

Sequential Observation and Selection with Rank-Dependent Payoffs: An Experimental Study

J. Neil Bearden
University of Arizona

Amnon Rapoport
University of Arizona
and
Hong Kong University of Science and Technology

Ryan O. Murphy
University of Arizona

May 3, 2005

Keywords: Secretary Problem, Optimal Stopping, Alternative Decision Rules, Sequential Search

JEL Classification: D83, C44, C61

Please address all correspondence to:

Before August 1, 2005

J. Neil Bearden
University of Arizona
Department of Management and Policy
405 McClelland Hall
Tucson, AZ 85721
Phone: 520-603-2092
Fax: 520-325-4171
Email address: jneilb@gmail.com

After August 1, 2005

J. Neil Bearden
Brain Imaging and Analysis Center
Duke University
163 Bell Research Building
Box 3918, Duke University Medical Center
Durham, NC 27710
Email address: jneilb@gmail.com

Abstract

We consider a class of sequential observation and selection decision problems in which applicants are interviewed one at a time, decision makers only learn the applicant's quality relative to the applicants that have been interviewed and rejected, only a single applicant is selected, and payoffs increase in the absolute quality of the selected applicant. Compared to the optimal decision policy, which we compute numerically, results from two experiments show that subjects terminated their search too early. We competitively test three behavioral decision rules and find that a multi-threshold rule, which has the same form as the optimal decision policy but is parameterized differently, best accounts for the data. Results from a probability estimation task show that subjects tend to overestimate the absolute quality of early applicants, and give insufficient consideration to the yet-to-be-seen applicants.

1. Introduction

Decision makers (DMs) must often choose among alternatives that are presented sequentially, one at a time. Consequently, in many situations, they must choose options without knowing the full choice set. Consider the problem of deciding when to sell an asset on an open market. On any given day, a trader observes a selling price and must decide whether to sell without knowing what the price will be on the following day. Similarly, in tight housing markets when properties may disappear from the market soon after they appear, potential buyers must quite often make irrevocable decisions without knowing what options will appear on the horizon. The same dilemma arises in decisions to hire job applicants in times of low unemployment, when passing up the opportunity to hire a good applicant may mean that one cannot find a better one in the future.

Sequential decision problems have been specified formally in a number of ways. Problems in which the DM is assumed to have complete information about the distribution from which the observations are sampled are referred to as *full-information problems*. There are also *partial-information* problems in which the DM is assumed to know certain features of the distribution from which the observations are drawn (e.g., that it is Gaussian) but not others (e.g., the mean and variance of the distribution). The class of problems requiring the weakest assumptions about the DM's state of knowledge regarding the distribution from which the observations are drawn are known as *no-information* problems. As models of real-world sequential decision problems, this latter class has a number of virtues. We will illustrate these by first noting the drawbacks of the full- and partial-information problems.

Consider the full-information sequential search problems. It seems doubtful that a DM trying to decide when to unload an asset has perfect information about the distribution of price changes for that asset. To model the asset-selling decision problem as one involving full-information is therefore unrealistic. One might argue that a partial-information formulation might be more defensible. One can assume that the DM is well informed and knows that the price changes are log-normally distributed, for example, but that she does not know the parameters of the distribution. This formulation may be realistic

for the asset-selling problem. However, both the full- and partial-information formulations break down as soon as one moves beyond decision alternatives that can be characterized by a single attribute.

In *some* sequential decision problems, the value of an option is easily decidable. For example, when selling an asset, the price of the asset on any given day *is* its value. The values of decision alternatives are less apparent in *many* other sequential decision problems. The values of houses and potential employees, for example, can be difficult to determine: It is hard to map these multi-attribute alternatives into a single measure of value. On the other hand, it is considerably easier to *rank* them in terms of quality. One can decide that one applicant is better than another without assigning each of them meaningful scalar measures of quality. In the no-information problems that we consider, one need only assume that the DM can rank decision alternatives; there is no need to assume that the DM knows the multivariate distribution from which the alternatives are drawn.

The *Secretary Problem* is a stylized no-information sequential observation and selection problem that first appeared in the February 1960 column of Martin Gardner in *Scientific American*. It was immediately taken up and developed by prominent statisticians and applied mathematicians. Since then, the problem has been extended and generalized in many different ways and given rise to a “field” of study in applied probability. For partial reviews see Ferguson (1989), Freeman (1983), and Samuels (1991). The secretary problem has also stimulated experimental, in addition to purely theoretical, research on optimal stopping in a class of sequential observation and selection problems with rank-dependent payoffs—i.e., where the payoff depends only on the rank of the selected observation and not on its actual value— by Corbin, Olson, and Abbondanza (1975), Seale (1996), Seale and Rapoport (1997, 2000), and Zwick, Rapoport, Lo, and Muthukrishnan (2003).

Although there are many variants of the secretary problem, in its simplest and original form the *Classical Secretary Problem* (CSP) is defined by the following assumptions:

1. There is only a single position to be filled.
2. There are n applicants for the position. The value of n is known before the search commences.
3. It is assumed that the DM can rank the applicants from best to worst without ties.

4. The applicants are interviewed sequentially, one at a time, in a random order with each of the $n!$ orderings being equally likely.
5. As each applicant is being interviewed (observed, evaluated), the DM must either accept the applicant for the position and thereby terminate the search, or reject the applicant and interview the next one, if any.
6. The decision to either accept or reject the current applicant must be based only on the relative ranks of the applicants interviewed so far.
7. Once rejected, an applicant cannot be recalled.
8. The DM's objective is to select the best applicant. With no loss of generality, this objective implies that the DM wins 1, if she selects the best applicant, and 0 otherwise.

Managerial, economic, and marketing sequential observation and selection problems that the CSP has been used to model include, among others, hiring an employee for a job, purchasing some product on the market, assigning a job to a single machine, choosing an investment alternative, and searching to rent an apartment (Zwick et al., 2003). These problems are characterized by decision alternatives (“applicants”) that are inspected sequentially. In these problems a balance is sought between the risk of stopping the search too soon and accepting an apparently desirable alternative when an even better one might be still to come, and the risk of searching too long and discovering that the best alternative was rejected earlier (Sardelis & Valahas, 1999).

It has long been recognized that the CSP is much too restrictive for most applications. Therefore, almost every assumption of the CSP has been relaxed. For example, the assumption that the value of n is finite and known has been generalized to include the more realistic case where n is uncertain and only its distribution function is known (Pressman & Sonin, 1972). The assumption that an applicant, once rejected, cannot be recalled has been generalized by allowing backward solicitation (or recall), with a probability of success depending on how far in the past the applicant was observed (e.g., Yang, 1974, Smith, 1975, Choe & Bai, 1983). Perhaps the most limiting restriction of the CSP is Assumption 8, that the DM is satisfied with nothing but the best. This is quite restrictive, since in all of the applications

mentioned earlier the DM presumably derives positive utility from selecting some applicant who is not the overall best. Different and more realistic objectives have been proposed in an attempt to generalize Assumption 8. Gilbert and Mosteller (1966) considered the case where the DM's objective is to select one of the k best applicants ($k > 1$) without differentiating among them. Bartoszynski and Govindarajulu (1978) considered a very special case where the DM receives payoffs of a , b , or 0 , $a \geq b > 0$, if she stops and selects the overall best, second best, or any other applicant, respectively. Chow, Moriguti, Robbins, and Samuels (1964) considered yet another case where the DM's objective is to minimize the expected rank of the chosen applicant, where the top applicant has rank 1, the second best rank 2, etc.

