

THE ODDS ALGORITHM BASED ON SEQUENTIAL UPDATING AND ITS PERFORMANCE

F. THOMAS BRUSS* AND

GUY LOUCHARD,** *Université Libre de Bruxelles*

Abstract

Let I_1, I_2, \dots, I_n be independent indicator functions on some probability space $(\Omega, \mathcal{A}, \mathbb{P})$. We suppose that these indicators can be observed sequentially. Furthermore, let T be the set of stopping times on (I_k) , $k = 1, \dots, n$, adapted to the increasing filtration (\mathcal{F}_k) , where $\mathcal{F}_k = \sigma(I_1, \dots, I_k)$. The odds algorithm solves the problem of finding a stopping time $\tau \in T$ which maximises the probability of stopping on the last $I_k = 1$, if any. To apply the algorithm, we only need the odds for the events $\{I_k = 1\}$, that is, $r_k = p_k/(1 - p_k)$, where $p_k = \mathbb{E}(I_k)$, $k = 1, 2, \dots, n$. The goal of this paper is to offer tractable solutions for the case where the p_k are unknown and must be sequentially estimated. The motivation is that this case is important for many real-world applications of optimal stopping. We study several approaches to incorporate sequential information. Our main result is a new version of the odds algorithm based on online observation and sequential updating. Questions of speed and performance of the different approaches are studied in detail, and the conclusiveness of the comparisons allows us to propose always using this algorithm to tackle selection problems of this kind.

Keywords: Optimal stopping; sequential estimation; algorithmic performance; simulation; asymptotic studies; clinical trials; secretary problem

2000 Mathematics Subject Classification: Primary 60G40

1. Introduction

Let I_1, I_2, \dots be independent indicator functions on a probability space $(\Omega, \mathcal{A}, \mathbb{P})$ with $p_k = \mathbb{E}(I_k)$. Furthermore, let

$$q_k = 1 - p_k, \quad r_k = \frac{p_k}{q_k},$$

that is, r_k presents the odds of the event $\{I_k = 1\}$. We may observe the indicators sequentially and we may stop on at most one, but only *online*, that is, at the moment of observation. We win if we stop on the *last* $I_k = 1$ (if any) and lose otherwise (including not stopping at all). Formally, let T be the set of nonanticipative stopping rules defined by $T = \{\tau : \{\tau = k\} \in \mathcal{F}_k\}$, where \mathcal{F}_k is the σ -algebra generated by I_1, I_2, \dots, I_k . The odds theorem of optimal stopping [2] determines the rule which maximises the probability of stopping on the last indicator that takes the value 1 (if any). This solution is conveniently computed via the odds algorithm described in Algorithm 1.1, below.

Received 4 April 2008; revision received 14 January 2009.

* Postal address: Université Libre de Bruxelles, Département de Mathématique, Facultés des Sciences, CP 210, Boulevard du Triomphe, B-1050 Bruxelles, Belgium. Email address: tbruss@ulb.ac.be

** Postal address: Université Libre de Bruxelles, Département d'Informatique, Facultés des Sciences, CP 212, Boulevard du Triomphe, B-1050 Bruxelles, Belgium. Email address: louchard@ulb.ac.be

Algorithm 1.1. (*The odds algorithm.*)

- *Input.* Define

$$R_k := r_n + r_{n-1} + \dots + r_k, \quad k = n, \dots, 1,$$

$$Q_k := q_n q_{n-1} \dots q_k, \quad k = n, \dots, 1,$$

and precompute

$$s = s(n) = \begin{cases} 1 & \text{if } R_1 < 1, \\ \sup\{k: R_k \geq 1\} & \text{otherwise.} \end{cases} \tag{1.1}$$

- *Output: an optimal stopping rule.* The optimal stopping rule to stop on the last ‘1’ is: stop on the first indicator I_k with $I_k = 1$ and $k \geq s$. If none exists, stop on n and lose. This strategy is named the odds strategy.

We say that we *win* with the odds strategy if the first $I_k = 1$ with $k \geq s$ is the last 1.

The optimal win probability (as seen at time $1, 2, \dots, s - 1$) equals $R_s Q_s$. The odds algorithm is very convenient and allows for many interesting applications (e.g. secretary problems, selection problems for randomly arriving objects, timing problems, buying and selling problems, clinical trials, automated maintenance problems, and others). See [2], [4], [10], [12], and [17]. The algorithm can also be adapted to continuous-time decision processes with Poisson arrivals (see [2]). Related problems have been studied by Hsiau and Yang [11], Suchwalko and Szajowski [15], and Szajowski [16]. In these problems the odds algorithm is also of partial help. It would be tempting to try to extend the applicability of the odds algorithm further to the full-information problem of Kurushima and Ano [13] and Mazalov and Tomaki [14], as well as to the maximum duration problem of Ferguson *et al.* [9]. However, in the first case, the indicator functions expressing the payoff for stopping become dependent, whereas, in the second case, the payoff is a continuous random variable. Therefore, we see no way to adapt the use of the odds algorithm to such problems.

A particular feature of the odds algorithm is that the number of computational steps to find the optimal value of s is (sub)-linear in n . The algorithm is optimal itself, in the sense that clearly no algorithm can exist which would in general yield the rule with less than $\mathcal{O}(n)$ computations. Since it yields the optimal rule and the optimal win probability at the same time, and is optimal itself, it is—in colloquial terms of optimal stopping—a rare ‘triple ace’.

A related problem to the problem of stopping on the last event $\{I_k = 1\}$ is the problem of stopping with maximum probability on the k th last indicator which equals 1 (see [6]). The precise solution is more complicated, but a slight modification of the odds algorithm gives a good approximation. A harder related problem is the problem of stopping on a last specific pattern in an independent sequence of variables taken from some finite or infinite alphabet, as studied in [5]. In these problems, the p_k are supposed to be known.

1.1. Unknown odds

The applicability of the above odds algorithm is restricted, because in many practical applications, the decision maker would not know beforehand the values p_k , at least not precisely. The corresponding optimal stopping problem for unknown p_k is now in general much harder. In some cases, the precise solution can be given, and this also within the framework of the odds algorithm (see [18]), but these cases are very special. In this paper we study the problem in more generality.

Note that we cannot give too much freedom to the randomness of the p_k , because, if we allow, as we typically do, the p_k to be different from each other, they must still be estimable.

More precisely, the odds $r_{k+1}, r_{k+2}, \dots, r_n$ must be estimable from I_1, I_2, \dots, I_k . This means that the number of unknown parameters on which the p_k (and, thus, the r_k) may depend, must stay very small compared with n . Since n is, in many important applications, as for instance the compassionate-use clinical trial example (see [4]), not large (10 or 15, say) we focus our interest in this paper on only one unknown parameter p . Hence, the p_k are thought of as being deterministic functions of one unknown parameter p .

1.2. The model $p_k = pf_k$

This is our main model. The parameter p is unknown, but the factor f_k is supposed to be known. This is an adequate setting for many problems. In the abovementioned clinical trial example, for instance, p is considered as the unknown success probability for a medical treatment and f_k is a factor (between 0 and 1) which reduces the success probability for the k th patient according to his or her state of health.

The idea is to combine the convenience of the odds algorithm with the concurrent task of estimating the ‘future odds’ from preceding observations. We will study both the case of a Bayesian setting with a prior for the unknown parameter p as well as the case of a completely unknown p . Both cases are well motivated. If a new type of practical problem is encountered, we sometimes have so little information that we should not take the risk of introducing a bias by a prior distribution. However, with some confirmed prior information, the Bayesian setting typically has the advantage of leading to more efficient estimators.

Let $(f_k), k = 1, \dots, n$, be a sequence of known real nonnegative values. We set

$$p_k = pf_k, \quad p \in [0, 1], pf_k \leq 1.$$

Here it is understood that if we suppose a support $[a, b]$ for the distribution of p other than $[0, 1]$ then the f_k may range between 0 and $1/b$, that is, f_k is not necessarily a reducing factor, but may also increase the intrinsic success probability.

