Available online at www.sciencedirect.com
SCIENCE DIRECT
ELSEVIER
Journal of Mathematical Psychology 49 (2005) 410–422
Journal of Mathematical Psychology
www.elsevier.com/locate/jmp

A multi-attribute extension of the secretary problem: Theory and experiments

J. Neil Bearden*, Ryan O. Murphy, Amnon Rapoport

University of Arizona, Department of Management and Policy, 405 McClelland Hall, Tucson, AZ 85721, USA

Received 21 December 2004; received in revised form 25 April 2005

Abstract

We present a generalization of a class of sequential search problems with ordinal ranks, referred to as "secretary" problems, in which applicants are characterized by multiple attributes. We then present a procedure for numerically computing the optimal search policy and test it in two experiments with incentive-compatible payoffs. With payoffs dependent on the absolute ranks of the attributes, we test the optimal search model with both symmetric (Experiment 1) and asymmetric (Experiment 2) search problems. In both experiments we find that, relative to the optimal search policy, subjects stop the search too early. Our results show that this bias is largely driven by a propensity to stop prematurely on applicants of intermediate (relative) quality.

Keywords: Sequential search; Secretary problem; Dynamic programming

1. Introduction

Consider the problem of searching for an employee to fill an open position that requires both strong technical skills and good interpersonal skills. Ideally, one would hire an applicant who is outstanding on both; more likely, however, one will have to make trade-off decisions, perhaps by accepting an applicant who has remarkable technical skills but only average interpersonal skills. Decisions of this sort have received considerable attention in static contexts in which the decision maker (DM) must choose among a set of options presented *simultaneously* (see Payne, Bettman, & Johnson, 1993, for a comprehensive review). Here, we are interested in problems in which options are observed *sequentially*, and decisions to accept or reject an option must be made in the absence of full information about the multi-dimensional distribution of the attributes. Returning to the hiring example, when one decides to terminate the search by hiring an applicant, one then forgoes the opportunity of hiring another applicant, potentially better, who has yet to be interviewed. Likewise, in times of low unemployment, not hiring a seemingly excellent applicant on the spot may mean that one forgoes the opportunity to hire that applicant. As a result, one may be forced into a position of hiring a less qualified applicant later on.

Most previous research on sequential search problems has presupposed that options are represented by a single (scalar) value of quality or goodness. These problems fall into three general classes. *Full information problems* present DMs with options that are random variables drawn i.i.d. from a distribution assumed to be known to the DM before the search commences. In *partial information problems* the assumption that the DM knows the parameters of the distribution from which the options are sampled is relaxed by, for example, assuming that the DM knows that the distribution is normal, but that she must learn its mean and variance during the search process. *No information problems* suppose that the distribution from which the options are

*Corresponding author.
E-mail addresses: jneilb@gmail.com (J.N. Bearden), murphro@eller.arizona.edu (R.O. Murphy), amnon@u.arizona.edu (A. Rapoport).

0022-2496/$ - see front matter
doi:10.1016/j.jmp.2005.08.002

taken is unknown to the DM and cannot be learned during the search process. The most famous example of a no information problem is the "secretary problem" (e.g., Ferguson, 1989; Freeman, 1983; Gilbert & Mosteller, 1966; Samuels, 1991). In it, the DM is only informed about the relative rank of each encountered option, specifically, whether each option is the best observed up to that point.

Experimental work on sequential search problems has primarily focused on the full-information case (e.g., Cox & Oaxaca, 1989; Hey, 1981, 1982, 1987; Kogut, 1990; Rapoport, 1969; Rapoport & Tversky, 1970; Sonnemans, 1998, 2000). Some have examined the partial information case (e.g., Kahan, Rapoport, & Jones, 1967; Shapira, 1981). More recently, the no information case has received growing attention (e.g., Bearden, Rapoport, & Murphy, 2004; Corbin, Olson, & Abbondanza, 1975; Seale & Rapoport, 1997, 2000; Zwick, Rapoport, Lo, & Muthukrishnan, 2003). In all of these cases, the options are represented by a scalar value (either a ratio measure of quality or rank information). Often, however, as in the job search example, DM must search through options composed of *multiple attributes*. In the current paper, we describe a new class of sequential search problems that generalize the secretary problem to options composed of multiple attributes, present a method for computing the optimal policies, and describe results from two experiments in which we test the descriptive power of the optimal search model.

2. Secretary problems

In the *classical secretary problem* (CSP), a DM sequentially observes applicants randomly drawn from a pool of n applicants for a single position. When she observes the jth applicant in the sequence, she learns only the quality of that applicant with respect to those previously seen. Her objective is to select the one who is best overall—i.e., relative to all applicants, those seen and those not-yet-seen. The CSP can be formally stated as follows:

1. There is a fixed and known number n of applicants competing for a single position who can be ranked from best (1) to worst (n) with no ties.
2. The applicants are interviewed (observed) sequentially in a random order (with all $n!$ orderings occurring with equal probability).
3. For each applicant j, the DM can only ascertain the *relative rank* of the applicant, that is, how valuable or attractive the applicant is relative to the $j-1$ previously viewed applicants.
4. Once rejected, an applicant cannot be later recalled. If reached, the nth applicant must be accepted.
5. The DM earns a payoff of 1 for selecting the applicant with *absolute rank* 1 (i.e., the overall best applicant in the population of n applicants) and 0, otherwise.

The optimal (expected payoff maximizing) search policy is to interview and reject the first t^*-1 applicants and then to accept the first one thereafter with a relative rank of 1 (Gilbert & Mosteller, 1966). The optimal cutoff can be obtained by

$$t^* = \min\left\{t \geqslant 1 : \sum_{k=t+1}^{n} \frac{1}{k-1} \leqslant 1\right\}. \quad (1)$$

The cutoff t^* converges to ne^{-1}, where e is the base of the natural logarithm, and the optimal policy selects the best applicant with probability $e^{-1} \approx .3679$ as $n \to \infty$. Both t^* and the selection probability converge from above. When $n=4$, $t^*=2$ and the best applicant is selected with probability .4583. Already at $n=20$, $t^*=8$ and the probability of success is .3842. At $n=100$, $t^*=38$ and the success probability is .3710. A historical review of the CSP can be found in Ferguson (1989) and in Samuels (1991). Depending on the context, the problem is sometimes referred to by other names (e.g., the "Sultan's dowry" problem).

Seale and Rapoport (1997) had subjects play a large number of random instances of the CSP in two different experimental conditions: $n=40$ and 80. In both, they found that subjects tended to terminate their search too early relative to the dictates of the optimal policy. The authors proposed several different decision heuristics that DMs might have used in the CSP, and competitively tested them using their experimental data. They concluded that a threshold rule of the same form as the optimal policy best accounted for their data. The DMs' thresholds were simply shifted toward early applicants; more precisely, the thresholds tended to be positioned below the optimal (t^*th) position. Seale and Rapoport (1997) suggested that the bias to stop too early might result from endogenous search costs. Since search is costly in terms of time (Stigler, 1961), the optimal policy to which the subjects' behavior is compared may be inappropriate. One cannot rule out the possibility that the search policies used by the subjects are, in fact, net payoff maximizing (and therefore optimal) if endogenous search costs are taken into consideration.