The most general formulation of the DM's objective specifies that the DM's payoff increases monotonically in the *quality* of the selected applicant, so that the lower the ranking of the selected applicant the higher the payoff. Let w_j denote the DM's payoff if the applicant she selects has rank A_j among all the n applicants. Then, the payoff function for this problem, which we dub the *Generalized Secretary Problem* (GSP), has the form $w_1 \geq w_2 \geq \dots \geq w_n$. In the best choice problem (CSP), $w_1 = 1$, $w_2 = w_3 = \dots = w_n = 0$. In the case studied by Bartoszynski and Govindarajulu (1978), $w_1 \geq w_2 > 0$, $w_3 = \dots = w_n = 0$. If the objective is to select one of the top k applicants, then the payoff function assumes the form $w_1 = \dots = w_k > 0$, $w_{k+1} = \dots = w_n = 0$. And in the minimization of expected rank, the payoff function has the very special form $w_j = n - a_j$, where a_j is the overall (absolute) rank of the selected applicant.

The optimal policy for the GSP implies a particular form of the decision rule that is not immediately obvious. A major purpose of the present study is to test experimentally whether financially motivated subjects participating in the GSP adhere to this decision rule. If not, we wish to characterize the decision rules (heuristics) they actually use in deciding when to stop the search, and to understand the psychological basis for their stopping decisions. Section 2 briefly characterizes the optimal decision policies for the CSP, the minimization of expected rank case, the GSP, and the numerical algorithms used to compute them. Experimental results from several variants of the secretary problem that pertain to the present study are summarized in Section 3. Section 4 presents the method and results of our first experiment on the GSP. Section 5 presents results from a second GSP experiment with different

parameter values, as well as results from an associated probability estimation task. The findings are summarized and discussed in Section 6.

2. Optimal Decision Rules

In all the variants of the secretary problem, the observations are taken to be the *relative ranks* s_1, s_2, \dots, s_n , where s_j is the rank of the j^{th} applicant among the first j applicants to be interviewed, rank 1 being the best. The *absolute ranks* are denoted by a_1, a_2, \dots, a_n , where a_j is the rank of applicant j among all the n applicants including the ones that have yet to be interviewed. To illustrate the relationship between absolute and relative ranks, consider $n=9$ applicants with absolute ranks 2, 9, 5, 3, 6, 1, 8, 7, 4, who are interviewed *in this order*. The resulting relative ranks, which constitute the only information available to the DM, are shown (in bold) in the bottom row:

Period of search	1	2	3	4	5	6	7	8	9
Absolute Rank	2	9	5	3	6	1	8	7	4
Relative Rank	1	2	2	2	4	1	6	6	4

In the CSP, applicants with relative rank 1 are called *candidates*. Obviously, in the CSP one only selects a candidate. In the example above there are two candidates, namely, the applicants interviewed in periods 1 and 6. The optimal decision rule for the CSP has a very simple form:

Reject the first r^-1 applicants and then stop and select the next candidate.*

The optimal policy is obtained numerically from the equation

$$r^* = \min \left\{ r \geq 1 : \sum_{k=r+1}^n \frac{1}{k-1} \leq 1 \right\}.$$

Values of r^* and P_r^* , which is the probability of selecting the best overall applicant given r^* , for $n=1, 2, \dots, 9$ are presented in the following table (Gilbert & Mosteller, 1966):

N	1	2	3	4	5	6	7	8	9
r^*	1	1	2	2	3	3	3	4	4
P_r^*	1.000	0.500	0.500	0.458	0.433	0.428	0.414	0.410	0.406

Perhaps surprising for someone not familiar with the CSP is that as $n \rightarrow \infty$, $r^* \rightarrow ne^{-1}$ and $P_r^* \rightarrow e^{-1}$. In words, for a large n the DM should pass up the first 36.8% of the applicants without taking any action and

then stop the search by selecting the next encountered candidate, if any. The asymptotic probability of selecting the best of the n applicants via the optimal decision rule is $e^{-1} \approx 0.368$. As illustrated above, convergence is reached rather quickly. For example, if $n=20$, then already $P_r^* = 0.384$.

If the DM's objective is to minimize the expected rank of the selected applicant, then the optimal decision rule has the following form:

If at period j you interview an applicant of relative rank x , then stop the search and select this applicant if $j \geq r_x^$.*

We refer to decision policies of this form as *Multi-threshold Rules* (MTRs). They are constrained such that $r_1^* \leq r_2^* \leq \dots \leq r_n^* = n$. Following an MTR entails passing the first $r_1^* - 1$ applicants without stopping; from period r_1^* through period $r_2^* - 1$ select an applicant of relative rank 1; from period r_2^* through period $r_3^* - 1$ select an applicant of relative rank 1 or 2, and so on. Chow et al. (1964) derived the optimal decision rule for this case (see also Moriguti, 1993). To illustrate it, assume that $n=25$. Then, $r_1^* = 8$, $r_2^* = 14$, $r_3^* = 17$, $r_4^* = 19$, and so on. In this example, the DM should observe about 1/4 of the applicants without taking any action, then from period 8 to period 13 select an applicant only if he happens to be the best of all those observed so far, from period 14 to 16 select an applicant only if he is either the best or second best of those interviewed so far, and so on. As more applicants are evaluated, the criterion for accepting an applicant and stopping the search is relaxed.

When monotone payoffs $w_1 \geq w_2 \geq \dots \geq w_k$ are assigned to the best k applicants, Mucci (1973) showed that the MTR is optimal with $r_1 \leq r_2 \leq \dots \leq r_n = n$. Denoting by $\mathbf{r} = (r_1, \dots, r_k)$ the MTR threshold values for any monotone order k stopping rule, Yeo and Yeo (1994) showed that the choice of an optimal stopping rule, \mathbf{r}^* , is made by finding the vector \mathbf{r} that maximizes

$$Q(\mathbf{r}) = \sum_{a=1}^k w_a P(a | \mathbf{r}),$$

where $P(a|\mathbf{r})$ is the probability of selecting the a^{th} best applicant given the policy \mathbf{r} . It is computed from

$$P(a|\mathbf{r}) = \frac{1}{n \binom{n-1}{a-1}} \sum_{d=1}^k \sum_{j=r_d}^{r_{d+1}-1} \prod_{i=1}^d \frac{r_i-1}{j-1} \sum_{s=1}^{\min(d,a)} \binom{j-1}{s-1} \binom{n-j}{a-s}$$

for $a=1, 2, \dots, n-r_1$, where s denotes the relative rank of the j^{th} applicant. Direct computations over a restricted k -dimensional grid provide the numerical solutions.

