On Generalized Secretary Problems

J. Neil Bearden and Ryan O. Murphy

University of Arizona, Department of Management and Organizations,
405 McClelland Hall, Tucson, AZ 85721, USA
e-mail: jneilb@gmail.com

Abstract. This paper is composed of two related parts. In the first, we present a dynamic programming procedure for finding optimal policies for a class of sequential search problems that includes the well-known "secretary problem". In the second, we propose a stochastic model of choice behavior for this class of problems and test the model with two extant data sets. We conclude that the previously reported bias for decision makers to terminate their search too early can, in part, be accounted for by a stochastic component of their search policies.

Keywords: sequential search, secretary problem, optimization

1 Introduction and Overview

The secretary problem has received considerable attention by applied mathematicians and statisticians (e.g., Ferguson, 1989; Freeman, 1983). Their work has been primarily concerned with methods for determining optimal search policies, the properties and implications of those policies, and the effects of introducing constraints on the search process (e.g., by adding interview costs). More recently, psychologists and experimental economists have studied how actual decision makers (DMs) perform in these sorts of sequential search tasks (e.g., Bearden, Rapoport and Murphy, 2004; Corbin, et al. 1975; Seale and Rapoport, 1997, 2000; Zwick, et al. 2003).

The current paper is composed of two main parts. First, we present a procedure for computing optimal policies for a large class of sequential search problems that includes the secretary problem. It is hoped that the accessibility of this procedure will encourage additional experimental work with this class of search problems. Second, we present a descriptive model of choice for the search problems, describe some of its properties, and test the model with two extant data sets. We conclude with a cautionary note on the difficulties researchers may face in drawing theoretical conclusions about the cognitive processes underlying search behavior in sequential search tasks.

2 Secretary Problems

2.1 The Problems

The *Classical Secretary Problem* (CSP) can be stated as follows:

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

The payoff maximizing strategy for the CSP, which simply maximizes the probability of selecting the best applicant, is to interview and reject the first $t-1$ applicants and then accept the first applicant thereafter with a relative rank of 1 (Gilbert and Mosteller, 1966). Further, they proved that t converges to n/e as n goes to infinity. In the limit, as $n \to \infty$, the optimal policy selects the best applicant with probability $1/e$. The value of t and the selection probability converge from above.

Consider a variant of the secretary problem in which the DM earns a positive payoff $\pi(a)$ for selecting an applicant with absolute rank a, and assume that $\pi(1) \geq \ldots \geq \pi(n)$. Mucci (1973) proved that the optimal search policy for this problem has the same threshold form as that of the CSP. Specifically, 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. As she gets deeper into the applicant pool her standards relax and she is more likely to accept applicants of lower quality.

We obtain what we call the *Generalized Secretary Problem* (GSP) by replacing **5** in the CSP, which is quite restrictive, with the more general objective function:

5'. The DM earns a payoff of $\pi(a)$ for selecting an applicant with absolute rank a where $\pi(1) \geq \ldots \geq \pi(n)$.

Clearly, the CSP is a special case of the GSP in which $\pi(1) = 1$ and $\pi(a) = 0$ for all $a > 1$. Results for other special cases of the GSP have appeared in the literature. For example, Moriguti (1993) examined a problem in which a DM's objective is to minimize the expected rank of the selected applicant. This problem is equivalent to maximizing earnings in a GSP in which $\pi(a)$ increases linearly as $(n - a)$ increases.

2.2 Finding Optimal Policies for the GSP

We will begin by introducing some notation. The orderings of the n applicants' *absolute ranks* is represented by a vector $\mathbf{a} = \left(a^1, \ldots, a^n\right)$, which is just a random permutation of the integers $1, \ldots, n$. The *relative rank* of the jth applicant, denoted r^j, is the number of applicants from $1, \ldots, j$ whose absolute rank is smaller than or equal to a^j. A *policy* is a vector $\mathbf{s} = \left(s^1, \ldots, s^n\right)$ of nonnegative integers in which $s^j \leq s^{j+1}$ for all $1 \leq j < n$. The policy dictates that the DM stop on the first applicant for which $r^j \leq s^j$. Therefore, the probability that the DM stops on the jth applicant, conditional on reaching this applicant, is s^j/j; we will denote this probability by Q^j. A DM's *cutoff* for selecting an applicant with a relative rank of r, denoted t_r, is the smallest value j for which $r \leq s^j$. Hence, a policy $\mathbf{s}$ can also be represented by a vector $\mathbf{t} = (t_1, \ldots, t_n)$. Sometimes, the cutoff representation will be more convenient. Again, a DM's *payoff* for selecting an applicant with absolute rank a is given by $\pi(a)$.

Given the constraint on the nature of the optimal policy for the GSP proved by Mucci (1973), optimal thresholds can be computed straightforwardly by combining numerical search methods with those of dynamic programming. We will describe below a procedure for doing so. A similar method was outlined in Lindley (1961) and briefly described by Yeo and Yeo (1994).