In a subsequent paper, Seale and Rapoport (2000) relaxed assumption 1 of the CSP. At the beginning of each trial, subjects in their experiment were informed of the distribution from which n was sampled but not the actual value of n for the problem instance they played. These authors again concluded that the subjects tended to terminate their search too soon. Other variants of the CSP have been experimentally studied. For example, Zwick et al. (2003) relaxed assumption 4 by allowing

Table 1
Several GSPs, their optimal policies $\mathbf{t}^*$, expected payoffs $E(\pi|\mathbf{t}^*)$, expected search lengths $E(m|\mathbf{t}^*)$

GSP	n	$(\pi(1),\ldots,\pi(n))$	$\mathbf{t}^*=(t_1^*,\ldots,t_n^*)$	$E(\Pi\|\mathbf{t}^*)$	$E(m\|\mathbf{t}^*)$
1	40	$(1,0,\ldots,0)$	$(16,40,\ldots,40)$	.38	30.03
2	20	$(1,1,0,\ldots,0)$	$(8,14,20,\ldots,20)$	.69	14.15
3	40	$(15,7,2,0,\ldots,0)$	$(14,29,37,40,\ldots,40)$	6.11	27.21

DMs to recall previously interviewed applicants, with the success of recall being a probabilistic function of the time of recall and the position of the to-be-recalled applicant. When search was costless, they found that their subjects tended to search insufficiently; however, when the experimenters imposed a fixed search cost for each applicant, the pattern was reversed: the subjects tended to search for too long.

Bearden et al. (2004) questioned the generality of findings from experimental studies of the CSP (e.g., Seale & Rapoport, 1997). They worried that the CSP's payoff scheme (payoffs only for selecting the best applicant) was too artificial. In hiring an administrative assistant, for example, it is tautological to say that one is better off hiring a better applicant over a poorer applicant. Hence, the payoff scheme of the CSP misses an important feature of many actual search problems. To capture this feature, Bearden et al. proposed what they dubbed the *generalized secretary problem* (GSP) by replacing assumption 5 in the CSP with:

5′. The DM earns a payoff of $\pi(a)$ for selecting an applicant with absolute rank a where $\pi(1)\geqslant\cdots\geqslant\pi(n)$.

This formulation captures a number of interesting payoff schemes. Suppose, for example, that one's payoff increases linearly in the quality of the selected applicant. One can represent this in the GSP by setting $\pi(a)=\beta(n-a)$, when $\beta>0$. Under this scheme, the DM's utility increases at a rate of β in the quality of the selected applicant. For purposes of optimization, however, β can be omitted, as the optimal policy is invariant with respect to β. For all $\beta>0$, the problem reduces to one of simply minimizing the rank of the selected applicant (see Chow, Moriguti, Robbins, & Samuels, 1964). Multiple alternative payoff structures can be captured as well. Further, note that the CSP is a special case of the GSP in which $\pi(1)=1$ and $\pi(a)=0$ for all $a>1$.

The optimal policy for search problems with the payoff structure stated in assumption 5′ has the same threshold form as that of the CSP (Mucci, 1973); however, rather than a single threshold, the policy is represented by a vector of thresholds $\mathbf{t}^*=(t_1^*,\ldots,t_n^*)$. Under the optimal policy, the DM should interview and reject the first t_1^*-1 applicants, then between applicant t_1^* and applicant t_2^*-1 she should only accept applicants with relative rank 1; between applicant t_2^* and applicant t_3^*-1 she should accept applicants with relative ranks 1 or 2; and so on. Under this policy, the DM's standards relax as she plunges deeper into the applicant pool (and closer to the last applicant), and she is more apt to select lower quality applicants. Bearden and Murphy (2004) presented a dynamic programming procedure for computing optimal policies for the GSP.

Examples of several GSPs and their optimal policies are shown in Table 1. GSP1 is merely a CSP with $n=40$. For it, the DM should skip the first 15 applicants, and then take the first one thereafter with a relative rank of 1. Following this policy, the DM can expect to earn .38 and, on average, to search through about 30 applicants before making a selection. GSP2 corresponds to a case in which the DM's objective is to select the best or second best applicant, as both outcomes lead to a payoff of 1 and all others to a payoff of 0. When $n=20$, the DM should skip the first 7 applicants; then, between applicant positions 8 and 13 she should only accept applicants with relative rank 1; from position 14 on, she should take applicants with relative rank 1 or 2. Consequently, she can expect to earn .68, which is equivalent to selecting the best or second best applicant 68% of the time; and she will interview around 14 applicants before making a selection. The final GSP instance in Table 1 captures a situation in which the DM would like to get one of the top few applicants (of 40), and in which she earns more for selecting the best than the second best, and likewise for the third best. If she does not select one of the top three applicants, she earns nothing. GSP3 was studied experimentally by Bearden et al. (2004), whose work we turn to next.

Bearden et al. (2004) tested two different GSPs. In Experiment 1, they used payoffs $(25,13,6,3,2,1,0,\ldots,0)$ for a GSP with $n=60$; and, again, the payoffs for Experiment 2 are shown in Table 1 as GSP3. Despite the more plausible payoff structures, in both problems DMs tend to terminate their search too early relative to the dictates of the optimal policy. Bearden et al. offered an alternative to the endogenous search cost explanation proposed by Seale and Rapoport (1997). Using scoring rules, Bearden et al. had subjects estimate the probability of obtaining various payoffs for selecting applicants of different relative ranks in different applicant positions. Based on the estimation results, they then argued that the bias to terminate search too early in the GSP results from subjects overestimating the payoffs that would result from doing so. In fact, their subjects' estimates of obtaining positive payoffs were subadditive: for many values of the applicant positions j, the subjects' mean probability estimates for obtaining positive payoffs for stopping on j summed to more than 1. They explained this finding using Tversky and Koehler (1994) support theory, which is a descriptive theory of subjective probability. In short, they suggested that when evaluating early applicants the DMs do not give sufficient weight to the fact that a large number of applicants remain to be interviewed.

The current paper builds on the work of Bearden et al. (2004) and Bearden and Murphy (2004) by proposing a multi-attribute (or multi-dimensional) generalization of the GSP, presenting a method for computing its optimal policies, and testing it in two experiments with incentive-compatible payoffs. Since real-world search problems often involve trade-offs among attributes, we believe that this extension moves laboratory search problems, which provide us with an exceptional degree of control, closer to the types of problems faced by DMs in the wild.

3. A multi-attribute secretary problem

3.1. Statement of the problem

The *multi-attribute secretary problem* (MASP) is expressed by the following assumptions:

1. There is a fixed and known number n of applicants for a single position. The applicants differ along k different dimensions or attributes. Within a given attribute, the applicants can be ranked from best (1) to worst (n) with no ties. The k attributes are uncorrelated.
2. The applicants are interviewed sequentially in a random order (with all $n!$ orderings occurring with equal probability).
3. For each applicant j, the DM can only ascertain the *relative ranks* of the applicant's k attributes.
4. Once rejected, an applicant cannot be recalled. If reached, the nth applicant must be accepted.
5. For each attribute i of the selected applicant, the DM earns a payoff of $\pi^i(a^i)$, where a^i is the selected applicant's absolute rank on attribute i and $\pi^i(1) \geqslant \cdots \geqslant \pi^i(n)$.

The attributes might correspond to different characteristics of applicants for a job, such as education, work experience, degree of technical proficiency, interpersonal skills, etc. Keep in mind that the "no tie" assumption applies within an attribute. It is possible, for example, that in a two-attribute problem one applicant A could be best on attribute 1 and second best on attribute 2, while applicant B could be second best on attribute 1 and best on attribute 2. Assuming that the two attributes are equally important to the DM, A and B would, in a sense, be "tied." By assumption 1, however, we are only assuming that there are no ties *within* an attribute; it is still possible that there are ties *across* applicants.

Before describing the optimal policy for the MASP we must first introduce some notation. As in the single-attribute secretary problems (the CSP and the GSP), the DM's payoffs depend on the selected applicant's true or *absolute ranks*. For a given attribute, these represent each applicant's quality (value, worth, etc.) compared to each of the other $n-1$ applicants. More precisely, the absolute rank of the jth applicant on the ith attribute, denoted by a^i_j, is simply the number of applicants in the applicant pool, including j, whose ith attribute is at least as good as the jth applicant's. The jth applicant's set of absolute ranks can therefore be represented by a vector $\mathbf{a}_j = (a^1_j, \ldots, a^k_j)$.