To illustrate the optimal solution, consider the GSP with $n=50$, $k=5$, and payoffs $w_1=16$, $w_2=8$, $w_3=4$, $w_4=2$, and $w_5=1$. Then, the optimal threshold values are found to be $r_1^* = 17$, $r_2^* = 36$, $r_3^* = 44$, $r_4^* = 47$, $r_5^* = 49$. The probabilities of selecting the best applicant, second best, third best, fourth best, and fifth best under the optimal decision rule are 0.351, 0.227, 0.128, 0.067, and 0.033, respectively; the probability of no selection is 0.166; and the expected payoff for following the optimal policy is 8.11.

3. Previous Experimental Research on the Secretary Problem

Varying n , Seale and Rapoport (1997) studied the CSP using a between-subjects design with two conditions, namely, $n=40$ and $n=80$. These authors found that their subjects tended to stop the search too early. Further, they competitively tested decision rules (heuristics) that the subjects might be using and concluded that a threshold rule of the same form as the optimal cutoff rule best accounted for the data. However, rather than being set at r^* , the subjects' actual thresholds, r , tended to be smaller (i.e., $r < r^*$). In a subsequent study, Seale and Rapoport (2000) studied a variant of the CSP that included uncertainty about the exact value of n . Using a between-subject design, their study included two experimental conditions where the value of n was known to be randomly sampled with equal probability from either the set $\{1, 2, \dots, 40\}$ or the set $\{1, 2, \dots, 120\}$. They again reported a bias toward early stopping. A third study by Zwick et al. (2003) examined yet another variant of the CSP that allows for backward solicitation. Charging a fixed cost per observation in two of their four experimental conditions, Zwick et al. used a 2×2 between-subject design with two levels of number of applicants ($n=20$ vs. $n=60$) and two

levels of observation cost (0 vs. c , with $c=\$0.30$ if $n=20$, and $c=\$0.10$ if $n=60$, with a reward for selecting the best applicant of \$10.00). When search costs were set at zero, they replicated the earlier findings, but with positive search costs they found that their subjects searched too much.

It is possible that the nothing-but-the-best payoff scheme used in these previous studies of the CSP might be responsible for the biased search behavior. In addition, we are hard-pressed to find situations in which a DM derives positive utility for selecting the best applicant and nothing otherwise. More likely, DMs prefer better applicants to poorer ones. The GSP allows for more realistic payoff structures to test sequential search behavior. Using the GSP, we pose three questions:

1. In comparison to the optimal decision rule, do subjects search too little, too much, or just enough?
2. Do subjects adhere to a multi-threshold decision rule?
3. If not, what heuristics do they use to decide when to stop the search?

The answer to the first question seems to be consistent across the three studies of the CSP described above. If the search is costless, then, in comparison to the optimal policy subjects tend to stop their search too early (Seale & Rapoport, 1997, 2000; Zwick et al., 2003). The effect is rather strong. For example, Seale and Rapoport (1997) reported that 21 of their 25 subjects in condition $n=40$ and 21 of their 25 subjects in condition $n=80$ stopped the search too early.

The answer to the second question is ambiguous. Conducting their analyses on the individual level, Seale and Rapoport (1997, 2000) reported moderate support for a decision rule with a threshold value that differed across subjects and was, in general, smaller than the optimal value. In contrast, the results of Zwick et al. (2003) did not support the optimal decision rule for the CSP with backward solicitation and positive search cost.

With regard to the third question, Seale and Rapoport (1997) compared the optimal decision rule to three behavioral decision rules (heuristics). The first heuristic, already mentioned above, is a threshold decision rule with a possibly non-optimal threshold value, r , which may differ across subjects. The optimal decision rule is a special case of this rule. The second heuristic is a candidate decision rule stipulating that the DM stops the search on observing the g^{th} candidate. The third heuristic is a successive

non-candidate decision rule stipulating that the DM chooses the first candidate after observing no fewer than h non-candidates that follow the last candidate that was rejected. Each of these three heuristics has a single parameter (r, g, h) that allows for individual differences. Stein, Seale, and Rapoport (2003) have shown that the successive non-candidate decision rule can obtain near optimal payoffs in the CSP. Seale and Rapoport (1997, 2000) provided evidence that some of the subjects' patterns of behavior could be accounted for quite well by either the threshold or successive non-candidate decision rule. Zwick et al. (2003) provided evidence that subjects are sensitive to local patterns in the observed sequences of applicants (e.g., increasing relative ranks on successive periods), patterns that optimal DMs should ignore.

In the experiments described next, we test whether the bias for early stopping observed in previous studies of the CSP persists under more realistic payoff schemes. In addition, we examine the types of decision rules that are used by actual DMs in the GSP.

4. Experimental 1

Method

Subjects. Sixty-two subjects participated individually in Experiment 1. 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 30-40 minutes, was \$20 (minimum \$5, maximum \$50).

Procedure. The instructions (hard copy) explained the GSP in detail, placing special emphasis on the computation of relative ranks with the presentation of a new applicant. An example of $n=6$ applicants with absolute ranks 6, 3, 4, 1, 5, 2 presented in this order was given, and the updating of the relative ranks at each of the six periods (assuming no stopping) was illustrated and explained. Subjects were instructed that as long as the search continues absolute ranks would not be displayed, only relative ranks of all the applicants that had been observed and rejected. After reading the instructions, two practice problems were presented to verify the subject's understanding of the task. The experimental problems were presented once the subjects successfully completed these two practice problems.

Each subject completed 60 *trials* (replications) of the GSP. Each trial consisted of a single GSP with $n=60$ applicants. Each trial was structured in the same way. The relative rank of applicant 1 was first displayed on period 1, and then the subject was instructed to either select this applicant, thereby terminating the search, or proceed to the next interview *period* and observe a new applicant. If she continued the search to period j ($j=2, \dots, n$), then the relative ranks of the $j-1$ previously viewed applicants that had been observed and rejected were updated by a computer 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 n absolute ranks and the corresponding n relative ranks were displayed on the computer screen. In this way, subjects who stopped the search on different periods were provided with the same information about the actual sequences of absolute ranks of all the n applicants.

To allow comparison between subjects, we generated a *sample* of 60 different random sequences of applicants (out of a *population* of $60!$ sequences). Each subject was presented with the same 60 sequences in the sample, but the presentation order was randomly varied between subjects. There were six positive payoffs (in US dollars) that were set at $w_1=25$, $w_2=13$, $w_3=6$, $w_4=3$, $w_5=2$, $w_6=1$, and $w_g=0$, otherwise ($7 \leq g \leq 60$). The subjects were instructed that they would be paid for 2 (out of 60) randomly chosen trials. At the end of the experiment, the subjects each drew two integers (1 to 60) from a hat without replacement to determine their payoff trials.