The probability that the jth applicant out of n whose relative rank is r^j has an absolute (overall) rank of a is given by (Lindely, 1961):

$$Pr\left(A = a | R = r^j\right) = \frac{\binom{a-1}{r-1}\binom{n-a}{j-r}}{\binom{n}{j}}, \tag{1}$$

when $r^j \leq a \leq r^j + (n-j)$; otherwise $Pr(A = a|R = r^j) = 0$. Thus, the expected payoff for selecting an applicant with relative rank r^j is:

$$E\left(\pi^j | r^j\right) = \sum_{a=r^j}^{n} Pr\left(A = a | R = r^j\right) \pi(a)\,. \tag{2}$$

The expected payoff for making a selection at stage j for some stage j policy $s^j > 0$ is:

$$E\left(\pi^j | s^j\right) = \left(s^j\right)^{-1} \sum_{i=1}^{s^j} E\left(\pi^j | r^j = i\right)\;; \tag{3}$$

otherwise, when $s^j = 0$, $E\left(\pi^j | s^j\right) = 0$. Now, denoting the expected payoff for starting at stage $j+1$ and then following a fixed threshold policy $(s^{j+1}, \ldots, s^n)$ thereafter by v^{j+1}, the value of v^j for any $s^j \leq j$ is simply:

$$v^j = Q^j E\left(\pi^j | s^j\right) + \left(1 - Q^j\right) v^{j+1}\,. \tag{4}$$

Since the expected earnings of the optimal policy at stage n are $v^n = n^{-1}\sum_{a=1}^{n}\pi(a)$, we can easily find an s^j for each j ($j = n-1,\ldots,1$) that maximizes v^j by searching through the feasible s^j; the expected earnings of the optimal threshold s^{j*} we denote by v^{j*}. These computations can be performed rapidly, and the complexity of the problem is just linear in n[1]. From the monotonicity constraint on the s^j, the search can be limited to $0 \le s^j \le s^{j+1}$. Thus, given v^{n*}, starting at stage $n-1$ and working backward, the dynamic programming procedure for finding optimal policies for the GSP can be summarized by:

$$s^{j*} = \arg \max_{s\in\{0,\ldots,s^{j+1*}\}} v^j \,. \tag{5}$$

The expected payoff for following a policy $\mathbf{s}$, then, is:

$$E\left(\pi|\mathbf{s}\right) = \sum_{j=1}^{n}\left[\prod_{i=0}^{j-1}\left(1-Q^i\right)\right]Q^j E\left(\pi^j|s^j\right) = v^1, \tag{6}$$

where $Q^0 = 0$. The optimal policy $\mathbf{s}^*$ is the policy $\mathbf{s}$ that maximizes Eq. 6. Denoting the applicant position at which the search is terminated by m, the probability that the DM stops on the $(j < n)$th applicant is:

$$Pr\left(m = j\right) = \left[\prod_{i=0}^{j-1}\left(1-Q^i\right)\right]Q^j, \tag{7}$$

and the expected stopping position is (Moriguti, 1993):

$$E\left(m\right) = 1 + \sum_{j=1}^{n-1}\left[\prod_{i=1}^{j}\left(1-Q^i\right)\right]. \tag{8}$$

Optimal cutoffs for several GSPs are presented in Table 1. In the first column, we provide a shorthand for referring to these problems. The first one, GSP1, corresponds to the CSP with $n = 40$. The optimal policy dictates that the DM should search through the first 15 applicants without accepting any and then accept the first one thereafter with a relative rank of 1. GSP2 corresponds to another CSP with $n = 80$. In both, the DM should search through roughly the first 37% and then take the first encountered applicant with a relative rank of 1. These two special cases of the CSP have been studied experimentally by Seale and Rapoport (1997). GSPs 3 and 4 were discussed in Gilbert and Mosteller (1966), who presented numerical solutions for a number of problems in which the DM earns a payoff of 1 for selecting either the best or second best applicant and nothing otherwise. GSPs 5 and 6 correspond to those studied by Bearden, Papoport and Murphy (2004) in Experiments 1 and 2,

[1] More elegant solutions can be used for special cases of the GSP. The method described here can be easily implemented for all special cases of the GSP.

respectively. In the first, the DM searches through the first 13 applicants without accepting any; then between 14 and 28 she stops on applicants with relative rank of 1; between 29 and 36, she takes applicants with relative rank 1 or 2; etc. Finally, GSP7 corresponds to the rank-minimization problem studied by Moriguti (1993). The results of our method are in agreement with all of those derived by other methods.

When inexperienced and financially motivated decision makers are asked to play the GSP in the laboratory, they have no notion of how to compute the optimal policy. Why then should one attempt to test the descriptive power of the optimal policy? One major reason is that tests of the optimal policies for different variants of the GSP (e.g. Bearden, Rapoport and Murphy, 2004; Seale and Rapoport, 1997, 2000; Zwick, et al. 2003) may provide information on the question of whether DMs search too little, just enough, or too much. This question has motivated most of the research in sequential search in economics (e.g., Hey, 1981, 1982, 1987) and marketing (e.g., Ratchford and Srinivasan, 1993; Zwick, et al.). However, tests of the optimal policy do not tell us what alternative decision policies subjects may be using in the GSP. And because they prescribe the same fixed threshold values for all subjects, they cannot account for within-subject variability across iterations of the sequential search task or between-subject variability in the stopping behavior.

Seale and Rapoport (1997, 2000) have proposed and tested three alternative decision policies in their study of two variants of the CSP. These decision policies (descriptive models) are not generalizable in their present form to the GSP. Moreover, because all of them are deterministic, they cannot account for within subject variability in stopping times across trials. Rather than attempting to construct more complicated deterministic choice models for the GSP, with a considerable increase in the number of free parameters, we propose an alternative stochastic model of choice for the GSP. Next, we describe the model and discusses its main properties. Then we summarize empirical results from some previous studies of the GSP and use them to test the model. Finally, we conclude by discussing some problems that arise in drawing theoretical conclusions about choice behavior in the GSP and related sequential search tasks.

