

ON-LINE SELECTION OF AN ACCEPTABLE PAIR

J. PREATER,* *Keele University*

Abstract

A sequence of objects with independent, identically distributed qualities is presented to a selector who must choose two on-line, i.e. without anticipation or recall. The selector's aim is to obtain a satisfactory pair as quickly as possible. Two versions of the problem are considered, and optimal selection rules are derived and compared. An investigation is also made of a heuristic rule suitable for a selector who has no prior knowledge of the nature of the object sequence.

Keywords: Optimal stopping; multiple selection

2000 Mathematics Subject Classification: Primary 60G40

1. Introduction and model

We investigate a rather primitive version of sequential selection with two choices. Imagine that some objects are offered to us one by one and that our aim is to choose expeditiously two that are satisfactory. The usual rules of on-line selection apply: the decision to accept or reject the current object must be made before observing the next one and cannot be based on the characteristics of the objects yet to be revealed; in short, the choices are irrevocable and nonanticipatory. Other recent work on two-choice problems can be found in [1] and [6].

In addition to the substantial theoretical literature on sequential selection there are a number of established applications, the foremost of which is employee selection. Biologists have used single selection to model mate choice ([9] contains a useful bibliography); but imagine a zoologist searching for a promising mating pair of some species. Economists have interpreted on-line choice in terms of job search [7]; but suppose that both a job and a house are sought, either of which may be secured first as opportunities arise. More starkly, if the vertices from a known or unknown graph are presented to us by means of sampling with replacement, how should we go about obtaining an edge? In all these cases, what is deemed a satisfactory second object will be strongly influenced by the choice made for the first. Moreover, in many potential applications it is unlikely that the defining attributes of an acceptable pair could be stated explicitly; this raises the issue of selector ignorance.

We begin with the following simple model for the problem. Let Y_1, Y_2, \dots be independent random variables with values in the unit interval $[0, 1]$ and distribution F ; these variables represent the qualities of successive objects. Any given pair of object qualities is adjudged acceptable or unacceptable. We suppose for now that the set, $\mathcal{A} \subset [0, 1]^2$, of acceptable pairs is nonempty, *symmetric*, i.e.

$$(y, y') \in \mathcal{A} \implies (y', y) \in \mathcal{A},$$

Received 19 October 2005; revision received 17 February 2006.

* Postal address: Department of Mathematics, Keele University, Keele, Staffordshire ST5 5BG, UK.

Email address: j.preater@keele.ac.uk

and *monotone*, i.e.

$$(y, y') \in \mathcal{A} \quad \text{and} \quad (z, z') \geq (y, y') \quad \implies \quad (z, z') \in \mathcal{A},$$

with ' \geq ' being the usual coordinatewise order on $[0, 1]^2$. The confinement of the object qualities to $[0, 1]$ clearly incurs no real loss of generality; indeed, by transforming them and \mathcal{A} we could take F to be uniform, except that this turns out not to be expedient (see Lemma 2.2, below).

It is assumed that F and \mathcal{A} are known prior to the disclosure of any objects and are such that acceptable pairs do arise, i.e. $P((Y_1, Y_2) \in \mathcal{A}) > 0$. A nonanticipatory selection is, as usual, represented by a stopping time relative to the Y_t sequence: we let \mathcal{T} denote the set of admissible stopping time pairs (τ, τ') with either $1 \leq \tau < \tau'$ or $1 \leq \tau \leq \tau' = \infty$. A generalization of the model encompassing all of the applications mentioned above will be discussed in due course.

There are two natural objectives. First, the aim might be to maximize the probability of selecting an acceptable pair from the first N objects, where $N \geq 1$ is known; in this case the maximal *win probability* is

$$w_N = \sup\{E(\mathbf{1}_{\mathcal{A}}(Y_\tau, Y_{\tau'})\mathbf{1}_{\{\tau' \leq N\}}) : (\tau, \tau') \in \mathcal{T}\}, \quad (1.1)$$

with $\mathbf{1}_A$ the indicator of a set A . Here and elsewhere the convention is that $(Y_t, Y_{t'}) \notin \mathcal{A}$ when either t or t' is infinite.

Second, the intention may be to minimize the expected number of object observations; in this case we let

$$u = \inf\{E(\tau') : (Y_\tau, Y_{\tau'}) \in \mathcal{A}, (\tau, \tau') \in \mathcal{T}\}. \quad (1.2)$$

In particular, $u = \infty$ when $\tau' = \infty$.