Results

Preliminaries. Following Yeo and Yeo (1994), we devised a numerical procedure that uses a restricted 6-dimensional grid for determining the optimal threshold values for the GSP under investigation in the present study. The six optimal threshold values, r_x^* , are presented in the second row of Table 1. They imply that the optimal DM should pass over exactly 1/3 of the applicants without ever stopping, stop the search between periods 21 and 42 only when observing an applicant with relative rank 1, stop the search between periods 43 and 52 only when observing an applicant with relative rank 1 or 2, and so on up to the threshold r_6^* for selecting any of the applicants with relative rank $x \leq 6$.

--Insert Table 1 about here--

Since we only presented the subjects with a relatively small sample (60) of the possible problem instances, the expected earnings under the application of the optimal policy to these instances need not equal the expected earnings of the policy when applied to all $n!$ feasible instances. Table 2 (columns 2 and 3) shows the expected earnings under the application of the optimal policy for the limiting situation (second row)—when all $n!$ instances are observed—and also for the 60 problem instances used in Experiment 1 (third row). In addition to the expected earnings, Table 2 also displays the expected number of applicants that are interviewed before a selection is made (column 3) for the limiting and sample cases. The sample expectations for the earnings and stopping position will serve as the benchmark in all the analyses reported below. The bottom row of Table 2 displays mean experimental results; these will be discussed below.

--Insert Table 2 about here--

Throughout, analyses of earnings and stopping position use individual subjects as the unit of analysis. For instance, in the analyses of earnings below we use the mean earnings for each subject as the basic datum. Reported measures of variability correspond to the variability in these means across subjects.

Earnings. Subjects in Experiment 1 earned significantly less ($M=10.06$, $SD=1.99$) than predicted by the application of the optimal policy, which earns \$12.10, $t(61)= 8.04$, $p<0.001$.

Stopping Times. Figure 1 exhibits the cumulative distribution of stopping time per period under the optimal decision rule (dotted line) and compares it to the observed cumulative distribution (solid line). The two functions show the cumulative probability of stopping on the j^{th} applicant (period) or sooner. For example, the cumulative probability of stopping on period 30 is just under 0.40 for both the theoretical and observed functions.

--Insert Fig. 1 about here--

We compared the mean stopping time (period) for each subject across all 60 trials to the expected stopping time under the optimal decision rule. On average, the subjects stopped the search significantly earlier ($M=35.34$, $SD=6.16$; $t(61)=3.09$, $p<0.003$) than predicted by the optimal decision rule (37.90). Figure 1 shows that the propensity to stop searching too early is almost entirely due to early stopping

decisions in periods 1-20. About 15% of all stopping decisions occurred between periods 1 and 20 compared to the predicted value of 0. After period 23, the difference between the predicted and observed functions largely disappears.

In searching for evidence of learning, for each subject we separately computed the mean stopping times for the first half (block 1: trials 1-30) and second half of the session (block 2: trials 31-60). A paired-sample t -test was used to test the null hypothesis of no difference between mean stopping times on the two blocks. Mean stopping time in the first block ($M=34.39$, $SD=6.39$) was smaller than in the second block ($M=36.28$, $SD=6.39$). This result, which is highly significant ($t(61)=4.33$, $p<0.001$), suggests that the subjects' propensity to stop the search too early decreased, but did not fully disappear, with experience in playing the GSP.

Comparison of Alternative Behavioral Decision Rules. To investigate the nature of the subjects' decision policies in the GSP, we begin by first assuming that the policies are of the same form (MTR) as the optimal policy but with possibly non-optimal threshold values $r_1 \leq r_2 \leq \dots \leq r_6$. Although threshold values may vary between subjects, for each subject they are assumed to remain fixed across trials. We then estimate the threshold values for each subject separately that best account for her stopping decisions in all 60 trials. A best fitting MTR is one minimizing the number of trials that the rule wrongly predicts the stopping time. We refer to an incorrect prediction as a *violation*. For example, a violation is recorded if the decision rule dictates that a subject stops on the 25th applicant where, in fact, the subject continues past this applicant. Since each subject completed 60 trials, the number of violations can range from 0 to 60. Estimated MTRs using this criterion may not yield a unique vector of threshold values; there may be several such rules with different threshold values yielding the same number of violations. Our analysis resulted in unique MTRs for 28 of the 62 subjects. For a single subject, 5 decision rules minimized the same number of violations; this was the subject for whom the decision rule was most underdetermined. The mean number of best-fitting MTRs was smaller than two. Thus, the non-uniqueness problem was not of major concern, as the recovered strategies are not greatly underdetermined. When the estimated MTR was not unique, we chose the one closest to the optimal policy.

Two main features of the results are noteworthy. First, there is significant variability in the estimated threshold values across subjects. Second, and consistent with earlier observations, the threshold values are, on average, smaller than those predicted by the optimal policy. This latter feature is also consistent with the observation that subjects stop too early. Table 1 presents the optimal thresholds values (second row) and the mean observed threshold values (third row) under the MTR. The mean number of violations for the best-fitting rules is displayed in the final column of the table.

We cannot conclude that the subjects are actually using MTRs based on the results just presented. The most appropriate way to determine whether subjects are using the MTR is by comparing the MTR to alternative decision rules.¹ To do so, we formulated two alternative decision rules to the MTR, estimated their parameter values individually across all 60 trials using the same procedure as for the MTR, and then applied the criterion of number of violations of the best-fitting decision rule to competitively test the three decision rules. These two decision rules generalize the ones proposed and tested by Seale and Rapoport (1997, 2000).

We call the first alternative heuristic the *Horse Race Decision Rule* (HRR). A subject using this decision rule keeps track of the number of times she observed applicants of relative rank x ($x=1, 2, \dots, 6$). Thus, she maintains six separate counters. Denote by $c_{x,j}$ the value of counter x at period j ($c_{x,j=0} = 0$, for all x before the search starts, $c_{x=1,j=1} = 1$ as the first applicant necessarily has relative rank 1, $c_{x=1,j=2}$ is either 1 or 2, and so on). One of the six counters is increased by 1 only when a subject encounters an applicant of relative rank that does not exceed 6. The HRR assumes that the subject maintains a separate threshold value r_x for each x , and that she stops the search on period j and selects an applicant of relative rank x once $c_{x,j} \geq r_x$. Because a subject should only select an applicant that yields a positive payoff, we can represent the feasible search strategies for the HRR by a vector $\mathbf{r}=(r_1, r_2, \dots, r_6)$. Using again a restricted grid search algorithm, we computed for each subject separately the best fitting vector \mathbf{r} under the HRR,

¹ Of course, inferences based on this method are still underdetermined. By using this method of comparison, we can, however, increase our degree of confidence in a particular inference.

the one minimizing the number of violations of this heuristic across the 60 trials. Table 1 (fourth row) shows the mean threshold values and violations across the 62 subjects.