Although the DM's payoffs are determined on the basis of absolute ranks, for each attribute she only observes an applicant's *relative ranks*, that is, how the applicant compares to the previously observed applicants, not to the entire pool of n applicants. To be clear, the relative rank of the jth applicant on the ith attribute, r^i_j, is the number of applicants from 1 to j whose ith attribute is at least as good as the jth's. When making a selection decision for the jth applicant, the DM only observes $\mathbf{r}_j = (r^1_j, \ldots, r^k_j)$.

By assumption 5, the payoff for selecting the jth applicant can be expressed as

$$\Pi_j = \sum_{i=1}^{k} \pi^i(a^i_j). \tag{2}$$

Hence, the "importance" of an attribute to the DM is captured by the attribute's payoffs relative to those of the other attributes. Consider a simple case with payoffs for attribute 1: $\pi^1(1) = 3$, $\pi^1(2) = 2$, $\pi^1(3) = 1$; and for attribute 2: $\pi^1(1) = 1.5$, $\pi^1(2) = 1$, $\pi^1(3) = .50$. Effectively, these payoffs capture a situation in which the first attribute is twice as important to the DM as the second attribute.

Note that when $\pi^1(a) > 0$ for some a, and for $i > 1$ $\pi^i(a) = 0$ for all a, the MASP reduces to a GSP. Therefore, the methods we describe below for computing optimal policies for the MASP can also be used to obtain solutions for GSPs (and, therefore, for CSPs).

A number of problems related to the MASP have appeared in the literature. Gnedin (1981) presented the solution to a multi-attribute CSP in which the attributes are independent, and the DM's objective is to select an applicant who is best on at least one attribute. The optimal policy for this problem consists of two thresholds g^* and h^* with $g^* \leqslant h^*$, and works as follows: skip the first $g^* - 1$ applicants. Between g^* and $h^* - 1$ take only applicants with relative ranks of 1 on both attributes. From applicant h^* to n take any applicant with a relative rank of one on at least one attribute. For this problem, as $n \to \infty$, h^*/n and the associated success probability both converge to .50. Ferguson (1992) generalized the problem presented by Gnedin by allowing dependencies between the attributes, and showed that the optimal policy has the same threshold form as the standard single attribute CSP. Samuels and Chotlos (1987) extended the rank minimization problem of Chow et al. (1964). They sought an optimal policy for minimizing the sum of two ranks for independent attributes. The rank sum minimization problem they studied is equivalent to the MASP in which $\pi^1(a) = \pi^2(a) = n - a$. The MASP is considerably more general than these previous problems, as it only constrains the payoff functions to be nondecreasing in the quality of the selected applicant's attributes. Next, we describe a procedure for computing optimal policies for the MASP.

3.2. A procedure for computing optimal policies

The probability that the ith attribute of the jth applicant whose relative rank on that attribute is r_j^i has an absolute (overall) rank of a_j^i is given by (Lindley, 1961)

$$Pr(A_j^i = a_j^i | R_j^i = r_j^i) = \frac{\binom{a_j^i - 1}{r_j^i - 1}\binom{n - a_j^i}{j - r_j^i}}{\binom{n}{j}}, \tag{3}$$

where $r_j^i \leqslant a_j^i \leqslant r_j^i + (n - j)$; otherwise, $Pr(A_j^i = a_j^i | R_j^i = r_j^i) = 0$. We assume that the k attributes are pairwise independent; that is, $Pr(A_j^i = a \wedge A_j^{i'} = a') = Pr(A_j^i = a)\, Pr(A_j^{i'} = a')$ for any pair of attributes i and i'. Therefore, the expected payoff for selecting the jth applicant is

$$E(\Pi_j | \mathbf{r}_j) = \sum_{i=1}^{k} \sum_{a_j^i = r_j^i}^{n} Pr(A_j^i = a_j^i | R_j^i = r_j^i) \pi^i(a_j^i). \tag{4}$$

The stipulation in assumption 1 of the MASP that the attributes be uncorrelated may seem unduly restrictive. Without it, however, Ferguson (1992) showed that the problem would likely be intractable. Consider a case with just $k = 2$ attributes. The first problem is defining what one means by correlated attributes. One possible method for generating correlated absolute ranks is to sample n pairs of values (Z_j^1, Z_j^2) $(j = 1, \ldots, n)$, called *worths*, from a bi-variate distribution with mean μ, variance σ^2, and correlation coefficient ρ. One could then generate absolute ranks within an attribute using the sampled worths. Specifically, the absolute rank of the jth applicant's first attribute, a_j^1, would be the rank of its worth Z_j^1 among the worths $Z_1^1, \ldots, Z_n^1$. Likewise for the second attribute. And the relative ranks could then be generated on the basis of the absolute ranks. Under this scheme, the resulting vectors of relative ranks $\mathbf{r}_1, \ldots, \mathbf{r}_n$ are not necessarily independent, as they are when the independence of attributes assumption is met. Consequently, Ferguson showed that the probability that some applicant j with relative ranks $\mathbf{r}_j$ has particular absolute ranks (a_j^1, a_j^2) depends on the history of observed relative ranks, i.e., on $\mathbf{r}_1, \ldots, \mathbf{r}_{j-1}$. To complicate matters further, these probabilities depend on the particular distribution from which the worths are sampled. The bottom line, then, is that going beyond uncorrelated attributes quickly leads one to a problem that is generally intractable and would nearly be impossible to communicate to an experimental subject.

We desire a policy that maximizes expected payoff in the MASP. The expected payoff for following such a policy is denoted by V^*. Following convention, the expected payoff for following the optimal policy from stage j to n is denoted by V_j^*. Hence, $V^* = V_1^*$.

At each stage j of the decision problem, the DM must decide to accept or reject an applicant knowing only the applicant's relative ranks $\mathbf{r}_j$. We represent a decision policy for each stage j as a set of acceptable $\mathbf{r}_j$ for that stage, which we denote by $\mathbf{R}_j$. Under the stage policy $\mathbf{R}_j$, the DM stops on an applicant with relative ranks $\mathbf{r}_j$ if and only if $\mathbf{r}_j \in \mathbf{R}_j$. The global policy is just the collection of stage policies $\mathbf{R} = \{\mathbf{R}_1, \ldots, \mathbf{R}_n\}$. By Bellman's (1957) Principle of Optimality, for an optimal (global) policy $\mathbf{R}^*$, each sub-policy $\{\mathbf{R}_j, \ldots, \mathbf{R}_n\}$ from stage j to n must also be optimal. Given this property, we can find the optimal policy using dynamic programming methods by working backward from stage n to stage 1. A procedure for constructing optimal stage policies $\mathbf{R}_j^*$ follows from Proposition 1, which we present below. To simply exposition, we first make the following assumption:

Assumption 1. When the expected value of stopping at stage j equals the expected value of continuing to stage $j+1$ and behaving optimally thereafter, the optimal DM stops at j.

Proposition 1. $\mathbf{r} \in \mathbf{R}_j^* \Leftrightarrow E(\Pi_j | \mathbf{r}) \geqslant V_{j+1}^*$.

Proof. Suppose that $E(\Pi_j | \mathbf{r}) > V_{j+1}^*$ for some $\mathbf{r} \notin \mathbf{R}_j^*$. Therefore, rejecting this $\mathbf{r}$ entails moving to $j+1$ where the expected payoff, V_{j+1}^*, is strictly less than stopping on j. Hence, by the Principle of Optimality, this $\mathbf{r}$ must be in $\mathbf{R}_j^*$. Now, suppose that $E(\Pi_j | \mathbf{r}) < V_{j+1}^*$ for some $\mathbf{r} \in \mathbf{R}_j^*$. Then, the DM will stop on this $\mathbf{r}$ when continuing the search has a higher expected value, V_{j+1}^*. Thus, by the Principle of Optimality, $\mathbf{R}_j^*$ cannot be optimal if it contains this $\mathbf{r}$. By Assumption 1, when $E(\Pi_j|\mathbf{r}) = V_j^*$, this $\mathbf{r} \in \mathbf{R}_j^*$. Therefore, $\mathbf{r} \in \mathbf{R}_j^*$ if and only if $E(\Pi_j|\mathbf{r}) \geqslant V_{j+1}^*$. □

Proposition 2. $\mathbf{r} \in \mathbf{R}_j^* \Rightarrow \mathbf{r} \in \mathbf{R}_{j+1}^*$.

