TUTORIALS IN OPERATIONS RESEARCH
INFORMS—New Orleans 2005

 | ISBN 0000-0000
DOI 10.1287/educ.1053.0000

Chapter 1

Operations Research in Experimental Psychology

J. Neil Bearden
Department of Management and Policy, University of Arizona, Tucson, AZ, jneilb@gmail.com

Amnon Rapoport
Department of Management and Policy, University of Arizona, Tucson, AZ, amnon@u.arizona.edu

Abstract This chapter reviews some of the uses of operations research methods in experimental psychology. We begin by describing some basic methodological issues that arise in the study of human decision making. Next, we describe in more detail research in experimental psychology that has used methods common to operations research—such as dynamic programming—to understand sequential observation and selection behavior. We then suggest some ways in which experimental psychology and operations research can each provide the other with new research questions.

Keywords experimental psychology; experimental methodology; dynamic programming

---

Introduction

Theories of decision making are typically classified in one of three ways. *Normative* theories provide a framework for understanding how decisions should be made. Implicitly at least, these theories quite often rely on rational agents who have unlimited computational and cognitive capacities. *Descriptive* theories help explain how actual—rather than ideal—agents make decisions. Typically these theories emerge from experimental studies of human decision makers. *Prescriptive* studies of decision making are aimed at determining how actual decision makers (DMs) could behave more in accord with the dictates of normative theories with some systematic reflection (see, [42], for a classic treatment). In order to determine appropriate prescriptions it is necessary to understand how it is that decision making goes wrong, that is, how actual decision making departs from normative decision making [68].

The normative theories of decision making most often emerge from two fields: economic theory (including game theory) and operations research (OR) [34]. Utility theory, as developed by von Neumann and Morganstern [66], and later extended by Savage [54] to accomodate subjective probabilities, has received considerable attention by experimental psychologists and more recently by experimental economists. Just about any review of the field of behavioral decision theory—particulary those from the 1960s, 70s and 80s—dedicates considerable space to discussions of utility theories of different sorts and how actual human decision making compares to them (e.g., [50], [59]). Camerer [12] should be consulted for an up-to-date review of experimental tests of game-theoretic predictions of behavior. More complex decision problems that involve solving optimization problems of one sort or the other have received less experimental attention. In the area of dynamic decision making, normative theories that are standardly applied in operations research have been relatively neglected by experimental psychologists.

One might wonder why psychologists should be concerned with research in OR. If asked, a number of answers can be given. First, coming from an applied field, OR problems tend to have some correspondence to the kinds of problems likely to be faced by actual DMs. Second, OR problems tend to be clearly formulated, with assumptions and predictions fully specified, thereby increasing their testability. Quite often, little is required to turn an abstract OR problem into one that can be tested in the laboratory. Finally, optimality results from OR can be useful in understanding the decision policies employed by human DMs.

Before proceeding, a brief comment on the logic of using optimal normative theories to understand decision behavior is in order. Since the assumptions of a normative theory entail the predictions of the theory, comparing empirical decision data to normative predictions can be informative. First, consider the case in which the decision data are incompatible with the predictions. Supposing that one's experiment is designed and run properly (issues which we discuss in the next paragraph), simple *modus tollens* reasoning allows one to conclude that at least one assumption underlying the normative theory is violated in the human decision process. More formally, suppose a theory T is composed of a conjunction of assumptions $T = \{A_1 \wedge \ldots \wedge A_k\}$, and that T entails some outcome O. Then, if we observe $\neg O$, we know that $\neg A_i$ is true for at least one assumption A_i of the theory. One can then begin to get an idea of the nature of the difference between the normative model and the actual decision behavior by modifying the assumptions of the former in order to derive predictions consistent with the latter. One need not take the position that human decision behavior *ought* to be consistent with the dictates of normative theory in order for the theory to be useful. Without taking a stance on the ought issue, one can simply use the normative theory as a reference point—a benchmark—for evaluating behavior. Similar arguments have been offered elsewhere in favor of the use of optimality theory in theoretical biology (e.g., [30], [39]).

But there is an asymmetry: When behavior is consistent with normative predictions, fewer inferences are permissible. One cannot conclude that subjects are, for example, performing the calculations needed to arrive at the normative predictions—say by solving a linear programming problem; rather, one can only say that their decision behavior is consistent with the theory.

Of course, a number of auxiliary assumptions go into experimentally testing normative theories in the laboratory (or any theory, for that matter; see, e.g., [25], [41]). One must assume that the subjects fully understand the decision problems they face, that they are motivated, and that basic protocols of good experimental procedure are followed [26]. To ensure that subjects fully understand a task, instructions should clearly explain the task and provide a number of examples demonstrating the "rules of the game." By offering subjects non-negligible sums of money contingent on their performance, one can ensure that they approach laboratory problems with a level of seriousness at least close to that with which they approach problems outside the laboratory. If subjects behave sub-optimally because they do not really care about the outcomes of their decisions, then little is learned. If, however, they exhibit departures from normative theory when real money is on the line, something interesting may be going on. Good basic experimental protocols are too numerous to name and largely depend on the nature of the experimental task. We do not discuss these here.