Rather fancifully, but in keeping with [4], we shall call the case in which (1.1) holds the *primal* problem and the case in which (1.2) holds the *dual* problem. Any admissible pair of stopping times achieving either the supremum in (1.1) or the infimum in (1.2) is said to be *optimal* for the respective problem. In both cases the second choice is trivial – we simply wait for an acceptable partner for the object already chosen; the first choice is the crux. Thus, for any $\tau \geq 1$, we fix

$$\tau' = \inf\{t > \tau : (Y_\tau, Y_t) \in \mathcal{A}\}, \quad (1.3)$$

with $\tau' = \infty$ when either $\tau = \infty$ or the set is empty.

The primal problem is a special case of the general monotone two-choice problem

$$v_N = \sup\{E(\Psi(Y_\tau, Y_{\tau'})) : 1 \leq \tau < \tau' \leq N\},$$

where the function $\Psi: \mathbb{R}^2 \rightarrow [0, \infty)$ is nondecreasing and symmetric (i.e. $\Psi(x, y) = \Psi(y, x)$ for $x, y \in \mathbb{R}$). This was discussed in [8]. The question there was whether the selector should be progressively less particular as the object sequence unfolds, as is the case for the corresponding single-selection problem. This is deceptively plausible but seems difficult to prove for all Ψ ; indeed it is sometimes false when the Y_t are \mathbb{R}^2 -valued, for example. Our principal result is that this *time monotonicity* property does hold for the primal problem, in which $\Psi = \mathbf{1}_{\mathcal{A}}$.

Chen *et al.* [2] and Coffman *et al.* [4] studied an associated pair of on-line selection problems. In the former, m elements from an infinite Y_t sequence must be chosen so that their sum does not exceed a constant, $C > 0$; the objective is to minimize the expected number of observations required to achieve this. In the latter, the goal is to maximize the expected number of selections made from a finite Y_t sequence under the same constraint. The former, for $m = 2$, is essentially

included in our dual problem, whereas the latter is informed just a little by our primal problem, as illustrated later.

The paper is organized as follows. In Section 2 we study the primal problem, proving time monotonicity and giving some asymptotics. The dual problem, which is more straightforward, is considered in Section 3, and the optimal selection policy is compared with that for the primal problem. Finally, in the discursive Section 4, we examine the implications of broadening the problem formulation, and reflect on the matter of selector ignorance.

2. The primal problem

In this section we adopt the model from Section 1 and, after some preparatory work, address the primal problem, (1.1).

2.1. Preliminaries

A subset of $[0, 1]$ of the form $(x, 1]$ or $[x, 1]$ will be called an *upper interval* with *threshold* x . For any distribution G on $[0, 1]$, define the constant

$$\beta_G = \inf\{x \in [0, 1]: G(x) \geq 1 - x\}.$$

Let Y be a typical object quality. Define $p(x) = P((x, Y) \notin \mathcal{A})$, $x \in [0, 1]$, and let H denote the distribution function of $p(Y)$. Since \mathcal{A} is monotone, p is nonincreasing. Thus, the right-continuous inverse of p is given by

$$\bar{p}(y) = \inf\{x \in [0, 1]: p(x) \leq y\}, \quad y \in [0, 1],$$

with $\inf \emptyset = 1$; \bar{p} is also nonincreasing.

Now consider the sets

$$A_{\mathcal{A}} = \{y \in [0, 1]: (y, y) \in \mathcal{A}\}, \quad B_H = \{y \in [0, 1]: p(y) \leq \beta_H\},$$

each of which is an upper interval, and let $a^* = \inf A_{\mathcal{A}}$.

Lemma 2.1. *We have $A_{\mathcal{A}} \subset B_H$, and if p is strictly decreasing then*

$$a^* = \inf B_H = \bar{p}(\beta_H).$$

Proof. Let $y \in A_{\mathcal{A}}$. Then, because \mathcal{A} is monotone, $(y, Y) \notin \mathcal{A}$ implies that $Y < y$ and, so, $p(y) \leq F(y-)$. Consequently, if $x < p(y)$ then

$$x < p(y) \leq F(y-) \leq F(\bar{p}(x)-) = P(Y < \bar{p}(x)) = P(p(Y) > x) = 1 - H(x),$$

showing that $x \leq \beta_H$. It follows that $p(y) \leq \beta_H$, i.e. that $y \in B_H$, and we conclude that $A_{\mathcal{A}} \subset B_H$.