We omit the proof of Proposition 2 as it follows directly from Corollary 2.1b in Mucci (1973). Proposition 2 tells us that if it is optimal to stop at stage j when one observes $\mathbf{r}$, then it is optimal to stop when one observes $\mathbf{r}$ in the next stage; by induction, then, it is optimal to stop given $\mathbf{r}$ in *all* subsequent stages. This property will be useful below because it allows us to represent the optimal policies rather compactly.

Since the DM must accept the nth applicant, if reached,

$$V_n^* = n^{-1} \sum_{i=1}^{k} \sum_{a_j^i=1}^{n} \pi^i(a_j^i). \tag{5}$$

The expected payoff for the last applicant ($j = n$) under the optimal policy (or any other permissible policy) is simply the payoff one expects for selecting an applicant at random. The expected payoff for following the optimal stage $j<n$ policy and then following the optimal policy thereafter is expressed by the functional equation

$$V_j^* = Q(\mathbf{R}_j^*)E(\Pi_j|\mathbf{R}_j^*) + [1 - Q(\mathbf{R}_j^*)]V_{j+1}^*, \tag{6}$$

where $E(\Pi_j|\mathbf{R}_j^*) = |\mathbf{R}_j^*|^{-1} \sum_{\mathbf{r}\in\mathbf{R}_j^*} E(\Pi_j|\mathbf{r})$ is the expected payoff for stopping at stage j under the optimal stage j policy, and $Q(\mathbf{R}_j^*) = |\mathbf{R}_j^*|/k^j$ is the probability of stopping on j under the optimal stage j policy. (The numerator on the right-hand side of the $Q(\mathbf{R}_j^*)$ function corresponds to the number of relative rank profiles that produce a stopping decision at stage j, and the denominator to the number of feasible relative rank profiles at stage j.)

Given V_n^*, working backward from stage $n-1$ to stage 1 by alternating between the application of Proposition 1 and the computation of Eq. (6), the optimal global policy $\mathbf{R}^*$ is easily constructed.

Denoting the applicant position at which the search is terminated by m, the expected stopping position is

$$E(m|\mathbf{R}^*) = 1 + \sum_{j=1}^{n-1} \left(\prod_{h=1}^{j} [1 - Q(\mathbf{R}_h^*)] \right). \tag{7}$$

Eq. (7) will be useful below when we discuss the behavior of actual DMs in the MASP.

3.3. An example of a MASP and the application of its optimal policy

An example of an instance of the MASP for a case in which $n = 6$ and $k = 2$ is shown in Table 2. The payoffs for each a for each attribute i are shown in the top panel.

Table 2
An example of a MASP with $n = 6$ and $k = 2$

Payoff values						
a	1	2	3	4	5	6
$\pi^1(a)$	6	5	4	3	2	1
$\pi^2(a)$	5	4	3	2	0	0
Example applicant sequence						
Applicant (j)	1	2	**3**	4	5	6
a_j^1	2	4	**3**	6	5	1
a_j^2	5	2	**1**	3	6	4
r_j^1	1	2	**2**	4	4	1
r_j^2	1	1	**1**	3	5	4
Optimal policy and payoffs						
Applicant (j)	1	2	**3**	4	5	6
V_{j+1}^*	7.82	7.67	**7.37**	6.83	5.83	–
$E(\Pi_j\|\mathbf{r}_j)$	5.83	5.73	**7.55**	2.93	1.83	8.00
Π_j	5.00	7.00	**9.00**	4.00	2.00	8.00

See text for explanation.

The center panel displays the absolute and relative ranks of each applicant. Applicant 1 has absolute ranks of 2 and 5 on attributes 1 and 2, respectively; her relative ranks are 1 for both attributes. Applicant 2 has absolute ranks of 4 and 2, and therefore relative ranks of 2 and 1, for attributes 1 and 2, respectively, etc. The bottom panel displays the value of the optimal policy for each applicant position (stage) and the expected payoffs for selecting each applicant j. Under the optimal policy, the expected earnings are $V_1^* = 7.82$.

Consider how the optimal policy would be applied here. For applicant 1, the DM should stop only if the expected payoff for selecting the first applicant meets or exceeds 7.82. However, since the expected payoff for the first applicant will always be 5.83 because her relative ranks will *always* be 1, the DM will never stop on the first applicant. For applicant 2, the expected payoff for selection must not be less than 7.67 for the DM to make a selection; hence, the DM will stop only when the second applicant has relative ranks of 1 on both attributes (because $E(\Pi_2|(1,1)_2) = 8.26$; $E(\Pi_2|(1,2)_2) = 5.93$; $E(\Pi_2|(2,1)_2) = 5.73$; and $E(\Pi_2|(2,2)_2) = 3.40$). In this example, the optimal policy dictates that the DM stop on applicant 3 because $E(\Pi_3|(3,1)_3) = 7.55 > V_4^* = 7.37$. Since $\mathbf{a}_3 = (3,1)$, the DM earns $\Pi_3 = 9.00$ for her selection. Fortunately for her, in this instance she could not have earned more by selecting any other applicant.

Next we describe two experiments in which we tested the predictions of the optimal search policy with actual subjects. After describing the experiments and their results, we describe some implications and discuss future directions for this line of research.

4. Experiments 1 and 2

Given the similarity of Experiments 1 and 2, we will report them together. Experiment 1 used *symmetric* payoffs. That is, the attributes contribute equally ($\pi^1(a) = \pi^2(a)$, $\forall a$) to the selection payoff. The symmetric payoff scheme corresponds to scenarios in which the DM gives equal weight to the attributes of the options through which she is searching. For example, in hiring for a faculty position at some universities the search committees might give equal weight to applicants' teaching and research accomplishments. At other universities significantly more weight might be given to teaching than to research. The latter scenario captures the *asymmetric* payoff scheme used in Experiment 2. Under it, one of the attributes contributes more ($\pi^1(a) \geqslant \pi^2(a)$, $\forall a$) to the DM's payoff than the other. We ask: How well do actual DMs search in multi-attribute sequential search problems? And what are the basic properties of the search policies that actual DMs employ?

4.1. Method

4.1.1. Subjects

Thirty subjects participated individually in Experiment 1; the same number participated in Experiment 2. All of them were University of Arizona students recruited by advertisements asking for volunteers to participate in a decision making experiment with payoffs contingent on performance. The mean payoff per session, which typically lasted 40–60 min, was \$21 (minimum \$5, maximum \$50) for Experiment 1, and \$20 (minimum \$5, maximum \$40) for Experiment 2. In addition to the monetary payoff, subjects received class credit for their participation if they requested it.

4.1.2. Procedure

The instructions (hard copy) explained the MASP in detail, placing special emphasis on the computation of the relative ranks with the presentation of a new applicant. In the instructions, the subjects read through an example with $n = 6$ applicants, $k = 2$ attributes, and the same payoff scheme used in the experiment (symmetric payoffs for Experiment 1, and asymmetric payoffs for Experiment 2; see Table 3). The example explained and illustrated the updating of the relative ranks of each attribute for each applicant.

That attributes were uncorrelated was carefully communicated to the subjects, as the independence assumption is crucial in determining the optimal decision policy. Specifically, the subjects were told:

> In the task you will be asked to perform, **the two attributes are uncorrelated**; that is, knowing the value of one attribute for an applicant tells you nothing about the value of the applicant's other attribute. The two are not related. It is possible that one applicant will be the best on Attribute 1 and Attribute 2; it is equally possible that the applicant is best on Attribute 1 and worst on Attribute 2, or the 3rd best on Attribute 1 and the 17th best on Attribute 2, etc. [Bold in the original instructions.]