The paper is organised as follows. In Section 2 we consider the case of fixed p . In Section 3 we assume that p is distributed according to some distribution $P(p)$. In Section 4 we analyse the algorithm cost. Section 5 is devoted to asymptotics as $n \rightarrow \infty$. The Bayesian approach is developed in Section 6. The particular case $f_k = 1$ is analysed in Section 7, and in Section 8 we present the conclusion of the paper. A complete version of this paper, including more graphs, is available at <http://www.ulb.ac.be/di/mcs/louchard/genodds.pdf>. The code of computations and algorithms is available upon request.

2. Fixed p

Let $p_k = pf_k, q_k := 1 - pf_k$, and

$$r_k := \frac{pf_k}{1 - pf_k},$$

that is, r_k is the (unknown) odds for $\{I_k = 1\}$. If $I_k = 1$, we say that a success occurs at time k . Furthermore, let $I_k(p)$ be the indicator of the event that a success occurs at time k when the success probability equals p . It is easy to see that

$$E\left(\sum_1^s I_k(p)\right) = p \sum_1^s f_k$$

and

$$\mathbb{V}\left[\sum_1^s I_k(p)\right] = p \sum_1^s f_k(1 - pf_k) = V_1(s), \quad \text{say,}$$

where \mathbb{V} denotes the variance. The odds algorithm gives

$$s^* = \begin{cases} \sup\left\{s : \sum_s^n r_k \geq 1\right\} & \text{if } \left\{s : \sum_s^n r_k \geq 1\right\} \neq \emptyset, \\ 1 & \text{otherwise.} \end{cases}$$

We should write $s^*(n)$, but here and in the sequel we drop the n to simplify the notation when there is no ambiguity. Hence, s^* is the time index from which onwards it is optimal to stop on the first event $I_k = 1$, and the corresponding optimal win probability equals $R_{s^*}Q_{s^*}$ (see Algorithm 1.1). Here, of course, R_s and Q_s are functions of p and f_1, \dots, f_s . We think of the f_k as being fixed and write

$$s^* = \varphi(p)$$

and

$$\psi(s, p) = \prod_s^n q_l \sum_s^n r_l.$$

Hence, the optimal win probability for a given p is given by

$$\psi^*(p) = \psi(s^*, p). \tag{2.1}$$

2.1. Sequential estimation

We use

$$\hat{p}(s, p) = \frac{\sum_1^s I_k(p)}{\sum_1^s f_k}, \tag{2.2}$$

as an estimator of p , which we do for two reasons. First, $\hat{p}(s, p)$ is an unbiased estimator of p . Indeed,

$$E(\hat{p}(s, p)) = \frac{\sum_1^s E(I_k(p))}{\sum_1^s f_k} = p.$$

Similarly, it follows from the independence of the I_k that

$$\mathbb{V}[\hat{p}(s, p)] = \frac{V_1(s)}{(\sum_1^s f_k)^2}.$$

Secondly, this estimator is efficient for constant f_k , that is, it has the smallest possible variance, as one can readily show using the Fisher information and Rao–Cramer’s bound. We note, however, that (2.2) is in general not a maximum likelihood estimator of p , as one can easily check. This is why we also offer an alternative approach later on.

Let us consider the distribution for \hat{p} for index s and parameter p both fixed. We denote it by

$$\hat{P}(\rho | s, p) := P[\hat{p}(s, p) = \rho].$$

We can see that $\hat{P}(\rho | s, p)$ becomes the binomial distribution if the f_k are constant. In the general case, it can be numerically computed by extracting the coefficients from the generating function

$$G_s(z) := \prod_1^s (pf_i z + 1 - pf_i). \tag{2.3}$$

We obtain

$$P\left[\hat{p}(s, p) = \frac{\nu}{\sum_1^s f_k}\right] = [z^\nu]G_s(z),$$

where $[z^n]f(z)$ denotes the coefficient of z^n in the power expansion of $f(z)$.

2.2. The distribution of the number of successes

Let

$$\nu(s) := \sum_{k=1}^s I_k = \# \text{ of successes up to time } s.$$

We note that $\nu(s)$ follows no well-known distribution unless the f_k are constant. However, we can construct a tractable recurrence relation for the law of $\nu(s)$ from $G_s(z)$, as given by (2.3). Let $p_k = pf_k$ and $q_k = 1 - pf_k$. Then, from (2.3), limiting ourselves to $z \geq 0$,

$$G_s(z) = \prod_{k=1}^s (q_k + p_k z), \quad z \geq 0. \tag{2.4}$$

We note that we can confine our interest to the case in which $q_k > 0$ for all $k = 1, 2, \dots, s$, because, for each k with $q_k = 0$, we have $I_k = 1$ almost surely, so that we can reduce $\nu(s)$ to a new random variable $\tilde{\nu}(s)$ which counts only nontrivial successes after the last $q_k = 0$ (that is if $p_k = 1$). Hence, we can and do suppose that

$$G_s(0) > 0. \tag{2.5}$$

Now, from (2.4),

$$\ln G_s(z) = \sum_{k=1}^s \ln(q_k + p_k z),$$

and, by differentiating,

$$G'_s(z) = G_s(z) \sum_{k=1}^s \frac{p_k}{q_k + p_k z}, \quad z \geq 0. \tag{2.6}$$

Here we used the facts that $G_s(z) \geq G_s(0) > 0$ from (2.4) and $G_k(0) > 0$ from (2.5). Now let

$$R_s(z) = \sum_{k=1}^s \frac{p_k}{q_k + p_k z}. \tag{2.7}$$

Then (2.6) becomes

$$G'_s(z) = G_s(z)R_s(z), \quad z \geq 0. \tag{2.8}$$

We will use (2.8) to build a recurrence formula for arbitrarily high derivatives.

Let $G_s^{(m)}(z)$ denote the m th derivative of $G_s(z)$ with respect to z . Then we obtain, from (2.8) and the well-known formula for derivatives of higher order of products,

$$\begin{aligned} G_s^{(m)}(z) &= \frac{d^{m-1}}{dz^{m-1}} G_s^{(1)}(z) \\ &= \frac{d^{m-1}}{dz^{m-1}} (G_s(z)R_s(z)) \\ &= \sum_{k=0}^{m-1} \binom{m-1}{k} G_s^{(k)}(z)R_s^{(m-1-k)}(z). \end{aligned} \tag{2.9}$$

Note that the right-hand side of (2.9) uses only the derivatives up to order $(m - 1)$, so that we will have a convenient recurrence relation for the $G_s^{(m)}(z)$, $m = 1, 2, \dots$, as soon as we have a formula for $R_s^{(j)}(z)$.

Lemma 2.1. *We have*

$$R_s^{(j)}(z) = (-1)^j j! \sum_{k=1}^s \left(\frac{p_k}{q_k + p_k z} \right)^{j+1} \text{ for all } j \in \mathbb{N}.$$

Proof. The statement is true for $j = 0$, since

$$R_s^{(0)}(z) = R_s(z) = \sum_{k=1}^s \left(\frac{p_k}{q_k + p_k z} \right),$$

which coincides with definition (2.7). Suppose now that the statement is true for $j = 0, 1, \dots, l$, $l \in \mathbb{N}$. Then

$$R_s^{(l)}(z) = (-1)^l l! \sum_{k=1}^s \left(\frac{p_k}{q_k + p_k z} \right)^{l+1},$$

and, therefore,

$$\begin{aligned} R_s^{(l+1)}(z) &= (-1)^{l+1} l! \sum_{k=1}^s (l+1) \left(\frac{p_k}{q_k + p_k z} \right)^l \frac{-p_k^2}{(q_k + p_k z)^2} \\ &= (-1)^{l+1} (l+1)! \sum_{k=1}^s \left(\frac{p_k}{q_k + p_k z} \right)^{l+2}, \end{aligned}$$

which completes the proof.

