

A Multi-Attribute Extension of the Secretary Problem: Theory and Experiments

J. Neil Bearden ^{*,1}, Ryan O. Murphy, Amnon Rapoport

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

Abstract

We present a generalization of a class of sequential search problems with ordinal ranks (“secretary” problems) in which applicants are characterized by multiple attributes that are evaluated independently. 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.

Key words: 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 attempt to employ an applicant who is outstanding on both; most likely, however, one will have to make trade-off decisions, perhaps by

* Corresponding author.

Email addresses: jneilb@gmail.com (J. Neil Bearden),
murphro@eller.arizona.edu (Ryan O. Murphy), amnon@u.arizona.edu (Amnon Rapoport).

¹ 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.

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 28 DEC 2004		2. REPORT TYPE		3. DATES COVERED 00-00-2004 to 00-00-2004	
4. TITLE AND SUBTITLE A Multi-Attribute Extension of the Secretary Problem: Theory and Experiments			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Arizona, Department of Management and Policy, Tucson, AZ, 85721			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT We present a generalization of a class of sequential search problems with ordinal ranks (secretary problems) in which applicants are characterized by multiple attributes that are evaluated independently. 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.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 24	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

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 et al., 1993, for an thorough 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.

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 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 and 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 and Oaxaca, 1989; Hey, 1981, 1982, 1987; Kogut, 1990; Rapoport, 1969; Rapoport and Tversky, 1970; Sonnemans, 1998, 2000). Bearden et al. (2004) have criticized the full information case as being too restrictive. Some have examined the partial information case (e.g., Kahan et al., 1967; Shapira, 1981). More recently, the no information case has received growing attention (e.g., Bearden et al., 2004; Corbin, et al., 1975; Seale and Rapoport, 1997, 2000; Zwick et al., 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, decision makers must search through options composed of *multiple attributes*. In the current paper, we describe a new class of sequential search problems that generalizes 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 j th 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 in terms of their quality 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 n th 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 and Mosteller, 1966). The cutoff t converges to $\frac{n}{e}$ and the optimal policy selects the best applicant with probability $\frac{1}{e}$ as $n \rightarrow \infty$. Both t and the selection probability converge from above.

Seale and Rapoport (1997) had subjects play a large number of random instances of the CSP in two different experimental conditions: $n = 40$ and $n = 80$. 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 ($\frac{n}{e}$ th) position. Seale and Rapoport 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 factored in.

In a subsequent paper, Seale and Rapoport (2000), relaxed assumption 1 of the CSP; at the beginning of each trial, subjects in their experiment knew the distribution from which n was sampled but not the actual value of n for the problem instance they played. They also concluded that the DMs 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 DMs to recall previously interview 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 DMs tended to search insufficiently; however, when the experimenters imposed a fixed search cost for each applicant, the pattern was reversed: the DMs tended to search for too long.

Bearden et al. (2004) presented and tested a generalization of the CSP in which the DM earns payoffs that are nondecreasing in the quality of the selected applicant. They argued that this payoff structure better captures features of many real-world search problems than the nothing-but-the-best (0,1) payoff scheme of the CSP. 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 $\pi(1) \geq \dots \geq \pi(n)$.

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 $\pi(a) = \alpha + \beta(n - a)$, when $\beta > 0$. An infinite number of alternative schemes 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 search policy for problems with the payoff structure stated in assumption 5' has the same threshold form as that of the CSP (Mucci, 1973). 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.

