1993

The Vprt - A Sequential Testing Procedure Dominating The Sprt

Noel A. Cressie

University of Wollongong, ncressie@uow.edu.au

Peter Morgan

State University of New York

Follow this and additional works at: <https://ro.uow.edu.au/infopapers>

Part of the [Physical Sciences and Mathematics Commons](#)

Recommended Citation

Cressie, Noel A. and Morgan, Peter: The Vprt - A Sequential Testing Procedure Dominating The Sprt 1993, 431-450.

<https://ro.uow.edu.au/infopapers/3913>

The Vprt - A Sequential Testing Procedure Dominating The Sprt

Abstract

Under more general assumptions than those usually made in the sequential analysis literature, a variable-sample-size-sequential probability ratio test (VPRT) of two simple hypotheses is found that maximizes the expected net gain over all sequential decision procedures. In contrast, Wald and Wolfowitz [25] developed the sequential probability ratio test (SPRT) to minimize expected sample size, but their assumptions on the parameters of the decision problem were restrictive. In this article we show that the expected net-gain-maximizing VPRT also minimizes the expected (with respect to both data and prior) total sampling cost and that, under slightly more general conditions than those imposed by Wald and Wolfowitz, it reduces to the one-observation-at-a-time sequential probability ratio test (SPRT). The ways in which the size and power of the VPRT depend upon the parameters of the decision problem are also examined.

Keywords

sprt, dominating, sequential, procedure, vprt, testing

Disciplines

Physical Sciences and Mathematics

Publication Details

Cressie, N. A. & Morgan, P. (1993). The Vprt - A Sequential Testing Procedure Dominating The Sprt. *Econometric Theory*, 9 (3), 431-450.

THE VPRT: A SEQUENTIAL TESTING PROCEDURE DOMINATING THE SPRT

NOEL CRESSIE
Iowa State University

PETER B. MORGAN
SUNY at Buffalo

Under more general assumptions than those usually made in the sequential analysis literature, a variable-sample-size-sequential probability ratio test (VPRT) of two simple hypotheses is found that maximizes the expected net gain over all sequential decision procedures. In contrast, Wald and Wolfowitz [25] developed the sequential probability ratio test (SPRT) to minimize expected sample size, but their assumptions on the parameters of the decision problem were restrictive. In this article we show that the expected net-gain-maximizing VPRT also minimizes the expected (with respect to both data and prior) total sampling cost and that, under slightly more general conditions than those imposed by Wald and Wolfowitz, it reduces to the one-observation-at-a-time sequential probability ratio test (SPRT). The ways in which the size and power of the VPRT depend upon the parameters of the decision problem are also examined.

1. INTRODUCTION

Formally, the problem we consider is the fundamental issue of selecting between two competing models, viewed here as simple hypotheses. The framework within which the analysis is conducted is quite general and incorporates a number of special cases already established in the sequential analysis literature. The ideas in this paper have precedents in two literatures that can be surveyed only very briefly here. One literature contains the very influential results of those who have developed sequential probability ratio testing. The other contains explicitly economic models that allow strategies for trad-

We are grateful to G. Simons for perceptive comments that led to a restatement of Theorem 4.1. The article has benefited considerably from an anonymous referee's constructive remarks.

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 an ASA/NSF/Census Fellow.

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.

ing off augmentation of information sets with the cost of obtaining the new information.

The first literature commences with the outstanding works by Neyman and Pearson [17], Wald and Wolfowitz [25], and Arrow, Blackwell, and Girshick [1]. These and their antecedents find valuable daily application in many contexts. The use of sequential procedures in clinical trials is common in the pharmaceutical industry [4]. In the United States, the introduction of a new chemical, of unknown carcinogenic potency, is subjected to strict regulatory processes. These typically involve "impact" reports that are cost-benefit analyses considering, for example, the incremental net benefits of introducing a new pesticide. A recent example is given by Olson [18], who develops an optimal sequential screening rule that uses the probability-ratio-testing methodology described in this paper.

The second, explicitly economic, literature is more diverse. The literature describing sequential search rules for the economically efficient acquisition of data is summarized by Morgan and Manning [16]. A sophisticated sequential research and development rule is presented by Roberts and Weitzman [21]. Easley and Kiefer [8] and Kiefer and Nyarko [14,15] have examined the economically rational way in which a Bayesian agent can explore the location of a linear regression process with unknown parameters by manipulating the variables usually treated as exogenous.

Wald and Wolfowitz [25] added a sequential experimental design to the problem of choosing between two simple hypotheses and showed that the SPRT is optimal in the sense that, of all level α tests having the same power, the SPRT needs the fewest expected number of observations under either hypothesis; extensions can be found in [5,13]. The VPRT admits a wider class of experimental designs and it is shown here that this endows the VPRT with, first, an expected net gain (risk) always at least as large (small) as any member of a very large class of tests (including, for example, the Neyman-Pearson test and the SPRT) and, second, the expected (with respect to both data and prior) total sampling cost is always at least as small as that of any test with level no larger and power no smaller than those of the VPRT.

The design innovation added by Wald and Wolfowitz was to allow one-at-a-time sampling with a stop/continue decision point available after the acquisition of each new datum. Wald [24, pp. 101-104] recognized that restricting the design to an incremental sample size of unity could be relaxed and others (see [11, p. 224], [12], and [19]) have presented extensions to designs in which a fixed incremental sample size is either exogenously specified or else is chosen in an ad hoc manner. Our paper focuses attention on the design that *optimizes* on the sample-size-choice component of a sequential decision procedure. An initial examination of this problem is provided in [9,23,26] for special cases of the more general model examined here. These experimental design issues are related to search-rule developments described by Morgan and Manning [16].

Section 2 sets out the notation and the model structure. Section 3 establishes a sequence of sets of data-independent probabilities that completely describe the optimal decision procedure. In Section 4 the optimal decision procedure is shown to be equivalent to a VPRT that dominates the SPRT and reduces to the SPRT under conditions slightly more general than those imposed by Wald and Wolfowitz. Section 5 offers comments on the structure of the level and the power of the VPRT. Concluding remarks are given in Section 6 and longer proofs are collected in the Appendix.

2. MODEL STRUCTURE AND NOTATION

2.1. Sequential Decision Rules

Let $\Omega = \{\omega_0, \omega_1\}$ be a two-element set of the possible states of nature. The problem is to decide which of ω_0 or ω_1 is the true state. Let X be an observation space over which is defined a cumulative distribution function (c.d.f.) $F(x; \omega)$; $x \in X$ and $\omega \in \Omega$. Observations may be drawn from X to assist in the choice between ω_0 and ω_1 . Let f denote the Radon-Nikodym derivative of F with respect to a dominating measure (usually, either Lebesgue measure or the counting measure).

Throughout this article we shall consider the case of simple hypotheses (as did Wald and Wolfowitz [25]), since that is where the concepts are clearest. Even when two hypotheses are composite, "least favorable" choices from each of the hypothesized parameter spaces lead to testing simple hypotheses.

Let T denote the number of decision points available to the decision maker, $1 \leq T \leq \infty$. Let $t \in \{0, \dots, T\}$ denote the t th decision point. When $T < \infty$, we have a truncated procedure. In other words, a decision about the true state of nature can be made or put off at any of the $t = 0, 1, \dots, T - 1$ discrete time points, but a choice between ω_0 and ω_1 *must* be made by time T .

The decision maker's initial information, at $t = 0$, is: the value of the truncation time T , complete knowledge of Ω and $\{F(\cdot; \omega), \omega \in \Omega\}$, an initial information vector y_0 , and an associated prior probability p_0 for the event $\{\omega = \omega_0\}$.