Now assume that p is strictly decreasing. From above we know that $a^* \geq \inf B_H$. To prove the reverse inequality, suppose that $y < a^*$. Then $(y, y) \notin \mathcal{A}$ and, so,

$$p(y) \geq F(y-) = P(Y < y) = P(p(Y) > p(y)).$$

Thus, $p(y) \geq \beta_H$, which in turn entails that $y \leq \inf B_H$, by virtue of the strict monotonicity of p . Therefore, $a^* \leq \inf B_H$, completing the proof.

2.2. A real sequence

The analysis of the primal problem depends on the properties of a particular sequence of real numbers.

Lemma 2.2. *Let X be a random variable taking values in the interval $[0, 1]$ and having distribution G . If the sequence of numbers (b_n) is defined by*

$$b_1 = 1, \quad b_{n+1} = E([b_n \wedge X]^n)^{1/(n+1)}, \quad n = 1, 2, \dots, \tag{2.1}$$

then

$$\text{the } b_n \text{ are nonincreasing in } n \tag{2.2}$$

and

$$\lim_n b_n = \beta_G. \tag{2.3}$$

Proof. Let X be as prescribed. To prove (2.2) we proceed by induction. Clearly $b_1 \geq b_2$. Suppose that

$$b_1 \geq \dots \geq b_n. \tag{2.4}$$

Pick a $b \in (0, 1)$ and impose the constraint $b_n = b$. Since b is arbitrary it then suffices, for the induction step, to show that $E(b^n \wedge X^n) \leq b^{n+1}$, and in doing so we may assume that $0 \leq X \leq b$. This follows by an obvious continuity argument once we have observed that, on the one hand, moving all of the mass of G to b entails that $b_n = b^{(n-1)/n} \geq b$, while, on the other hand, $b_n \leq b$ is the result of moving to b any mass lying to the right thereof.

It remains to show that

$$E(X^n) \leq b^{n+1}. \tag{2.5}$$

Using (2.1) and (2.4) we obtain $E(X^{n-1}) = b^n$. Set $Z = X^{n-1}$. Then

$$0 \leq Z \leq b^{n-1}, \quad E(Z) = b^n, \quad E(h(Z)) = E(X^n), \tag{2.6}$$

where $h(z) = z^{n/(n-1)}$, $0 \leq z \leq 1$. Now, since h is convex, it follows from a standard dilation result that $E(h(Z))$ is maximized, subject to the first two properties in (2.6), only when Z takes the value 0 or the value b^{n-1} . Here, this entails that $P(Z = b^{n-1}) = b$. Thus, the X maximizing the left-hand side of (2.5) satisfies $P(X = b) = b = 1 - P(X = 0)$, and, hence, $E(X^n) = b^{n+1}$. This proves (2.5) and, consequently, (2.2).

Turning to the verification of (2.3), let $b_\infty = \lim_n b_n$ and $\beta \equiv \beta_G$. We have $b_1 = 1 \geq \beta$, and if $b_n \geq \beta$ then

$$b_{n+1}^{n+1} = E([b_n \wedge X]^n) \geq E([\beta \wedge X]^n) \geq \beta^n (1 - G(\beta-)) \geq \beta^{n+1},$$

giving $b_{n+1} \geq \beta$. Thus, $b_n \geq \beta$ for all n by induction and, hence, $b_\infty \geq \beta$.

To prove the reverse inequality, suppose that $b_\infty = \eta > \beta$. Then, for any $n \geq 1$, we have $b_n \geq \eta$ and

$$\begin{aligned} b_{n+1}^{n+1} &= b_n^n E([1 \wedge X/b_n]^n) \leq b_n^n E([1 \wedge X/\eta]^n) \\ &= b_n^n [(1 - G(\eta-)) + E((X/\eta)^n \mathbf{1}_{\{X < \eta\}})]. \end{aligned}$$

However, in the last expression $1 - G(\eta-) < \eta$ and, by the bounded convergence theorem, the expectation is vanishingly small for large n . Therefore, for some $\delta > 0$ and some positive integer N ,

$$b_{n+1}^{n+1} \leq b_n^n (\eta - \delta), \quad n = N, N + 1, \dots$$

Iterating, we obtain

$$b_{n+N}^{n+N} \leq b_N^N (\eta - \delta)^n, \quad n = 1, 2, \dots,$$

which requires that $b_\infty = \lim_n b_n \leq \eta - \delta < \eta$, a contradiction. We conclude that $b_\infty \leq \beta$, and (2.3) follows.

Statement (2.2) is the kernel of our work; the case where $b_2 \geq b_3$, i.e. $E(E(X) \wedge X^2) \leq E(X)^{3/2}$, is of mild intrinsic interest.