Below, we review some of the work in experimental psychology that has relied on computational procedures from operations research. In particular, we focus on multi-stage decision problems that can be solved by dynamic programming. We do so for two reasons. First, this will provide a certain cohesion. Second, much of our own work involves experimental investigations of behavior in dynamic decision problems. Hence, when discussing these problems, we can give an insider's view of ways in which experimentalists think about the OR methods and how we think these methods can inform us about human cognition.

Although we restrict this review to experimental studies of sequential observation and selection behavior, OR researchers should be aware that other areas of dynamic decision making have been brought into the laboratory. Toda [63] pioneered the study of multi-stage decision behavior more than forty years ago. He devised a one-person game called the "fungus-eater" game, in which subjects were asked to control the sequential search of a robot that attempts to maximize the amount of a valuable resource ("uranium") it collects while ensuring that it has sufficient fuel. All of this happens on a hypothetical planet on which uranium and fuel are distributed randomly. Other dynamic decision problems that were studied experimentally in the last forty years, in which dynamic programming has been used to compute the optimal decision policy, include control of explosive systems (Rapoport, [43] and [44]); portfolio selection over discrete time and multi-stage gambling (Ebert [19]; Rapoport, Jones, and Kahan [47]); vehicle navigation (Anzai [1]; Jagacinkski and Miller [31]); health management (Kleinmuntz and Thomas [33]); inventory control (Rapoport [45]; Sterman [61]); and fire-fighting (Brehmer and Allard [10]). For cumulative reviews of this literature, see Rapoport [46], Brehmer [9], Sterman [62], Kerstholt and Raaijmakers [32], and more recently Busemeyer [11] and Diederich [17].

Optimal and Empirical Results on Optimal Stopping

The problems that we will discuss in this chapter involve sequential search, selection, and assignment. Problems of this sort are referred to in different literatures by different names. Sometimes they are dubbed *optimal* stopping problems; at other times one sees them referred to as *optional* stopping problems. Depending on their precise formulation and on the literature in which they appear (e.g., OR, economics, psychology), the problems are sometimes referred to as job-search problems, rent-seeking problems, secretary problems, etc. All of the problems in this class share an important feature with a number of interesting real-world choice problems, namely the choice alternatives are encountered *sequentially.* Given their multi-stage structure, dynamic programming (DP) methods are often invoked to find optimal decision policies for these problems.

Full-Information Problems

Suppose the DM can observe as many as n observations X_j $(j = 1, \ldots, n)$, but for each observation she must pay a fixed cost $c \geq 0$. The observations are drawn independently according to the density function $f(x)$, which is known to the DM. The DM's objective is to maximize the expected value of her selected observation minus her total costs; she cannot return to an observation once she has rejected it; and if she reaches the nth one, she must accept it. Because the DM knows the distribution from which the observations are sampled, this problem is known as a *full-information* optimal stopping problem.

Under the optimal policy, the DM should stop on observation j whenever the value of the draw x_j exceeds the expected value of moving to the next stage (i.e., to $j+1$) and behaving optimally thereafter, denoted V^*_{j+1}. Since the DM is forced to accept the nth observation, if reached, $V^*_n = \int_{-\infty}^{\infty} x f(x) dx - c$. Hence, at stage $n-1$ the optimal DM sets her *cutoff* to $s^*_{n-1} = V^*_n$. The cutoffs for each stage determine the observation values that the optimal DM finds acceptable at each stage; specifically, the DM stops on observation j whenever $x_j \geq s^*_j$. Thus, more generally, $V^*_{j+1} = s^*_j$. The optimal cutoffs for each stage $j = n-1, \ldots, 1$ can be obtained by computing the recurrence (Sakaguchi, [51]):

$$s^*_j = \int_{s^*_{j+1}-c}^{\infty} \left[x - \left(s^*_{j+1} - c\right)\right] f(x) dx + s^*_{j+1} - c. \tag{1}$$

Optimal policies for observations from a standard normal distribution with $n = 24$ are shown in Figure 1 for various values of c. These well-known results form the classical basis for optimal stopping problems. We present them here merely to set the stage for a discussion of experimental studies of optimal stopping, which we turn to next.

FIGURE 1. Optimal cutoffs s_j^* for various costs for the standard normal distribution.