We refer to the second alternative heuristic as the *Successive Undesirable Applicant Decision Rule*, and denote it by SUAR. A subject adhering to the SUAR keeps track of the number of successive applicants of relative rank $x > 6$ since she last observed an applicant of relative rank $x \leq 6$. Put differently, she keeps track of the number of successively observed applicants with relative rank entailing a payoff of zero (hence the term “undesirable”). Once she observes an applicant with relative rank equal to or smaller than 6, she sets the counter to zero and starts counting again. Denote the value of the counter at period j by y_j . For each applicant with relative rank equal or smaller than 6, the subject sets a (possibly different) threshold value r_x . She then stops the search and selects an applicant of relative rank x , if $y_j \geq r_x$. Note that under this decision rule the subject groups the applicants into two exclusive sets: applicants who yield a payoff of zero ($x > 6$) and applicants who yield a positive payoff ($x \leq 6$). The differences among applicants in the second set are reflected in the threshold values r_x . Because under the SUAR the subject would only select an applicant of relative rank $x \leq 6$, we can again represent the feasible search strategies by a vector $\mathbf{r} = (r_1, r_2, \dots, r_6)$. As with the two previous decision rules (MTR and HRR), we used a restricted grid search algorithm to find the vector \mathbf{r} that minimizes the number of violations of the best-fitting SUAR. Once again, this was accomplished separately for each subject across all 60 trials. Table 1 (bottom line) presents the mean threshold values and violations taken over subjects.

The HRR and SUAR generated, on average, more than two to three times the number of violations of the MTR. Comparison of the three decision rules on the individual level shows that the MTR substantially outperforms the other two rules for *each* of the 62 subjects. Our results suggest that in selecting an applicant in the GSP subjects consider both the applicants’ relative ranks and their position in the sequence, rather than only the number of applicants with relative ranks equal to or smaller than 6 (HRR), or the successive number of undesirable applicants who follow a desirable applicant (SUAR). A major feature of the MTR is that the threshold for selecting an applicant is relaxed as the sequence is nearing its end. This feature is not shared by the HRR and SUAR.

Discussion

The results of Experiment 1 suggest that subjects adhere to the MTR that is sub-optimally parameterized. Specifically, the subjects' thresholds are placed such that they tend to terminate their search too early. Although there is strong support for the MTR on the aggregate and individual levels, two issues remain unresolved. First, the results of Experiment 1 pertain to a single set of parameter values. More general conclusions could be drawn were we to change the number of trials, number of positive payoffs, form of the payoff function, or any combination of the above. Second, the results do not tell us *why* DMs stop the search too early. In trying to answer this question, Seale and Rapoport (1997) suggested that DMs stop the search too early because of an implicit cost of search. However, they proposed no procedure to test this hypothesis directly. In contrast, we hypothesize that subjects stop the search too early because of misperception of the probabilities of receiving positive payoffs for early applicants. Experiment 2 was designed in two parts both to generalize Experiment 1 to a different set of parameter values and to test the following hypothesis: DMs terminate their search too early in the GSP because they overestimate the probability of obtaining positive payoff for selecting early applicants.

5. Experiment 2

Overview. Experiment 2 consisted of two parts. Part 1 included a sequential observation and selection task of the same type studied in Experiment 1, which allows us to assess the generalizability of the results reported in Experiment 1. Part 2 was introduced to test the probability overestimation hypothesis.

Method

Subjects. Thirty subjects recruited in the same fashion as in Experiment 1 participated in Experiment 2. The mean payoff for the 60-minute session was \$17 (minimum \$10, maximum \$37).

Procedure. The general procedure for Part 1 (the sequential search task) was the same as in Experiment 1 with the following exceptions. The number of periods on each trial was reduced from 60 to 40, the number of positive payoffs was decreased from 6 to 3, and the payoffs (in US dollars) were set at: $w_1=12$, $w_2=7$, $w_3=2$, and $w_g=0$ ($g=4, \dots, 40$). Table 2 (columns 4 and 5) displays the expected earnings and stopping position for this problem. Whereas the payoffs in Experiment 1 decreased exponentially as the

quality of the selected applicant decreased, those in Experiment 2 decrease linearly. The rather rapid decrease in the payoff function used in Experiment 1 more closely approximates the (0, 1) payoff function used in studies of the CSP, most of which also found that DMs terminate the search too early. Hence, the current payoff function in Experiment 2 allows us to assess the generalizability of the results. The subjects were told that they would be paid for a single (rather than 2) randomly chosen trial.

Part 2: A Probability Estimation Task. After completing the 60 sequential search trials in Part 1, all the subjects performed a probability estimation task that constituted Part 2. They were instructed that, rather than making hiring decisions as in Part 1, they would serve as consultants for a person making hiring decisions by estimating and providing her with the probability that applicants with certain relative ranks they know have absolute ranks that they do not know. The instructions explained the task and provided several examples of the types of estimates that would be requested. For example, the subjects were given an example in which they observed that the third applicant out of six applicants had a relative rank of 2. They were then asked to make estimates of the following sort: “*Please estimate the probability that Applicant Number 3, whose relative rank is 2, has an absolute rank of 2.*” For another example, “*Please estimate the probability that Applicant Number 3, whose relative rank is 2, has an absolute rank of 3.*” The subjects were instructed that they would be paid according to the accuracy of their estimates and were shown a table of the payoffs they would receive based on the error (the absolute difference between the true and estimated probabilities) of their estimate. To motivate the subjects to state their true subjective probabilities, we used a logarithmic scoring rule (see Winkler, 1969) to determine the payoffs. Denoting the absolute difference between the true and estimated probability by x , subjects received \$25 for $x < 0.01$, $\$5[-\ln(x)]$ for $0.01 \leq x < 0.50$, and \$0, otherwise. Under this scoring rule, the optimal response is to report one’s true subjective probability. Subjects were told that at the end of Part 2 they would be paid for one of their estimates chosen at random.

In Part 2, the subjects were presented with 30 trials of 40 applicants each and were asked to assign probability estimates for only a subset of the applicants. To keep the task similar to Part 1, on each trial they were required to sequentially view all 40 applicants, even those for which they were not asked to

provide estimates. We elicited estimates only for applicants whose relative ranks could entail positive payoffs—i.e., only for applicants with relative ranks of 1, 2, and 3. To determine the applicants for which estimates are required, we partitioned the 40 applicants on each trial into 8 groups of 5, corresponding to the first 5 applicants, second 5 applicants, and so on. Then, within each partition, for each subject over the course of the 30 trials, we asked for a total of six estimates for applicants with relative rank 1, two estimates for applicants with relative rank 2, and a single estimate for applicants for relative rank 3. In total, we elicited 72 probability estimates for each subject. For estimates of applicants with relative rank 1, with probabilities $1/3$, $1/3$, and $1/3$ we asked the subject to estimate the probability that the applicant had an absolute rank of 1, 2, or 3, respectively. To be clear, for each applicant for whom an estimate was requested, a subject was only asked to estimate the probability of one particular absolute rank. For applicants with relative rank 2, estimates for absolute ranks of 2 and 3 were requested with equal probability. For those with relative rank 3, subjects were asked to only estimate the probability that the applicant’s absolute rank was 3. Rather than asking the subjects to input numerical probabilities, all the estimates were elicited by using a slider on the computer screen that allowed for estimates (in percent metric) between 0 and 100 in units of 1. The slider was always initially positioned at 0.