2.3. Main theorem

We introduce the sets

$$A_n = \{y \in [0, 1]: p(y) \leq (1 - w_n)^{1/n}\}, \quad n = 1, 2, \dots, \tag{2.7}$$

where w_n is the primal problem win probability, (1.1), for n objects. Each A_n is an upper interval with threshold

$$a_n = \bar{p}((1 - w_n)^{1/n}), \quad n = 1, 2, \dots \tag{2.8}$$

Note that $a_1 = 0$.

Theorem 2.1. *The primal problem has an optimal stopping pair (τ_p, τ'_p) with*

$$\tau_p = \inf\{t: Y_t \in A_{N-t}\}, \tag{2.9}$$

where the stopping sets A_n are as given in (2.7), and τ'_p is as defined in (1.3). There is time monotonicity in that

$$[0, 1] = A_1 \supset A_2 \supset \dots, \tag{2.10}$$

and, thus, the thresholds a_n defined in (2.8) are nondecreasing in n . Furthermore,

$$a_\infty := \lim_n a_n = \bar{p}(\beta_H), \tag{2.11}$$

where, recall, H is the distribution of $p(Y)$.

Remark 2.1. The rather elaborate description in (2.9) and (2.10) of a time-monotone threshold rule, along with ‘stopping set’ terminology, is used simply because the A_n may be upper intervals of either type.

Proof of Theorem 2.1. For any τ , $1 \leq \tau \leq N - 1$, and τ' defined by (1.3), we have

$$E(\mathbf{1}_{\mathcal{A}}(Y_\tau, Y_{\tau'}) \mathbf{1}_{\{\tau' \leq N\}}) = 1 - E(p(Y_\tau)^{N-\tau}).$$

Hence, on writing $u_n = 1 - w_n$ and $X_t = p(Y_t)$, the two-choice problem (1.1) reduces to a single-choice problem for the X_t sequence:

$$u_N = \inf\{E(X_\tau^{N-\tau}): 1 \leq \tau \leq N - 1\}.$$

According to standard stopping time theory [3], the reduced problem admits the Bellman equation

$$u_1 = 1, \quad u_{n+1} = E(u_n \wedge X^n), \quad n = 1, 2, \dots, N - 1, \tag{2.12}$$

and has optimal stopping time

$$\tau = \inf\{t \geq 1: X_t^{N-t} \leq u_{N-t}\} = \inf\{t \geq 1: p(Y_t) \leq u_{N-t}^{1/(N-t)}\}.$$

This establishes (2.9).

By virtue of (2.12), the numbers $b_n = u_n^{1/n}$ satisfy (2.1). Property (2.10) is therefore a direct consequence of (2.2) and (2.7).

Finally, since X has distribution H , (2.3) implies that $\lim_n u_n^{1/n} = \beta_H$. The limit (2.11) then follows from (2.8) and the right continuity of \bar{p} .

The following corollary, which is immediate from Lemma 2.1, shows that the limit of the thresholds depends typically on \mathcal{A} but not on F .

Corollary 2.1. *We have $a_\infty \leq a^*$, with equality when p is strictly decreasing.*

The condition to have $a_\infty = a^*$ in Corollary 2.1 can be weakened; for example, it is enough that p be strictly decreasing in an open neighbourhood of a_∞ . However, it is easy to construct cases where $a_\infty < a^*$; for instance, if $\mathcal{A} = [0, 1]^2 - [0, a]^2$, $0 < a < 1$, then $a_n = 0$ for every n and $a^* = a$, whatever the distribution F is.

2.4. An illustration

Consider the following instance of Coffman *et al.*'s [4] aforementioned problem: maximize the expected number of selections from an independent, identically distributed sequence Z_1, \dots, Z_N , when Z_1 is uniformly distributed on $(\frac{1}{3}, 1)$, subject to their sum not exceeding 1. Since at least one selection is always made, whereas three or more are not feasible, an equivalent problem is to maximize the probability of making two selections. By working with transformed variables $Y_t = \frac{4}{3} - Z_t$, we reduce this to our primal problem with

$$\mathcal{A} = \{(y, y') : y + y' \geq \frac{5}{3}\}.$$

Theorem 2.1 then tells us that the initial selection should be at the first time t for which $Y_t \geq a_{N-t}$, i.e. for which

$$Z_t \leq \frac{4}{3} - a_{N-t},$$

where the thresholds a_n are nondecreasing. This monotonicity proved difficult to establish for the general problem in [4]. Furthermore, as p is strictly decreasing where it matters, the thresholds have limit $a_\infty = a^* = \frac{5}{6}$.