3 A Stochastic Model of Choice in Secretary Problems

3.1 Background

Stochastic models have a long history in psychological theories. As early as 1927, L.L. Thurstone posited that observed responses are a function of an underlying (unobservable) component together with random error (Thurstone, 1927a, 1927b). For reviews of the consequences of Thurstone's ideas, see Bock and Jones (1968) and Luce (1977, 1994).

Table 1. Optimal policies for several GSPs

GSP	n	$\pi = (\pi(1), \ldots, \pi(n))$	$\mathbf{t}^* = (t_1^*, \ldots, t_n^*)$	$E(\pi \mid \mathbf{s}^*)$	$E(m)$
1	40	$(1, 0, \ldots, 0)$	$(16, 40, \ldots, 40)$	.38	30.03
2	80	$(1, 0, \ldots, 0)$	$(30, 80, \ldots, 80)$	.37	58.75
3	20	$(1, 1, 0, \ldots, 0)$	$(8, 14, 20, \ldots, 20)$	.69	14.15
4	100	$(1, 1, 0, \ldots, 0)$	$(35, 67, 100, \ldots, 100)$	.58	68.47
5	40	$(15, 7, 2, 0, \ldots, 0)$	$(14, 29, 37, 40, \ldots, 40)$	6.11	27.21
6	60	$(25, 13, 6, 3, 2, 1, 0, \ldots, 0)$	$(21, 43, 53, 57, 58, 59, 60, \ldots, 60)$	12.73	41.04
7	25	$(25, 24, 23, \ldots, 1)$	$(8, 14, 17, 19, 21, 22, 23, 23, 24, 24, 24, 25, \ldots, 25)$	22.88	14.46

More recently, theorists have shown that unbiased random error in judgment processes can produce seemingly biased judgments. For example, Erev, et al. (1994) have shown that symmetrically distributed random error can produce confidence judgments consistent with overconfidence even when the underlying (unperturbed) judgments are well-calibrated (see also, Juslin, et al. 1997; Pfeifer, 1994; Soll, 1996).

In related work, Bearden, Wallsten and Fox (2004) have shown that unbiased random error in the judgment process is sufficient to produce subadditive judgments. Suppose we have an event X that can be partitioned into k mutually exclusive and exhaustive subevents $X = \bigcup_{i=1}^{k} X_i$. Denote a judge's underlying (or true) probability estimate for X by $C(X)$ and her overt expression of the probability of X by $R(X)$. Bearden et al. assumed that $R(X) = f(C(X), e)$, where e is a random error component that is just as likely to be above as below $C(X)$. They proved that under a range of conditions $R(X)$ is regressive, i.e., it will be closer than $C(X)$ to .50. As a result, the overt judgment for X can be smaller than the sum of the judgments for the X_i, even when $C(X) = \sum_i C(X_i)$. Put differently, the overt judgments can be subadditive even when the underlying judgments are themselves additive. A considerable body of research has focused on finding high-level explanations such as availability for subadditive judgments (e.g., Rottenstreich and Tversky, 1997; Tversky and Koehler, 1994). Bearden et al. simply demonstrated that unbiased random error in the response process is sufficient to account for the seemingly biased observed judgments. One need not posit higher-level explanations. We follow this line of research and look at the effects of random error in the GSP.

Empirical research on the GSP has consistently shown that DMs exhibit a bias to terminate their search too soon (Bearden, Rapoport and Murphy, 2004; Seale and Rapoport, 1997, 2000). At the level of description, this observation is undeniable. However, researchers have gone beyond this observation by offering psychological explanations to account for the bias. In a paper on the CSP, Seale and Rapoport (1997) suggested that the bias results from an endogenous search cost: Because search is inherently costly (see, Stigler, 1961), the DM's payoff increases in the payoff she receives for selecting the best applicant but decreases in the amount of time spent searching. Therefore, early stopping may be the result of a (net) payoff maximizing strategy. Bearden, Papoport and Murphy (2004) offered a different explanation. They had DMs estimate the probability of obtaining various payoffs for selecting applicants of different relative ranks in different applicant positions. Based on their findings, they argued that the bias to terminate the search too soon in a GSP results from DMs overestimating the payoffs that would result from doing so.

In Sect. 3.2 we present a simple stochastic model of search in the secretary problem and show that it produces early stopping behavior even when DMs use decision thresholds that are *symmetrically* distributed about the optimal thresholds.

3.2 The Model

Recall that under the optimal policy for the GSP, the DM stops on some applicant j if and only if the applicant's relative rank does not exceed the DM's threshold for that stage (i.e., when $r^j \leq s^{j*}$). Experimental results, however, conclusively show that DMs do not strictly adhere to a deterministic policy of this sort. Rather, we posit that DMs' thresholds can be modelled as random variables. Each time the DM experiences an applicant with a relative rank r, she is assumed to sample a threshold from her distribution of thresholds for applicants with relative rank r; then, using the sampled threshold, she makes a stopping decision[2]. Denoting the sampled threshold σ_r, she stops on an applicant with relative rank r^j if and only if $r^j \leq \sigma_r$. (Note that at each stage j, the DM samples from a distribution that depends on the relative rank of the applicant observed at that stage. The distribution *is not* conditional only on the stage; it is only conditional on the relative rank of the observed applicant at that stage.) We assume that the probability density function for the sampled threshold is given by:

$$f\left(\sigma_r\right) = \frac{e^{-(\sigma_r - \mu_r)/\beta_r}}{\beta_r \left[1 + e^{-(\sigma_r - \mu_r)/\beta_r}\right]^2} . \tag{9}$$

Consequently, conditional on being reached, the probability that an applicant with relative rank r^j is selected is:

$$Pr\left(r^j \leq \sigma_r\right) = \frac{1}{1 + e^{-(j-\mu_r)/\beta_r}} . \tag{10}$$

We assume that $\mu_1 \leq \ldots \leq \mu_n$ and $\beta_1 \geq \ldots \geq \beta_n$. This is based on the constraint of the GSP that payoffs are nonincreasing in the absolute rank of the selected applicant. Hence, it seems reasonable to assume that $Pr\left(r^j \leq \sigma_r\right) \geq Pr\left(r'^j \leq \sigma_{r'}\right)$ whenever $r \leq r'$. That is, the DM should be more likely to stop on any given j whenever the relative rank of the observed applicant decreases. The constraints on the ordering of μ and β do not guarantee this property but do encourage it[3].