It follows that

$$R_s^{(j)}(0) = (-1)^j j! \sum_{k=1}^s r_k^{j+1}, \tag{2.10}$$

where, as before,

$$r_k = \frac{p_k}{q_k} = \frac{p f_k}{1 - p f_k}.$$

It follows, from (2.9) and (2.10), that

$$\begin{aligned} P[v(s) = m] &= \frac{1}{m!} G_s^{(m)}(0) \\ &= \frac{1}{m!} \sum_{k=0}^{m-1} \binom{m-1}{k} G_s^{(k)}(0) (-1)^{m-1-k} (m-1-k)! \sum_{j=1}^s r_j^{m-k} \\ &= \frac{1}{m} \sum_{k=0}^{m-1} \frac{1}{k!} G_s^{(k)}(0) (-1)^{m-1-k} \sum_{j=1}^s r_j^{m-k}. \end{aligned} \tag{2.11}$$

This yields the recurrence to compute $\{P[v(s) = m]\}_{m=0,1,\dots,s}$, namely, from (2.11),

$$P[v(s) = m] = \frac{1}{m} \sum_{k=0}^{m-1} P[v(s) = k] (-1)^{m-1-k} \sum_{j=1}^s r_j^{m-k}$$

with initial condition

$$P[v(s) = 0] = q_1 q_2 \cdots q_s.$$

2.3. The future number of successes

The value of our recurrence relation is that it immediately lends itself to the distribution of

$$\mu(s) = v(n) - v(s),$$

that is, $\mu(s)$ counts the number of successes from time s onwards. Indeed, it suffices to replace the p_1, p_2, \dots by p_s, p_{s+1}, \dots and to replace the second sum $\sum_{j=1}^s r_j^{m-k}$ by $\sum_{j=s}^n r_j^{m-k}$, that is,

$$P[\mu(s) = m] = \frac{1}{m} \sum_{k=0}^{m-1} P[\mu(s) = k] (-1)^{m-1-k} \sum_{j=s}^n r_j^{m-k}, \quad m = 0, 1, \dots, n - s + 1,$$

with initial condition

$$P[\mu(s) = 0] = q_s q_{s+1} \cdots q_n.$$

2.4. A stopping rule based on the sequential maximum likelihood

Since at time s the history of all successes and failures is known up to time s , we can plug in the corresponding maximum likelihood estimators (at time s), denoted by $p_{s+1}^{ML}(s), p_{s+2}^{ML}(s), \dots, p_n^{ML}(s)$. The rule to stop on a success if

$$P[\mu(s) = 1 \mid I_s = 1, p_{s+1}^{ML}(s), p_{s+2}^{ML}(s), \dots, p_n^{ML}(s)] \geq P[\mu(s+1) = 1 \mid p_{s+1}^{ML}(s), p_{s+2}^{ML}(s), \dots, p_n^{ML}(s)] \tag{2.12}$$

is then optimal in the class of one-step-look-ahead stopping rules (see [8, Chapter 5]).

This is the alternative approach, which we announced before. However, the computation of the maximum likelihood estimators always needs the full history of the processes I_1, I_2, \dots and cannot be based on a simpler sufficient statistic of the history. Hence, this approach is computationally much more involved than our approach based on (2.2) and the odds algorithm. We do not study the maximum likelihood approach performance as such. The computations are much more complicated and the few easy examples we went through showed no convincing overall improvement which would justify a more detailed study of this approach.

Finally, even this more complicated approach would not always yield the overall optimal strategy because (2.12) does not imply that the stopping problem is a monotone case problem (see, e.g. the section on monotone rules in [7]). Hence, the overall optimal rule need not be in the class of one-step-look-ahead stopping rules.

We therefore focus our interest on the odds algorithm using estimator (2.2).

2.5. Qualitative assessment

Let us now discuss the intrinsic weakness of any approach based on sequential estimation.

If $\hat{p}(s, p)$ is small at the beginning (no events $\{I_k = 1\}$ at the beginning), the stopping threshold s is also small and we could consequently stop too early. It is true that we can only stop on a success, so that \hat{p} jumps up at such instances. This reduces the risk of under-estimating p and, hence, the q_k . However, it does not exclude it.

Similarly, if we wait some time before we compute and use $\hat{p}(s, p)$, and if p is small, we could stop too late. As an alternative, we may decide to use some fixed learning samples and to never stop on the first $s_d - 1$ values, that is, we start the algorithm at $s = s_d$. Here $s_d = 1$ corresponds to the classical algorithm with no delay. The question of an optimal delay s_d will be analysed later on.

The odds algorithm for the stopping threshold s leads to the equation

$$\varphi(\hat{p}(s, p)) \leq s.$$

The threshold computation procedure is given in Algorithm 2.1, below.

Algorithm 2.1. (*The odds algorithm with sequential estimation of odds.*)

- *Input.* Precompute the optimal delay s_d (if we use a delay).
- *Output:* an optimal stopping threshold s .
 $s := s_d$
 cont := true
while cont **do**
 $v := \sum_1^s I_k, \hat{p}(s) = v / \sum_1^s f_k$
 if $\sum_{s+1}^n r_k(\hat{p}(s)) < 1$ **then** cont := false
 else
 $s := s + 1$
 if $s = n$ **then** cont := false
 end if
 end if
end while
return s

2.6. Winning probability

Here s is a random variable with some distribution $\phi(s, p)$, say.

Fix p . For each time s , the possible values of the random variable $v := \sum_1^s I_k$ satisfying

$$\varphi\left(\frac{v}{\sum_1^s f_k}\right) \leq s$$

are constrained to stay in an interval denoted by $[0, \gamma[s]]$. In order to stop at any case not later than n , we just set $\gamma[n] = n$. In the case of delaying we just set $\gamma[s] = -1, s = 1, \dots, s_d - 1$. Here v is represented by a Markov chain. In the following we drop the p parameter to ease the notation. Let

$$\Pi[s, \mu] := P[v = \mu, \text{ no stopping threshold before } s].$$

Then,

$$\Pi[1, 1] = pf_1, \quad \Pi[1, 0] = 1 - pf_1, \quad \phi[1, p] = \sum_{\mu=0}^{\gamma[1]} \Pi[1, \mu],$$

and, for $s \geq 2$,

$$\Pi[s, \mu] = \Pi[s - 1, \mu - 1]pf_s \mathbf{1}_{\{\mu \neq 0 \wedge \mu - 1 > \gamma[s-1]\}} + \Pi[s - 1, \mu](1 - pf_s) \mathbf{1}_{\{\mu > \gamma[s-1]\}}.$$

The stopping threshold probability distribution is now given by

$$\phi(s, p) = \sum_{\mu=0}^{\gamma[s]} \Pi[s, \mu].$$

Finally,

$$P[\text{win}] = P[\text{algorithm stops on the last } 1 \mid p] = \sum_{s=1}^n \phi(s, p)\psi(s, p) = \Theta(p), \quad \text{say.}$$

2.7. Choice of f_k and n

In the examples given in this paper we use two different choices for the sequence (f_k) . One choice is $f_k = 1$ for all k . This is a natural choice for the case when all the I_k are independently and identically distributed Bernoulli random variables. We could also have used $f_k = C$ for some fixed constant $0 < C \leq 1$. Our second and most frequent choice is $f_k = 1/k$. One reason is that we want to cover the case when all the odds are different. Besides this, there is nothing really special about this choice except that it solves a new version of a well-known best-choice problem, that is, the secretary problem with *unknown* availability probability. Indeed, suppose that in a sequence of candidates, all arrival orders of different ranks are equally probable, and that the k th candidate is available, independently of his rank, with probability p . Then this candidate is best so far *and* available with probability p/k . See also [1].