Observations $x_t = (x_{t1}, \dots, x_{tin_t})'$ are asked for at time t , collected between t and $t + 1$, and received by the decision maker at time $t + 1$, at a cost of $c_t(n_t)$ paid at time t . The case $n_t = 0$ corresponds to no observations being taken. Assume that for each $t \geq 0$, $c_t(0) \geq 0$, $c_t(n)$ is monotonic strictly increasing with n , and $c_t(\cdot)$ is unbounded above. The information vector y_t at time t is updated to include any new information x_t :

$$y_{t+1} = \begin{cases} y_t; & n_t = 0 \\ (y_t', x_t')'; & n_t \geq 1, \end{cases} \quad \text{for } t = 0, \dots, T - 1. \quad (2.1)$$

Define Y_t to be the range associated with the information vector y_t ; $t = 0, \dots, T$.

A (sequential) decision rule to choose between ω_0 or ω_1 consists of a stopping rule, a sample-size rule, and a terminal decision rule.

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

DEFINITION 2. A payoff function at time t is a mapping $U_t: \Omega \times \{0, 1\} \rightarrow R$, with

$$U_t(\omega, \delta(y_t)) = \begin{cases} u_{t00}, & \text{if } \delta(y_t) = 0 \text{ and } \omega = \omega_0 \\ u_{t01}, & \text{if } \delta(y_t) = 0 \text{ and } \omega = \omega_1 \\ u_{t10}, & \text{if } \delta(y_t) = 1 \text{ and } \omega = \omega_0 \\ u_{t11}, & \text{if } \delta(y_t) = 1 \text{ and } \omega = \omega_1; \end{cases} \quad t = 0, \dots, T. \quad (2.2)$$

Typically $u_{t00} > u_{t10}$ and $u_{t11} > u_{t01}$, since correct decisions are usually rewarded more generously than incorrect decisions. Each u_{tij} is assumed to be finite. ■

DEFINITION 3. A sample-size rule ν is a sequence $\{\nu_t\}_{t=0}^T$ where $\nu_t: Y_t \rightarrow Z^+ \cup \{0\}$, the set of non-negative integers; $\nu_T(\cdot) \equiv 0$. Call N the set of sample-size rules. ■

DEFINITION 4. A stopping rule S is a sequence $\{S_t\}_{t=0}^T$ where $S_t: Y_t \rightarrow \{0, 1\}$; sampling continues at t if $S_t(\cdot) = 0$ and a terminal decision is made at t if $S_t(\cdot) = 1$; $S_T(\cdot) \equiv 1$. Call Σ the space of stopping rules. ■

DEFINITION 5. A sequential decision rule d is an ordered triple (S, ν, δ) composed of a stopping rule S , a sample-size rule ν , and a terminal decision rule δ . Call $D \equiv \Sigma \times N \times \Delta$ the space of sequential decision rules. ■

In the decision-theory literature, payoffs are often expressed as negative losses, by way of a prespecified loss function. Which approach one uses is a matter of taste; in what follows, we shall look to maximize expected payoff net of costs (see Definition 6), which is clearly equivalent to minimizing the Bayes risk. Raiffa and Schlaifer [20, pp. 79–86] give a detailed discussion of this matter and are particularly helpful in developing the idea of costs as the opportunity losses that are an integral component of any optimization problem. Further, it is allowed that future payoffs and costs may be discounted by a discount factor $\tau \in [0, 1]$, such as in (2.8) below.

In this section the goal is to establish the existence and form of a decision rule $d^* \in D$ that yields the highest expected net gain amongst all rules $d \in$

D for choosing between ω_0 and ω_1 . As has been noted, such a rule is in fact a *Bayes rule*. The considerable breadth of D is best illustrated by considering some of its members. Purely sequential (one-at-a-time sampling) decision rules correspond to $\nu_t(\cdot) \equiv 1$; $t = 0, \dots, T-1$, and group-sequential decision rules correspond to $\nu_t(\cdot) \equiv n$; $t = 0, \dots, T-1$. The class of one-stage test procedures is defined by restricting $S_t(\cdot) \equiv 1$; $t \geq 1$ and the class of two-stage procedures is defined by restricting $S_t(\cdot) \equiv 1$; $t \geq 2$. Clearly, these are just a few of the many types of rules found in D .

2.2. Optimal Decision Rules

As part of the notational development, we now give the decision rule that turns out to be optimal when T is finite. Let p_t be the decision maker's posterior probability at t of the event $\omega = \omega_0$. By Bayes' theorem, for $t = 0, \dots, T-1$ and $n_t > 0$,

$$\begin{aligned} p_{t+1} &= \frac{f(x_t; \omega_0)p_t}{f(x_t; \omega_0)p_t + f(x_t; \omega_1)(1 - p_t)} \\ &= \frac{f(y_{t+1}; \omega_0)p_0}{f(y_{t+1}; \omega_0)p_0 + f(y_{t+1}; \omega_1)(1 - p_0)}, \end{aligned} \quad (2.3)$$

where $f(\cdot)$ represents the density (or probability mass) function of its argument. If $n_t = 0$ then $p_{t+1} = p_t$, for any $t \in \{0, \dots, T-1\}$. Also, the unconditional c.d.f. of x_t at t is

$$F_t(x_t) = p_t F(x_t; \omega_0) + (1 - p_t) F(x_t; \omega_1).$$

A dynamic-programming approach is used in the following analysis; the classical references in its vast literature are the books [2,3]. An expected-net-gain maximization rule is achieved by choosing the stopping rule, the sample-size rule, and the terminal decision rule that maximizes the expected net gain *at each stage* of the sequential procedure. Backward induction (e.g., see [20, p. 7]) is the technique that allows the optimal rule to be specified.

At time t , the expected net gain from a terminal decision procedure δ_t is

$$\begin{aligned} &\{u_{t00}(1 - \delta_t(y_t)) + u_{t10}\delta_t(y_t)\}p_t \\ &+ \{u_{t01}(1 - \delta_t(y_t)) + u_{t11}\delta_t(y_t)\}(1 - p_t). \end{aligned} \quad (2.4)$$

This expected net gain is maximized by choosing the terminal decision procedure

$$\delta_t^*(y_t) = \begin{cases} 0, & \text{if } u_{t00}p_t + u_{t01}(1 - p_t) > u_{t10}p_t + u_{t11}(1 - p_t) \\ \gamma, & \text{if } u_{t00}p_t + u_{t01}(1 - p_t) = u_{t10}p_t + u_{t11}(1 - p_t) \\ 1, & \text{if } u_{t00}p_t + u_{t01}(1 - p_t) < u_{t10}p_t + u_{t11}(1 - p_t), \end{cases} \quad (2.5)$$

where $\gamma \in [0, 1]$. Without loss of generality, we can select $\gamma = 0$. The maximal expected net gain is

$$E_\omega[U_t(\omega, \delta_t^*(y_t)) | y_t, p_t] \\ = \max\{u_{t00}p_t + u_{t01}(1 - p_t), u_{t10}p_t + u_{t11}(1 - p_t)\}, \quad (2.6)$$

where E_ω denotes the expectation with respect to the measure on Ω defined by the posterior probability p_t . Should the decision maker sample until $t = T$ is reached, then a terminal decision *must* be made. Thus, the expected gain at $t = T$ of a decision rule d incorporating δ_T^* is

$$V_T^T(y_T, d, p_T) \equiv E_\omega[U_T(\omega, \delta_T^*(y_T)) | y_T, p_T]. \quad (2.7)$$

At time $t = T - 1$, the decision maker can either choose between ω_0 and ω_1 , and expect a gain of $E_\omega[U_{T-1}(\omega, \delta_{T-1}^*(y_{T-1})) | y_{T-1}, p_{T-1}]$, or collect $n_{T-1} \geq 0$ additional observations and expect a net gain of

$$-c_{T-1}(n_{T-1}) + \tau E_{F_{T-1}}[V_T^T(y_T, d, p_T) | y_{T-1}, p_{T-1}, n_{T-1}], \quad (2.8)$$

where $\tau \in [0, 1]$ denotes the discount factor and $E_{F_{T-1}}$ denotes expectation with respect to the unconditional c.d.f. of the n_{T-1} observations collected at $t = T - 1$.