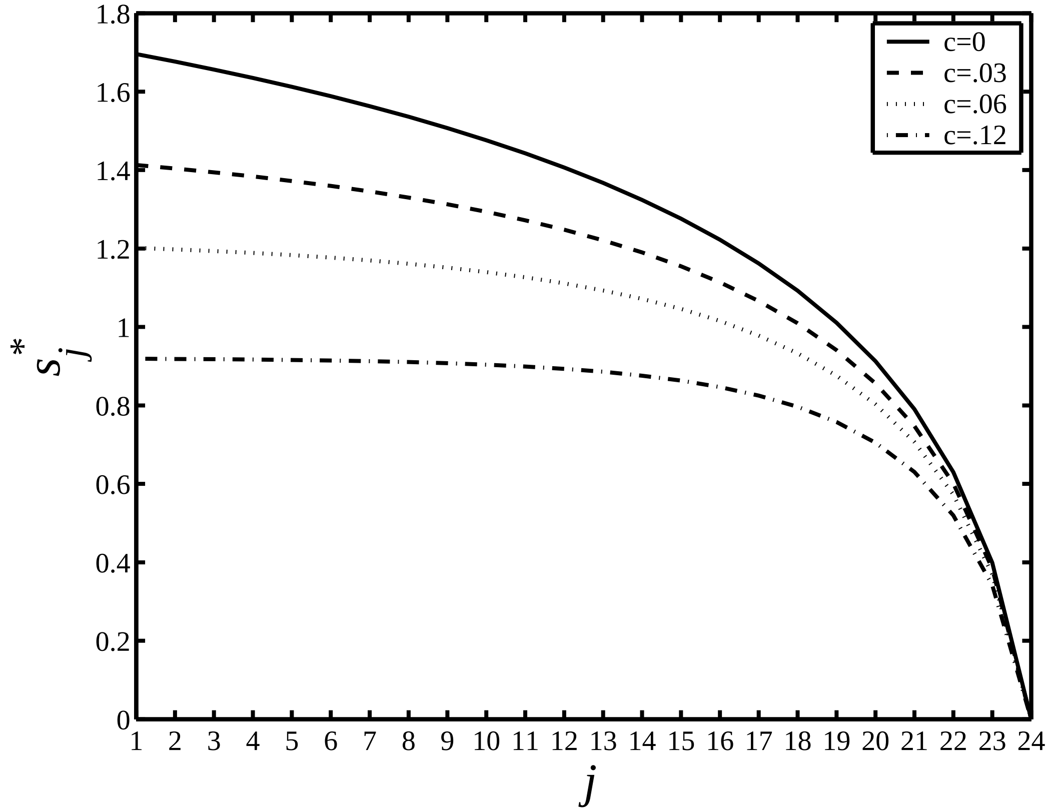

Experimental Studies of Full-Information Optimal Stopping

Rapoport and Tversky [49] conducted an experimental test of full-information optimal stopping. One difficulty in testing the full-information problem is that it requires that the DM know the distribution from which the observations are taken (cf. [48]). Simply telling a naïve subject that observations are taken from, for example, "a normal distribution with mean 100 and standard deviation 10" is obviously problematic, as the subject is unlikely to have a good grasp of precisely what this means. In order to get a strong test of the predictions of the optimal policy, one must ensure that the subjects have a good sense of the distribution. To guarantee that their subjects understood the nature of the distribution from which observations were taken, Rapoport and Tversky had them first perform a signal detection task in which observations were taken from two normal distributions A and B with a common standard deviation (167) but with different means (1630 and 1797, respectively). The subjects were required to judge whether a sample was from A or B. Over a six week period, five times a week, subjects observed a total of 7800 samples from each of the two distributions. Hence, presumably, the subjects ultimately had a good sense of the nature of the distributions.

Once the distributions were learned, the subjects performed optimal stopping tasks over several sessions. In each session, the subjects were told whether the values would be sampled from A or from B, and that $n = 24$. They were also told the cost c for each sampled observation, which was held constant within a trial. The cost values used in the study were 0, 5, 10, or 20. (Rapoport and Tversky studied optimal stopping problems both with and without recall. But for brevity, we will only discuss results from the latter.)

The main results are summarized in Table 1. For each cost condition, Rapoport and Tversky observed that, on average, subjects tended to sample fewer observations than predicted

Table 1. Mean number of observations for each cost condition from Rapoport and Tversky [49].

	$c = 0$	$c = 5$	$c = 10$	$c = 20$
Average Number of Draws	9.51	9.61	8.44	4.11
Expected Number of Draws	10.42	11.25	10.80	4.50

Note. The expected number of observations are not monotonically decreasing in cost because the values shown are based on the application of the optimal policy to the empirically sampled observation values, not on the limiting expectation under random sampling.

by the application of the optimal policy. Rapoport and Tversky suggested that the stopping results were reasonably consistent with subjects using a fixed threshold cutoff rule. Under this decision rule, the subject stops on an observation x_j whenever $x_j \geq s$, where s is fixed for all j ($j = 1, \ldots, n$). Thus, assuming this is the correct model for the choice behavior, the subjects were insensitive to the search horizon, and sought a single "target" value.

Other experimental studies of full-information optimal search have examined the effects of allowing the recall of previously encountered observations [48]; others (e.g., [27], [28], [29]) have more closely examined the factors that influence decisions to stop, such as the influence of the history of observation values (e.g., whether the sequence has been increasing or decreasing).