When the n applicants appear in a random order, the true probability that the j^{th} of n applicants whose relative rank is s has an absolute rank of a is given by:

$$P(A = a | R = s; j) = \frac{\binom{a-1}{s-1} \binom{n-a}{j-s}}{\binom{n}{j}}.$$

For example, the probability that the 15th out of 40 applicants whose relative rank is 1 has absolute ranks 1, 2, or 3 are 0.38, 0.24, and 0.15, respectively. The probability that the 15th applicant whose relative rank is 3 has an absolute rank of 3 is 0.05.

Part 1 Results

Earnings. The experimental subjects earned significantly less ($M=4.62$, $SD=0.55$) than they would have following the optimal policy (5.05), $t(29)=4.24$, $p<.001$.

Stopping Times. Using the same format as Fig. 1, the aggregate (cumulative) stopping times under the optimal policy and for the experimental subjects are displayed in Fig. 2. Figure 2 shows that the empirical stopping times are shifted to the left of the optimal stopping times, revealing that the subjects in Experiment 2 had the same tendency to terminate their search too early. This was statistically confirmed, as we again find that the subjects stopped the search significantly earlier ($t(61)=4.83, p<0.001; M=26.54, SD=3.09$) than expected under the optimal decision rule (29.27). Consistent with the results from Experiment 1, the mean stopping time in the first block of 30 trials ($M=25.89, SD=3.12$) was significantly smaller than in the second block ($M=27.20, SD=3.37, t(29)=3.60, p<0.001$).

--Insert Fig. 2 about here--

Comparison of Alternative Behavior Decision Rules. We again estimated parameter values for each of the three decision rules tested in Experiment 1. Using the same format as Experiment 1, the MTR again (overwhelmingly) best accounts for the stopping data. The results for each of the three decision rules are displayed in Table 3. Consistent with the early stopping results, we see that the best-fitting thresholds for the MTR are shifted to the left of the optimal thresholds.

--Insert Table 3 about here--

Given the remarkable consistency of the search results across Experiments 1 and 2, in the rest of this section we focus on the probability estimation data and their relation to the stopping data.

Overall Accuracy of Probability Estimates. Figure 3 exhibits the mean estimated probabilities (averaged over the 30 subjects) as a function of the true (objective) probability. Over virtually the entire range of true probabilities, the subjects overestimated the true probabilities. Analysis of the individual subject estimates reveals the same pattern for every subject.

--Insert Fig. 3 about here--

Probability Estimates over Periods. One explanation for the early stopping in Part 1 is that the subjects misperceive the probability that hiring early applicants will produce positive payoffs. Recall that subjects in Part 1 earned positive payoffs for hiring applicants with absolute ranks 1, 2, or 3, and nothing otherwise. Denoting a subject's estimate that applicant j whose relative rank is s has an absolute rank of a

by $sp(j,s,a)$, we can infer that a subject's estimate that hiring applicant j will result in positive payoff by

$\pi(j,s) = \sum_{a=s}^3 sp(j,s,a)$. Likewise, we denote the true probability of positive payoff by $\pi^*(j,s)$. Figure 4

exhibits the values of $\pi^*(j,s)$ (top panel) and $\pi(j,s)$ (bottom panel) across periods ($j=1,\dots,40$) for $s=1, 2$, and 3. Both $\pi^*(j,s)$ and $\pi(j,s)$ are seen to increase in j and decrease in s . However, the discrepancy between the true and (derived) estimated values is remarkable: The subjects strongly overestimate $\pi^*(j,s)$. Even more interesting is the finding that the estimates are often *supercertain*—i.e., in many cases $\pi(j,s) > 1$.

--Insert Fig. 4 about here--

The vertical lines in Fig. 4 represent the optimal (top panel) and mean estimated observed (bottom panel) thresholds under the MTR. Recall (Table 3) that the optimal thresholds are $r^*_1=14$, $r^*_2=29$, $r^*_3=37$, and the mean estimated thresholds are $r_1=13$, $r_2=22$, $r_3=30$. As in Experiment 1, the estimated thresholds are shifted to the left of the optimal thresholds, consistent with the observed early stopping behavior. Consider period 13, where the average r_1 threshold is 13; here, we observe (Fig. 4) that $\pi(j,1)$ is about 1.25, whereas $\pi^*(j,1)$ is 0.70. Put differently, the actual probability of earning a positive payoff for a 13th applicant with a relative rank of 1 is 0.70, but the subjects' mean probability estimates for the 13th applicant having absolute ranks of 1, 2, or 3 sum to more than 1. Obviously, the subjects do not believe that it is more than certain that hiring the 13th applicant in this case will result in positive payoff; however, it is safe to infer that they do believe that the probability of positive payoff is considerably greater than it actually is. The error of the probability estimates was greatest in the early periods where most true probabilities were well below 1. In the later periods—between 30 and 40—the subjects' estimates tended to become more accurate, and we actually observe that $\pi(j,1)$ and $\pi(j,2)$ tend to decrease in this range.

Relation Between Probability Estimates and Stopping Results. The results just exhibited in Fig. 4 are at the *aggregate* level. Next, we examine the results at the *individual* level. We ask: Can the estimation results be used to predict the stopping results for an individual subject? To address this question, we computed a measure of the degree to which subjects' probability estimates deviated from the true

underlying probabilities. Specifically, we computed the mean difference between the subjects' estimates and the true probabilities, \bar{d} . When positive, \bar{d} reflects a tendency to overestimate probabilities, and when negative to underestimate them. Thus, if the subjects' overconfidence in obtaining positive payoffs drives their stopping decisions, we should find that \bar{d} is negatively related to their mean stopping position \bar{m} . Since we again observe learning during early trials of the stopping task in Experiment 2, we used the last 40 trials in Part 1 to compute \bar{m} separately for each subject. The Spearman rank-order correlation between \bar{d} and \bar{m} was, indeed, negative and significant, $\rho_s(28)=-0.34, p=0.03$. (Using all 60 trials from the stopping task to compute \bar{m} leads to the same qualitative conclusion, though with the noisier estimates the effect is attenuated, $\rho_s(28)=-0.24, p=0.10$.)