We usually use the sample size $n = 15$, but again there is nothing special about this choice and most graphs would look similar for n not too small. Clearly, small n , say $n \leq 6$, lead to unreliable odds estimates and, hence, to bad performance.

2.7.1. *Examples.* In Figure 1 we plot $\Theta(p)$ as a function of p for $s_d = 1, \dots, 5$. We have chosen $n = 15$ and $f_k = 1/k$ (these parameters will always be used in the sequel). The circles denote $\psi^*(p)$ and the horizontal line represents $1/e$. The relevance of a comparison with $1/e$ will be explained below.

Note that $\Theta(p)$ possesses a local maximum and a local minimum for some values of s_d . This can be explained as follows. When p is small, the chance of having 1s is small, and, hence, the total win probability is small for any strategy. Since the estimated odds are very likely to be small as well for small p , the risk of a wrong decision by the odds strategy is also small simply because stopping on the very first 1 (if any) is best. But, for growing p , this risk increases in




FIGURE 1: Plots of $\Theta(p)$ as a function of p for $s_d = 1, \dots, 5$ (in order from top to bottom at $p = 0.1$), $n = 15$, and $f_k = 1/k$, $k = 1, \dots, n$. The circles denote $\psi^*(p)$ and the horizontal line denotes $1/e$.




FIGURE 2: Plots of $\Theta(p)$ as a function of p for the choice $f_k = 1/k$ and $s_d = 1$, where the solid curve displays the exact solution and the dotted curve displays the results of the simulation.

the middle range of p so that the total win probability goes somewhat down before getting the full benefit of large success probabilities. The difference between the local maximum and the local minimum is of course also dependent of the choice of the f_k . A more detailed approach is given in Section 5.

There is a good reason why the comparison of the performance of this algorithm with the value $1/e$ is the most adequate one. Indeed, it was shown that if the odds are known and if their sum is at least 1, then $1/e$ is the exact lower bound for the win probability over all such sequences p_1, p_2, \dots, p_n (see [3]). But, moreover, if n becomes large and if $\sum_{k=1}^n p_k^2 / \sum_{k=1}^n p_k \rightarrow 0$ as $n \rightarrow \infty$, then the win probability converges to $1/e$ (see [2]). Hence, in particular, if the sum of all the odds is at least 1 then it suffices that $p_k \rightarrow 0$ as $k \rightarrow \infty$.

We finally observe that, for any value of p , there is an optimal value of s_d . This will be useful later on.

We have also made a simulation of $\phi(s, p)$ for 1000 trials, and computed $\Theta(p)$ for $s_d = 1$ with the simulated matrix and $\psi(s, p)$. This is given in Figure 2, showing that the fit is good.

3. Unknown p according to a density $P(p)$

We now suppose that the unknown parameter p follows a density $P(p)$, which is unknown to the decision maker. Let $\Theta(p)$ denote, as before, the conditional probability of winning for a given p . The absolute win probability using our algorithm is then given by

$$P_w := P[\text{win}] = \int_0^1 P(p)\Theta(p) dp.$$

There is no statistical inference on p other than using the sequential estimator (2.2). The only focus is the impact of delaying as a function of the distribution $P(p)$. Statistical inference based on a (known) prior distribution of p will be used in the Bayesian approach in Section 6.

3.1. Examples

(i) As a first example, we let $P(p)$ be given by a parabola on $[0, 1]$, with maxima occurring respectively at $p_m \in [\frac{1}{16}, \frac{1}{8}, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, \frac{7}{8}]$. The parabola starts at the origin for $p_m = \frac{1}{16}, \frac{1}{2}, \frac{3}{4}, \frac{7}{8}$,




FIGURE 3: A plot of P_w as a function of s_d for $p_m = \frac{7}{8}$.

and ends at 0 for $p = 1$ and $p_m = \frac{1}{8}, \frac{1}{4}$. Again, we use $f_k = 1/k$. For $p_m = \frac{1}{8}$, P_w as a function of the delay parameter s_d is strictly decreasing. Nothing is gained by delaying stopping. The situation is the same for $p_m = \frac{1}{16}, \frac{1}{4}$. However, for $p_m \in [\frac{1}{2}, \frac{3}{4}, \frac{7}{8}]$, we see that it is better to ignore the first event. We see this, in Figure 3, for $p_m = \frac{7}{8}$, where P_w is plotted as function of the delay parameter s_d .

The optimal s_d values, for our six parabolae, are given by $[1, 1, 1, 2, 2, 2]$. We see that these optimal values are rather robust: minimal information about the shape of $P(p)$ is enough to choose s_d .

(ii) As an example of large s_d , we have computed P_w with a linear $P(p) = 2p$. This leads to an optimal delay $s_d = 4$.

In the case of sequential updating, we will denote by $P_w(p_m)$ the success probability without delay and by $P_{w, \text{opt}}(p_m)$ the success probability with optimal delay for our six parabolae distributions.

If we know p beforehand, we must use $\psi^*(p)$. Hence,

$$P_w^* = \int_0^1 P(p)\psi^*(p) dp.$$

4. Algorithm cost

The computational cost of Algorithm 2.1 depends essentially on the computation of $\hat{p}(s)$ and on the instruction: if $\sum_{s+1}^n r_k(\hat{p}(s)) < 1$. Assuming, for simplicity, that each numerical operation costs 1 unit, we have, at time s , a cost of

$$C(s, p) = \sum_1^s (n - v + 1) = \left(n + \frac{1}{2}\right)s - \frac{s^2}{2}, \tag{4.1}$$

and $C'_s(s, p) = n + \frac{1}{2} - s \geq 0$. The mean cost $M(p)$ is given by

$$M(p) = \sum_1^n \phi(s, p) \left(\left(n + \frac{1}{2}\right)s - \frac{s^2}{2} \right) = \left(n + \frac{1}{2}\right)\bar{s} - \frac{\bar{s}^2}{2},$$

where

$$\bar{s}^i := \sum_1^n \phi(s, p) s^i.$$

Similarly, the second moment of the cost becomes

$$M^{(2)}(p) = \sum_1^n \phi(s, p) \left(\left(n + \frac{1}{2} \right) s - \frac{s^2}{2} \right)^2 = \frac{\bar{s}^4}{4} - \left(n + \frac{1}{2} \right) \bar{s}^3 + \left(n + \frac{1}{2} \right)^2 \bar{s}^2,$$

and the variance is given by

$$\mathbb{V}[p] = M^{(2)}(p) - M(p)^2.$$

The cost distribution is itself computed as follows. For fixed p , we have

$$\begin{aligned} C(1, p) &= n, \\ C(2, p) &= 2n - 1, \\ C(3, p) &= 3n - 3, \\ &\vdots \\ C(n, p) &= \frac{n(n+1)}{2}. \end{aligned}$$

For each possible value of the function C , the corresponding value of s is given, from (4.1), by

$$s = \frac{(2n + 1) - \sqrt{(2n + 1)^2 - 8C}}{2}.$$

This allows, with $\phi(s, p)$, the computation of the cost distribution $H(p)$. For instance, for $H(\frac{1}{2})$ with our usual parameters, only three values of s lead to nonnull values of $\phi(s, p)$, and only three values of the cost determine the distribution.

4.1. Asymptotic behaviour of the cost as $n \rightarrow \infty$

To study the asymptotic behaviour of the cost as $n \rightarrow \infty$, we must distinguish between two cases.