Table 3
Payoff schemes used in Experiments 1 and 2

Experiment 1						
a	1	2	3	4	5	6–30
$\pi^1(a)$	25	12	8	4	2	0
$\pi^2(a)$	25	12	8	4	2	0
Experiment 2						
a	1	2	3	4	5	6–30
$\pi^1(a)$	25	12	8	4	2	0
$\pi^2(a)$	15	8	4	2	1	0

Payoffs are in US dollars.

Once the subjects understood the instructions, they were seated at individual computers. They then performed two practice problems to acquaint them with the computer-controlled task. The experimental problems were presented once the subjects completed these two practice problems.

Each subject completed 100 trials (replications) of the MASP with $n = 30$ applicants and $k = 2$ attributes. The orderings of the absolute ranks for each attribute were generated randomly and independently for each subject and each trial. The payoff structure (fully described below) was fixed over all trials and each trial was structured in the same way: the relative ranks of applicant j on two attributes were displayed, and then the subject was allowed to either select the applicant, thereby terminating the search, or proceed and observe a new applicant. If she decided to continue the search on applicant j ($j = 1, \ldots, n-1$), then the relative ranks of all j applicants that had been observed and rejected were updated and displayed. If she opted not to stop the search, then she was forced to accept the nth applicant. When the subject stopped the search, thereby terminating the trial, all the n absolute ranks for both attributes and their corresponding relative ranks were displayed on the computer screen. In this way, subjects who stopped the search on different periods were provided with full information about the actual sequences of absolute ranks of all the n applicants.

4.1.3. Payoff structure

Experiment 1. Experiment 1 (symmetric payoffs) had the same payoff structure for both attributes. Specifically, selecting an applicant with an absolute rank of 1 on either attribute contributed \$25 to a DM's payoff; selecting ones with 2, 3, 4, or 5 contributed \$12, \$8, \$4, \$2, respectively. For absolute ranks 6–30, they earned nothing.

Experiment 2. The payoff structure in Experiment 2 (asymmetric payoffs) was not the same for both attributes. The DM's payoff was more heavily influenced by the selected applicant's absolute rank on attribute 1 than her rank on attribute 2. The payoffs used for the attribute 1 were the same as in Experiment 1. For attribute 2, however, the subjects earned \$15 for selecting an applicant whose absolute rank on that attribute was 1; \$8 for selecting an applicant whose second attribute had an absolute rank of 2; and \$4, \$2, and \$1 for selecting applicants whose absolute ranks on attribute 2 were 3, 4, and 5, respectively. For absolute ranks greater than 5 on attribute 2 they earned nothing.

Subjects were paid for a single randomly selected trial, which they determined for themselves by drawing a number from a hat. Hence, the subjects could earn as much as \$50 (\$40) and as little as \$0 (\$0) in Experiment 1 (Experiment 2).

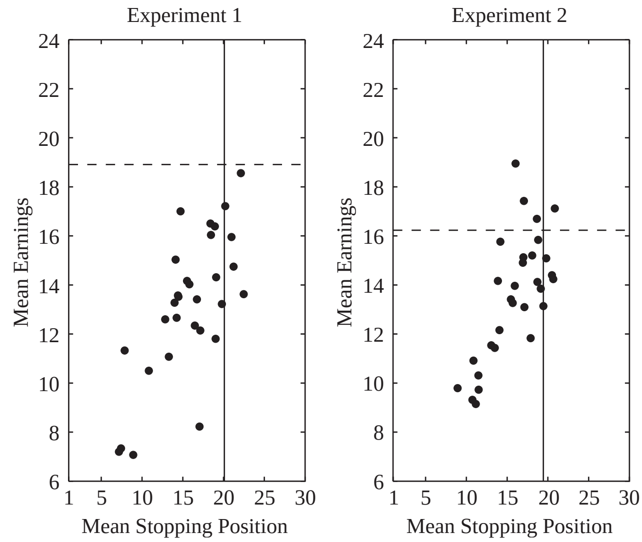

Fig. 1. Mean earnings as a function of mean stopping position for Experiments 1 and 2. The vertical (solid) line in each plot corresponds to the expected stopping position under the optimal policy. The horizontal (dotted) line in each represents the expected earnings under the optimal policy.

4.2. Results

4.2.1. Earnings

Experiment 1. Under the optimal policy, a DM expects to earn $V_1^* = 18.91$. (All payoffs are in US dollars. We omit the dollar signs.) Taking the mean earnings for each subject over all 100 trials ($M = 13.13$, $SD = 2.97$) and comparing these to the expected payoff under optimal play, we find that the actual payoffs are significantly smaller, $t(29) = 10.57$, $p < .001$.

Experiment 2. The optimal DM can expect to earn $V_1^* = 16.23$. Computing the mean earnings for each subject ($M = 13.53$, $SD = 2.53$), we find that the mean empirical payoff is significantly smaller than the expected optimal payoff, $t(29) = 5.83$, $p < .001$.

4.2.2. Stopping position

Experiment 1. We computed the mean stopping positions over the 100 trials of the MASP and compared them to the expected stopping position under the optimal policy, $E(m|\mathbf{R}^*) = 20.09$, which results from the application of Eq. (7). The mean observed stopping position ($M = 15.89$, $SD = 4.29$) was significantly smaller than expected under the optimal policy, $t(29) = 5.48$, $p < .001$. On average, the subjects stopped the search about four observations shorter than expected under the optimal policy.

The linear correlation between the subjects' mean stopping position and their mean earnings was positive and significant ($r = .72$, $p < .001$). A scatterplot of the relationship is presented in the left panel of Fig. 1. Subjects who tended to search longer also tended to earn higher payoffs; however, in all cases, the mean earnings are below those expected under the optimal policy. Hence, a reasonable inference is that even those subjects who tended to search, on average, about the same amount as expected used policies that differed from the optimal policy. Below, we discuss in more detail the nature of the subjects' policies.

Experiment 2. The expected stopping position under the optimal policy is $E(m|\mathbf{R}^*) = 19.45$. As in Experiment 1, the mean stopping position ($M = 15.90$, $SD = 3.38$) was significantly smaller than expected under optimal search, $t(29) = 5.75$, $p < .001$. Once again, the correlation between observed mean stopping position and mean payoff was positive and significant ($r = .71$, $p < .001$). As the subjects tended to search less, they earned less (see Fig. 1, right panel).

4.2.3. Evidence of learning

Experiment 1. We first searched for evidence of learning by regressing the mean earnings for each trial onto the trial numbers. The slope of the regression line was positive and significant, $b = .0145$, $R^2 = .31$, $p = .049$, indicating that earnings increased with experience. However, the increase in earnings is quite mild across trials, and still well below the expected earnings in the final trials (Fig. 2, left panel). The subjects also tended to search longer with experience. Regressing the mean stopping position on each trial onto the trial numbers, we find that the slope of the regression line is

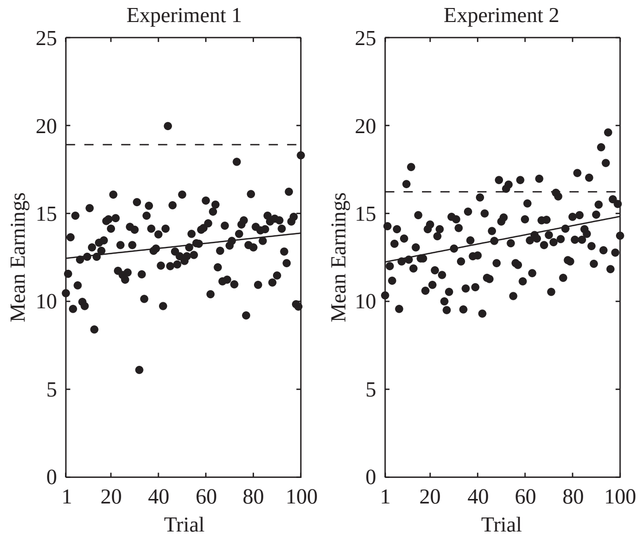

Fig. 2. Mean earnings (over subjects) as a function of trial for Experiments 1 and 2. The solid lines are based on OLS best fits. The horizontal (dotted) line in each represents the expected earnings under the optimal policy.