In two experimental studies of the GSP, Bearden et al. (2004) concluded that, as in the CSP, DMs tend to terminate their search too early relative to the dictates of the optimal policy. They offered an alternative to the endogenous search cost explanation proposed by Seale and Rapoport (1997). Using scoring rules, Bearden et al. had DMs estimate the probability of obtaining various payoffs for selecting applicants of different relative ranks in different applicant positions. They then argued that the bias to terminate search too early in the GSP results from DMs 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's (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*

Suppose that we have a set of n applicants who differ from one another along k different dimensions or attributes. In the domain of job search, the attributes might represent education, work experience, degree of technical proficiency, interpersonal skills, etc. The applicants are interviewed in a random order, with all $n!$ orderings obtaining with equal probability. For each attribute i ,

the applicants can be ranked from best to worst with no ties.²

The *absolute rank* of the j th applicant on the i th attribute, denoted by a_i^j , is simply the number of applicants in the applicant pool, including j , whose i th attribute is at least as good as the j th applicant's. Hence, an absolute rank of 1 means that there is no other applicant in the pool with an attribute as good; an absolute rank of 2 means that there is exactly one other applicant with a better score on that attribute; etc. The j th applicant's absolute ranks can, therefore, be represented by a vector $\mathbf{a}^j = (a_1^j, \dots, a_k^j)$.

The *relative rank* of the j th applicant on the i th attribute, denoted by r_i^j , is the number of applicants from 1 to j whose i th attribute is at least as good as the j th's. The first applicant ($j = 1$) will always have a relative rank of 1 on all k attributes; applicant $j = 2$ will have a relative rank of either 1 or 2 on each attribute; and so on. When the DM observes the j th applicant, she observes the relative ranks of the applicant along each of the k attributes. She observes $\mathbf{r}^j = (r_1^j, \dots, r_k^j)$ and must make her selection decision on the basis of this information. As in the GSP, once she passes (rejects) an applicant she cannot return to that applicant—there is no recall.

Although she only observes relative ranks $\mathbf{r}^j$, the DM's payoff for selecting the j th applicant, denoted Π^j , is based on the applicant's absolute ranks $\mathbf{a}^j$; specifically,

$$\Pi^j = \sum_{i=1}^k \pi_i(a_i^j), \quad (1)$$

where π_i is a function that maps absolute ranks on the i th attribute into real payoffs. We constrain $\pi_i(x) \geq \pi_i(y)$ for all i and all $x \leq y$; that is, *ceteris paribus*, one never earns less for selecting an applicant with a better attribute. An optimal policy for this problem, which we dub the *Multi-attribute Secretary Problem* (MASP), is one that maximizes the expected value of the selected applicant.

A few related problems 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. Ferguson (1992) generalized the problem presented by Gnedin by allowing dependencies between the attributes, and showed that the

² The no tie assumption considerably simplifies matters with little loss of generality, assuming that the attribute values fall on a continuum and that there will be at least infinitesimal differences in the attributes of any two applicants. Of course, this rules out nominal valued attributes such as sex; but most attributes can be represented on a continuum.

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 i th attribute of the j th applicant whose relative rank on that attribute is r has an absolute (overall) rank of a is given by (Lindley, 1961):

$$Pr(A = a | R = r) = \frac{\binom{a-1}{r-1} \binom{n-a}{j-r}}{\binom{n}{j}}, \quad (2)$$

where $r \leq a \leq r + (n - j)$; otherwise, $Pr(A = a | R = r) = 0$. We assume that the k attributes are pairwise independent; that is, $Pr(a_i = a \wedge a_{i'} = a') = Pr(a_i = a)Pr(a_{i'} = a')$ for any pair of attributes i and i' . Therefore, the expected payoff for selecting the j th applicant is:

$$E(\Pi^j | \mathbf{r}^j) = \sum_{i=1}^k \sum_{a=r_i^j}^n Pr(A = a | R = r_i^j) \pi_i(a). \quad (3)$$

Again, 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 V^{j*} . Hence, $V^* = V^{1*}$.

At each stage j of the decision problem, the DM must decide either 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 $\mathbf{R}^j = \{\mathbf{r}^j\}$. 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 $\mathbf{R} = (\mathbf{R}^1, \dots, \mathbf{R}^n)$. By Bellman's (1957) *Principle of Optimality*, for an optimal (global) policy $\mathbf{R}^*$, each sub-policy $(\mathbf{R}^j, \dots, \mathbf{R}^n)$ from stage j to n must also be optimal. Given this property, we can find the optimal policy

using straightforward 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 simplify 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 $\mathbf{r} \in \mathbf{R}^{j*} \Leftrightarrow E(\Pi^j|\mathbf{r}) \geq V^{j+1*}$.