The expected net gain of continuing to sample at $t = T - 1$ is maximized by choosing the sample-size rule ν_{T-1}^* , defined as

$$\nu_{T-1}^*(y_{T-1}) = n_{T-1}^* \in \operatorname{argmax}_{n_{T-1} \geq 0} \{-c_{T-1}(n_{T-1}) \\ + \tau E_{F_{T-1}}[V_T^T(y_T, d, p_T) | y_{T-1}, p_{T-1}, n_{T-1}]\}. \quad (2.9)$$

(The notation, $\operatorname{argmax}_{n \geq 0}\{\dots\}$, is used to denote the value(s) of the argument n that achieves a maximum for the expression in braces.) Now, n_{T-1}^* exists and is finite since $c_{T-1}(\cdot)$ is unbounded above, and a unique value is obtained by choosing the smallest value in the argmax set. Thus, the expected net gain at $t = T - 1$ is maximized by choosing the stopping rule,

$$S_{T-1}^*(y_{T-1}) = \begin{cases} 1, & E_\omega[U_{T-1}(\omega, \delta_{T-1}^*(y_{T-1})) | y_{T-1}, p_{T-1}] \\ & \geq \{-c_{T-1}(\nu_{T-1}^*(y_{T-1})) \\ & + \tau E_{F_{T-1}}[V_T^T(y_T, d, p_T) | y_{T-1}, p_{T-1}, \nu_{T-1}^*(y_{T-1})]\} \\ 0, & \text{otherwise,} \end{cases} \quad (2.10)$$

and the maximal expected net gain is

$$V_{T-1}^T(y_{T-1}, d, p_{T-1}) \\ \equiv \max\{E_\omega[U_{T-1}(\omega, \delta_{T-1}^*(y_{T-1})) | y_{T-1}, p_{T-1}], -c_{T-1}(\nu_{T-1}^*(y_{T-1})) \\ + \tau E_{F_{T-1}}[V_T^T(y_T, d, p_T) | y_{T-1}, p_{T-1}, \nu_{T-1}^*(y_{T-1})]\}. \quad (2.11)$$

Backward induction yields the sequential decision rule $d^* \equiv (S^*, \nu^*, \delta^*)$:

(a) $\delta^* \equiv \{\delta_t^*\}_{t=0}^T$, where

$$\delta_t^*(y_t) = \begin{cases} 0, & \text{if } u_{t00}p_t + u_{t01}(1-p_t) \geq u_{t10}p_t + u_{t11}(1-p_t) \\ 1, & \text{if } u_{t00}p_t + u_{t01}(1-p_t) < u_{t10}p_t + u_{t11}(1-p_t). \end{cases} \quad (2.12a)$$

(b) $\nu^* \equiv \{\nu_t^*\}_{t=0}^T$, where

$$\nu_t^*(y_t) \in \underset{n_t \geq 0}{\operatorname{argmax}} \{-c_t(n_t) + \tau E_{F_t}[V_{t+1}^T(y_{t+1}, d^*, p_{t+1}) | y_t, p_t, n_t]\};$$

$$t = 0, \dots, T-1, \quad (2.12b)$$

and

$$\nu_T^*(y_T) \equiv 0. \quad (2.12c)$$

Recall that E_{F_t} denotes expectation with respect to the unconditional c.d.f. of the n_t observations collected at time t .

(c) $S^* \equiv \{S_t^*\}_{t=0}^T$, where

$$S_t^*(y_t) = \begin{cases} 1, & \text{if } \max\{u_{t00}p_t + u_{t01}(1-p_t), u_{t10}p_t + u_{t11}(1-p_t)\} \\ & \geq \{-c_t(\nu_t^*(y_t)) \\ & \quad + \tau E_{F_t}[V_{t+1}^T(y_{t+1}, d^*, p_{t+1}) | y_t, p_t, \nu_t^*(y_t)]\} \\ 0, & \text{otherwise;} \quad t = 0, \dots, T-1, \end{cases} \quad (2.12d)$$

and

$$S_T^*(y_T) \equiv 1. \quad (2.12e)$$

The sequence of maximal expected net gain functions implied by d^* is

$$V_t^T(y_t, d^*, p_t) = \max\{E_\omega[U_t(\omega, \delta_t^*(y_t)) | y_t, p_t], -c_t(\nu_t^*(y_t)) \\ + \tau E_{F_t}[V_{t+1}^T(y_{t+1}, d^*, p_{t+1}) | y_t, p_t, \nu_t^*(y_t)]\};$$

$$t = 0, \dots, T-1 \quad (2.13a)$$

and

$$V_T^T(y_T, d^*, p_T) = E_\omega[U_T(\omega, \delta_T^*(y_T)) | y_T, p_T]. \quad (2.13b)$$

DEFINITION 6. A decision rule $d' \in D$ is optimal if, for any given $y_t \in Y_t$ and any $p_t \in [0, 1]$,

$$V_t^T(y_t, d', p_t) \geq V_t^T(y_t, d, p_t), \quad \text{for all } d \in D; \quad t = 0, \dots, T, \quad (2.14)$$

where, for $d = (S, \nu, \delta)$,

$$V_t^T(y_t, d, p_t) = \{E_\omega[U_t(\omega, \delta_t(y_t)) | y_t, p_t]\}I(S_t(y_t) = 1) \\ + \{-c_t(\nu_t) + \tau E_{F_t}[V_{t+1}^T(y_{t+1}, d, p_{t+1}) | y_t, p_t, \nu_t]\} \\ \times I(S_t(y_t) = 0); \quad t = 0, \dots, T-1,$$

and

$$V_T^T(y_T, d, p_T) = E_\omega[U_T(\omega, \delta_T(y_T)) | y_T, p_T].$$

The problem of discovering an optimal decision procedure d' can be viewed as a special case of a search-theoretic problem analyzed by Morgan and Manning [16]. They provide a proof of the *existence* of an optimal procedure under conditions satisfied here and show (their Theorem 2, p. 397) the decision rule d^* described by (2.12) to be an optimal procedure.

3. PROPERTIES OF THE OPTIMAL DECISION RULE

The important result of this section is that for any $t \in \{0, \dots, T-1\}$ there exists a pair of *data-independent* probabilities that completely define the optimal decision rule. These probabilities, p_{iL}^T and p_{iU}^T , partition $[0,1]$ into three intervals: an interval where ω_0 is chosen, an interval where ω_1 is chosen, and an interval where sampling is continued. Continued sampling occurs when p_t is in the interval (p_{iL}^T, p_{iU}^T) , which is further divided into subintervals over which particular sample sizes are optimal. These results have been noted by Ehrenfeld [9] and proved by Spahn and Ehrenfeld [23] for the special case $T = \infty$, $\tau = 1$, time-invariant payoffs (with $u_{00} = u_{11} = 0$), and time-invariant costs.

Formula (2.12) shows that the elementary components $S_t^*(y_t)$, $\nu_t^*(y_t)$, $\delta_t^*(y_t)$ of d^* and the maximum-expected-net-gain functions $V_t^T(y_t, d^*, p_t)$ depend on the data only through p_t . Therefore, from now on denote these as $S_t^*(p_t)$, $\nu_t^*(p_t)$, $\delta_t^*(p_t)$, and $V_t^T(p_t, d^*)$.

The following three lemmas present results for the optimal decision rule d^* that are similar to those established for the optimal decision rule within the smaller class of pure sequential (one-at-a-time sampling) decision rules.

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

Proof. The proof is very similar to that of Ferguson [10, Lemma 1, p. 362], and follows from the fact that the maximum of convex functions is convex, and therefore continuous. ■

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

Proof. The result is easy to establish upon realizing that $V_t^T(\cdot) \equiv V_{t+1}^{T+1}(\cdot)$ and that $V_{t+1}^{T+1}(\cdot)$ is a maximum over $V_{t+1}^{T+1}(\cdot)$ and other terms. ■

LEMMA 3.3. *$\{V_t^T\}_{T=t}^\infty$ is uniformly convergent to V_t^∞ , for all $t \geq 0$.*

Proof. The result follows immediately from Dini's theorem: Monotonicity is established in Lemma 3.2, and continuity of the functions V_t^T is established in Lemma 3.1. ■

At any time t there is a value of p_t such that, if sampling ceased at this time, the decision maker would be indifferent between choosing ω_0 or ω_1 . This value, denoted by $\hat{p}_t$, is the value for which $u_{t00}p_t + u_{t01}(1 - p_t)$ equals $u_{t10}p_t + u_{t11}(1 - p_t)$, and hence it is data independent.