Discussion

The probability estimation results provide insight into the early stopping behavior observed in Experiments 1 and 2. We have reported strong evidence that subjects overestimate the probability of obtaining positive payoffs across the range of applicant positions (periods). In fact, we find that the sum of the subjects' estimates often exceeded 1. This result is consistent with research showing that probability estimates are often subadditive; that is, the sum of probability estimates assigned to mutually exclusive sub-events often exceeds the probability assigned to the event that is their union. This finding can be explained by Tversky and Koehler's (1994) support theory (see also Rottenstreich & Tversky, 1997). Under support theory, the (subjective) probability assigned to the event "A rather than not A" is given by

$$P(A, B) = \frac{\phi(A)}{\phi(A) + \phi(\neg A)},$$

where $\phi(A)$ returns the *support* one accrues for the hypothesis A . Support is taken to be a measure of the strength of evidence in favor of the evaluated hypothesis. Research has shown that the focal hypothesis, in this case A , typically receives greater consideration than its alternative and thereby has its support increased to a greater degree. Whenever a DM in the GSP encounters an applicant with relative rank s and

has to make a stopping decision, she should consider the expected earnings of selecting that applicant, which requires that she—in some fashion (not necessarily consciously)—estimate the probabilities of the applicant’s possible absolute ranks given his relative rank. Presumably, the absolute ranks about which she must make the estimations that are most salient are those corresponding to positive payoffs; those absolute ranks that do not produce positive payoffs are perhaps lumped together as “all other absolute ranks.” As a result, then, the positive payoff absolute ranks, which are focal, will receive disproportionate weight in the decision to stop. This overoptimism, which can result from insufficiently considering that hiring the applicant *may not result in a positive payoff*, is sufficient to bias the DMs to terminate their search too early. Consistent with this explanation, our probability estimation data show that subjects strongly overestimate the probability that an applicant’s relative rank is one with a corresponding positive payoff; furthermore, at the individual level, the overestimation is negatively correlated with the mean length of search.

6. Discussion and Conclusions

We investigated individual decision behavior in a class of observation and selection problems with a finite number of applicants and rank-dependent payoffs in which payoffs increase monotonically in the quality of the selected applicant. In agreement with the results reported by Seale and Rapoport (1997, 2000), who tested experimentally two simpler variants of the CSP, our results show that subjects stop their search too early. This finding holds for two different payoff functions. In Experiment 1, the payoff function was *convex* in the quality of the selected applicant, with the payoffs dropping off at a rapid but diminishing rate in the quality of the selected applicant. Experiment 2 used a (piecewise) *linear* payoff scheme for the first three applicants; thus, the payoff difference for selecting applicants with true ranks of 1 versus 2 was the same as the difference for 2 versus 3. Seale (1996) experimentally studied an extension of the CSP in which the DM earned a payoff of 1 for selecting the best or second best applicant, and nothing otherwise. Similarly, Zwick et al. (unpublished) studied a CSP variant with payoffs of 1 for selecting any of the top three applicants. Note that the payoff structures used by Seale and by Zwick et al. were *concave*. Early stopping was also observed in these two studies. Thus, it seems unlikely that the

tendency to search insufficiently reported in the current study is simply an artifact of our two different payoff schemes. This bias has been reported under many payoffs schemes, and even in different countries.

Another concern might be that the early stopping bias is a result of using rank-dependent payoffs, and that the same effect might not be observed in full-information problems. Kogut (1990) and Rapoport and Tvserky (1970) compared sequential search behavior in full-information problems and also found that subjects tended to stop searching earlier than predicted by the optimal policy.² We can, therefore, rule out the possibility that our findings are due solely to the use of rank-dependent payoffs.

There are some qualifications to the early stopping finding. First, our results show that as more experience is gained in selecting applicants, the discrepancy between observed and optimal stopping decisions decreases. Decision makers experienced in this sort of sequential search may no longer exhibit this bias. Second, as Zwick et al. (2003) have shown in a considerably different variant of the CSP with backward solicitation, when fixed cost per search is charged the tendency to stop too early is reversed. It is yet to be determined experimentally whether this latter finding generalizes to the GSP. Third, the results reported by Seale and Rapoport, as well as the results of Experiments 1 and 2 reported here, apply to relatively large values of n , larger than sometimes encountered in practice. Caution should be exercised in extrapolating them to smaller values of n .

In summary, the current work extends previous work on the well-known secretary problem by presenting and empirically testing a generalization, the Generalized Secretary Problem (GSP), that better captures properties of many real-world sequential search tasks. Most often in sequential search, a DM derives utility that is monotonically increasing in the quality of her selection; the Classical Secretary Problem (CSP) does not capture this property. However, as in previous studies of the CSP, we again observed that DMs in the GSP tend to terminate their search too early and, consequently, fail to maximize expected earnings. Our results suggest that this bias may result from the DMs overestimating the quality

² It is difficult to assess optimality in partial-information search problems. To do so requires making strong assumptions about the mechanism by which DMs learn the characteristics of the distribution from which observations are sampled. That the distributions need not be learned—and cannot be learned—is another (methodological) virtue of no-information search problems.

of early applicants by failing to give sufficient weight to the prospect that better applicants are among those yet-to-be-seen.

References

- Bartoszynski, R. and Govindarajulu, Z. (1978). The secretary problem with interview cost. Sankhyà: The Indian Journal of Statistics, **40**, 11-28.
- Choe, K. I. and Bai, D. S. (1983). A secretary problem with backward solicitation and uncertain employment. Journal of Applied Probability, **20**, 891-896.
- 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. Organizational Behavior and Human Performance, **14**, 207-216.
- Ferguson, T. S. (1989). Who solved the secretary problem? Statistical Science, **4**, 282-296.
- 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.
- Kogut, C. A. (1990). Consumer search behavior and sunk costs. Journal of Economic Behavior and Organization, **14**, 381-392.
- Moriguti, S. (1993). Basic theory of selection by relative rank with cost. Journal of the Operations Research Society of Japan, **36**, 46-61.
- Mucci, A. G. (1973). Differential equations and optimal choice problems. Annals of Statistics, **1**, 104-113.
- Pressman E. L. and Sonin, I. M. (1972). The best choice problem for a random number of objects. Theory of Probability and Its Applications, **17**, 657-668.
- Rapoport, A. and Tversky, A. (1970). Choice behavior in an optional stopping task. Organizational Behavior and Human Performance, **5**, 105-120.
- Rottenstreich, Y. and Tversky, A. (1997). Unpacking, repacking, and anchoring: Advances in support theory. Psychological Review, **104**, 406-415.