Let us now compute the thresholds. From (2.8) and (2.12), $a_n = \bar{p}(b_n)$, where the b_n satisfy (2.1) with $X = p(Y)$. Straightforward calculations show that

$$\bar{p}(y) = \begin{cases} 0, & y = 1, \\ \frac{4}{3} - \frac{2}{3}y, & \frac{1}{2} \leq y < 1, \\ 1, & 0 \leq y < \frac{1}{2}; \end{cases}$$

that H , the distribution of X , is given by

$$H(x) = \begin{cases} 1, & x = 1, \\ x - \frac{1}{2}, & \frac{1}{2} \leq x < 1, \\ 0, & 0 \leq x < \frac{1}{2}; \end{cases}$$

and that recursion (2.1) reads

$$b_1 = 1, \quad b_{n+1} = \left[\left(\frac{3}{2} - \frac{n}{n+1} b_n \right) b_n^n - \frac{1}{n+1} \left(\frac{1}{2} \right)^{n+1} \right]^{1/(n+1)}, \quad n = 1, 2, \dots$$

From this we obtain $a_1 = 0, a_2 = 0.7097, a_3 = 0.7344, a_4 = 0.7507, \dots$, with slow convergence to the limit.

3. The dual problem

In this section we examine the dual problem, in which (1.2) holds. The second selection is again governed by (1.3).

In seeking to attain the infimum in (1.2) we can obviously restrict attention to stopping times τ with finite expectation. Conditional upon $\tau = t$ and $Y_t = y$, the waiting time, $\tau' - \tau$, until the second selection has a geometric distribution with parameter $1 - p(y)$. Thus,

$$\begin{aligned} E(\tau') &= E(\tau) + E(\tau' - \tau) = E(\tau) + E(E(\tau' - \tau \mid \mathcal{F}_\tau)) \\ &= E(\tau) + E((1 - p(Y_\tau))^{-1}) \\ &= E(\tau + (1 - p(Y_\tau))^{-1}), \end{aligned}$$

where \mathcal{F}_τ is the pre- τ σ -field relative to the Y_t sequence. Therefore, writing

$$Z_t = -(1 - p(Y_t))^{-1}$$

(which equals $-\infty$ if $p(Y_t) = 1$), we obtain

$$u = \inf_{\tau} \{E(\tau + (1 - p(Y_\tau))^{-1}) : E(\tau) < \infty\} = -\sup_{\tau} \{E(Z_\tau - \tau) : E(\tau) < \infty\}.$$

The supremum represents the classical problem of maximizing the net gain from the sale of an asset subject to independent, identically distributed offers Z_t and unit observation cost. This problem was solved in [3] and discussed further in, for example, [5]. (In the cited work $E(|Z_1|) < \infty$, but a simple truncation argument shows that the results also apply in our case, where $Z_1 \leq 0$ and $P(Z_1 > -\infty) > 0$.) The following stopping time, where $\zeta = -u$, the maximal net gain, is the unique solution to

$$E((Z_1 - \zeta)^+) = 1, \tag{3.1}$$

is optimal:

$$\tau = \inf\{t \geq 1 : Z_t \geq \zeta\}. \tag{3.2}$$

The next result summarizes these findings and appends a complement roughly to the effect that the dual selector is no more particular than the primal selector when N is large. Let

$$B = \{y \in [0, 1] : p(y) \leq 1 - u^{-1}\}, \tag{3.3}$$

which is an upper interval with threshold

$$b = \bar{p}(1 - u^{-1}).$$

Theorem 3.1. *The dual problem has an optimal stopping pair (τ_d, τ'_d) with*

$$\tau_d = \inf\{t \geq 1 : Y_t \in B\}, \tag{3.4}$$

where B is as given in (3.3), and τ'_d is as defined in (1.3). The minimal expected number of observations, u , needed to procure an acceptable pair is the unique solution to

$$E((u - (1 - p(Y))^{-1})^+) = 1. \tag{3.5}$$

Moreover, the dual threshold, b , is no larger than the asymptotic primal threshold, a_∞ .