DEFINITION 7. For $1 \leq T \leq \infty$,

$$\hat{p}_t \equiv \frac{u_{t11} - u_{t01}}{(u_{t11} - u_{t01}) + (u_{t00} - u_{t10})}; \quad t = 0, \dots, T. \quad (3.1)$$

The major significance of $\hat{p}_t$ in the following two theorems is that $\hat{p}_t$ locates the smallest possible expected value of stopping sampling at time t ; that is,

$$\begin{aligned} u_{t00}\hat{p}_t + u_{t01}(1 - \hat{p}_t) &= u_{t10}\hat{p}_t + u_{t11}(1 - \hat{p}_t) \\ &= \min_{p_t \in [0, 1]} \max\{u_{t00}p_t + u_{t01}(1 - p_t), u_{t10}p_t \\ &\quad + u_{t11}(1 - p_t)\}. \end{aligned} \quad (3.2)$$

THEOREM 3.1. For given $1 \leq T \leq \infty$ and $t \in \{0, \dots, T - 1\}$, if

$$\begin{aligned} -c_t(v_t^*(\hat{p}_t)) + \tau E_{F_t}[V_{t+1}^T(p_{t+1}, d^*) | \hat{p}_t, v_t^*(\hat{p}_t)] &> u_{t00}\hat{p}_t + u_{t01}(1 - \hat{p}_t) \\ &= u_{t10}\hat{p}_t + u_{t11}(1 - \hat{p}_t), \end{aligned} \quad (3.3)$$

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

$$\begin{cases} \text{stop sampling and choose } \omega_1, & \text{if } p_t \in [0, p_{tL}^T] \\ \text{defer a terminal decision and collect} \\ \quad v_t^*(p_t) \text{ further observations,} & \text{if } p_t \in (p_{tL}^T, p_{tU}^T) \\ \text{stop sampling and choose } \omega_0, & \text{if } p_t \in [p_{tU}^T, 1]. \end{cases}$$

Proof. The proof follows a similar argument to that of Ferguson [10, pp. 363–364]. The inequality condition in the statement guarantees the intersection of $-c_t(v_t^*(p_t)) + E_{F_t}[V_{t+1}^T(p_{t+1}, d^*) | p_t, v_t^*(p_t)]$ with $u_{t00}p_t + u_{t01}(1 - p_t)$ at $p_t = p_{tU}^T$, and with $u_{t10}p_t + u_{t11}(1 - p_t)$ at $p_t = p_{tL}^T$. Crucial to the argument is the convexity proved in Lemma 3.1. Finally, each (p_{tL}^T, p_{tU}^T) pair does not depend upon p_t and is therefore data-independent. ■

By convention, put $p_{TL}^T = \hat{p}_T = p_{TU}^T$. Whenever the inequality condition (3.3) fails to hold there is no value of p_t for which sampling is the best action; the expected net gain of stopping exceeds that of continuing for every $p_t \in [0, 1]$ and so p_{tL}^T and p_{tU}^T do not exist.

The sequence $\{(p_{tL}^T, p_{tU}^T)\}_{t=0}^{T-1}$ can be computed before any observations are collected, because each continue-sampling interval is data-independent. Less

obvious is that the sample-size rule $\{\nu_i^*(p_i)\}_{i=0}^{T-1}$ can also be specified prior to any data being collected. Section 4.2 shows why this is so.

Theorem 3.1 could be expressed equivalently in terms of likelihood ratios rather than posterior probabilities. This we do in Section 4 where we define the VPRT.

THEOREM 3.2. *For given $t \geq 0$, $(p_{iL}^T, p_{iU}^T) \subseteq (p_{iL}^{T+1}, p_{iU}^{T+1})$ for every $T \geq t$, and (p_{iL}^T, p_{iU}^T) converges uniformly to $(p_{iL}^\infty, p_{iU}^\infty)$ as $T \rightarrow \infty$.*

Proof. Since the straight lines $y = u_{i00}p_t + u_{i01}(1 - p_t)$ and $y = u_{i10}p_t + u_{i11}(1 - p_t)$ do not depend upon T , Lemma 3.2 gives the nesting and increasing properties. Lemma 3.3 gives the convergence property. ■

For given $t \leq T$, the continue-sampling intervals (p_{iL}^T, p_{iU}^T) widen in a nested manner as T increases because the expected net gain of continuing to sample is an increasing function of the number of intervals remaining to the truncation time T . Now suppose that $T = \infty$ and the payoff and cost functions are the same in each period; that is,

$$\begin{aligned} u_{tij} &= u_{ij}; & i, j &= 0, 1, & t &= 0, \dots, T, & \text{and} \\ c_t(\cdot) &\equiv c(\cdot); & t &= 0, \dots, T-1. \end{aligned} \quad (3.4)$$

Then, $p_{iL}^\infty \equiv p_L^\infty$ and $p_{iU}^\infty \equiv p_U^\infty$ are independent of t because the passage of time does not reduce the number of periods remaining for the decision maker. This property, of time-invariant boundaries for the continue-sampling intervals when $T = \infty$ and (3.4) holds, is well known for the special case of the optimal one-at-a-time sequential decision procedure (e.g., see Ferguson [10, pp. 355–356]); the difference here is the one-at-a-time procedure's continue-sampling region, $(p_L', p_U') \subseteq (p_L^\infty, p_U^\infty)$. The containment will be strict in many circumstances. A class of problems for which $(p_L', p_U') = (p_L^\infty, p_U^\infty)$ is described in Theorem 4.3.

4. THE VARIABLE-SAMPLE-SIZE SEQUENTIAL PROBABILITY RATIO TEST (VPRT)

This section extends the pioneering results of Wald and Wolfowitz [25] and Arrow, Blackwell, and Girshick [1] from the class of purely sequential (one-at-a-time sampling) decision rules to the much larger class of decision rules described in Section 3. A brief review of the relevant literature is given in Section 1.

4.1. Wide-Sense Sequential Probability Ratio Tests

Given $T \leq \infty$ and $y_t \in Y_t$, define the likelihood ratio

$$\lambda_t \equiv \frac{f(y_t; \omega_0)}{f(y_t; \omega_1)} \equiv \frac{(1 - p_0)p_t}{p_0(1 - p_t)}; \quad t = 1, \dots, T, \quad (4.1a)$$

and

$$\lambda_0 \equiv 1. \quad (4.1b)$$

DEFINITION 8. Given $1 \leq T \leq \infty$, a wide-sense sequential probability ratio test (WSPRT) is a sequence $\{(S_t, \nu_t, \delta_t, P_t, N_t)\}_{t=0}^T$, where S_t, ν_t, δ_t are given by Definitions 4, 3, and 1, respectively, and

$$P_t \equiv \{P_{t0}, \dots, P_{tm_t}\}; \quad t = 0, \dots, T-1, \quad (4.2a)$$

$$0 \leq A_t \equiv P_{t0} < P_{t1} < \dots < P_{tm_t} \equiv B_t \leq \infty; \quad t = 0, \dots, T-1, \quad (4.2b)$$

$$0 \leq A_T = B_T \leq \infty \quad \text{and} \quad m_T = 0, \quad (4.2c)$$

$$N_t \equiv \{n_{t1}, \dots, n_{tm_t}\}, \quad n_{tj} \geq 0 \quad \text{for } j = 1, \dots, m_t; \\ t = 0, \dots, T-1, \quad \text{and} \quad N_T = \emptyset. \quad (4.2d)$$