In our view, the main shortcoming of full-information optimal stopping problems—when taken as models of actual human search problems—is the assumption that $f(x)$ is known to the DM. This assumption is critical for testing the optimal search model because the values of s_j^* are located at the right tail of the distribution $f(x)$, and are, therefore, very sensitive to deviations from them. Perhaps in many situations DMs do have a "good sense" of the operative distribution; in many others, we suspect, this condition is not met. Next, we turn to *no-information* search problems that do not require that the DM have any knowledge about the distribution from which observations are sampled.

No-Information Problems

The standard no-information stopping problem is the "Secretary Problem." To contrast it with other no-information problems that we discuss later, we will refer to the most common formulation of the problem as the *Classical Secretary Problem* (CSP). It 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 from best to worst 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 [23]. The optimal cutoff can be obtained by:

$$t^* = \min\left\{ t \geq 1 : \sum_{k=t+1}^{n} \frac{1}{k-1} \leq 1 \right\}. \tag{2}$$

Interestingly, t^* converges to ne^{-1} as $n \to \infty$, and the policy selects the best applicant with probability e^{-1}.

An historical review of the CSP can be found in Ferguson [20] and in Samuels [52]. Numerous generalizations of the CSP have been proposed. For example, Corbin [15] and Yang [69] presented procedures for computing optimal policies for secretary problems in which options can be recalled with probabilistic success; Pressman and Sonin [40] discussed problems in which the number of applicants n is itself unknown, but the DM does know the distribution from which n is sampled.

Experimental Studies of the CSP

Seale and Rapoport [55] tested two versions of the CSP, one with $n = 40$ and another with $n = 80$. Using a between-subjects experimental design, each subject in their experiment was either in the $n = 40$ or $n = 80$ condition. Each subject played a total of 100 independent instances (trials) of the CSP in a computer-controlled environment. The subjects were first given a cover story that described their task as one of hiring applicants for a position. Each of the trials proceeded as follows. The *relative rank* of the first applicant was displayed (which, by definition, was always 1). The subjects could then choose to select (hire) the applicant or proceed (interview) to the next applicant. Whenever the subject chose to proceed from applicant j to applicant $j + 1$, the computer displayed the relative ranks of all applicants up to $j + 1$. Once an subject selected an applicant, the *absolute ranks* of all n applicants were displayed. If the subject selected the best overall applicant, the computer informed her that she had made a correct selection and added her payoff for that trial to her cumulative earnings. For the $n = 40$ condition, a subject earned \$.30 each time she selected an applicant with absolute rank 1; for the $n = 80$ condition, the corresponding payoff was \$.50.

The probability of stopping on applicant j or sooner under the optimal policy if $n = 80$ is displayed in Figure 2. Figure 2 also shows the proportion of times that subjects in Seale and Rapoport's $n = 80$ condition stopped on applicant j or sooner. Note that the empirical curve is shifted considerably to the left of the optimal curve, demonstrating a propensity of the subjects to stop earlier than is predicted by the optimal policy.

From the results displayed in Figure 2 we can see that the subjects are behaving sub-optimally; however, we cannot infer how it is that the subjects are making their decisions: we cannot infer their actual decision policies. Seale and Rapoport tried to get a handle on the DM's underlying (unobservable) decision policies by competitively testing three different single parameter decision policies or heuristics. The particular policies studied were chosen for their psychological plausibility: *a priori* they seemed like policies that people might reasonably use. We will consider each of these in turn and will describe some of the properties of the policies that were later derived by Stein, Seale, and Rapoport [60].

The Cutoff Rule (CR): Do not accept any of the first $t - 1$ applicants; thereafter, select the first encountered *candidate* (i.e., an applicant with relative rank 1). This rule has as a special case the optimal policy for the CSP for which t^* is obtained by Equation 1. In addition to the optimal policy, the CR can be set to begin accepting candidates earlier in the sequence $(t < t^*)$ or later $(t > t^*)$.

Candidate Count Rule (CCR): Select the hth encountered candidate. Note that this rule does not necessarily skip any applicants; it only considers how many candidates have been observed, not how deep the DM is in the applicant sequence.

Successive Non-Candidate Rule (SNCR): Select the first encountered candidate after observing g successive non-candidates (i.e., applicants with relative rank > 1).

Each of these three heuristic decision policies can be represented by a single parameter (t, h, and g). Before returning to the results of Seale and Rapoport [55], let us examine the theoretical performance of the heuristics. We know that no heuristic can outperform the CR with $t = t^*$. But how well can the others do?

FIGURE 2. Optimal and empirical cumulative stopping probabilities from the $n = 80$ condition in Seale and Rapoport [55].