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 $E(\Pi^j|\mathbf{r}) = V^{j+1*}$, this $\mathbf{r} \in \mathbf{R}^{j*}$. Therefore, $\mathbf{r} \in \mathbf{R}^{j*}$ if and only if $E(\Pi^j|\mathbf{r}) \geq V^{j+1*}$. $\square$

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 n th applicant, if reached,

$$V^{n*} = (n)^{-1} \sum_{i=1}^k \sum_{a=r_i^j}^n \pi_i(a). \quad (4)$$

The expected payoff for the last applicant 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

$$V^{j*} = Q(\mathbf{R}^{j*}) E(\Pi^j|\mathbf{R}^{j*}) + [1 - Q(\mathbf{R}^{j*})] V^{j+1*}, \quad (5)$$

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. 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. 5, the optimal global policy $\mathbf{R}^*$ is easily constructed.

Denoting the applicant position at which the search is terminated by m , the probability that the DM stops on the $(j < n)$ th applicant under the optimal policy is:

$$Pr(m = j) = \left(\prod_{h=0}^{j-1} [1 - Q(\mathbf{R}^{h*})] \right) Q(\mathbf{R}^{j*}), \quad (6)$$

where $Q(\mathbf{R}^{0*}) = 0$. The expected stopping position is, then:

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

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 1. The payoffs for each a for each attribute i are shown in the top panel. 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 $V^{1*} = 7.82$ in this example.

Let us look at 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

Table 1

An example of a MASP with $n = 6$ and $k = 2$. See text for explanation.

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_1^j	2	4	3	6	5	1
a_2^j	5	2	1	3	6	4
r_1^j	1	2	2	4	4	1
r_2^j	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

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 DMs. In the first, the attributes are equally weighted, that is, $\pi_1(a) = \pi_2(a)$ for all a . The second experiment examines a case in which the attributes are unequally weighted. After describing the experiments and their results, we describe some implications and discuss future directions for this line of research.

4 Experiment 1: Equally Weighted Attributes

4.1 Method

4.1.1 Subjects

Thirty subjects participated individually in the experiment. 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, that typically lasted 40-60 minutes, was \$21 (minimum \$5, maximum \$50). 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. The example explained and illustrated the updating of the relative ranks of each attribute for each applicant. Once the subjects understood the instructions, they were then seated at individual computers. Then they performed two practice problems to verify their understanding of the task. The experimental problems were presented once the subjects successfully 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 (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 ($j = 1, \dots, n - 1$), 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 n th 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

The DMs earned positive payoffs for selecting applicants with attributes whose absolute rank did not exceed 5; for absolute ranks 6 through 30, they earned nothing. Selecting an applicant with an absolute rank of 1 contributed \$25 to a DM's payoff; selecting ones with 2, 3, 4, or 5 contributed \$12, \$8, \$4, \$2, respectively. Hence, the subjects could earn as much as \$50 on a given trial and as little as \$0. Subjects were paid for a single randomly selected trial, which they determined for themselves by drawing a number from a hat. The payoff schemes used in the current experiment and also in Experiment 2 are presented in Table 2. The optimal policy for the problems studied in both experiments is presented in the next section.

Table 2

Payoff schemes used in Experiments 1 and 2. Payoffs are in US dollars.

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

4.2 Results

4.2.1 Earnings

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, $t(29) = 10.57$, $p < .001$.

4.2.2 Stopping Position

We computed the mean stopping positions over the 100 trials of the MASP and compared these to the expected stopping position under the optimal policy,

$E(m) = 20.09$, which results from the application of Eq. 7. The mean observed stopping position ($M = 15.89$, $SD = 4.29$) was significantly smaller than that expected under the optimal policy, $t(29) = 5.48$, $p < .001$. On average and across all 100 trials, 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 Figure 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.