- 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.
- Sardelis, D. A. and Vahahas, T. M. (1999). Decision making: A golden rule. Mathematical Monthly, **106**, 215-226.
- Seale, D. A. (1996) Sequential observation and selection with relative ranks: An empirical investigation of the secretary problem. Unpublished Ph.D. dissertation, University of Arizona.
- 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.
- Smith, M. H. (1975). A secretary problem with uncertain employment. Journal of Applied Probability, **12**, 620-624.
- Stein, W. E., Seale, D. A., and Rapoport, A. (2003). Analysis of heuristic solutions to the best choice problem. European Journal of Operational Research, **51**, 140-152.
- Tversky, A., & Koehler, D. (1994). Support theory: A nonextensional representation of subjective probability. Psychological Review, **101**, 547-567.
- Winkler, R. L. (1969). Scoring rules and the evaluation of probability assessors. Journal of the American Statistical Association, **64**, 1073-1078.
- Yang, M. C. K. (1974). Recognizing the maximum of a random sequence based on relative rank with backward solicitation. Journal of Applied Probability, **11**, 504-512.
- Yeo, A. J. and Yeo, G. F. (1994). Selecting satisfactory secretaries. Australian Journal of Statistics, **36**, 185-198.
- Zwick, R., Murphy, R., and Rapoport, A. (unpublished). Sequential search with recall and cost: An experimental study. Working paper. Department of Marketing, The Hong Kong University of Science and Technology.

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.

Acknowledgements

We gratefully acknowledge financial support by a contract F49620-03-1-0377 from the AFOSR/MURI to the Department of Industrial Engineering and the Department of Management and Policy at the University of Arizona. We would also like to thank the AE and three anonymous reviewers whose critical feedback helped us think more clearly about these problems.

Table 1. Optimal (r_x^*) and mean observed estimated threshold values (r_x) across subjects and trials for three different decision rules for Experiment 1.

x	1	2	3	4	5	6	Mean Violations
r_x^*	21	43	53	57	58	59	--
Mean r_x :MTR	12	22	28	35	40	44	13
Mean r_x :HRR	4	7	5	6	7	7	44
Mean r_x :SUAR	16	15	15	14	14	13	35

Table 2. Expected and mean empirical payoffs and stopping position for Experiments 1 and 2.

	Experiment 1		Experiment 2	
	Payoffs	Stopping Position	Payoffs	Stopping Position
<i>Limiting</i>	12.73	41.04	6.11	27.21
<i>Sample</i>	12.10	37.90	5.05	29.27
<i>Empirical</i>	10.06	35.34	4.24	26.54

Table 3. Optimal (r_x^*) and mean observed estimated threshold values (r_x) across subjects and trials for three different decision rules for Experiment 2.

x	1	2	3	Mean Violations
r_x^*	14	29	37	--
Mean r_x :MTR	13	22	30	12
Mean r_x :HRR	4	3	2	44
Mean r_x :SUAR	5	7	6	49

Figure 1. Observed and predicted cumulative distributions of stopping time across subjects by period of search for Experiment 1.

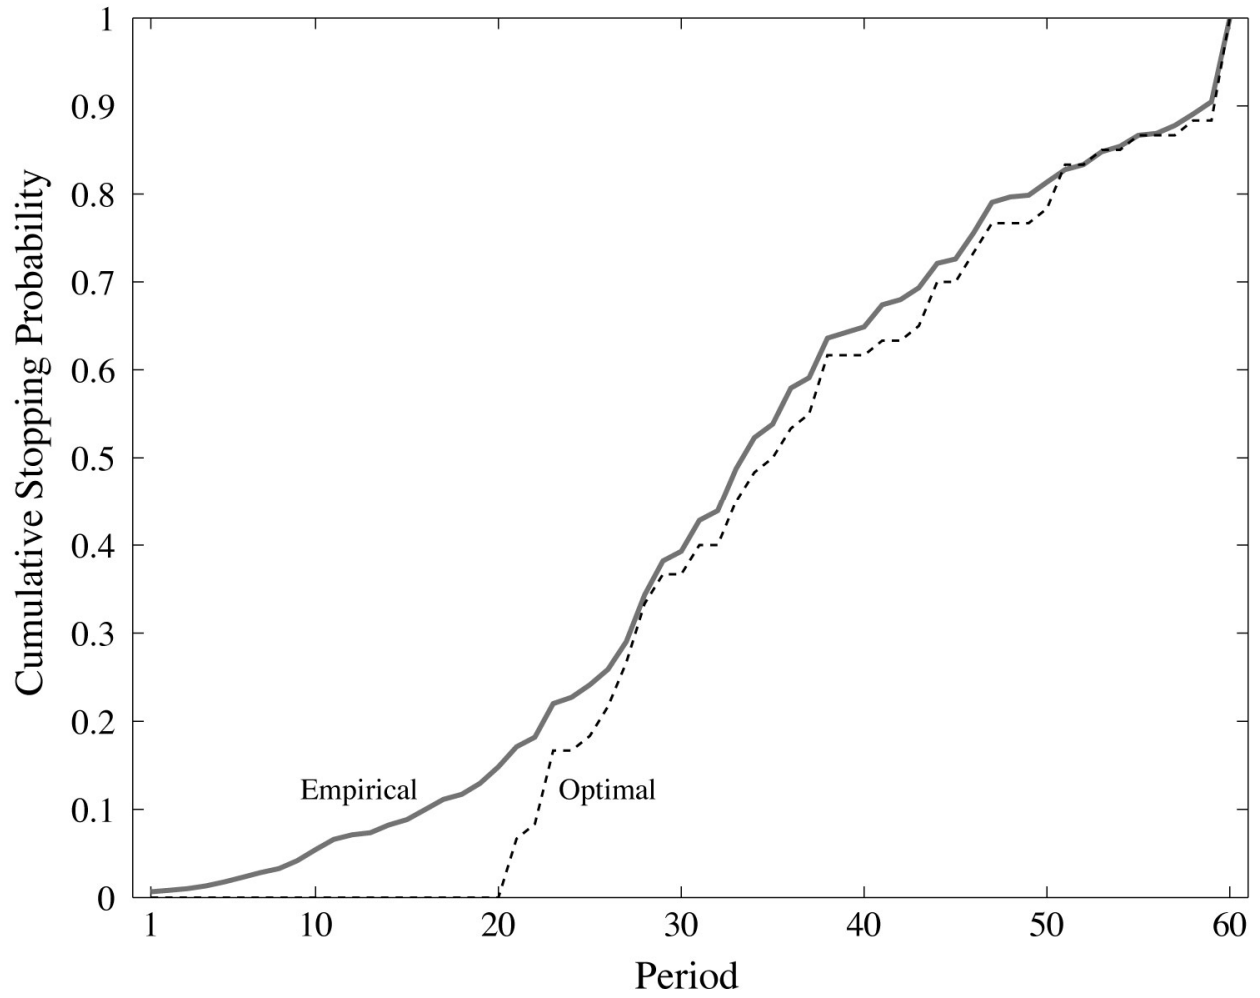


Figure 2. Observed and predicted cumulative distributions of stopping time across subjects by period of search for Experiment 2.




Figure 3. Mean estimated probability as a function of true probability for the estimate task (Task 2) in Experiment 2.




Figure 4. True (top panel) and derived (bottom panel) estimates of probability of obtaining positive payoffs for stopping in period j for an applicant with relative rank of s .