The associated decision rule is, for $y_t \in Y_t$:

$$\left\{ \begin{array}{ll} S_t(y_t) = 1 \quad \text{and} \quad \delta_t(y_t) = 1, & \text{if } 0 \leq \lambda_t \leq A_t \\ S_t(y_t) = 0 \quad \text{and collect:} & \\ \quad \nu_t(y_t) = n_{tj} \text{ extra observations,} & \text{if } P_{t,j-1} < \lambda_t \leq P_{tj}, \\ & \text{for } j \in \{1, \dots, m_t - 1\}, \\ \quad \nu_t(y_t) = n_{tm_t} \text{ extra observations,} & \text{if } P_{t,m_t-1} < \lambda_t < P_{tm_t} \\ S_t(y_t) = 1 \quad \text{and} \quad \delta_t(y_t) = 0, & \text{if } B_t \leq \lambda_t \leq \infty; \\ & t = 0, \dots, T-1. \end{array} \right. \quad (4.3a)$$

For $t = T$,

$$S_T(\cdot) \equiv 1, \quad \text{with } \delta_T(y_T) = 1 \quad \text{if } 0 \leq \lambda_T < A_T = B_T \quad \text{and} \\ \delta_T(y_T) = 0 \quad \text{otherwise.} \quad (4.3b)$$

Let $\mathcal{W} \subset \mathcal{D}$ denote the class of wide-sense sequential probability ratio tests. Obviously, the usual one-at-a-time sequential probability ratio tests belong to $\mathcal{W}$, as do the group-sequential procedures. Less obviously, the Neyman-Pearson test [17] based on a fixed sample size, and the two-stage tests defined by likelihood-ratio test statistics, also belong to $\mathcal{W}$. The class $\mathcal{W}$ is very broad.

Due to Wald and Wolfowitz's optimality result for the SPRT, we use it as the "yardstick" against which to measure the performance of any member of $\mathcal{W}$. In what follows, we demonstrate that the maximum-expected-net-gain decision rule d^* described in Section 3 (i.e., the VPRT) is contained in $\mathcal{W}$, and hence dominates every member of $\mathcal{W}$, including the SPRT.

4.2. The VPRT

At any time $t \in \{0, \dots, T-1\}$ the VPRT's continue-sampling interval (p_{iL}^T, p_{iU}^T) may be partitioned as

$$(p_{i0}^*, p_{i1}^*] \cup \dots \cup (p_{i, m_t^* - 2}^*, p_{i, m_t^* - 1}^*] \cup (p_{i, m_t^* - 1}^*, p_{i, m_t^*}^*), \quad (4.4)$$

where $p_{iL}^T = p_{i0}^* < p_{i1}^* < \dots < p_{i, m_t^* - 1}^* < p_{i, m_t^*}^* = p_{iU}^T$. On these intervals,

$$\nu_t^*(p_t) = \begin{cases} n_{ij}^*; & p_t \in (p_{i, j-1}^*, p_{ij}^*], \quad \text{for } j = 1, \dots, m_t^* - 1 \\ n_{i, m_t^*}^*; & p_t \in (p_{i, m_t^* - 1}^*, p_{i, m_t^*}^*), \end{cases} \quad (4.5)$$

where $n_{i1}^* \neq n_{i2}^*, n_{i2}^* \neq n_{i3}^*, \dots, n_{i, m_t^* - 1}^* \neq n_{i, m_t^*}^*$. That is, for $t \in \{0, \dots, T-1\}$ the optimal sample-size rule is piecewise constant on (p_{iL}^T, p_{iU}^T) and there are $m_t^* - 1$ changes in the value of the sample-size rule. Consider a typical value p_{ij}^* at which a change occurs; at that value, the expected net gain of sampling by collecting n_{ij}^* additional observations is the same as for collecting $n_{i, j+1}^*$ additional observations. This value of p_t is thus data-independent in the same way as is $\hat{p}_t$ defined by (3.1). Therefore, the change points and the sample-size rule can be completely specified in advance of data collection.

Define

$$P_{ij}^* \equiv \frac{p_{ij}^*(1 - p_0)}{(1 - p_{ij}^*)p_0}; \quad j = 0, \dots, m_t^*, \quad t = 0, \dots, T-1, \quad \text{and} \\ A_t^* \equiv P_{i0}^*, \quad B_t^* \equiv P_{i, m_t^*}^*. \quad (4.6)$$

Further, define

$$\eta_t^*(\lambda_t) \equiv \nu_t^* \left(\frac{\lambda_t p_0}{\lambda_t p_0 + 1 - p_0} \right). \quad (4.7)$$

From Theorem 3.1 the optimal (expected-net-gain-maximizing) decision at time t is to choose ω_0 if $p_t \in [p_{iU}^T, 1]$, which is equivalent to

$$\text{choose } \omega_0 \quad \text{if } \lambda_t \geq \frac{p_{iU}^T(1 - p_0)}{(1 - p_{iU}^T)p_0} \equiv B_t^*; \quad t = 0, \dots, T-1. \quad (4.8)$$

Similarly, it is optimal to

$$\text{choose } \omega_1 \quad \text{if } \lambda_t \leq \frac{p_{iL}^T(1 - p_0)}{(1 - p_{iL}^T)p_0} \equiv A_t^*; \quad t = 0, \dots, T-1. \quad (4.9)$$

It is likewise optimal to

$$\text{collect} \begin{cases} n_{ij}^* = \eta_t^*(\lambda_t) \text{ extra observations} & \text{if } P_{i, j-1}^* < \lambda_t \leq P_{ij}^*; \\ & j = 1, \dots, m_t^* - 1 \\ n_{i, m_t^*}^* = \eta_t^*(\lambda_t) \text{ extra observations} & \text{if } P_{i, m_t^* - 1}^* < \lambda_t \leq P_{i, m_t^*}^*; \end{cases} \quad (4.10)$$

$t = 0, \dots, T-1$. And, for $t = T$, it is optimal to

$$\text{choose } \omega_1 \quad \text{if } \lambda_T < A_T^* = B_T^* \equiv \frac{\hat{p}_T(1 - p_0)}{(1 - \hat{p}_T)p_0}; \quad \text{otherwise choose } \omega_0. \quad (4.11)$$

Therefore, the optimal decision rule d^* can be rewritten as a WSPRT.

DEFINITION 9. *Given T , $1 \leq T \leq \infty$, the variable-sample-size sequential probability ratio test (VPRT) is a sequence $\{(S_t^*, \eta_t^*, \delta_t^*, P_t^*, N_t^*)\}_{t=0}^T$, where S_t^* and δ_t^* are defined by (2.12); η_t^* is defined by (4.7) and (2.12); $P_t^* \equiv \{P_{t0}^*, \dots, P_{t m_t^*}^*\}$, m_t^* , and P_{ij}^* are defined by (4.4), (4.5), and (4.6); and $N_t^* \equiv \{n_{t1}^*, \dots, n_{t m_t^*}^*\}$ is defined by (4.5). ■*

Because the payoffs and cost functions of the VPRT may be time-dependent, the components $(S_t^*, \eta_t^*, \delta_t^*, P_t^*, N_t^*)$ of the VPRT are time-dependent even in the limiting case of $T = \infty$; $(S_t^*, \eta_t^*, \delta_t^*, P_t^*, N_t^*)$ are time-invariant if $T = \infty$ and (3.4) holds. Even in this special case, however, the VPRT is more general than a group-sequential procedure since at any time t , $\eta_t^*(\lambda_t)$ is not necessarily a constant for every $\lambda_t \in (A^*, B^*)$.