Note that the model approaches a deterministic model as $\beta_r \to 0$ for each r. Further, the optimal policy for an instance of a GSP obtains when β_r is small (near 0) and $t_r^* - 1 < \mu_r < t_r^*$ for each r.

Examples of the distributions of thresholds and resulting stopping probabilities for a possible DM are exhibited in Fig. 1 for the GSP2 (i.e., for

---

[2] The thresholds are, of course, unobservable. The model specified here as an as-if one: We are merely suggesting that the DM's observed behavior is in accord with her acting as if she is randomly sampling thresholds subject to the constraints of the model we propose.

[3] Adding the strong constraint that $Pr\left(r^j \leq \sigma_r\right) \geq Pr\left(r'^j \leq \sigma_{r'}\right)$ for all $r \leq r'$ makes dealing with the model too difficult. The numerical procedures used below to derive maximum likelihood estimates of the model's parameters from data would be infeasible under the strong constraint.

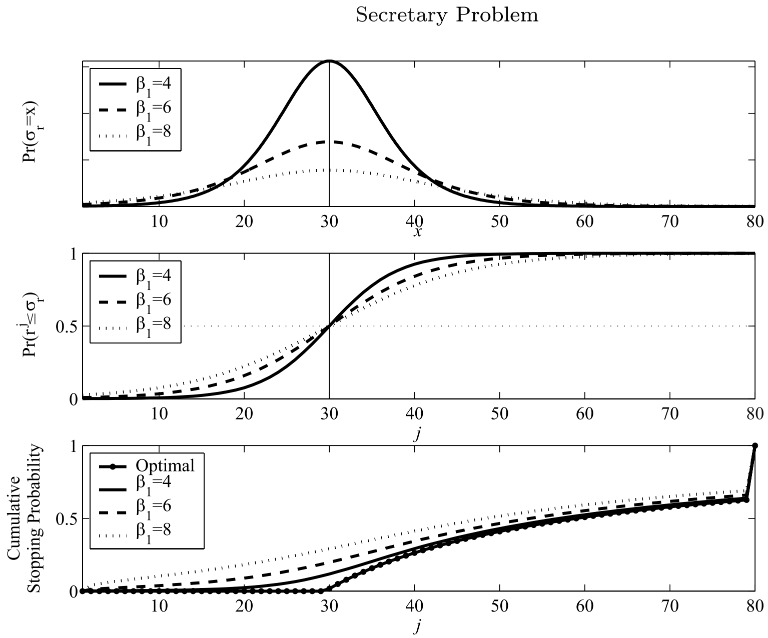

Fig. 1. Hypothetical threshold distributions and resulting stopping probabilities (conditional and cumulative) for one of the GSPs (GSP2) studied by Seale and Rapoport (1997) for various values of β. These results are based on $\mu_1 = t_1^*$. The *cumulative stopping probabilities* under the optimal policy are also shown in the bottom panel

a CSP with $n = 80$). In all cases shown in the figure, $\mu_1 = t_1^*$; that is, all of the threshold distributions are centered at the optimal cutoff point for the problem. The top panel shows the pdf of the threshold distribution. The center panel shows that for $j < \mu_1$ the probability of selecting a candidate (i.e., an applicant with a relative rank of 1) increases as β increases; however, for $j > \mu_1$, the trend is reversed. The bottom panel shows the probability of stopping on applicant j or sooner for the model and also for the optimal policy. Most importantly, in this example we find that the propensity to stop too early increases as the variance of the threshold distribution (β) increases, and in none of the model instances do we observe late stopping.

Under the model, the probability that the DM stops on the $(j < n)$th applicant, given that she has reached him, is:

$$\hat{Q}^j = \sum_{r^j=1}^{j} \frac{1}{j} Pr\left(r^j \leq \sigma_r\right) . \qquad (11)$$

Replacing Q^j in Eq. 8 with $\hat{Q}^j$, we can easily compute the *model expected stopping position*. Some examples of model expected stopping positions for various values of μ_1 and β_1 for the GSP2 are presented in Table 2. Several features of the $E(m)$ are important. First, whenever $\mu_1 < t_1^*$, the expected stopping position under the model is smaller than the expectation under the optimal policy. Second, even when $\mu_1 \geq t_1^*$ and β_1 is non-negligible, we find that the model tends to stop sooner than the optimal policy. Also, when $t_1^* - 1 < \mu_1 < t_1^*$ (that is, when the mean of the model threshold distribution is just below the optimal cutoff), the expected stopping position under the model is *always* less than under the optimal. Finally, as β increases, the expected stopping position decreases. In other words, as the variance of the threshold distribution increases, the model predicts that stopping position move toward earlier applicants. This general pattern of results obtains for the other GSPs as well.