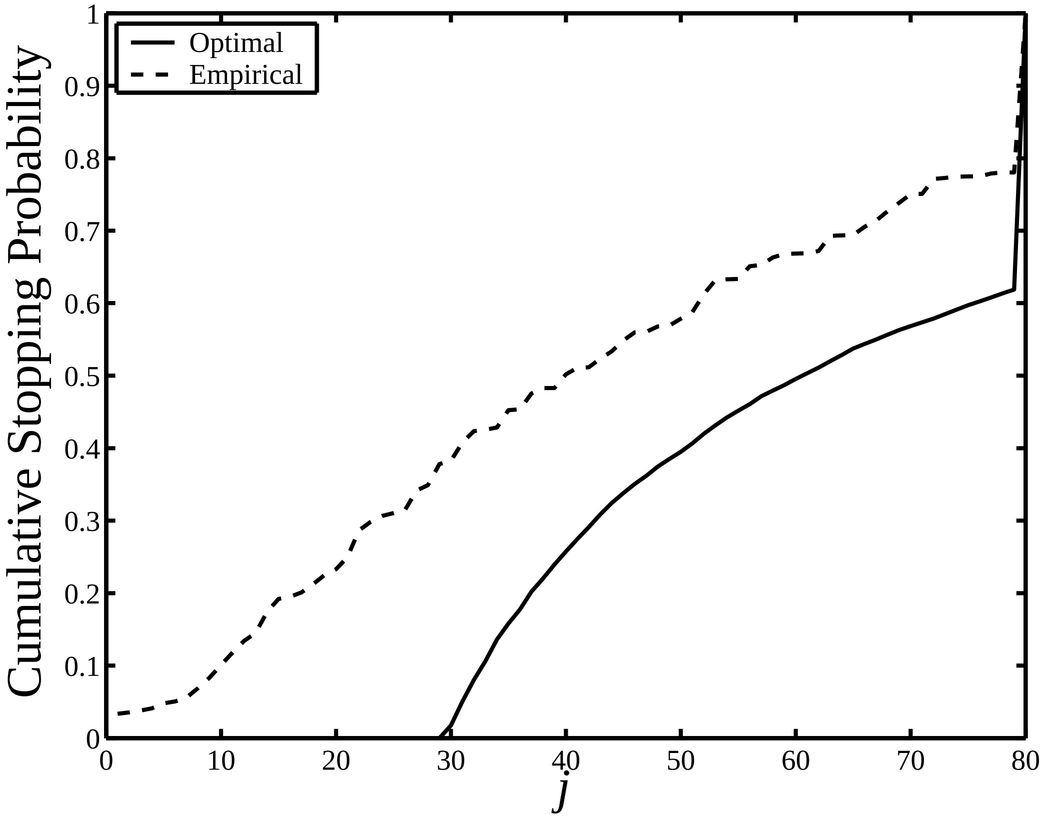

Results for each the heuristics for $n = 80$ CSP are displayed in Figure 3. The horizontal axis corresponds to the values of each heuristic's single parameter (t, h, and g for the CR, CCR, and SNCR, respectively). There are a number of interesting features of these results. First, the CR and SNCR heuristics strongly outperform the CCR heuristic. In addition, when properly tuned, the SNCR can perform nearly as well as the optimal policy (i.e., the CR with $t = t^*$), with the former earning about 95% of what is earned by the latter. Finally, the CR is relatively robust to mis-parameterization. Whenever $.63t^* \leq t \leq 1.47t^*$ the DM can expect to earn at least 90% of what is earned under the optimal policy. From sensitivity analyses, the flatness of payoff functions for a number of problems studied by behavioral decision theorists has often been noted (see, e.g., [46], [67]).

Seale and Rapoport [55] fit each of the three heuristics just described to each subject's choice data. For each subject and each heuristic, they used a *brute force* search procedure to find the heuristic's parameter that minimized the number of incorrect stops (or violations) predicted by the heuristic. Put differently, if a subject stopped on applicant j on a given trial and the heuristic predicted that the subject should stop on applicant $j' \neq j$ then a violation was obtained. Hence a particular parameterization of a heuristic could produce between 0 and 100 (the total number of experimental trials) violations for a given subject.

The results for both the $n = 40$ and the $n = 80$ conditions were consistent; for brevity we will restrict discussion to those from the latter condition. The most important result is that the subjects' decision data were best characterized by the CR heuristic. For 21 of 25 subjects, the CR best fit the data; the SNCR best fit the choices for 8 of the 25 subjects; and the CCR fit best for only 1 subject. (The number of best fitting heuristics sums to more than 25 because of ties. For some subjects the best-fitting CR and SNCR heuristics

Figure 3. Probability of success for three CSP heuristics.