4.2.3 Evidence of Learning

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, $b = .0145$, $p = .049$, indicating that earnings increased with experience. (A Durbin-Watson test revealed that the independence assumption necessary for the inference was not violated, $d = 2.07$, $p > .05$) However, the increase in earnings is quite mild across trials, and still well below the expected earnings in the final trials (Figure 2). 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 significantly positive, $b = .025$, $p < .001$; $d = 2.07$, $p > .05$. The mean stopping positions over trials are displayed in Figure 3. The subjects seem to have learned that searching longer improves payoffs.

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 ($r_1^j = x, r_2^j = y$) if and only if the cutoff for (x, y) has been reached, i.e., if $c_{x,y}^* \geq j$. The $c_{x,y}^*$ are ordered such that $c_{x,y}^* \leq c_{x,y+1}^*$ and $c_{x,y}^* \leq c_{x+1,y}^*$. 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 sets the empirical cutoffs $\hat{\mathbf{c}}$ for each




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.




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.

experimental 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 $\hat{c}_{x,y} \geq j$; 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

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

subject to $\hat{c}_{x,y} \leq \hat{c}_{x,y+1}$ and $\hat{c}_{x,y} \leq \hat{c}_{x+1,y}$, 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 violations (Eq. 8), we used the threshold accepting (TA) algorithm developed by Dueck and 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 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. We refer the reader to Dueck and Scheuer 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 36 rather than 900 cutoffs without any loss. The median estimated $\hat{t}_{x,y}$ ($x, y = 1, \dots, 6$) are displayed in the center panel of Table 3. 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 subjects' decisions.

The difference panel (bottom panel) in Table 3 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 ($r \geq 6$), 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 $(1, r \geq 6)$.

4.3 Conclusion

Consistent with previous experimental studies of secretary problems (e.g., Bearden et al., 2004; Seale and Rapoport, 1997, 2000; Zwick et al., 2003), we find that DMs in the MASP tend to terminate their search too early relative to the optimal policy. 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 ($r \geq 6$) 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. To investigate this finding further, we conducted a second experiment with a payoff scheme that gives considerably more weight to one attribute relative to the

Table 3

Optimal and empirical (estimated) cutoffs for Experiment 1. 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).

		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

other. Under the unequal weighting payoff scheme, relying on a heuristic policy of the sort observed in Experiment 1 is less efficient. In addition, since, in many real-world search scenarios, DMs often assign more weight to certain attributes over others, our second experiment allows us to test the generality of the results of Experiment 1 in a broader class of sequential search problems.

5 Experiment 2: Unequally Weighted Attributes

5.1 *Method*

5.1.1 *Subjects*

Thirty subjects were recruited in the same manner as those in Experiment 1. The mean payoff for the session was around \$20 (minimum \$5, maximum \$37).

5.1.2 *Procedure*

The procedure was identical to that of Experiment 1 in all but one way: The payoff scheme used in the example in the instructions was changed to correspond to the payoff scheme used in the current experiment.

5.1.3 *Payoff Scheme*

The payoffs used for the first attribute 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 \$0. Hence, in this experiment the subjects could earn between \$0 and \$37 on a single trial. Again, each subject was paid for a single randomly chosen trial by drawing a random number from a hat. The payoff scheme is summarized in Table 2.

5.2 Results

5.2.1 Earnings

The optimal DM can expect to earn $V^{1*} = 16.23$ in Experiment 2. Computing the mean earnings for each subject and taking the mean of these individual results ($M = 13.53$, $SD = 2.53$), our results show that the mean observed payoff is significantly smaller than the expected payoff, $t(29) = 5.83$, $p < .001$.

5.2.2 Stopping Position

The expected stopping position under the optimal policy is $E(m) = 19.45$. Similarly to Experiment 1, the mean stopping position ($M = 15.90$, $SD = 3.38$) was significantly smaller than that expected under optimal search, $t(29) = 5.75$, $p < .001$. As observed in Experiment 1, 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, in turn, earned less (see Figure 1).