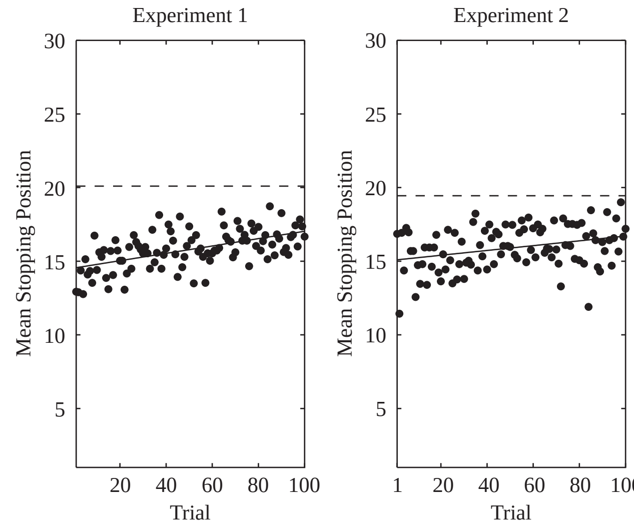

Fig. 3. Mean stopping position (over subjects) as a function of trial for Experiments 1 and 2. The solid lines are based on OLS best fits. The horizontal (dotted) line in each represents the expected stopping position under the optimal policy.

significantly positive, $b = .025$, $p < .001$, though the fit is rather poor $R^2 = .04$.[1] The mean stopping positions over trials are displayed in Fig. 3. The subjects seem to have learned that searching longer improves payoffs.

Experiment 2. Earnings tended to increase with experience, as evidenced by the significant positive slope of the earnings regressed onto trial number, $b = .0262$, $R^2 = .12$, $p < .001$. The trend is greater than observed in Experiment 1. In fact, the earnings in late trials are quite close to the optimal expected earnings (see Fig. 2, right panel). The increase in earnings over trials seems to have been driven, at least in part, by the subjects learning to search deeper into the options before making a selection. The slope of the regression line of mean stopping position onto trial is positive and significant, $b = .017$, $R^2 = .10$, $p < .001$. The mean stopping position over trials is exhibited in the right panel of Fig. 3.

4.2.4. *Estimated policies*

By Proposition 2, the optimal policy can be represented by a set of cutoffs for each feasible pair of relative ranks. The cutoffs dictate at which applicant positions it becomes optimal to select applicants with different sets of relative ranks. Specifically, under this representation, the optimal DM stops on a pair of relative ranks ($r_j^1 = x, r_j^2 = y$) if and only if the cutoff for (x, y) has been reached, i.e., if $c^*_{x,y} \geqslant j$. The $c^*_{x,y}$ are ordered such that $c^*_{x,y} \leqslant c^*_{x,y+1}$ and $c^*_{x,y} \leqslant c^*_{x+1,y}$. That is, under the optimal policy, represented by the set of all optimal cutoffs $\mathbf{c}^*$, the threshold for a pair of relative ranks cannot be below the threshold for another pair of ranks that is strictly better. In this section, we describe a procedure for estimating the empirical cutoffs $\hat{\mathbf{c}}$ for each subject from the experimental choice data.

We treat the estimation as a minimization problem. For each subject and each trial, we have a set of continue and stop decisions. On trial t, the decision on applicant j with relative ranks (x, y) to stop is denoted by $\delta(t, j, x, y) = 1$ and to continue by $\delta(t, j, x, y) = 0$. For a set of cutoffs $\hat{\mathbf{c}}$, we denote the corresponding predicted decisions by $\hat{\delta}(t, j, x, y)$. Precisely, $\hat{\delta}(t, j, x, y) = 1$ for some policy $\hat{\mathbf{c}}$ if and only if $\hat{c}_{x,y} \geqslant j$; otherwise, $\hat{\delta}(t, j, x, y) = 0$. For each subject, our objective is to find a set of cutoffs $\hat{\mathbf{c}}$ that minimizes

$$\frac{\sum_{t=1}^{100} \sum_{j=1}^{m_t} |\delta(t, j, x, y) - \hat{\delta}(t, j, x, y)|}{\sum_{t=1}^{100} m_t}, \tag{8}$$

subject to $\hat{c}_{x,y} \leqslant \hat{c}_{x,y+1}$ and $\hat{c}_{x,y} \leqslant \hat{c}_{x+1,y}$, and where m_t denotes the applicant position on which the subject stopped on trial t. In words, Eq. (8) simply returns the proportion of decisions made by a subject that are incompatible with the cutoffs $\hat{\mathbf{c}}$. For shorthand, we refer to this objective function as the *violation function*.

To minimize the number of violations (Eq. (8)), we used the threshold accepting (TA) algorithm developed by Dueck & Scheuer (1990), which is an extension of the better-known simulated annealing algorithm. The set of feasible cutoff sets is enormous, making enumeration (brute-force) infeasible. Further, our problem is not convex and more traditional optimization procedures

[1]Due to the stochastic nature of the problem instances, the R^2 values will tend to be depressed. As we report these regressions merely to demonstrate the general trends in the data, and not as a serious model-fitting exercise, the precise model fits are of minor importance.

are not applicable. TA allows us to efficiently search the space and also to avoid local minima, which tend to plague many combinatorial optimization problems (Kirkpatrick, Gelatt, & Vecchi, 1983). We refer the reader to Dueck & Scheuer (1990) for a complete description of the algorithm.

Since all relative ranks greater than 5 entailed 0 payoffs, we converted all of these to relative rank of 6. This allows us to estimate considerably fewer cutoffs.

Experiment 1. Since the payoffs are symmetric in Experiment 1, we further constrained the search to policies such that $\hat{c}_{x,y} = \hat{c}_{y,x}$, i.e., we imposed symmetry on the estimated cutoffs. In sum, then, the procedure required the estimation of 21 cutoffs.

The median estimated $\hat{c}_{x,y}$ $(x, y = 1, \ldots, 6)$ are displayed in the center panel of Table 4. The mean value of the violation function for the estimated cutoffs was rather small $M = .034$ $(SD = .017)$; on average, the estimated cutoffs predicted more than 96% of the subjects' decisions.

The difference panel (bottom panel) in Table 4 is quite telling. First, note that most of the differences between the observed and optimal cutoffs are negative, indicating that the subjects' cutoffs were generally shifted toward stopping too early. The differences are most negative for intermediately small pairs of relative ranks (e.g., $(2,2)$, $(2,3)$, $(3,4)$, etc.), indicating a strong bias to stop early on these pairs. The estimated cutoff for applicants whose relative ranks are both 1 is neither too early nor too late. In contrast, the subjects' tended to pass up applicants with one good attribute $(r = 1)$ and one poor one $(r \geqslant 6)$, when stopping had a higher expected payoff. Therefore, the observed early stopping seems to be largely driven by the subjects' strong tendency to stop early on "middle quality" pairs of relative ranks. This observation is confirmed by an analysis of the actual probabilities of stopping on applicants with each pair of relative ranks. The subjects tended to stop considerably more often on applicants with intermediately small relative ranks than is dictated by the optimal policy. Further, they stopped less often than they should have (by the optimal policy) for pairs of $(1, r \geqslant 6)$.