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

4.3. Optimality Properties of the VPRT

The optimality criterion employed so far, that is, expected-net-gain-maximization (equivalently, Bayes-risk minimization), appears to be different from that used in establishing the optimality of the SPRT. Wald and Wolfowitz [25] show that the SPRT is optimal in the sense that, over *all* test procedures with type I and type II error probabilities no larger than those of the SPRT, it minimizes the expected total number of observations needed to choose from $\Omega = \{\omega_0, \omega_1\}$, irrespective of whether ω_0 or ω_1 is the true state of nature. It is important to notice that, since Wald and Wolfowitz chose a constant cost c per observation, minimizing the expected number of observations is equivalent to minimizing the expected total sampling cost. It will be shown that it is the latter notion that generalizes readily to other sampling cost functions and to variable sample-size rules. Can optimality in the sense of expected-net-gain-maximization imply optimality in the sense of expected-total-cost-minimization? Theorem 4.1 shows that for $\tau = 1$, the VPRT minimizes the expected (with respect to both data and prior) total sampling cost over *all* test procedures in W (including the SPRT) with type I and type II error probabilities no larger than those of the VPRT. This result may appear to conflict with Wald and Wolfowitz's result, but it does not. Theorem 4.2 shows that if $T = \infty$, $\tau = 1$, $c(0) = 0$, and $c(n)$ is nonnegative and superadditive (e.g., $c(n) \equiv c \cdot n$ satisfies these conditions), then the VPRT and the SPRT coincide.

The development of the expected-net-gain-maximizing rule d^* (i.e., the VPRT) in Sections 2 and 3 is more general than earlier developments (Cressie

and Morgan [6]) in that it admits *time-dependent* payoff functions and time-dependent cost functions. However, to prove the VPRT's expected-total-sampling-cost optimality and its coincidence with the SPRT, we must assume in the two theorems below that (3.4) holds (i.e., time-invariant payoff and cost function).

THEOREM 4.1. *Let α^* and β^* denote the type I and type II error probabilities, respectively, for the VPRT with given payoffs $u_{00}, u_{01}, u_{10}, u_{11}$, $1 \leq T \leq \infty$, discount factor $\tau = 1$ and an unbounded non-negative strictly increasing cost function $c(n)$. Let $D(\alpha^*, \beta^*)$ denote the set of all WSPRT's with type I error probability no larger than α^* and type II error probability no larger than β^* (equivalently, power no smaller than $1 - \beta^*$). Let T_s denote the stopping time variable of the WSPRT. The VPRT minimizes the expected total sampling cost $E[\sum_{i=0}^{T_s-1} c(n_i)]$ over $D(\alpha^*, \beta^*)$.*

Proof. See the Appendix. ■

Notice that when $c(n) \equiv c \cdot n$; $c > 0$, then minimizing the expected total sampling cost is equivalent to minimizing the expected total sample size. However, Theorem 4.1 does not generalize Wald and Wolfowitz's optimality result given in Section 1; see Schmitz [22]. It is a weaker result in the sense that the expectation (of the expected total sampling cost) is taken with respect to both the data and the prior. Wald and Wolfowitz [25] proved that, under somewhat restrictive assumptions, the SPRT minimizes the expected total sample size *given H_0 or given H_1* , where the expectation is taken with respect to the data.

THEOREM 4.2. *The VPRT and the SPRT have identical (optimal) expected net gain functions under the conditions of Theorem 4.1 with $T = \infty$, $c(0) = 0$, $c(n) > 0$ for $n > 0$, and $c(n) - c(n-1)$ a nondecreasing function for $n \geq 1$.*

Proof. See the Appendix. ■

Theorem 4.2 shows that the SPRT is an expected-net-gain-maximizing procedure under conditions more general than the classical ones ($c(n) \equiv n$, $u_{00} = u_{11} = 0$) used by Wald and Wolfowitz [25] and Arrow, Blackwell, and Girshick [1]. We have been able to construct numerical examples that show the conditions listed in Theorem 4.2 are minimal in the sense that relaxing any one of them allows the construction of a numerical counterexample to the conclusion.

5. POWER AND LEVEL CONSIDERATIONS

Let $\Phi(\omega) \equiv \sum_{t=0}^T \text{pr}(\text{choose } \omega_1 \text{ at } t | \omega)$. Then the level and power of a decision rule are, respectively,

$$\alpha = \Phi(\omega_0) \tag{5.1}$$

and

$$\pi = \Phi(\omega_1). \quad (5.2)$$

For the VPRT given in Section 4.2, and for $\omega \in \{\omega_0, \omega_1\}$,

$$\begin{aligned} \Phi(\omega) = & \Pr\left(\lambda_0 \leq \frac{p_{0L}^T(1-p_0)}{(1-p_{0L}^T)p_0} \mid \omega\right) \\ & + \sum_{i=1}^{T-1} \Pr\left(\bigcap_{i=0}^{i-1} \left\{ \frac{p_{iL}^T(1-p_0)}{(1-p_{iL}^T)p_0} < \lambda_i < \frac{p_{iU}^T(1-p_0)}{(1-p_{iU}^T)p_0} \right\} \right. \\ & \quad \left. \cap \left\{ \lambda_i \leq \frac{p_{iL}^T(1-p_0)}{(1-p_{iL}^T)p_0} \right\} \mid \omega\right) \\ & + \Pr\left(\bigcap_{i=0}^{T-1} \left\{ \frac{p_{iL}^T(1-p_0)}{(1-p_{iL}^T)p_0} < \lambda_i < \frac{p_{iU}^T(1-p_0)}{(1-p_{iU}^T)p_0} \right\} \right. \\ & \quad \left. \cap \left\{ \lambda_T \leq \frac{\hat{p}_T(1-p_0)}{(1-\hat{p}_T)p_0} \right\} \mid \omega\right). \end{aligned} \quad (5.3)$$

Notice that α and π 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 decision maker's prior probability p_0 , and the cost function $c(n)$. Therefore, the usual hypothesis testing procedure of choosing a test of a certain level α that achieves a power π for testing $\omega = \omega_0$ against $\omega = \omega_1$ could be viewed as a selection of a subset of values for $u_{00}, u_{01}, u_{10}, u_{11}, \tau, T, p_0$, and $c(n)$ that satisfy the restrictions (5.1) and (5.2). Thus, 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 maximize expected-net-gain). For example, it is possible that choosing the uninformative prior $p_0 = \frac{1}{2}$, as well as a value for the type I error probability α , is inconsistent with optimal decision making whenever (as is often the case) the payoffs, discount rate, truncation time, and sampling cost function are exogenously determined for the decision maker. Cressie and Morgan [7] investigate the (classical one-stage) case, $T = 1$, in some detail.

Notice further that when $T = \infty$ the time-invariance of the continue-sampling interval (p_L^∞, p_U^∞) makes the sequence of probabilities summed in (5.3) a monotonic decreasing sequence. This reveals that α is "spent" and π is "generated" most rapidly at early decision points; stated informally, first impressions matter most when T is large.

6. CONCLUDING REMARKS

Under assumptions more general than those usually imposed in the sequential analysis literature, we have found a sequential decision procedure that

is optimal among all procedures that choose between two possible states of nature. "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 instead of payoffs leads to minimizing the Bayes risk and identical results.) We call the optimal decision rule the "variable-sample-size sequential probability ratio test" (VPRT); it is not restricted to collecting a fixed number (such as one) of observations per time period. We show in Section 4 that, under appropriate conditions, the VPRT also minimizes expected (with respect to both data and prior) total sampling cost.

The determination of the VPRT is generally a computationally intensive task. Ehrenfeld [9] initiated a discussion of computing the optimal procedure for cases where the likelihood ratio takes a special form, and provided a computationally simple example of an independent and identically distributed sequence of Bernoulli trials. Spahn and Ehrenfeld [23] explain how a sequence of linear programming solutions may be used to compute the optimal VPRT procedure when $T = \infty$, $\tau = 1$, and $u_{00} = u_{11} = 0$. They use their algorithm to define six approximate (suboptimal) procedures, and show for a Bernoulli example that these approximations are very good. As well, they show that the SPRT's Bayes risk can be very much larger than that of the VPRT. Their study reports Bayes risk values and (constant) group sequential sample size values but does not reveal type I or type II error probabilities, expected total sample sizes, expected total sampling costs, the distribution of the stopping time variable, and so on. We have developed a computer program that performs these tasks not only for VPRT's but, as well, for best (within their class) one-sample, two-sample, group-sequential, and pure sequential probability ratio tests. A numerical comparison of these procedures will be reported elsewhere. The outcomes of a preliminary numerical study are given in [6].