Proof. All but the final statement of the theorem comes from the preceding discussion, (3.4) and (3.5) being transliterations of (3.2) and (3.1), respectively. To show that $b \leq a_\infty$, we first define the function K by

$$K(x) = \int_1^x H((1 - y^{-1})-) \, dy, \quad x > 1,$$

where H is the distribution of $p(Y)$, as before. Since the integrand is nondecreasing and left continuous, K is nondecreasing and convex, with left-hand derivative, $'K$ say, given by

$$'K(x) = H((1 - x^{-1})-), \quad x > 1.$$

We have $K(1) = 0$. Moreover, using (3.5), we obtain

$$\begin{aligned} 1 = E((u - (1 - p(Y))^{-1})^+) &= \int_0^\infty P(u - (1 - p(Y))^{-1} > z) \, dz \\ &= \int_0^{u-1} H((1 - (u - z)^{-1})-) \, dz \\ &= \int_1^u H((1 - y^{-1})-) \, dy \quad (y = u - z) \\ &= K(u). \end{aligned}$$

Therefore,

$$\begin{aligned} 1 - H(1 - u^{-1}) &\leq 1 - H((1 - u^{-1})-) \\ &= 1 - 'K(u) \\ &\leq 1 - \frac{K(u) - K(1)}{u - 1} \\ &= 1 - (u - 1)^{-1} \\ &< 1 - u^{-1}, \end{aligned}$$

the second inequality owing to the convexity of K . It follows that $\beta_H \leq 1 - u^{-1}$, whence $b = \bar{p}(1 - u^{-1}) \leq \bar{p}(\beta_H) = a_\infty$, as claimed.

4. Complements

In this section we broaden the model and tackle selector ignorance.

4.1. Generalization

In the simple model, object qualities were numerical and \mathcal{A} was monotone. However, this setting cannot capture, for example, the male–female selection problem because, even if the individuals within each gender can be scored numerically, they must occupy disjoint subsets of $[0, 1]$; thus, the monotonicity of \mathcal{A} will almost certainly be lost. In general, object qualities may not even be ordered. To construct a model in these circumstances, let (S, \mathcal{S}) be a measurable space and let the object qualities Y_i be independent, S -valued random variables, again with distribution F . The set $\mathcal{A} \subset S \times S$ of acceptable pairs is assumed to be nonempty, symmetric, and $\mathcal{S} \times \mathcal{S}$ measurable (which was guaranteed previously by monotonicity). The remainder of the formulation in Section 1 carries through unchanged.

Let us review the consequences for the primal and dual problems. Beginning with the former, note first that the update

$$p(x) = P((x, Y) \notin \mathcal{A}), \quad x \in S,$$

is well defined by virtue of \mathcal{A} 's measurability, but is no longer necessarily monotonic. Nevertheless, with (2.7) replaced by

$$A_n = \{y \in S : p(y) \leq (1 - w_n)^{1/n}\}, \quad n = 1, 2, \dots,$$

we have the following reworking of Theorem 2.1.

The generalized primal problem has an optimal stopping pair (τ_p, τ'_p) with τ_p given by (2.9) and τ'_p defined by (1.3). There is time monotonicity in that

$$S = A_1 \supset A_2 \supset \dots$$

Furthermore, the sets A_n decrease to $A_\infty = \{y \in S : p(y) \leq \beta_H\}$.

The proof of Theorem 2.1 (apart from the last sentence) holds also for this version. There is no counterpart to Corollary 2.1.

Turning to the dual problem, replace (3.3) by

$$B = \{y \in S : p(y) \leq 1 - u^{-1}\}.$$

Then Theorem 3.1 and its proof carry through essentially unchanged, as follows.

The generalized dual problem has an optimal stopping pair (τ_d, τ'_d) with τ_d given by (3.4) and τ'_d defined by (1.3); u is determined by (3.5). Moreover, $A_\infty \subset B$.

Here, the final claim of Theorem 3.1 which is the analogue of $b \leq a_\infty$, is immediate because $\beta_H \leq 1 - u^{-1}$ still applies.

4.2. Ignorance

Our solutions to the primal and dual problems, in both original and generalized forms, presuppose that \mathcal{A} , F , and N are all known to the selector; notwithstanding some special cases, this is clearly necessary for achieving exact optimality. Partial knowledge is arguably more realistic, and could be modelled in a variety of ways; indeed there are many precedents in the sequential selection literature, for instance [10]. However, being rather primitive, we wish to consider briefly the plight of a selector who, ignorant of \mathcal{A} , F , and N , times the choices solely on the basis of knowing which pairs of observed objects are acceptable. More formally, the stopping times used must be adapted to the filtration (\mathcal{G}_t) , where

$$\mathcal{G}_t = \sigma(\mathbf{1}_{\{(Y_r, Y_s) \in \mathcal{A}\}} : 1 \leq r \leq s \leq t), \quad t = 1, 2, \dots$$

Notice that we have generously granted the selector a modicum of imagination, in that $r = s$ is permitted, i.e. the acceptability, or otherwise, of an observed object paired with an exact replica is assumed known.