Case (i). If $\sum_1^\infty f_k$ converges and $\sum_1^\infty f_k / (1 - f_k) > 1$ (otherwise we always stop at $s = 1$), we have, for each p , a maximum $s^*(p)$ such that $\sum_{s^*}^\infty r_k(p) \geq 1$. Here $\varphi(p)$ is asymptotically independent of n and $\phi(s, p)$ also becomes independent of n , and we have a cost given by (4.1), which is linear in n . Also, setting

$$\hat{s} := \sup \left\{ j : \sum_j^\infty \frac{f_k}{1 - f_k} \geq 1 \right\},$$

we have $s^*(p) \leq \hat{s}$ and

$$C(s, p) \leq \left(n + \frac{1}{2} \right) \hat{s} - \frac{\hat{s}^2}{2}.$$

Case (ii). If $\sum_1^\infty f_k$ diverges, $\varphi(p)$ is close to n , $\phi(s, p)$ gives a maximum weight in the neighbourhood of n , and the cost is now of the order of n^2 .

For instance, if $f_k = 1$, the odds algorithm gives $s \sim n - q/p$ and

$$C(s, p) \sim \frac{n^2}{2},$$

if $f_k = 1/k$, we have $s \sim ne^{-1/p}$ and

$$C(s, p) \sim \left(e^{-1/p} - \frac{e^{-2/p}}{2} \right) n^2.$$

5. The asymptotic behaviour of $\psi^*(p, n)$ for $f_k = 1/k$

In this section we explicitly introduce the parameter n into $\psi(s, p)$ and $\psi^*(p)$: we will use $\psi(s, p, n)$ and $\psi^*(p, n)$. All asymptotics in this section are limit relations as n tends to ∞ .

We will use a continuous variable approach: our tool will be the Euler–Maclaurin formula with one error term (to simplify). In our previous expressions we replace sums by integrals, using

$$\sum_a^b f(k) \sim \int_a^b f(x) dx + \frac{1}{2}f(a) + \frac{1}{2}f(b).$$

Thus, if we equate the approximation with the odds-sum threshold 1, we obtain

$$\int_{s^*(p,n)}^n \frac{p/v}{1-p/v} dv + \frac{p/s^*(p,n)}{2(1-p/s^*(p,n))} + \frac{p/n}{2(1-p/n)} = 1.$$

Solving for $s^*(p, n)$, we obtain accordingly

$$s^*(p, n) = ne^{-1/p} - pe^{-1/p} + p + \frac{1}{2} + \frac{1}{2}e^{-1/p} + \mathcal{O}\left(\frac{1}{n}\right). \tag{5.1}$$

Also,

$$s^{*'}(p, n) = \frac{ne^{-1/p} - pe^{-1/p} - p^2e^{-1/p} + p^2 + 1/2e^{-1/p}}{p^2} + \mathcal{O}\left(\frac{1}{n}\right). \tag{5.2}$$

A critical value $p^*(n)$ is given by

$$s^*(p^*(n), n) = 1,$$

that is,

$$\int_1^n \frac{p^*(n)/v}{1-p^*(n)/v} dv + \frac{p^*(n)}{2(1-p^*(n))} + \frac{p^*(n)/n}{2(1-p^*(n)/n)} = 1,$$

which leads to

$$p^*(n) = \frac{1}{\ln(n)} - \frac{1}{2\ln(n)^2} - \frac{5}{4\ln(n)^3} + \frac{9}{8\ln(n)^4} + \mathcal{O}\left(\frac{1}{\ln(n)^5}\right).$$

5.1. Uniqueness of the maximum of $\psi^*(p, n)$

In the upper range for p , $p^*(n) \leq p \leq 1$, we have

$$\begin{aligned} \psi(s, p, n) = \exp & \left[\int_s^n \ln\left(1 - \frac{p}{u}\right) du + \frac{1}{2} \ln\left(1 - \frac{p}{s}\right) + \frac{1}{2} \ln\left(1 - \frac{p}{n}\right) \right] \\ & \times \left(\int_s^n \frac{p/v}{1-p/v} dv + \frac{p/s}{2(1-p/s)} + \frac{p/n}{2(1-p/n)} \right) \end{aligned}$$

and

$$\begin{aligned} \psi^*(p, n) &= \psi(s^*(p, n), p, n) \\ &= \exp\left[\int_{s^*(p, n)}^n \ln\left(1 - \frac{p}{u}\right) du + \frac{1}{2} \ln\left(1 - \frac{p}{s^*(p, n)}\right) + \frac{1}{2} \ln\left(1 - \frac{p}{n}\right)\right] \end{aligned} \tag{5.3}$$

with

$$\begin{aligned} \int_s^n \ln\left(1 - \frac{p}{u}\right) du &= -p \ln(n) + \ln\left(\frac{n-p}{n}\right) - p \ln\left(\frac{n-p}{n}\right) + p \ln(s) \\ &\quad - s \ln\left(\frac{s-p}{s}\right) + p \ln\left(\frac{s-p}{s}\right). \end{aligned}$$

To show that the limiting extremum for $\psi^*(p, n)$ is unique in the upper range, we differentiate $\psi^*(p, n)$ with respect to p , and set it to 0. This leads to the equation

$$\begin{aligned} 0 &= -\ln\left(1 - \frac{p}{s^*(p, n)}\right) s^{*'}(p, n) - \frac{1}{p} \left(1 - \frac{p/s^*(p, n)}{2(1 - p/s^*(p, n))} - \frac{p/n}{2(1 - p/n)}\right) \\ &\quad + \frac{-1/s^*(p, n) + ps^{*'}(p, n)/s^{*2}(p, n)}{2(1 - p/s^*(p, n))} - \frac{1/n}{2(1 - p/n)}. \end{aligned}$$

Using (5.1) and (5.2), we obtain, in a first approximation,

$$0 = \frac{1}{n} \left(-\frac{1}{2} e^{1/p} (2pe^{-1/p} - 2p + 1)\right) + \mathcal{O}\left(\frac{1}{n^2}\right).$$

This yields

$$p_{\min}^*(n) = \frac{1}{\text{Lambert}(-2e^{-2} + 2)} + \mathcal{O}\left(\frac{1}{n}\right),$$

where

$$\text{Lambert}(x)e^{\text{Lambert}(x)} = x.$$

This is related to the tree function $T(x)$: $T(x) = -\text{Lambert}(-x)$. We have

$$\frac{1}{\text{Lambert}(-2e^{-2} + 2)} = 0.627\,500\,487\,4\dots$$

Also, an asymptotic series for $\psi^*(p, n)$ is given from (5.3) by

$$\psi^*(p, n) = e^{-1} - \frac{1}{n} \left[\frac{1}{2} e^{-1+1/p} p^2 (-1 + e^{-1/p})\right] + \mathcal{O}\left(\frac{1}{n^2}\right),$$

and we can show that the extremum is actually a minimum with

$$\psi^*(p_{\min}^*(n), n) = \frac{0.284\,028\,432\,2\dots}{n} + \mathcal{O}\left(\frac{1}{n^2}\right).$$

At the critical point, we have

$$\psi^*(p^*(n), n) = \exp\left[\int_1^n \ln\left(1 - \frac{p^*(n)}{u}\right) du + \frac{1}{2} \ln(1 - p^*(n)) + \frac{1}{2} \ln\left(1 - \frac{p^*(n)}{n}\right)\right],$$

which is

$$\psi^*(p^*(n), n) = e^{-1} + \frac{3e^{-1}}{4 \ln(n)^2} + \mathcal{O}\left(\frac{1}{\ln(n)^3}\right). \tag{5.4}$$

Let us now turn to the lower range for p , $0 \leq p \leq p^*(n)$. Recall that $s^*(p, n) = 1$ in this range. We obtain

$$\begin{aligned} \psi^*(p, n) &= \exp\left[\int_1^n \ln\left(1 - \frac{p}{u}\right) du + \frac{1}{2} \ln(1 - p) + \frac{1}{2} \ln\left(1 - \frac{p}{n}\right)\right] \\ &\times \left(\int_1^n \frac{p/v}{1 - p/v} dv + \frac{p}{2(1 - p)} + \frac{p/n}{2(1 - p/n)}\right). \end{aligned} \tag{5.5}$$

Now use

$$\int_1^n \ln\left(1 - \frac{p}{u}\right) du = -p \ln(n) + \ln\left(\frac{n - p}{n}\right) - p \ln\left(\frac{n - p}{n}\right) - \ln(1 - p) + p \ln(1 - p)$$

and

$$\int_1^n \frac{p/v}{1 - p/v} dv = p \ln(n - p) - p \ln(1 - p).$$

This leads to

$$\begin{aligned} \psi^*(p, n) &= \exp\left[-p + p \ln(1 - p) - \frac{1}{2 \ln(1 - p)}\right] \\ &\times \frac{p \ln(n) - p \ln(1 - p) + p/(2(1 - p))}{n^p} \\ &+ \mathcal{O}\left(\frac{\ln(n)}{n^{1+p}}\right), \end{aligned} \tag{5.6}$$

so that

$$\psi^*(p^*(n), n) = e^{-1} + \frac{3e^{-1}}{4 \ln(n)^2} + \mathcal{O}\left(\frac{1}{\ln(n)^3}\right).$$

Note that this agrees with (5.4).