REFERENCES

1. Arrow, K.J., D. Blackwell & M.A. Girshick. Bayes and minimax solutions of sequential decision problems. *Econometrica* 17 (1949): 213-244.
2. Bellman, R.E. *Dynamic Programming*. Princeton, NJ: Princeton University Press, 1957.
3. Bellman, R.E. & S.E. Dreyfus. *Applied Dynamic Programming*. Princeton, NJ: Princeton University Press, 1962.
4. Berry, D.A. & C.H. Ho. One-sided sequential stopping boundaries for clinical trials: A decision-theoretic approach. *Biometrics* 44 (1988): 219-227.
5. Bhat, B.R. & H.J. Vaman. Discount optimality of Wald SPRT for iid and certain Markov dependent sequences. *Journal of the Indian Statistical Association* 22 (1984): 13-21.
6. Cressie, N. & P.B. Morgan. The VPRT: Optimal sequential and nonsequential testing. In S.S. Gupta and J.O. Berger (eds.), *Statistical Decision Theory and Related Topics, IV*, vol. 2, pp. 107-118. New York: Springer, 1988.
7. Cressie, N. & P.B. Morgan. Design considerations for Neyman-Pearson and Wald hypothesis testing. *Metrika* 36 (1989): 317-325.
8. Easley, D. & N.M. Kiefer. Controlling a stochastic process with unknown parameters. *Econometrica* 56 (1986): 1045-1064.

9. Ehrenfeld, S. On group sequential sampling. *Technometrics* 14 (1972): 167–174.
10. Ferguson, T.S. *Mathematical Statistics: A Decision Theoretic Approach*. New York: Academic Press, 1967.
11. Ghosh, B.K. *Sequential Tests of Statistical Hypotheses*. Reading, MA: Addison-Wesley, 1970.
12. Gupta, S.S. & K.J. Miescke. Sequential selection procedures—A decision theoretic approach. *The Annals of Statistics* 12 (1984): 336–350.
13. Irle, A. Extended optimality of sequential probability ratio tests. *The Annals of Statistics* 12 (1984): 380–386.
14. Kiefer, N.M. & Y. Nyarko. Control of a linear regression process with unknown parameters. In W.A. Barnett, E.R. Berndt & H. White (eds.), *Dynamic Economic Modeling: Proceedings of the Third Econometric Symposium in Economic Theory and Econometrics*, pp. 105–120. Cambridge: Cambridge University Press, 1988.
15. Kiefer, N.M. & Y. Nyarko. Optimal control of an unknown linear process with learning. *International Economic Review* 30 (1989): 571–586.
16. Morgan, P.B. & R. Manning. Optimal search. *Econometrica* 53 (1985): 923–944.
17. Neyman, J. & E.S. Pearson. On the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society, A* 231 (1933): 140–185.
18. Olson, L.J. The search for a safe environment: The economics of screening and regulating environmental hazards. *Journal of Environmental Economics and Management* 19 (1990): 1–18.
19. Pocock, S.J. Group sequential methods in the design and analysis of clinical trials. *Biometrika* 64 (1977): 191–199.
20. Raiffa, H. & R. Schlaifer. *Applied Statistical Decision Theory*. Cambridge: The M.I.T. Press, 1961.
21. Roberts, K. & M. Weitzman. Funding criteria for research, development, and exploration projects. *Econometrica* 49 (1981): 1261–1288.
22. Schmitz, N. Wald-Wolfowitz optimality of sequentially planned tests: Remarks and conjectures. *Annals of Operations Research* 32 (1991): 205–213.
23. Spahn, M. & S. Ehrenfeld. Optimal and suboptimal procedures in group sequential sampling. *Naval Research Logistics Quarterly* 21 (1974): 53–68.
24. Wald, A. *Sequential Analysis*. New York: Wiley, 1947.
25. Wald, A. & J. Wolfowitz. Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics* 19 (1948): 326–339.
26. Whittle, P. Some general results in sequential design (with discussion). *Journal of the Royal Statistical Society, B* 27 (1965): 371–394.

APPENDIX

This Appendix contains the proofs of Theorems 4.1 and 4.2.

Proof of Theorem 4.1. Define a sequence of functions $\{\psi_t^*\}_{t=0}^T$ where $\psi_t^* : [0, 1]^{t+1} \rightarrow \{0, 1\}$ and

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

where $p_t \equiv (p_0, \dots, p_t)$. Then the stopping time variable T_s^* of the VPRT is $T_s^* \equiv \sum_{t=0}^{T-1} t\psi_t^*$. Furthermore, the expected net gain of the VPRT, given $\omega = \omega_0$, is

$$\begin{aligned} V_0^T(p_0, d^*; \omega_0) &= u_{00} \left[\psi_0^*(p_0)(1 - \delta_0^*(p_0)) + \sum_{t=1}^{T-1} \int_{Y_t} \psi_t^*(p_t)(1 - \delta_t^*(p_t)) \right. \\ &\quad \times \prod_{i=1}^{t-1} (1 - \psi_i^*(p_i)) f(x_i; \omega_0) dx_i \\ &\quad \left. + \int_{Y_T} (1 - \delta_T^*(p_T)) \prod_{i=1}^{T-1} (1 - \psi_i^*(p_i)) f(x_i; \omega_0) dx_i \right] \\ &+ u_{10} \left[\psi_0^*(p_0)\delta_0^*(p_0) + \sum_{t=1}^{T-1} \int_{Y_t} \psi_t^*(p_t)\delta_t^*(p_t) \right. \\ &\quad \times \prod_{i=1}^{t-1} (1 - \psi_i^*(p_i)) f(x_i; \omega_0) dx_i \\ &\quad \left. + \int_{Y_T} \delta_T^*(p_T) \prod_{i=1}^{T-1} (1 - \psi_i^*(p_i)) f(x_i; \omega_0) dx_i \right] \\ &- \sum_{t=0}^{T-1} \int_{Y_{t+1}} \sum_{k=1}^t c(\nu_k^*(p_k)) \psi_{t+1}^*(p_{t+1}) \\ &\quad \times \left\{ \prod_{i=1}^t (1 - \psi_i^*(p_i)) f(x_i; \omega_0) dx_i \right\} f(x_{t+1}; \omega_0) dx_{t+1} \quad (\text{A.2}) \end{aligned}$$

$$= u_{00}(1 - \alpha^*) + u_{10}\alpha^* - E \left[\sum_{k=0}^{T_s^*-1} c(\nu_k^*(p_k)) | \omega_0 \right]. \quad (\text{A.3})$$

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

$$V_0^T(p_0, d^*; \omega_1) = u_{01}\beta^* + u_{11}(1 - \beta^*) - E \left[\sum_{k=0}^{T_s^*-1} c(\nu_k^*(p_k)) | \omega_1 \right]. \quad (\text{A.4})$$

Morgan and Manning [16, Theorems 2 and 3] show that $|V_0^T(p_0, d^*; \omega_\ell)| < \infty$, for $\ell = 0, 1$. Repeating the arguments above for the expected net gain of any other WSPRT $d \in D(\alpha^*, \beta^*)$ shows this expected net gain (possibly, $-\infty$) is either

$$V_0^T(p_0, d; \omega_0) = u_{00}(1 - \alpha) + u_{10}\alpha - E \left[\sum_{k=0}^{T_s-1} c(\nu(p_k)) | \omega_0 \right]; \quad \alpha \leq \alpha^*, \quad (\text{A.5})$$

or

$$V_0^T(p_0, d; \omega_1) = u_{01}\beta + u_{11}(1 - \beta) - E \left[\sum_{k=0}^{T_s-1} c(\nu(p_k)) | \omega_1 \right]; \quad \beta \leq \beta^*, \quad (\text{A.6})$$

where α and β are, respectively, the WSPRT's type I and type II error probabilities. The VPRT maximizes expected net gain over $D(\alpha^*, \beta^*)$; so, from (A.3), (A.4), (A.5), (A.6), and recalling that $u_{00} \geq u_{10}$, $u_{11} \geq u_{01}$, we obtain

$$\begin{aligned}
 0 \geq & \left\{ E \left[\sum_{k=0}^{T_s^*-1} c(\nu_k^*(p_k)) | \omega_0 \right] - E \left[\sum_{k=0}^{T_s-1} c(\nu_k(p_k)) | \omega_0 \right] \right\} p_0 \\
 & + \left\{ E \left[\sum_{k=0}^{T_s^*-1} c(\nu_k^*(p_k)) | \omega_1 \right] - E \left[\sum_{k=0}^{T_s-1} c(\nu_k(p_k)) | \omega_1 \right] \right\} (1 - p_0). \quad (\text{A.7})
 \end{aligned}$$

Proof of Theorem 4.2. The VPRT's sampling rule necessarily satisfies $\nu_t^* \geq 1$ at any decision point t for which the VPRT's action of highest expected net gain is to continue sampling. The essence of the proof is to show that if, for some t , $\nu_t^*(p_t) \geq 2$, then the VPRT d^* can be improved.