The optimal policies for the GSP are represented by integers, but we are proposing a model in which the thresholds are real valued (and can even be negative); hence, some justification is in order. Using Eq. 10 to model choice probabilities has a number of desirable features. First, we can allow for shifts in both the underlying thresholds (or the means of the threshold distributions) by varying μ_r, and we can control the steepness of the response function about a given μ_r by β_r. As stated above, this can (in the limit) allow us to model both deterministic policies and noisy policies. The logistic distribution was chosen for its computational convenience (its CDF can be written in closed form); we have tried other symmetric distributions (e.g., the normal) and reached roughly the same conclusions that we report here for the logistic. (Actually, the tails of the normal distribution tend to be insufficiently fat to well-account for the empirical data.) Again, we desire a distribution with a symmetric PDF to model the thresholds in order for the thresholds to be *unbiased*. Empirical data show that DMs in secretary search tasks tend to terminate their search too early. We wish to demonstrate that this may result from an essentially unbiased stochastic process.

Table 2. Expected stopping times under the model for the GSP2 for different values of β_1 and μ_1. Keep in mind that $E(m) = 58.75$ under the optimal policy and $t_1^* = 30$. The average value of m for this problem in Seale and Rapoport (1997) is 43.61

β_1	$E(m \mid \mu_1 = 25)$	$E(m \mid \mu_1 = 29.5)$	$E(m \mid \mu_1 = 30)$	$E(m \mid \mu_1 = 35)$
.01	53.83	58.74	59.24	63.80
1	53.72	58.65	59.15	63.73
2	53.29	58.32	58.83	63.48
4	51.08	56.72	57.29	62.33
8	41.46	48.54	49.27	55.94
10	36.63	43.55	44.29	51.27
12	32.62	39.03	39.74	46.59
16	26.86	32.04	32.63	38.57

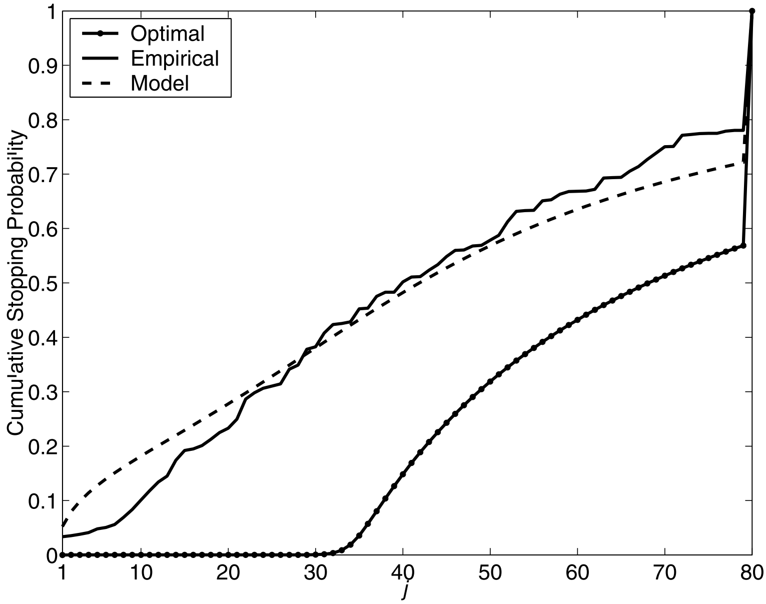

Fig. 2. Cumulative stopping probabilities for the GSP2 for the optimal and stochastic model policies and also for the empirical data reported by Seale and Rapoport (1997). The model probabilities are based on $\mu_1 = t_1^* = 30$ and $\beta_1 = 10$

Fig. 2 portrays optimal, empirical, and model cumulative stopping probabilities for the instance of the GSP that was studied empirically by Seale and Rapoport (1997). First, note that the empirical curve is shifted to the left of the optimal one. This indicates that DMs tended to stop earlier than dictated by the optimal policy. The model stopping probabilities are based on $\mu_1 = t_1^* = 30$, that is, the mean of the distribution from which the thresholds were sampled is set equal to the value of the optimal threshold. However, the model stopping probabilities are also shifted in the direction of stopping early. This is an important observation: In this example, the stochastic thresholds are distributed symmetrically about the optimal threshold and stopping behavior is biased toward early stopping. For example, it is just as likely that a DM's threshold will be 4 units above as below the optimal threshold, corresponding to too early and a too late thresholds, respectively; yet stopping behavior is biased toward early stopping.

The reason for early stopping under the stochastic model can be stated quite simply. First, there is a nonzero probability that a DM will stop sometime before it is optimal to do so; as a consequence, she will not have the opportunity to stop on time or stop too late. Secondly, though the threshold distribution itself is symmetric, the unconditional stopping probabilities are not. The probability of observing a given relative rank $r^j \leq j$ decreases in j. Consider $r^j = 1$. When $j = 1$, the probability of observing a relative rank of 1 is 1; when $j = 2$, the probability is 1/2; and in general it is $1/j$. Thus, for a given σ_r, the probability of stopping on applicant j is strictly decreasing in j. Therefore, properties of the problem itself can entail early stopping under the model. Researchers should, therefore, be cautious in attributing early stopping to general psychological biases.

Thus far we have only discussed the theoretical consequences of the stochastic model. Next, we evaluate the model using some of the empirical data reported in Seale and Rapoport (1997) and in Bearden, Papoport and Murphy (2004). We ask: Can the observed early stopping in these experiments be explained by unbiased stochastic thresholds?