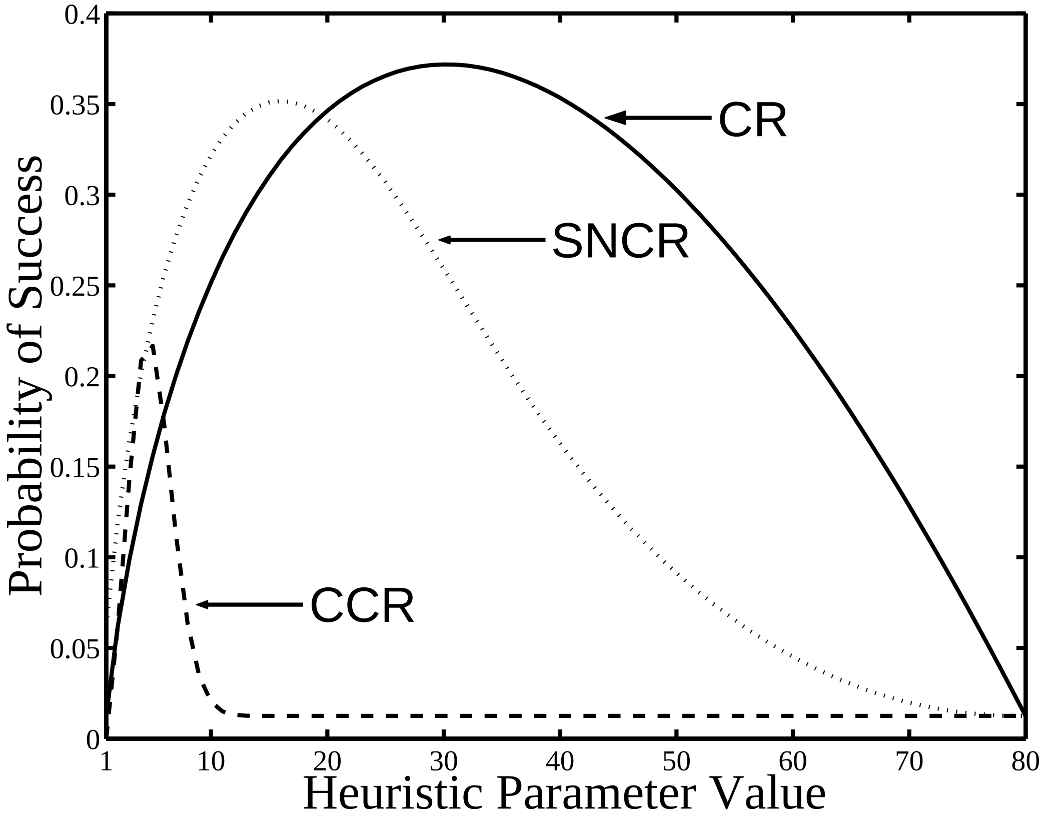

Note. These results are based on the derivations presented in Stein, Seale, and Rapoport [60].

produced the same number of violations.) The authors compared the subjects' CR cutoffs (t values) to the optimal cutoff, and found that $t < t^*$ for 21 subjects, $t > t^*$ for 3 subjects, and $t = t^*$ for only 1 subject.

This illustrates one logic that can be used to draw inferences from experimental data about the decision policies employed by actual decision makers. One might be tempted to argue that the three heuristics tested by Seale and Rapoport were arbitrarily chosen, and that any number of other heuristics *could* have been tested. The latter is certainly true. However, it is up to the critic to propose specific alternatives in order for the former argument to have force. Further, all explanations are going to be underdetermined (see, e.g., van Fraassen [65]), even those that come from our most successful scientific theories. We believe, however, that there is another criticism of the procedures used by Seale and Rapoport that should be carefully considered. Although the CR heuristic best accounted for most subjects' data, it did not do so perfectly. For most subjects, the best-fitting heuristic could only explain about 60-70% of the stopping decisions. What drove the other 30-40% of the stopping decisions? This question led Bearden and Murphy [3] to formulate a stochastic version of the CR and fit it to Seale and Rapoport's data; however, they remained agnostic on the source of the threshold variability. A fuller account of the decision behavior in the CSP should address this issue.

Several extensions of the CSP have been studied experimentally. Seale and Rapoport [56], for example, looked at a problem in which the DMs do not know n but only its distribution. Compared with the optimal policy, the subjects tended not to search enough. Zwick et al. [71] examined the effect of search cost and probabilistic recall of previously seen applicants.

The most important finding from this study is that the subjects tended to search for too long with positive search costs, whereas they did not search enough when search costs were set at 0.

The nothing-but-the-best payoff structure of the CSP seems patently unrealistic. One can imagine few situations in which selecting the best option from a pool yields positive utility, and selecting any other options yields no utility. What, then, might we have learned about actual human decision making in environments in which options are encountered sequentially from the experimental studies of the CSP? Might it be the case that the observed early stopping is simply a consequence of the CSP's special payoff function? That is, perhaps we now know that actual DMs cannot perfectly solve the CSP probability puzzle, but we have learned nothing about how people might actually go about selecting secretaries. To address this issue, we studied an extension of the CSP in which the DM's objective is to find a *good*—not necessarily the best—alternative. Next, we describe the formal problem, how to compute its optimal policy, and some experimental studies of decision making in this more realistic context.

The Generalized Secretary Problem

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 [38] 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. The values of the t^* depend on the operative payoffs and on the number of applicants n.

We obtain what we call the *Generalized Secretary Problem* (GSP) by replacing **5** in the CSP with the more general payoff 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 [37] 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.