Assume that there is a time t and a history vector y'_t generating a value of $p_t(y'_t)$ given by (2.3), for which $\nu_t^*(p_t(y'_t)) = n' \geq 2$. Denote the VPRT's history vector at time $(t+1)$ as $y_{t+1} = (y'_t : x_{t1}, x_{t2}, \dots, x_{tn'})$. (Recall that $(y : x)$ denotes the concatenation of vectors y and x .) The infinite truncation time allows us to compare the VPRT to a WSPRT that is identical to the VPRT up to time t but takes only one extra datum at time t and, if sampling continues at $(t+1)$, takes another $(n' - 1)$ data. From $t+2$ onward, the WSPRT is defined to be again identical to the VPRT. Thus, the WSPRT's history vector at $(t+1)$ is $y_{t+1} = (y'_t : x_{t1})$, and at $(t+2)$ is $y_{t+2} = ((y'_t : x_{t1}) : x_{t+1,1}, \dots, x_{t+1,n'-1})$, a vector of the same dimension as the VPRT's history vector at $(t+1)$.

The expected net gain of the VPRT at time t is

$$\begin{aligned}
 V_t^\infty(p_t(y'_t), d^*) \\
 = & \max\{E_\omega[U(\omega, \delta_t^*(p_t(y'_t))) | p_t(y'_t)], -c(n') \\
 & + E_{F_t}[\max\{E_\omega[U(\omega, \delta_{t+1}^*(p_{t+1}(y'_t : x_{t1}, \dots, x_{tn'}))) | p_{t+1}(y'_t : x_{t1}, \dots, x_{tn'})], \\
 & -c(\nu_{t+1}^*(p_{t+1}(y'_t : x_{t1}, \dots, x_{tn'}))) \\
 & + E_{F_{t+1}}[\max\{E_\omega[U(\omega, \delta_{t+2}^*(p_{t+2}(y_{t+2}))) | p_{t+2}(y_{t+2})], \\
 & \dots\} | p_{t+1}(y'_t : x_{t1}, \dots, x_{tn'})]\} | p_t(y'_t)]\}. \quad (\text{A.8})
 \end{aligned}$$

Since $c(n) - c(n-1)$ is nondecreasing,

$$c(1) + c(n' - 1) \leq c(n'). \quad (\text{A.9})$$

Substituting (A.9) into (A.8) gives,

$$\begin{aligned}
 V_t^\infty(p_t(y'_t), d^*) \\
 \leq & \max\{E_\omega[U(\omega, \delta_t^*(p_t(y'_t))) | p_t(y'_t)], -c(1) - c(n' - 1) \\
 & + E_{F_t}[\max\{E_\omega[U(\omega, \delta_{t+1}^*(p_{t+1}(y'_t : x_{t1}, \dots, x_{tn'}))) | p_{t+1}(y'_t : x_{t1}, \dots, x_{tn'})], \\
 & -c(\nu_{t+1}^*(p_{t+1}(y'_t : x_{t1}, \dots, x_{tn'}))) \\
 & + E_{F_{t+1}}[\max\{E_\omega[U(\omega, \delta_{t+2}^*(p_{t+2}(y_{t+2}))) | p_{t+2}(y_{t+2})], \\
 & \dots\} | p_{t+1}(y'_t : x_{t1}, \dots, x_{tn'})]\} | p_t(y'_t)]\} \quad (\text{A.10})
 \end{aligned}$$

$$\begin{aligned}
&\leq \max\{E_\omega[U(\omega, \delta_t^*(p_t(y'_t)))]|p_t(y'_t)], -c(1) - c(n' - 1) \\
&\quad + E_{F_t}[\max\{E_\omega[U(\omega, \delta_{t+1}^*(p_{t+1}(y'_t : x_{t1})))]|p_{t+1}(y'_t : x_{t1})], \\
&\quad E_{F_{t+1}}[\max\{E_\omega[U(\omega, \delta_{t+2}^*(p_{t+2}((y'_t : x_{t1}) : x_{t+1,1}, \dots, x_{t+1,n'-1}))) \\
&\quad |p_{t+2}((y'_t : x_{t1}) : x_{t+1,1}, \dots, x_{t+1,n'-1})], \\
&\quad -c(\nu_{t+2}^*(p_{t+2}((y'_t : x_{t1}) : x_{t+1,1}, \dots, x_{t+1,n'-1}))) \\
&\quad + E_{F_{t+2}}[\max\{E_\omega[U(\omega, \delta_{t+3}^*(p_{t+3}(y_{t+2} : x_{t+2})))]|p_{t+3}(y_{t+2} : x_{t+2})], \\
&\quad \dots\}|p_{t+2}((y'_t : x_{t1}) : x_{t+1,1}, \dots, x_{t+1,n'-1})]|p_{t+1}(y'_t : x_{t1})]|p_t(y'_t)]\}.
\end{aligned} \tag{A.11}$$

Since $c(n' - 1) > 0$, form (A.11),

$$\begin{aligned}
&V_t^\infty(p_t(y'_t), d^*) \\
&\leq \max\{E_\omega[U(\omega, \delta_t^*(p_t(y'_t)))]|p_t(y'_t)], \\
&\quad -c(1) + E_{F_t}[\max\{E_\omega[U(\omega, \delta_{t+1}^*(p_{t+1}(y'_t : x_{t1})))]|p_{t+1}(y'_t : x_{t1})], -c(n' - 1) \\
&\quad + E_{F_{t+1}}[\max\{E_\omega[U(\omega, \delta_{t+2}^*(p_{t+2}((y'_t : x_{t1}) : x_{t+1,1}, \dots, x_{t+1,n'-1}))) \\
&\quad |p_{t+2}((y'_t : x_{t1}) : x_{t+1,1}, \dots, x_{t+1,n'-1})], \\
&\quad -c(\nu_{t+2}^*(p_{t+2}((y'_t : x_{t1}) : x_{t+1,1}, \dots, x_{t+1,n'-1}))) \\
&\quad + E_{F_{t+2}}[\max\{E_\omega[U(\omega, \delta_{t+3}^*(p_{t+3}(y_{t+2} : x_{t+2})))]|p_{t+3}(y_{t+2} : x_{t+2})], \\
&\quad \dots\}|p_{t+2}((y'_t : x_{t1}) : x_{t+1,1}, \dots, x_{t+1,n'-1})]|p_{t+1}(y'_t : x_{t1})]|p_t(y'_t)]\}.
\end{aligned} \tag{A.12}$$

The right-hand side of (A.12) is the expected net gain of the WSPRT described at the beginning of this proof. Similar inequalities (or equalities) hold for all other values of t . Since $T = \infty$, the steps of the proof can be repeated to obtain a new WSPRT that has $n' - 2$ data at time $t + 2$, and has an expected net gain at least as large as that of the VPRT, and so forth. But the VPRT is optimal. Hence, necessarily, $\nu^*(p_t(y'_t)) = 1$; $t = 0, 1, \dots$, is an optimal sample-size rule. ■