Of course, it is possible that the subjects used policies of a different form. One reviewer suggested that they might be averaging the relative ranks of each applicant and using a cutoff rule for the averaged rank: select applicant j if and only if $(r_j^1 + r_j^2)/2 \leqslant s_j$, where s_j is the stage j cutoff. To test this possibility, we compared the probability of stopping on applicants with different average ranks. The results showed conclusively that the probability of stopping on applicants with average ranks of $x = (r^1 + r^2)/2$ decreased as the (absolute) difference in r^1 and r^2 decreased. Consider relative rank pairs with an average of 3: $(1,5)$, $(2,4)$, and $(3,3)$. Under the averaging hypothesis, the probability of stopping on a applicant with these relative rank pairs from position 5 onward should be the same. (Pairs before 5 must be disregarded because $(1,5)$ pairs are never observed before position 5.) The stopping probabilities are, however, .21, .14, and .12 for pairs $(1,5)$, $(2,4)$, and $(3,3)$, respectively. And for *all* average ranks, this decreasing pattern obtained. Thus, the subjects were sensitive to the particular values of *each* of the applicants' relative ranks and not just to the mean of their ranks.

Table 4
Optimal and empirical (estimated) cutoffs for Experiment 1

		Optimal cutoffs r^2					
		1	2	3	4	5	6
r^1	1	7	12	14	15	16	16
	2	12	19	22	24	25	26
	3	14	22	25	27	27	28
	4	15	24	27	28	29	29
	5	16	25	27	29	30	30
	6	16	26	28	29	30	30
		Empirical cutoffs r^2					
		1	2	3	4	5	6
r^1	1	7	9	10	12	13	18
	2	9	11	12	16	18	25
	3	10	12	16	20	22	27
	4	12	16	20	22	25	29
	5	13	18	22	25	29	30
	6	18	25	27	29	30	30
		Empirical–optimal cutoffs r^2					
		1	2	3	4	5	6
r^1	1	0	−4	−5	−3	−3	2
	2	−4	−9	−10	−9	−7	−1
	3	−5	−10	−9	−8	−5	−2
	4	−3	−9	−8	−6	−4	0
	5	−3	−7	−5	−4	−1	0
	6	2	−1	−2	0	0	0

The estimated cutoffs are based on the median cutoff taken over subjects. The bottommost panel shows the difference in the median empirical and optimal cutoff for each pair of relative ranks. Note that a negative difference obtains when the empirical cutoff is placed *before* the optimal cutoff (too early); the difference is positive when the empirical cutoff is located *after* the optimal cutoff (too late).

Experiment 2. We used the TA procedure described above to estimate the policy cutoffs $\hat{\mathbf{c}}$ for each subject in Experiment 2. However, since the payoffs are asymmetric, we did not constrain the thresholds to be symmetric, thereby requiring us to estimate 36 thresholds for each subject. The median cutoffs are displayed in Table 5. First, note that the subjects' cutoffs reveal that their policies are sensitive to the attribute weights.

Table 5
Optimal and empirical (estimated) cutoffs for Experiment 2

		Optimal cutoffs r^2					
		1	2	3	4	5	6
r^1	1	7	11	12	13	13	13
	2	14	19	22	23	24	24
	3	17	23	25	26	27	27
	4	19	25	27	28	29	29
	5	20	26	28	29	30	30
	6	21	27	29	30	30	30
		Empirical cutoffs r^2					
		1	2	3	4	5	6
r^1	1	7	7	12	12	21	21
	2	10	13	17	20	25	25
	3	14	20	22	25	25	27
	4	17	24	25	25	28	28
	5	22	26	28	29	30	30
	6	25	30	30	30	30	30
		Empirical– optimal cutoffs r^2					
		1	2	3	4	5	6
r^1	1	0	−4	0	−1	8	8
	2	−4	−6	−5	−3	1	1
	3	−3	−3	−3	−1	−2	0
	4	−2	−1	−2	−3	−1	−1
	5	2	0	0	0	0	0
	6	4	3	1	0	0	0

The estimated cutoffs are based on the median cutoff taken over subjects. The bottommost panel shows the difference in the median empirical and optimal cutoff for each pair of relative ranks. Note that a negative difference obtains when the empirical cutoff is placed *before* the optimal cutoff (too early); the difference is positive when the empirical cutoff is located *after* the optimal cutoff (too late). Recall that r^1 corresponds to the more heavily weighted attribute.

In no case is the cutoff for a given pair of relative ranks $(r^1, r^2) = (x, y)$ for $x \leqslant y$ greater than the cutoff for (y, x), and in most cases the cutoffs for (x, y) are smaller. This provides strong evidence that the subjects were giving more weight to the more important (higher payoff) attribute, as they ought to. The violation function for the estimated cutoffs was again quite small $M = .042$ ($SD = .014$), but slightly larger than in Experiment 1.

The general pattern of departures of the estimated cutoffs from the optimal cutoffs is quite similar to the one in Experiment 1 (compare Tables 4 and 5). Much of the early stopping in Experiment 2 is driven by the subjects' tendency to stop on intermediately small pairs of relative ranks. Once again, we find that the subjects' cutoffs for applicants with one good attribute ($r = 1$) and one poor one ($r \geqslant 6$) are shifted considerably toward later applicants. Taken together, the estimated cutoffs again suggest that the subjects are strongly biased to select applicants whose relative ranks may both entail positive payoffs; conversely, the subjects tend to be biased against selecting applicants for whom at least one attribute will certainly result in zero payoff. A comparison of the cutoffs for relative ranks $(1, 6)$ and $(6, 1)$ suggests that the subjects assigned disproportionate weight to the less important attribute. They should accept applicants with relative ranks $(1, 6)$ starting with the 13th applicant but do not tend to do so until the 21st applicant—eight applicant positions too late. The bias is much smaller for $(6, 1)$: the subjects should take these applicants starting on the 21st applicant and begin to do so just four applicants later.

4.3. Discussion

Consistent with previous experimental studies of secretary problems (e.g., Bearden et al., 2004; Seale & Rapoport, 1997, 2000; Zwick et al., 2003), in both experiments, we find that DMs in the MASP tend to terminate their search too early relative to the optimal policy. But our results allow us to say more than this. We find that the tendency to terminate the search too early is mostly driven by the DMs stopping prematurely on intermediately small relative ranks. Taken together with the finding that the DMs tend to search beyond applicants with one good ($r = 1$) attribute and one poor ($r \geqslant 6$) one when they ought not, it seems that they are giving considerable (disproportionate) weight to selecting an applicant who is "acceptable" on both attributes, where acceptable is defined as contributing a nonzero amount to the selection payoff.

This pattern of behavior is consistent with the use of a modified satisfying rule (Simon, 1955). The subjects seem to be searching for applicants who are acceptable on both attributes (i.e., both attributes can lead to positive payoffs); however, they do not seem to have a strict set of aspiration levels: they do tend to stop sooner on applicants with smaller pairs of relative ranks. When the relative ranks are both below 6 and can therefore both entail positive payoffs, the subjects do make trade-offs and behave in a way consistent with a form of optimization (though the behavior is still suboptimal with respect to the optimal policy). As soon as one relative rank entails zero payoffs for that attribute, the decision rule seems to become non-compensatory—subjects do not tend to make the same sorts of trade-offs in these cases. The estimated cutoff policies account for the data remarkably well. In Experiment 1, the estimated policies captured around 96% of the subjects' decisions; in Experiment 2, the estimated policies captured around 95% of the decisions.

The pattern of results suggests that the subjects are not following some sort of simple heuristic in which they use the average of relative ranks to make their decisions.

Their policies seem to have been more complicated, relying on the particular relative rank profiles of the observed applicants. Without resorting to more complicated decision rules (e.g., ones with additional free parameters), we think it is doubtful that we can better account for these results.