Next, we want to show that, asymptotically, $\psi^*(p, n)$ is increasing in the lower range. So we differentiate $\psi^*(p, n)$ as given by (5.6). Setting $p = \alpha / \ln(n)$, this gives

$$\psi^{*'}(p, n)n^p \sim \ln(n)(1 - \alpha) \geq 0, \quad \alpha \in [0, 1].$$

This shows that $\psi^*(p, n)$ is increasing for sufficiently large n .

Finally, we have, asymptotically, a unique maximum of $\psi^*(p, n)$ for

$$p = p^*(n) = p_{\max}^*(n), \quad \text{say.}$$

Using (5.3), (5.5), and $n = 15$, in Figure 4 we plot $\psi^*(p, n)$ on the whole range $p \in [0, 1]$. This function is continuous, but its derivative is not. We also compare it with the discrete expression for $\psi^*(p, n)$ given by (2.1).

The fit is reasonable, given that n is small and that we used Euler–Maclaurin with only one error term and a continuous $s^*(p, n)$ instead of the discrete one given in (1.1). Note that the plot in Figure 4 has a similar behaviour to that in Figure 1 for $\Theta(p)$ in the sequential updating approach. In Figure 1, the difference between the maximum and minimum is even more pronounced.




FIGURE 4: A plot of $\psi^*(p, n)$ versus $\psi^*(p, n)$ for $n = 15$ and $f_k = 1/k$, where the solid curve displays the continuous version and the dotted curve displays the discrete version.

For large n , the difference between exact (continuous) expressions and first-order asymptotics (neglecting $\mathcal{O}(1/n)$ errors) becomes negligible. We confine our interest to the difference between the discrete and the continuous approach.

For $n = 100$, the fit is excellent. Let us note that the optimal s_d values, for our six parabolae for $n = 100$ are given by [1, 3, 4, 6, 13, 13]. Again, these optimal values are rather robust.

We note that it would be hard to prove existence and unicity of min and max in the discrete case as well as for $\Theta(p)$.

6. Bayesian approach

6.1. The theory

We use in this approach the work of Van Lokeren [18] (mémoire de DEA under the supervision of F. T. Bruss, unpublished). The problem is as before, that is, maximising the probability of stopping on the last success. Allowing the different success parameters p_1, \dots, p_n to vary independently of each other leads to an ill-posed problem. Therefore, we make the following assumptions.

Let p be a random variable taking values in $[0, 1]$, and let $\Psi: [0, 1] \times \mathbb{N} \rightarrow [0, 1]$ be a deterministic (known) function. We assume that the success parameter p_k is given by

$$p_k = \Psi(p, k).$$

Furthermore, we suppose the following.

- (I1) The conditional law of I_k , given $p = x$, is a Bernoulli law with known success parameter $\Psi(x, k)$.
- (I2) The random variables I_1, I_2, \dots, I_n are conditionally independent, given $p = x$.

6.1.1. *The general solution.* Suppose that the prior distribution of p is given by a density $P(x)$. We first determine the form of the posterior distribution. Suppose that we have observed I_1, I_2, \dots, I_k yielding the outcomes a_1, a_2, \dots, a_k with $a_i \in \{0, 1\}$ for all i . The posterior

distribution of p at time k is defined by

$$P_k(x \mid a_1, \dots, a_k) := \frac{d}{dx} P[p \leq x \mid I_1 = a_1, \dots, I_k = a_k] \quad \text{for all } x \in [0, 1].$$

Now Bayes' rule gives

$$P_k(x \mid a_1, \dots, a_k) = \frac{P[I_1 = a_1, \dots, I_k = a_k \mid p = x]P(x)}{P[I_1 = a_1, \dots, I_k = a_k]}. \tag{6.1}$$

By conditional independence of the I_k given p , we have

$$\begin{aligned} P[I_1 = a_1, \dots, I_k = a_k \mid p = x] &= \prod_{i=1}^k P[I_i = a_i \mid p = x] \\ &= \prod_{i=1}^k (\Psi(x, i))^{a_i} (1 - \Psi(x, i))^{1-a_i}. \end{aligned}$$

Applying the law of total probability to the denominator of (6.1) yields

$$P[I_1 = a_1, \dots, I_k = a_k] = \int_0^1 \prod_{i=1}^k (\Psi(x, i))^{a_i} (1 - \Psi(x, i))^{1-a_i} P(x) dx,$$

so that, from (6.1),

$$P_k(x \mid a_1, \dots, a_k) = \frac{(\prod_{i=1}^k (\Psi(x, i))^{a_i} (1 - \Psi(x, i))^{1-a_i}) P(x)}{\int_0^1 (\prod_{i=1}^k (\Psi(x, i))^{a_i} (1 - \Psi(x, i))^{1-a_i}) P(x) dx}. \tag{6.2}$$

Similarly as in the case of the maximum likelihood approach (see (2.12)), the posterior distribution now depends on the full history up to time k .

We denote by \mathbf{a}_k the history up to time k , that is, $\mathbf{a}_k := (a_1, \dots, a_k)$, where a_i is the outcome of observation I_i . Write $(\mathbf{a}_k, 1)$ for $(a_1, \dots, a_k, 1)$. Formally, we set $\mathbf{a}_0 = 0$ to indicate that no observations have been taken. We also introduce the function

$$\tilde{C}(k, \mathbf{a}_k) := \int_0^1 \left(\prod_{i=1}^k (\Psi(x, i))^{a_i} (1 - \Psi(x, i))^{1-a_i} \right) P(x) dx. \tag{6.3}$$

Let $V_{n-k}(\mathbf{a}_k)$ be the optimal value obtained by applying the optimal strategy from time k onwards—that is, if there are still $n - k$ observations to come—given that no stop has yet been made strictly before time k and that the history up to time k is given by \mathbf{a}_k . We distinguish two cases: $k = n - 1$ and $k < n - 1$.