3.3 Parameter Estimation

We estimated the model parameters for the stochastic choice model for individual subjects from two previous empirical studies of the GSP. Seale and Rapoport (1997) had 25 subjects play the GSP2 for 100 trials under incentive-compatible payoffs. They reported that their subjects exhibited a tendency to terminate their searches too early, and explained this by a deterministic cutoff rule of the same form as the optimal policy but whose cutoff was shifted to the left of the optimal cutoff. They evaluated alternative deterministic decision policies and concluded that the alternatively parameterized cutoff rule best accounted for the data. To determine a subject's cutoff – t_1, in our notation – they found the value of $1 \leq t_1 \leq 80$ that maximized the number of selection decisions compatible with the cutoff. For the GSP2, $t_1^* = 30$; Seale and Rapoport estimated that the modal cutoff for their subjects was 21.

Bearden, Rapoport and Murphy (2004) had 61 subjects perform the GSP6 for 60 trials under incentive-compatible payoffs. They, too, concluded that their subjects terminated search too early, and that the stopping behavior was most compatible with a threshold stopping rule. For the GSP6, $t_1^* = 21$, $t_2^* = 43$, $t_3^* = 53$, $t_4^* = 57$, $t_5^* = 58$, and $t_6^* = 59$; for their subjects, they estimated that the mean thresholds were $t_1 = 12$, $t_2 = 22$, $t_3 = 28$, $t_4 = 35$, $t_5 = 40$, and $t_6 = 44$. In both Seale and Rapoport and Bearden et al., the authors (implicitly) assumed that the subjects used deterministic or fixed thresholds. Hence, for a given subject, they could not account for stopping decisions inconsistent with that subject's estimated threshold.

In the current paper, we assume that the subjects' thresholds are random variables (whose pdf is given by Eq. 9) and use maximum likelihood procedures to estimate the parameters of the distribution from which the thresholds are sampled. For each set of data that we examine, the researchers reported learning across early trials of play, but in both, the choice behavior seems to have stabilized by the 20th trial. Hence, for the tests below, we shall eliminate the first 20 trials from each data set from the analyses, and we will assume that the choice probabilities are i.i.d.

For a given trial of a GSP problem, the DM observes a sequence of applicants and their relative ranks, and for each applicant she decides to either accept or continue searching. Denoting a decision function for applicant j by $\delta\left(r^j\right)$, we let $\delta\left(r^j\right) = 0$ if the DM *does not stop* on applicant j and $\delta\left(r^j\right) = 1$ if she *does stop*. Hence, decisions for a particular trial k can be represented by a vector $\Delta^k = \left(\delta\left(r^1\right), \ldots, \delta\left(r^m\right)\right) = (0, 0, \ldots, 1)$, where m denotes the position of the selected applicant. Under the stochastic model, if $m < n$, the likelihood of Δ^k can be written as:

$$L\left(\Delta^k | \mu, \beta\right) = \left[\prod_{i=1}^{m-1} Pr\left(r^i > \sigma_r\right)\right] Pr\left(r^m \leq \sigma_r\right). \tag{12}$$

When $m = n$ (i.e., when the DM reaches the last applicant, which she *must* accept), we simply omit the final term in Eq. 12 since the DM's choice is determined. Assuming independence, the likelihood of a DM's choice responses over K trials of the GSP is just:

$$L\left[\left(\Delta^1, \ldots, \Delta^K\right) | \mu, \beta\right] = \prod_{k=1}^{K} L\left(\Delta^k | \mu, \beta\right). \tag{13}$$

Due to the small numbers involved, it is convenient to work with the log of the likelihood, rather than the likelihood itself. Taking the log of Eq. 13, we get:

$$\ell\left[\left(\Delta^1, \ldots, \Delta^K\right) | \mu, \beta\right] = \sum_{k=1}^{K} \ln\left[L\left(\Delta^k | \mu, \beta\right)\right]. \tag{14}$$

For each subject we computed the parameters μ and β that maximized Eq. 14 under different constraints. We only estimated the parameters for relative ranks that can entail positive payoffs. For the GSP2, we restrict estimates to $r = 1$, and for the GSP6 to $1 \leq r \leq 6$. Therefore, we omit from the analyses trials on which the DM chose to stop on the applicants whose relative rank could not entail a positive payoff. Very likely these were errors. Fewer than 2% of the trials were omitted.

We are primarily interested in testing the following:

Optimal but stochastic threshold hypothesis: $\mu_r = t_r^*$ for all r.

If this hypothesis is supported, the bias toward early stopping behavior could be the result of the stochastic nature of the thresholds. We evaluate the optimal but stochastic threshold hypothesis (OBSTH) using standard likelihood ratio tests. Under the *constrained model*, we impose that $\mu_r = t_r^*$ for all r and allow the β_r to freely vary; under the *unconstrained model* we allow both the μ_r and β_r to freely vary. Denoting the maximum log-likelihood of the constrained model ℓ_c (based on Eq. 15) and of the unconstrained model ℓ_u, the likelihood ratio is:

$$LR = \left(\ell_c - \ell_u\right). \tag{15}$$

The statistic $-2LR$ is χ^2 distributed with degrees of freedom (df) equal to the number of additional free parameters in the unconstrained model. Hence, for the Seale and Rapoport (1997), $df = 1$; and for Bearden, Rapoport and Murphy (2004), $df = 6$.

A few words about estimating the model parameters are in order. To estimate the model parameters we used a constrained optimization procedure (fmincon) in Matlab. We imposed the constraint that $\mu_r \leq \mu_{r'}$ whenever $r \leq r'$, and imposed the corresponding constraint on the β parameters. For each subject, we used a large number of initial starting values. We are confident that the estimated parameters provide globally optimal results for each subject.