Consider the simple model of Section 1 again, and let us focus first on the primal problem. Corollary 2.1 suggests making the first selection at time

$$\tau_0 = \inf\{t \geq 1 : Y_t \in A_{\mathcal{A}}\}, \tag{4.1}$$

which is the principal reason for our aforementioned generosity. As above, the time of the second selection is prescribed by (1.3), and we use τ'_0 to denote it. Let w_0 be the associated probability of obtaining an acceptable pair. It may well happen that $\tau_0 \geq N$, in which case $w_0 = 0$. We have the following result.

Theorem 4.1. *For any monotone \mathcal{A} , any F , and any $N \geq 2$,*

$$w_N - w_0 \leq \frac{1}{4};$$

moreover, the bound is tight.

Proof. Fix a monotone \mathcal{A} , an F , and an $N \geq 2$, and let J_p and J_0 be the events that an acceptable pair is obtained using τ_p and τ_0 , respectively. Write

$$M = \text{card}\{1 \leq t \leq N : Y_t \in A_{\mathcal{A}}\}, \quad a = P(Y \in A_{\mathcal{A}}), \tag{4.2}$$

and note that, since M has a binomial(N, a) distribution,

$$P(M = 1) = Na(1 - a)^{N-1} \leq \frac{1}{2}, \tag{4.3}$$

the bound being attained when $N = 2$ and $a = \frac{1}{2}$.

Next, when $M = 1$ there is just one object, say \mathcal{O} , having the largest quality among the first N objects. Denote by C and D the events that $M = 1$ and an acceptable partner to \mathcal{O} appears among the first N objects before and, respectively, after \mathcal{O} arrives. For the choice at τ_p to win (i.e. to result in an acceptable pair) and that at τ_0 to lose, C must occur and D must not occur; thus, $J_p \cap J_0^c \subset C \cap D^c$. Furthermore, by symmetry, the exclusive events $C \cap D^c$ and $C^c \cap D$ are equiprobable, conditional on $M = 1$. Therefore,

$$P(J_p \mid M = 1) - P(J_0 \mid M = 1) \leq P(J_p \cap J_0^c \mid M = 1) \leq P(C \cap D^c \mid M = 1) \leq \frac{1}{2}. \tag{4.4}$$

In view of the monotonicity of \mathcal{A} , the choice at τ_p loses when $M = 0$ and the choice at τ_0 wins when $M \geq 2$. We conclude that

$$\begin{aligned} w_N - w_0 &= P(J_p) - P(J_0) = \sum_{m=0}^N [P(J_p \mid M = m) - P(J_0 \mid M = m)] P(M = m) \\ &\leq [P(J_p \mid M = 1) - P(J_0 \mid M = 1)] P(M = 1) \\ &\leq \frac{1}{2} \times \frac{1}{2} \\ &= \frac{1}{4}, \end{aligned}$$

using (4.3) and (4.4).

Finally, when $\mathcal{A} = [0, 1]^2 - [0, \frac{1}{2}]^2$, F is uniform on $[0, 1]$, and $N = 2$, it is easy to see that $w_2 = \frac{3}{4}$, $w_0 = \frac{1}{2}$, and, thus, that the bound in the theorem is attained.

The application of (4.1) to the dual problem is less well motivated. Nevertheless, for this problem there is a counterpart to Theorem 4.1. Let u_0 denote the expected number of objects observed when using the stopping times (τ_0, τ'_0) .

Theorem 4.2. *For any monotone \mathcal{A} and any F ,*

$$\frac{u_0 - u}{u} \leq \frac{1}{4};$$

moreover, the bound is tight.

Proof. Recall from (4.1) and (3.4) that the stopping sets for τ_0 and τ_d are $A_{\mathcal{A}}$ and B , respectively. In the proof of Theorem 3.1 we learned that $\beta_H \leq 1 - u^{-1}$ and, therefore, that $A_{\mathcal{A}} \subset B$.

By monotonicity, an acceptable pair must include at least one object quality in $A_{\mathcal{A}}$. Thus, $\tau'_d \geq \tau_0$ and, so,

$$u = E(\tau'_d) \geq E(\tau_0) = 1/a, \tag{4.5}$$

τ_0 having a geometric(a) distribution with a as defined in (4.2).