Case (i): $k = n - 1$. We have

$$\begin{aligned} V_1(\mathbf{a}_{n-1}) &= P[I_n = 1 \mid \mathbf{a}_{n-1}] \\ &= \int_0^1 P[I_n = 1 \mid p = x] P_{n-1}(x \mid \mathbf{a}_{n-1}) dx \\ &= \int_0^1 \Psi(x, n) P_{n-1}(x \mid \mathbf{a}_{n-1}) dx. \end{aligned}$$

By the explicit form (6.2) of the posterior distribution for $k = n - 1$ and the definition of $\tilde{C}(n - 1, \mathbf{a}_{n-1})$, we obtain

$$\begin{aligned} V_1(\mathbf{a}_{n-1}) &= \int_0^1 \Psi(x, n) \frac{(\prod_{i=1}^{n-1} (\Psi(x, i))^{a_i} (1 - \Psi(x, i))^{1-a_i}) P(x)}{\tilde{C}(n - 1, \mathbf{a}_{n-1})} dx \\ &= \frac{\tilde{C}(n, (\mathbf{a}_{n-1}, 1))}{\tilde{C}(n - 1, \mathbf{a}_{n-1})}. \end{aligned}$$

Case (ii): $k < n - 1$. We obtain, from the optimality principle, the equation

$$\begin{aligned} V_{n-k}(\mathbf{a}_k) &= P[I_{k+1} = 0 \mid \mathbf{a}_k] V_{n-k-1}((\mathbf{a}_k, 0)) \\ &\quad + P[I_{k+1} = 1 \mid \mathbf{a}_k] \\ &\quad \times \max\{P[I_{k+2} = \dots = I_n = 0 \mid (\mathbf{a}_k, 1)], V_{n-k-1}((\mathbf{a}_k, 1))\}. \end{aligned}$$

The first term yields

$$\begin{aligned} P[I_{k+1} = 0 \mid \mathbf{a}_k] &= \int_0^1 P[I_{k+1} = 0 \mid p = x] P_k(x \mid \mathbf{a}_k) dx \\ &= \int_0^1 (1 - \Psi(x, k + 1)) P_k(x \mid \mathbf{a}_k) dx \\ &= \int_0^1 (1 - \Psi(x, k + 1)) \frac{(\prod_{i=1}^k (\Psi(x, i))^{a_i} (1 - \Psi(x, i))^{1-a_i}) P(x)}{\tilde{C}(k, \mathbf{a}_k)} dx \\ &= \frac{\tilde{C}(k + 1, (\mathbf{a}_k, 0))}{\tilde{C}(k, \mathbf{a}_k)}, \end{aligned}$$

and the second term is given by

$$P[I_{k+1} = 1 \mid \mathbf{a}_k] = \frac{\tilde{C}(k + 1, (\mathbf{a}_k, 1))}{\tilde{C}(k, \mathbf{a}_k)}.$$

Finally, by conditioning on the parameter p ,

$$\begin{aligned} &P[I_{k+2} = \dots = I_n = 0 \mid (\mathbf{a}_k, 1)] \\ &= \int_0^1 P[I_{k+2} = \dots = I_n = 0 \mid p = x] P_{k+1}(x \mid (\mathbf{a}_k, 1)) dx \\ &= \int_0^1 \prod_{i=k+2}^n (1 - \Psi(x, i)) P_{k+1}(x \mid (\mathbf{a}_k, 1)) dx \\ &= \int_0^1 \prod_{i=k+2}^n (1 - \Psi(x, i)) \\ &\quad \times \frac{(\prod_{i=1}^k (\Psi(x, i))^{a_i} (1 - \Psi(x, i))^{1-a_i}) \Psi(x, k + 1) P(x)}{\tilde{C}(k + 1, (\mathbf{a}_k, 1))} dx. \end{aligned}$$

If we write $\mathbf{a}_n^{(k)} := (a_1, \dots, a_k, 0, 0, \dots, 0)$ and $\mathbf{a}_n^{(k+)} := (a_1, \dots, a_k, 1, 0, \dots, 0)$, we can rewrite the last probability as

$$P[I_{k+2} = \dots = I_n = 0 \mid (\mathbf{a}_k, 1)] = \frac{\tilde{C}(n, \mathbf{a}_n^{(k+)})}{\tilde{C}(k + 1, (\mathbf{a}_k, 1))}.$$

The recurrence relation for $V_{n-k}(\mathbf{a}_k)$ then becomes

$$V_{n-k}(\mathbf{a}_k) = \frac{\tilde{C}(k+1, (\mathbf{a}_k, 0))}{\tilde{C}(k, \mathbf{a}_k)} V_{n-k-1}((\mathbf{a}_k, 0)) + \frac{\tilde{C}(k+1, (\mathbf{a}_k, 1))}{\tilde{C}(k, \mathbf{a}_k)} \max \left\{ \frac{\tilde{C}(n, \mathbf{a}_n^{(k+)})}{\tilde{C}(k+1, (\mathbf{a}_k, 1))}, V_{n-k-1}(\mathbf{a}_k, 1) \right\} \tag{6.4}$$

with boundary condition

$$V_1(\mathbf{a}_{n-1}) = \frac{\tilde{C}(n, (\mathbf{a}_{n-1}, 1))}{\tilde{C}(n-1, \mathbf{a}_{n-1})}. \tag{6.5}$$

We have thus proven the following theorem.

Theorem 6.1. *Let the prior distribution of p and the function Ψ be given. An optimal strategy is*

- (i) *to stop at the first I_k for which $I_k = 1$ and $\tilde{C}(n, \mathbf{a}_n^{(k)})/\tilde{C}(k, \mathbf{a}_k) \geq V_{n-k}(\mathbf{a}_k)$, where the right-hand side satisfies the recurrence relation (6.4) and its boundary condition (6.5);*
- (ii) *to stop at I_n if the above conditions are not fulfilled for $1 \leq k \leq n - 1$.*

The optimal value is given by $V_n(\mathbf{a}_0)$.

This represents an *exact solution* to the stopping problem.

If a discrete distribution of p is chosen, it suffices to replace the integral in (6.3) by a sum running over the possible values of p . All other expressions remain unchanged.

We are now ready to give the algorithm.

6.2. The algorithm for the Bayesian approach

The algorithm deals with a vector $a[1, \dots, n]$ of bits. We convert this vector into an integer $l = \sum_1^n a[i]2^{i-1}$ with the procedure $l := \text{conv}_1(a)$. Similarly, for any l , we compute the corresponding vector a with a procedure $a := \text{conv}_2(l)$. Then, according to [18], we compute the two n -by- n matrices $C[0..n, 0..2^n - 1]$ and $V[1..n, 0..2^n - 1]$ with the following formulae:

```

C[0, 0] := 1;
for i to 2^n - 1 do C[0, i] := 0 od;
for k to n do
for l from 0 to 2^n - 1 do
a := conv2(l);
    
```

where ‘od’ is the code for the end of the programme loop. Then, in general,

$$C[k, l] := \int_0^1 \prod_1^k (x f_i)^{a[i]} [1 - (x f_i)]^{1-a[i]} P(x) dx; \text{ od; od;}$$

$$\text{for } l \text{ from } 0 \text{ to } 2^n - 1 \text{ do } V[1, l] := C[n, l + 2^{n-1}] / C[n - 1, l]; \text{ od;}$$

$$\text{for } k \text{ from } n - 2 \text{ by } -1 \text{ to } 0 \text{ do}$$

$$\text{for } l \text{ from } 0 \text{ to } 2^k - 1 \text{ do}$$

$$A := \frac{C[n, l + 2^k]}{C[k + 1, l + 2^k]}; \quad B := V[n - k - 1, l + 2^k]; \quad T := \max(A, B);$$

$$V[n - k, l] := \frac{C[k + 1, l]}{C[k, l]V[n - k - 1, l]} + \frac{C[k + 1, l + 2^k]}{C[k, l]T}; \text{ od; od.}$$




FIGURE 5: Plots of $P_w^*(p_m)$, $P_{w_{opt}}(p_m)$, $P_w^B(p_m)$, and $P_w(p_m)$ (in order from top to bottom at $p = 0.8$) for $p_m \in [\frac{1}{16}, \frac{1}{8}, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, \frac{7}{8}]$, $n = 15$, and $f_k = 1/k$. The horizontal line at height $1/e$ serves as a comparison.

Finally, the Bayesian optimal value is given by

$$P_w^B = V[n, 0].$$

The practical procedure is given in Algorithm 6.1, below.

Algorithm 6.1. (The optimal strategy.)