Seale and Rapoport Data

Based on the likelihood ratio test with $df = 1$, the OBSTH could not be rejected for 12 of the 25 experimental subjects at the $\alpha = .01$ level. Seale and Rapoport concluded that 21 of their 25 subjects had thresholds below the optimal cutoff. Our analyses suggest that they overestimated the number of subjects with biased thresholds. Fig. 3 shows a distribution of thresholds

J. Neil Bearden

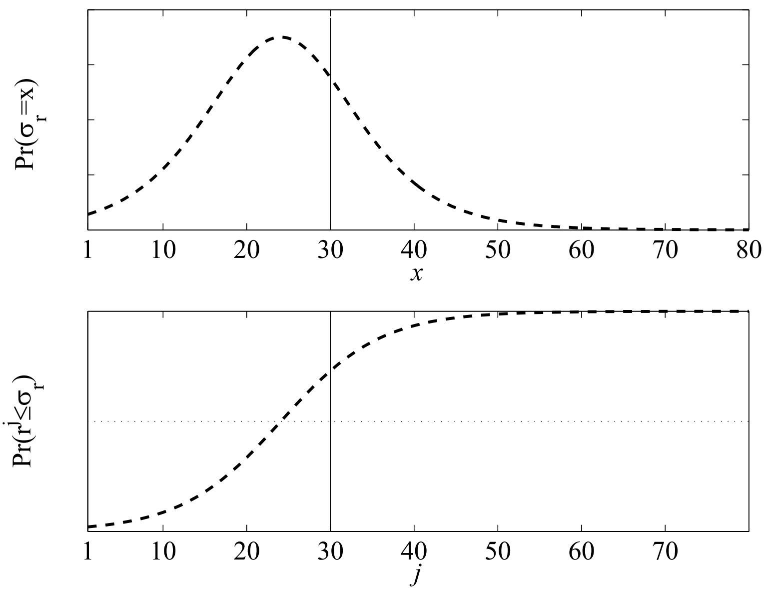

Fig. 3. Estimated threshold distribution and resulting stopping probabilities for the $n = 80$ CSP studied by Seale and Rapoport (1997) based on median estimated μ_1 and σ_1. The horizontal line is located at the optimal cutoff point ($t_1^* = 30$. The vertical line in the bottom panel corresponds to a probability of .50

(σ_1) that is based on the median estimated values of μ_1 and β_1 from the 25 experimental subjects. We find that the distribution of thresholds (based on the aggregate data) is, indeed, shifted to the left of the optimal cutoff, consistent with the observed early stopping behavior. Further, we find that the variance of the threshold distribution is considerably greater than 0. Thus, early stopping in Seale and Rapoport may be due both to thresholds that tend to be biased toward early stopping and also to stochastic variability in placement of the thresholds. Summary statistics from the MLE procedures are displayed in Table 3.

Bearden, Rapoport, and Murphy Data

The corresponding thresholds from Bearden, Rapoport and Murphy (2004) are displayed in Fig. 4. For these data, the OBSTH could not be rejected for 23 of the 61 subjects (i.e., for 37%). We find that the distribution of thresholds for $r = 1$ tends to be centered rather close to the optimal cutoff. Likewise for the $r = 6$ threshold. For $r = 2, \ldots, 5$, the thresholds tend to be shifted toward early stopping. The variances of the threshold distributions tend to decrease quite rapidly in r, but are all away from 0. Thus, as with the Seale and Rapoport (1997) data, the early stopping in the GSP6 seems to be driven by biased thresholds as well as the stochastic nature of those thresholds. Summary results are presented in Table 3.

Table 3. Summary of MLE results for Seale and Rapoport (n=80) condition and Bearden, Rapoport, and Murphy Experiment 1 data. Note: OBSTH compatible tests are based on $\alpha = .01$

Seale & Rapoport (1997) Data	
Number of subjects	25
Median μ	(24.08)
Median β	(5.97)
Median LR	4.19
Test *df*	1
OBSTH compatible	48%
Bearden, Rapoport, & Murphy (2004) Data	
Number of subjects	62
Median μ	$(23.16, 34.71, 43.96, 48.70, 54.49, 58.53)$
Median β	$(4.13, 3.56, 2.49, 1.24, 0.69, 0.43)$
Median LR	18.38
Test *df*	6
OBSTH compatible	37%

