c(\nu^*_t(p_t)-1)+E_{\omega}[U(\omega,\delta^*_{t+1}(p_{t+1}))|y_{t+1}=(y_t:x_{t1}),p_t],$$

$$-c(n_{t+1}=\nu^*_t(p_t)-1)+E_{F_{t+1}}[\max\{E_{\omega}[U(\omega,\delta^*_{t+2}(p_{t+2}))|y_{t+2}=(y_t:x_t),p_t],$$

$$-c(\nu^*_{t+2}(p_{t+2}))+E_{F_{t+2}}[\max\{\cdots\}|y_{t+2}=(y_t:x_t),p_t]\}|y_{t+1}=(y_t:x_{t1}),p_t]\}|p_t]\}$$

$$\leq \max\{E_{\omega}[U(\omega,\delta^{*}_{t}(p_t))|p_t],$$

$$- c(1) + E_{F_t}[\max\{E_{\omega}[U(\omega,\delta^{*}_{t+1}(p_{t+1}))|y_{t+1}=(y_t:x_{t1}),p_t],$$

$$- c(\nu^{*}_{t}(p_t)-1) + E_{F_{t+1}}[\max\{E_{\omega}[U(\omega,\delta^{*}_{t+2}(p_{t+2}))|y_{t+2}=(y_t:x_t),p_t],$$

$$-c(\nu^{*}_{t+2}(p_{t+2}))+E_{F_{t+2}}[\max\{\ldots\}|y_{t+2}=(y_t:x_t),p_t]\}|y_{t+1}=(y_t:x_{t1}),p_t]\}|p_t]\}. \tag{A.39}$$

Comparing (A.37) to the last part of (A.39) shows that if sampling is continued at t then the optimal procedure of taking $\nu^{*}_{t}(p_t)$ further observations at t is (weakly) dominated by the procedure of taking exactly one further observation at t and the remaining $\nu^{*}_{t}(p_t)-1$ observations at t+1. If $\nu^{*}_{t}(p_t) \geq 2$ then $\nu^{*}_{t}(p_t)-1 \geq 1$ and so $c(\nu^{*}_{t}(p_t)-1) > 0$; but then the last inequality of (A.39) is strict and contradicts the optimality of drawing $\nu^{*}_{t}(p_t)$ observations at t. Therefore $\nu^{*}_{t}(p_t) = 1$; that is, if the VPRT continues sampling at t then exactly one additional observation is taken at t, which means the VPRT is indistinguishable from the SPRT.

Q.E.D.

REFERENCES

Armitage, P. (1985). The search for optimality in clinical trials. International Statistical Review 53, 15-24.

Arrow, K.J., Blackwell, D. and Girshick, M.A. (1949). Bayes and minimax solutions of sequential decision problems. Econometrica 17, 213-244.

Cressie, N., and Morgan, P.B. (1986). Improving upon the Neyman-Pearson approach to testing hypotheses. Biometrika, submitted.

Ghosh, B.K. (1970). Sequential Tests of Statistical Hypotheses. Addison-Wesley, Reading.

Gupta, S.S., and Miescke, K.J. (1984). Sequential selection procedures - a decision theoretic approach. The Annals of Statistics 12, 336-350.

Irle, A. (1984). Extended optimality of sequential probability ratio tests. The Annals of Statistics 12, 380-386.

Morgan, P.B., and Manning, R. (1985). Optimal search. Econometrica 53, 923-944.

Neyman, J. and Pearson, E. S. (1933). On the problem of the most efficient tests of statistical hypotheses. Philosophical Transactions of the Royal Society, A 231, 140-185.

Pocock, S.J. (1977). Group sequential methods in the design and analysis of clinical trials. Biometrika 64, 191-199.

Quang, P.X. (1985). Robust sequential testing. The Annals of Statistics 13, 638-649.

Rudin, W. (1976). Principles of Mathematical Analysis, (3rd edition). McGraw-Hill.

Spahn, M., and Ehrenfeld, S. (1974). Optimal and suboptimal procedures in group sequential sampling. Naval Research Logistics Quarterly 21, 53-68.

Wald, A., and Wolfowitz, J. (1948). Optimum character of the sequential probability ratio test. The Annals of Mathematical Statistics 19, 326-339.