Now let $Z = Y_{\tau_0}$ be the quality of the first object that is eligible when using τ_0 (with $Z = 1$, say, on the null event $\tau_0 = \infty$). For $z \in A_{\mathcal{A}}$, define

$$B_z = \{y \in B : y \notin A_{\mathcal{A}}, (y, z) \in \mathcal{A}\}, \quad b_z = P(B_z),$$

and

$$C_z = \{\text{at least one quality from } B_z \text{ occurs before time } \tau_0\}.$$

Then

$$P(C_z) = \frac{b_z}{a + b_z}. \tag{4.6}$$

On C_z^c , the dual selection policy will procure precisely one object by time τ_0 , and the quality of this object will not be greater than Z . Thus,

$$E(\tau'_0 - \tau'_d \mid C_z^c, Z) \leq 0. \tag{4.7}$$

On C_z , $\tau'_0 - \tau'_d$ is at most the waiting time after τ_0 until an acceptable partner (determined by comparison of its quality with Z) appears, and this waiting time is independent of \mathcal{F}_{τ_0} . It follows that

$$E(\tau'_0 - \tau'_d \mid C_z, Z) \leq \frac{1}{P((Y, Z) \in \mathcal{A} \mid Z)} \leq \frac{1}{a + b_z}. \tag{4.8}$$

Therefore, employing (4.6), (4.7), and (4.8), we have

$$\begin{aligned} u_0 - u &= E(\tau'_0) - E(\tau'_d) \\ &= E(E(\tau'_0 - \tau'_d \mid Z)) \\ &= E(P(C_z \mid Z) E(\tau'_0 - \tau'_d \mid C_z, Z) + P(C_z^c \mid Z) E(\tau'_0 - \tau'_d \mid C_z^c, Z)) \\ &\leq E\left(\frac{b_z}{a + b_z} \times \frac{1}{a + b_z} + \frac{a}{a + b_z} \times 0\right) \\ &= E\left(\frac{b_z}{(a + b_z)^2}\right). \end{aligned}$$

From this and (4.5) we obtain

$$\frac{u_0 - u}{u} \leq E\left(\frac{abz}{(a + bz)^2}\right) = E\left(\frac{\frac{1}{4}((a + bz)^2 - (a - bz)^2)}{(a + bz)^2}\right) \leq \frac{1}{4}.$$

When \mathcal{A} is not monotone, or in the context of the generalized model, τ_0 has nothing to recommend it. One alternative is then

$$\tau^* = \inf\{t \geq 1 : (Y_s, Y_t) \in \mathcal{A} \text{ for some } s, 1 \leq s < t\},$$

which we think of, facetiously, as ‘stable door’ stopping. The properties of τ^* merit a separate study.

Acknowledgement

The author would like to thank the referee for a thorough reading and for a simplification of the proof of (2.2).

References

- [1] ASSAF, D., GOLDSTEIN, L. AND SAMUEL-CAHN, E. (2004). Two-choice optimal stopping. *Adv. Appl. Prob.* **36**, 1116–1147.
- [2] CHEN, R. W., NAIR, V. N. AND VARDI, Y. (1984). Optimal sequential selection of N random variables under a constraint. *J. Appl. Prob.* **21**, 537–547.
- [3] CHOW, Y. S., ROBBINS, H. AND SEIGMUND, D. (1971). *Great Expectations: The Theory of Optimal Stopping*. Houghton Mifflin, Boston, MA.
- [4] COFFMAN, E. G., JR., FLATTO, L. AND WEBER, R. R. (1987). Optimal selection of stochastic intervals under a sum constraint. *Adv. Appl. Prob.* **19**, 454–473.
- [5] KENNEDY, D. P. AND KERTZ, R. P. (1992). Comparisons of optimal stopping values and expected suprema for i.i.d. r.v.'s with costs and discounting. *Contemp. Math.* **125**, 217–230.
- [6] KÜHNE, R. AND RÜSCHENDORF, L. (2002). On optimal two-stopping problems. In *Limit Theorems in Probability and Statistics*, Vol. II, János Bolyai Mathematical Society, Budapest, pp. 261–271.
- [7] LIPPMAN, S. A. AND MCCALL, J. J. (1976). The economics of job search: a survey: part I. *Econom. Enquiry* **14**, 155–189.
- [8] PREATER, J. (1993). A note on monotonicity in optimal multiple stopping problems. *Statist. Prob. Lett.* **16**, 407–410.
- [9] RAMSEY, D. M. (1994). Models of evolution, interaction and learning in sequential decision processes. Doctoral Thesis, University of Bristol.
- [10] TAMAKI, M. (1988). A Bayesian approach to the best-choice problem. *J. Amer. Statist. Soc.* **83**, 1129–1133.