5.2.3 Evidence of Learning

Earnings tended to increase with experience, as evidenced by the significant positive slope of the earnings regressed onto trial number, $b = .0262$, $p < .001$. (As before, we used a Durbin-Watson test to verify that the independence assumption of the regression was not violated; it was not, $d = 2.00$, $p > .05$.) 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 Figure 2). 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, $b = .017$, $p < .001$; $d = 2.44$, $p > .05$. The mean stopping position over trials is exhibited in Figure 3.

5.2.4 Estimated Policies

We used the procedure described in Experiment 1 to estimate the decision policy of each subject in Experiment 2. The median cutoffs are displayed in Table 4. First, note that the subjects' cutoffs reveal that their policies are sensitive to the attribute weights. In no case is the cutoff for a given pair of relative ranks $(r_1, r_2) = (x, y)$ for $x \leq y$ 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. 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 ($r \geq 6$) 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 a 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.

5.3 Conclusion

The results of Experiments 1 and 2 are quite similar. In both experiments, the subjects tend not to search deeply enough into the set of applicants: They stop too soon. The bias to stop searching too early, however, does not obtain for all combinations of relative ranks. It occurs for applicants for whom both attributes are relatively good (and therefore may have positive payoffs), but the subjects are biased against selecting applicants with one relatively very good attribute (e.g., $r = 1$) and one poor attribute ($r \geq 6$), even when doing so is advantageous. This pattern of behavior is consistent with the use of a modified satisficing 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

Table 4

Optimal and empirical (estimated) cutoffs for Experiment 2. 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 is the more heavily weighted attribute.

		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

subjects’ decisions; the Experiment 2, the estimated policies captured around 95% of the decisions. Without resorting to more complicated decision rules (e.g., ones with additional free parameters), we think it doubtful that we can better account for these results.

6 Summary

We began this paper by presenting a realistic extension of the secretary problem in which the DM searches through applicants who vary on several independent 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 DMs 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 and Rapoport, 1997, 2000; Zwick et al., 2003), DMs 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 former 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 other scenarios in which one would expect the attributes to be correlated. Proofreading and writing abilities, for

example, are presumably related. This is another generalization of the MASP that we intend to pursue.

References

- Bearden, J. N., and Murphy, R. O. (2004). On the generalized secretary problem. Unpublished manuscript
- Bearden, J. N., Rapoport, A., and Murphy, R. O. (2004). Sequential observation 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., and 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., and M. Abbondanza, M. (1975). Context effects in optimal stopping rules. *Behavior and Human Performance*, 14, 207-216.
- Cox, J. C. and Oaxaca, R. L. (1989). Laboratory experiments with a finite horizon job-search model. *Journal of Risk and Uncertainty*, 2, 301-330.
- Dueck, G. and 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. and 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 & 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., and Jones, L. V. (1967). Decision making in a sequential search task. *Perception and Psychophysics*, 2, 374-376.
- 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., and 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 & Psychophysics*, 6, 234-240.
- Rapoport, A., and Tversky, A. (1970). Effects of search cost on sequential search behavior. *Perception & Psychophysics*, 6, 234-240.
- Samuels, S. M. (1991). Secretary problems. In B. K. Gosh and P. K. Sen (Eds.), *Handbook of Sequential Analysis*. New York: Marcel Dekker, pp. 381-405.
- Samuels, S. M., and 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*. Upton, NY: Brookhaven National Laboratory, pp. 62-78.
- Seale, D. A. and 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. and 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 & 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., and Koehler, D. (1994). Support theory: A nonextensional representation of subjective probability. *Psychological Review*, 101, 547-567.
- Zwick, R, Rapoport, A., Lo, A. K. C., and Muthukrishnan, A. V. (2003). Consumer sequential search: Not enough or too much? *Marketing Science*, 22, 503-519.