- *Input.* Precompute C and V .
- *Output:* an optimal strategy.
 - (i) Set $a[k] := I_k$, $k = 1, \dots, n$.
 - (ii) Stop at the first I_k for which $I_k = 1$ and, with $l_k := \text{conv}_1(a[1..k])$,

$$\frac{C[n, l_k]}{C[k, l_k]} \geq V[n - k, l_k].$$

- (iii) Stop at I_n if the above conditions are not fulfilled for any $1 \leq k \leq n - 1$.

As an illustration, we have computed the success probability given by the Bayesian approach, $P_w^B(p_m)$ for our five parabolae distributions.

Figure 5 gives ($n = 15$) $P_w^*(p_m)$, $P_w^B(p_m)$, $P_w(p_m)$, and $P_{w_{opt}}(p_m)$. Here P_w^* naturally gives the best result. The other ones are comparable, with a slight advantage for P_w^B and $P_{w_{opt}}$, but P_w is rather close. Note that, for some values of p_m , the value P_w^B is better than $P_{w_{opt}}$, but the opposite is true for other values of p_m .

7. The case $f_k = 1$

The case $f_k = 1$ for all $k = 1, 2, \dots, n$ is the simplest interesting special case. In Figure 6 we plot $\Theta(p)$ as a function of p for $s_d = 1, \dots, 14$ and $n = 15$. The circles denote $\psi^*(p)$




FIGURE 6: Plots of $\Theta(p)$ as a function of p for $s_d = 1, \dots, 14$ (in order from bottom to top at $p = 0.3$), $n = 15$, and $f_k = 1, k = 1, \dots, n$. The dotted curve displays $\psi^*(p)$.




FIGURE 7: Plots of $P_w^*(p_m)$, $P_{w_{opt}}(p_m)$, $P_w^B(p_m)$, and $P_w(p_m)$ (in order from top to bottom at $p = 0.8$) for $p_m \in [\frac{1}{16}, \frac{1}{8}, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, \frac{7}{8}]$, $n = 15$, and $f_k = 1$.

and the horizontal line denotes $1/e$. If we compare this graph with Figure 1, we see that the maximum and minimum are more pronounced (at least for small values of s_d).

The delay analysis shows that, for $p_m = \frac{1}{16}$, no delay is necessary, but, for $p_m = \frac{1}{8}$, we already have an optimal $s_d = 10$. The optimal s_d values for our six parabolae are given by [1, 10, 10, 11, 12, 12]. Again, these optimal values are rather robust: minimal information about the shape of $P(p)$ is enough to choose s_d .

Figure 7 displays $P_w^*(p_m)$, $P_w^B(p_m)$, $P_w(p_m)$, and $P_{w_{opt}}(p_m)$. Again, P_w^* gives the best result, but its advantage is less pronounced. Here P_w^B and $P_{w_{opt}}$ are rather close to each other, and P_w is definitively bad.

Of course, we could compute an equivalent continuous analysis of $\psi^*(p, n)$, as we did in Section 5, but we will not pursue this matter in this paper.

8. Conclusion

The solution of the problem of maximising the probability of stopping on a last success in a sequence of independent indicators has many real-world applications, ranging from best-choice problems (secretary problems) over buying-selling strategies to applications in sequential search, maintenance problems, and clinical trials. If the odds are known in advance, the odds algorithm provides this solution in a straightforward way, and this algorithm is itself optimal. If the odds are unknown and must be estimated from preceding observations, then the optimal rule is not obvious and can be made explicit in special cases only. The objective of this work was to examine the question of whether there are good approximations for the optimal rule. We have proposed an algorithm which is based on the odds algorithm and on a simple unbiased sequential estimator of the success probabilities $p_k = P[I_k = 1]$. Although we have no precise estimates for how much this algorithm misses optimality, we have established several important facts.

Firstly, it is asymptotically optimal, because, as $n \rightarrow \infty$, the sequential estimators of the odds will converge, in our model, to the true odds, and we know that for the true odds the odds algorithm gives the optimal solution.

Secondly, its cost seems lower than that of the more complicated decision rule obtained by maximal likelihood estimates. The maximal likelihood algorithm should not be better, of course, than the optimal algorithm, leading to P_w^* , and we have seen that our algorithm compares favourably with it.

Thirdly, a comparison is given with decision rules based on the Bayesian model. Here, again, the computational cost is incomparably higher, but the result is not uniformly better.

We can now summarise our conclusions. Taking all arguments together, we would suggest to always use the odds algorithm with sequential updating based on the estimator defined in (2.2). With some additional information, we may somewhat improve on this by a slight delay factor s_d , as explained before. Note also that this algorithm, working with a number of computations which are at most quadratic in n , stands out from a computational point of view.

Acknowledgement

We would like to thank the anonymous referee for his/her critical comments and suggestions.

References

- [1] ANO, K., TAMAKI, M. AND HU, M. (1996). A secretary problem with uncertain employment when the number of offers is restricted. *J. Operat. Res. Soc. Japan* **39**, 307–315.
- [2] BRUSS, F. T. (2000). Sum the odds to one and stop. *Ann. Prob.* **28**, 1384–1391.
- [3] BRUSS, F. T. (2003). A note on the odds theorem of optimal stopping. *Ann. Prob.* **31**, 1859–1861.
- [4] BRUSS, F. T. (2006). The art of a right decision: why decision makers may want to know the odds-algorithm. *Newsletter Europ. Math. Soc.* **62**, 14–20.
- [5] BRUSS, F. T. AND LOUCHARD, G. (2003). Optimal stopping on patterns in strings generated by independent random variables. *J. Appl. Prob.* **40**, 49–72.
- [6] BRUSS, F. T. AND PAINDAVEINE, D. (2000). Selecting a sequence of last successes in independent trials. *J. Appl. Prob.* **37**, 389–399.
- [7] CHOW, Y. S., ROBBINS, H. AND SIEGMUND, D. (1991). *The Theory of Optimal Stopping*. Corrected republication of the 1971 Dover edition. Dover, New York.
- [8] FERGUSON, T. S. (2006). *Optimal Stopping and Applications*. Available at <http://www.math.ucla.edu/tom/Stopping/Contents.html>.
- [9] FERGUSON, T., HARDWICK, J. AND TAMAKI, M. (1992). Maximizing the duration of owning a relatively best object. In *Strategies for Sequential Search and Selection in Real Time* (Amhurst, MA, 1990; Contemp. Math **125**), American Mathematical Society, Providence, RI, pp. 35–57.

- [10] HSIAU, S.-R. AND YANG, J.-R. (2000). A natural variation of the standard secretary problem. *Statistica Sinica* **10**, 639–646.
- [11] HSIAU, S.-R. AND YANG, J.-R. (2002). Selecting the last success in Markov-dependent trials. *J. Appl. Prob.* **39**, 271–281.
- [12] IUNG, B., LEVRAT, E. AND THOMAS, E. (2007). Odds-algorithm-based opportunistic maintenance task execution for preserving product conditions. *CIRP Ann.* **56**, 13–16.
- [13] KURISHIMA, A. AND ANO, K. (2003). A note on the full-information Poisson arrival selection problem. *J. Appl. Prob.* **40**, 1147–1154.
- [14] MAZALOV, V. V. AND TAMAKI, M. (2006). An explicit formula for the optimal gain in the full-information problem of owning a relatively best object. *J. Appl. Prob.* **43**, 87–101.
- [15] SUCHWALKO, A. AND SZAJOWSKI, K. (2003). On Bruss' stopping problem with general gain function. *Game Theory Appl.* **9**, 161–171.
- [16] SZAJOWSKI, K. (2007). A game version of the Cowan–Zabczyk–Bruss problem. *Statist. Prob. Lett.* **77**, 1683–1689.
- [17] TAMAKI, M. (2001). Optimal stopping on trajectories and the ballot theorem. *J. Appl. Prob.* **38**, 946–959.
- [18] VAN LOKEREN, M. (2007). DEA mémoire en statistique. Université Libre de Bruxelles. Unpublished manuscript.