5. Conclusion

We began this paper by presenting an extension of the secretary problem in which the DM searches through applicants who vary on several dimensions. In the standard (single-attribute) versions of the secretary problem (the CSP and the GSP), the DM is not faced with the dilemma—inherent to many decisions—of making trade-offs among the attributes of the decision alternatives. When hiring an administrative assistant, for example, it is not unlikely that applicants who are good in one domain (e.g., using a database) are less qualified in another domain (e.g., proofreading complex documents). As a result, the person making hiring decisions must make trade-offs among the attributes of the applicants. The multi-attribute secretary problem (MASP) that we introduced here captures important properties of these kinds of search problems.

Results from the two experiments suggest that subjects facing similar problems may behave suboptimally and exhibit predictable biases. Most notably, consistent with findings from the study of behavior in single-attribute secretary search problems (e.g., Bearden et al., 2004; Seale & Rapoport, 1997, 2000; Zwick et al., 2003), subjects tend to search insufficiently through applicants. Further, in problems like the MASP, they may make poor trade-off decisions within applicants, preferring applicants who are mediocre on all attributes to those who excel on one and are poor on others, even when the expected reward for the latter is greater. Perhaps in a number of real-world situations, however, this bias would actually be beneficial. It would make little sense to hire a database genius who introduces errors into legal documents that could result in considerable cost to a company. In future work, we intend to generalize the MASP to allow for noncompensatory payoff functions that capture these situations. One possibility is to render the payoff a function of the product of attribute values. We are currently developing methods for computing optimal policies for this and other extensions of the MASP.

As formulated here, the attributes in the MASP are pairwise uncorrelated. There are many situations in which this assumption is not likely to be violated. Intentionally, the examples used throughout this paper have involved attributes that we suspect are at most weakly correlated, such as technical and interpersonal skills. Of course, there are many other scenarios in which one would expect the attributes to be correlated. Proofreading and writing abilities, for example, are presumably related. Unfortunately, based on the results in Ferguson (1992), it looks like it would be difficult to generalize the MASP to allow for correlated attributes, since computing the joint probabilities of absolute rank profiles given relative rank profiles (analogous to what is computed in Eq. (3) for independent attributes) is quite challenging. The main problem is that the joint probabilities can depend on the entire history of observed relative ranks, which is not true when the attributes are uncorrelated. One possible way around this is to formulate a memoryless version of the problem in which we assume that the DM only knows the relative ranks of the jth applicant, and not those of the previous $j-1$ applicants. Perhaps this constraint on the MASP will make problems with correlated attributes more tractable.

Another useful way to extend the MASP is to impose a cost on learning the value of the applicant's attributes. When hiring for a position, it can be costly to acquire information about each applicant. One must pay for background checks, personality tests, computer skills tests, and so on. Presumably it makes little sense to pay for a computer skills test after one learns that an applicant is psychotic. How, then, should one go about deciding when to gather more information on an applicant? By adding a search problem within applicants, as well as across applicants (as in the standard MASP), we may be able to bring these abstract decision problems closer to the ''real world.''

Acknowledgments

We gratefully acknowledge financial support by a contract F49620-03-1-0377 from the AFOSR/MURI to the Department of Systems and Industrial Engineering and the Department of Management and Policy at the University of Arizona.

References

Bearden, J. N., & Murphy, R. O. (2004). On the generalized secretary problem. Unpublished manuscript.

Bearden, J. N., Rapoport, A., & Murphy, R. O. (2004). Sequential oberservation and selection with rank-dependent payoffs: An experimental test. Unpublished manuscript.

Bellman, R. (1957). *Dynamic programming*. Princeton, NJ: Princeton University Press.

Chow, Y. S., Moriguti, S., Robbins, H., & Samuels, S. M. (1964). Optimal selection based on relative rank (the ''secretary problem''). *Israel Journal of Mathematics*, *2*, 81–90.

Corbin, R. M., Olson, C. R., & Abbondanza, M. (1975). Context effects in optimal stopping rules. *Behavior and Human Performance*, *14*, 207–216.

Cox, J. C., & Oaxaca, R. L. (1989). Laboratory experiments with a finite horizon job-search model. *Journal of Risk and Uncertainty*, *2*, 301–330.

Dueck, G., & Scheuer, T. (1990). Threshold accepting: A general purpose optimization algorithm appearing superior to simulated annealing. *Journal of Computational Physics*, *90*, 161–175.

Ferguson, T. S. (1989). Who solved the secretary problem? *Statistical Science*, *4*, 282–296.

Ferguson, T. S. (1992). Best-choice problems with dependent criteria. *Contemporary Mathematics*, *125*, 135–151.

Freeman, P. R. (1983). The secretary problem and its extensions: A review. *International Statistical Review*, *51*, 189–206.

Gilbert, J., & Mosteller, F. (1966). Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, *61*, 35–73.

Gnedin, A. V. (1981). A multicriteria problem of optimal stopping of a selection process. *Automation and Remote Control*, *42*, 981–986.

Hey, J. D. (1981). Are optimal search rules reasonable? And vice versa? *Journal of Economic Behavior and Organization*, *2*, 47–70.

Hey, J. D. (1982). Search for rules of search. *Journal of Economic Behavior and Organization*, *3*, 65–81.

Hey, J. D. (1987). Still searching. *Journal of Economic Behavior and Organization*, *8*, 137–144.

Kahan, J. P., Rapoport, A., & Jones, L. V. (1967). Decision making in a sequential search task. *Perception and Psychophysics*, *2*, 374–376.

Kirkpatrick, S., Gelatt, C. D., & Vecchi, M. P. (1983). Optimization by simulated annealing. *Science*, *220*, 671–680.

Kogut, C. A. (1990). Consumer search behavior and sunk costs. *Journal of Economic Behavior and Organization*, *14*, 381–392.

Lindley, D. V. (1961). Dynamic programming and decision theory. *Applied Statistics*, *10*, 39–51.

Mucci, A. G. (1973). Differential equations and optimal choice problems. *Annals of Statistics*, *1*, 104–113.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. New York: Cambridge University Press.

Rapoport, A. (1969). Effects of observation cost on sequential search behavior. *Perception and Psychophysics*, *6*, 234–240.

Rapoport, A., & Tversky, A. (1970). Effects of search cost on sequential search behavior. *Perception and Psychophysics*, *6*, 234–240.

Samuels, S. M. (1991). Secretary problems. In: B. K. Gosh, P. K. Sen (Eds.), *Handbook of sequential analysis* (pp. 381–405). New York: Marcel Dekker.

Samuels, S. M., & Chotlos, B. (1987). A multiple criteria optimal selection problem. In: J. Van Ryzin (Ed.), *Adaptive statistical procedures and related topics*: *Proceedings of a symposium in honor of Herbert Robbins* (pp. 62–78). Upton, NY: Brookhaven National Laboratory.

Seale, D. A., & Rapoport, A. (1997). Sequential decision making with relative ranks: An experimental investigation of the secretary problem. *Organizational Behavior and Human Decision Processes*, *69*, 221–236.

Seale, D. A., & Rapoport, A. (2000). Optimal stopping behavior with relative ranks: The secretary problem with unknown population size. *Journal of Behavioral Decision Making*, *13*, 391–411.

Shapira, Z. (1981). Optional stopping on nonstationary series. *Organizational Behavior and Human Performance*, *27*, 32–49.

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, *69*, 99–118.

Sonnemans, J. (1998). Strategies of search. *Journal of Economic Behavior and Organization*, *35*, 309–332.

Sonnemans, J. (2000). Decisions and strategies in a sequential search experiment. *Journal of Economic Psychology*, *21*, 91–102.

Stigler, G. L. (1961). The economics of information. *Journal of Political Economy*, *69*, 213–225.

Tversky, A., & Koehler, D. (1994). Support theory: A nonextensional representation of subjective probability. *Psychological Review*, *101*, 547–567.

Zwick, R., Rapoport, A., Lo, A. K. C., & Muthukrishnan, A. V. (2003). Consumer sequential search: Not enough or too much? *Marketing Science*, *22*, 503–519.