Finding Optimal Policies for the GSP

We begin by introducing some notation. The orderings of the n applicants' *absolute ranks* is represented by a vector $\mathbf{a} = (a_1, \ldots, a_n)$, which is some 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 $\mathbf{s} = (s_1, \ldots, s_n)$ 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 $Q(s_j) = s_j/j$. 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.

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. Lindley [36] described a similar method, and another was briefly sketched in [70]. A fuller treatment is presented in Bearden and Murphy [3].

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:

$$Pr\left(A=a|R=r_j\right)=\frac{\binom{a-1}{r-1}\binom{n-a}{j-r}}{\binom{n}{j}}, \tag{3}$$

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:

$$E\left(\pi_j|r_j\right)=\sum_{a=r_j}^{n} Pr\left(A=a|R=r_j\right)\pi(a). \tag{4}$$

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

$$E\left(\pi_j|s_j\right)=\left(s_j\right)^{-1}\sum_{i=1}^{s_j} E\left(\pi_j|r_j=i\right); \tag{5}$$

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:

$$V_j=Q(s_j)E\left(\pi_j|s_j\right)+\left[1-Q(s_j)\right]V_{j+1}. \tag{6}$$

As with the optimal policy for the full-information optimal stopping problem described above, the optimal policy for the GSP entails stopping on an applicant with relative rank r_j at stage j whenever the expected payoff for doing so exceeds the expected payoff for proceeding and playing optimally thereafter. This follows directly from the Principle of Optimality [8]. Given that 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 $(j=n-1,\ldots,1)$ by backward induction. Since the last applicant must be selected, $s_n^*=n$; then, for $j=n-1,\ldots,1$,

$$s_j^*=\min\left\{s\in\left\{0,\ldots,s_{j+1}^*\right\}:V_j\geq V_{j+1}^*\right\}. \tag{7}$$

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

$$E\left(\pi|\mathbf{s}\right)=\sum_{j=1}^{n}\left[\prod_{i=0}^{j-1}\left[1-Q(s_i)\right]\right]Q(s_j)E\left(\pi_j|s_j\right)=V_1, \tag{8}$$

where $Q(s_0)=0$. Denoting the applicant position at which the search is terminated by m, the expected stopping position under the policy is:

$$E\left(m\right)=1+\sum_{j=1}^{n-1}\left[\prod_{i=1}^{j}\left[1-Q(s_i)\right]\right]. \tag{9}$$

Equation 9 can be useful in evaluating the results of actual DMs in GSPs (see, e.g., [6]).

Optimal cutoffs for several GSPs are presented in Table 2. 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. GSPs 3 and 4 were discussed in Gilbert and Mosteller [23], 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

Table 2. Several GSPs and their optimal policies.

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	60	$(25, 13, 6, 3, 2, 1, 0, \ldots, 0)$	$(21, 43, 53, 57, 58, 59, 60, \ldots, 60)$	12.73	41.04
6	40	$(15, 7, 2, 0, \ldots, 0)$	$(14, 29, 37, 40, \ldots, 40)$	6.11	27.21

studied by Bearden, Rapoport, and Murphy [6] in Experiments 1 and 2, respectively. In the first, the DM searches through the first 20 applicants without accepting any; then between 21 and 42 she stops on applicants with relative rank of 1; between 43 and 52, she stops on applicants with relative rank 1 or 2; etc.

Experimental Studies of the GSP

Recall that Seale and Rapoport [55] found that subjects playing the CSP (GSPs 1 and 2, actually) tended to search insufficiently into the applicants before making a selection—i.e., on average they stopped too soon. Above, we suggested that the CSP might be quite contrived and that this result could possibly be an artifact of the narrow payoff scheme of the CSP. Bearden, Rapoport, and Murphy [6] experimentally tested this using two variants of the GSP, GSPs 5 and 6. In Experiment 1, each subject played 60 instances of GSP5. (They were paid in cash for two randomly selected trials.) It is difficult to say what a realistic payoff function is, but it seems certain that the one for GSP5 is more realistic than the one for the CSPs. For the GSP5, the DM gets a substantial payoff for selecting the best applicant ($25 in the experiment), considerably less—but still significant—for the second best ($13), and so on. This captures the payoff properties of problems in which the DM would like to get a "really good" applicant, which seems like a desire that that DMs might often have. Using this more plausible payoff scheme we find that the early stopping result from studies of the CSP persists. Figure 4 shows the cumulative stopping results for the optimal policy and also those from the experimental subjects for the GSP5. We found that the subjects tend to terminate their searches too soon relative to the optimal policy. Comparing Figures 2 and 4, we can see that the tendency to stop searching too soon is less pronounced in the GSP than in the CSP, but drawing strong inferences from data gathered in two different experimental settings is problematic.

Bearden, Rapoport, and Murphy competitively tested multi-parameter generalizations of the three heuristics examined by Seale and Rapoport [55]. The generalized Cutoff Rule, for example, had thresholds for each of the relative ranks that could entail positive payoffs ($r \leq 6$). Under this rule, the DM stops on applicant j with relative rank r_j whenever $j \geq t_{r_j}$. The Candidate Count Rule and the Successive Non-Candidates Rule were extended in the same fashion. Consistent with the findings of Seale and Rapoport [55], the results were best captured by the Cutoff Rule with cutoffs shifted toward early stopping.