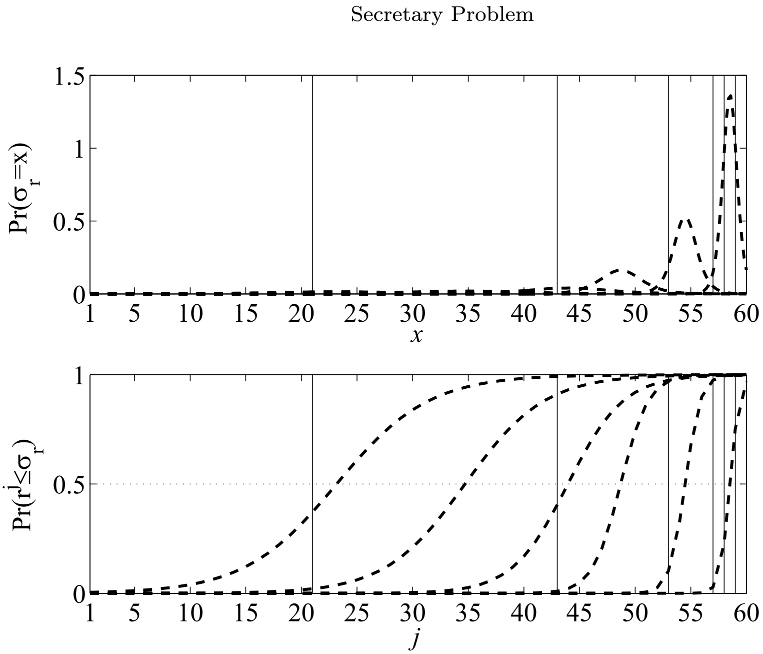

Fig. 4. Estimated threshold distribution and resulting stopping probabilities for the GSP6 studied by Bearden, Rapoport, and Murphy (2004). The curves are based on median estimated μ_r and σ_r $(r = 1, \ldots, 6)$, and are ordered from left $(r = 1)$ to right $(r = 6)$. Note that the variances of the $Pr\,(\sigma_r = x)$ distributions for $r = 1, 2, 3$ relative to the variance of the $r = 6$ distribution are quite small, making the resulting distributions rather flat and difficult to see

The estimation results suggest that researchers should be cautious in drawing conclusions about the underlying causes of early stopping in GSPs without taking random error into account. A straightforward question must be addressed before any claims are made: What does it mean for subjects to be biased to stop early? Is the statement merely an empirical one that describes that observed stopping behavior or does it have some theoretical import? Does the "bias" refer to a property of the choice process? Seale and Rapoport (1997) suggested that the subjects in their task seemed to follow cutoff policies that were of the same form as the optimal policy but were parameterized differently. Specifically, the cutoffs for the experimental subjects tended to be positioned earlier than the optimal cutoff. They suggested that the shift might be a compensation for endogenous search costs. Our results suggest, however, that the threshold may not have been biased toward early stopping for nearly 50% of the subjects in their $n = 80$ condition. For these subjects, stochastic thresholds centered at the optimal cutoff can account for the early stopping. Likewise, for roughly 37% of the subjects in Experiment 1 of Bearden, Rapoport and Murphy (2004), we can account for early stopping by the OBSTH.

We do not argue that early stopping is not driven by some genuine choice or judgment bias (e.g., by overestimating the probability of obtaining good payoffs for selecting early applicants). Rather, we simply wish to demonstrate that the effects of random error should be taken into consideration before drawing sharp conclusions about the magnitude of the effects of these potential biases on the stopping behavior.

4 Conclusions

We began this paper by presenting a simple dynamic programming procedure for computing optimal policies for a large class of sequential search problems with rank-dependent payoffs. The generality of the permissible payoff schemes allows a number of realistic (especially in contrast to the CSP, which has an only-the-best payoff scheme) search problems to be modelled.

Next, we described a simple stochastic model of choice behavior for the GSP and described some previous experimental results. The empirical results show that DMs tend to terminate their search too early relative to the stopping positions dictated by the optimal policy. Previous explanations for this finding have invoked endogenous search costs (Seale and Rapoport, 1997) and probability overestimation (Bearden, Rapoport and Murphy, 2004) as explanations. Our results suggest that at least part of the observed early stopping *can* be explained by unbiased stochastic variability in stopping thresholds.

Future research should contrast the endogenous search cost and probability overestimation explanations of early stopping in generalized secretary problems. Importantly, in such tests, researchers should be cautious of the contribution of random error to the apparently biased search behavior.

Acknowledgement. We thank Amnon Rapoport for his comments and Darryl Seale for providing us with his data. 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 Organizations at the University of Arizona.

References

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

Bearden, J.N., Wallsten, T.S., Fox, C.R. (2004), A stochastic model of subadditivity. Unpublished manuscript.

Bock, R.D., Jones, L.V. (1968), *The Measurement and Prediction of Judgment and Choice*. San Francisco: Holden-Day.

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

Erev, I., Wallsten, T.S., Budescu, D.V. (1994), Simulataneous over- and underconfidence: The role of error in judgment processes. *Psychological Review*, 101, 519–528.

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

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.

Juslin, P., Olsson, H., Björkman, M. (1997), Brunswikian and Thurstonian origins of bias in probability assessment: On the interpretation of stochastic components of judgment. *Journal of Behavioral Decision Making*, 10, 189–209.

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

Luce, R.D. (1977), Thurstone's discriminable process fifty years later. *Psychometrika*, 42, 461–489.

Luce, R.D. (1994), Thurstone and sensory scaling: Then and now. *Psychological Review*, 101, 271–277.

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.

Pfeifer, P.E. (1994), Are we overconfident in the belief that probability forecasters are overconfident. *Organizational Behavior and Human Decision Processes*, 58, 203–213.

Ratchford, B.T., Srinivasan, N. (1993), An empirical investigation of return to search. *Marketing Science*, 12, 73–87.

Rottenstreich, Y., Tversky, A. (1997), Unpacking, repacking, and anchoring: Advances in support theory. *Psychological Review*, 104, 406–415.

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

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

Stein, W.E., Seale, D.A., Rapoport, A. (2003), Analysis of heuristic solutions to the best choice problem. *European Journal of Operational Research*, 51, 140–152.

Soll, J.B. (1996), Determinants of overconfidence and miscalibration: The roles of random error and ecological structure. *Organizational Behavior and Human Decision Processes*, 65, 117–137.

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

Thurstone, L.L. (1927a), A law of comparative judgment. *Psychological Review*, 34, 273–286.

Thurstone, L.L. (1927b), Psychophysical analysis. *American Journal of Psychology*, 38, 368–389.

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

Yeo, A.J., Yeo, G.F. (1994), Selecting satisfactory secretaries. *Australian Journal of Statistics*, 36, 185–198.

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