In a second experiment, Bearden, Rapoport, and Murphy experimentally studied the GSP6. In all regards the choice results captured the qualitative properties of those found in the data from the first experiment, including the superiority of the Cutoff Rule in accounting for the data. After subjects performed 60 trials of the GSP6, they were then asked to perform a probability estimation task. They were shown the relative ranks r_j of applicants in different positions j and asked to estimate the probability that the applicant had various absolute ranks a. For example, a subject might have been asked: "What is the probability that applicant 10 (of 40), whose relative rank is 2, has an absolute rank of 2?" Or: "What is the probability that applicant 10 (of 40), whose relative rank is 2, has an absolute rank of 3?" (It is .06 for $a = 2$ and .09 for $a = 3$.) We only asked the subjects about absolute ranks

FIGURE 4. Optimal and empirical cumulative stopping probabilities from Experiment 1 in Bearden, Rapoport, and Murphy [6].

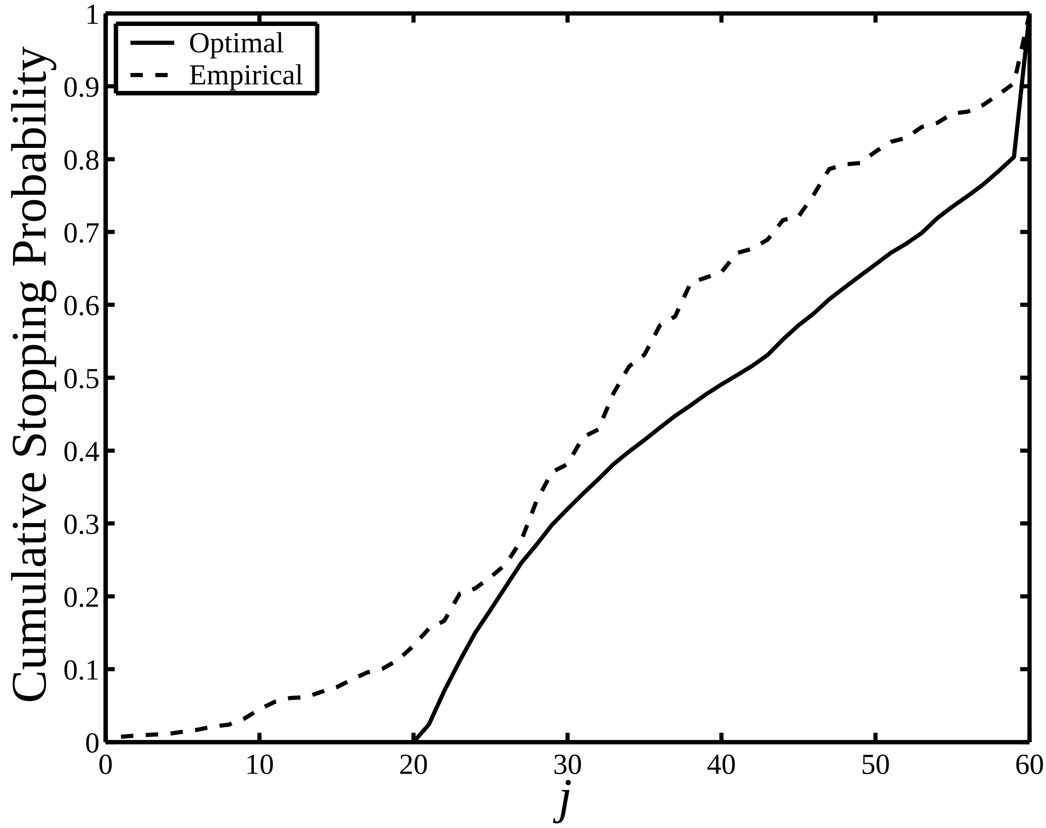

that could entail positive payoffs, i.e., about $a \leq 3$. We used a strictly proper scoring rule to encourage the subjects to give accurate estimates. Put simply, the subjects made more money as their estimates were closer to the true probabilities, which can be obtained from Equation 3, and there was no incentive for them to misrepresent their true estimates (i.e., to respond strategically).

The results were informative. Overwhelmingly, the subjects tended to overestimate the true probabilities, particularly for applicants early in the sequence. In fact, the estimates were quite often subadditive: For a given applicant position j and relative rank r_j, the sum of the probability estimates exceeded 1. Taking these results seriously, the subjects were often *supercertain* that they would obtain positive payoffs for selecting applicants. These findings are consistent with Tversky and Koehler's [64] *support theory*, an empirically successful theory of subjective probability. Under support theory, the subjective probability assigned to an event E, $p(E)$, is a function of the support (evidence) one gives to E and its complement $\neg E$; specifically,

$$p(E) = \frac{s(E)}{s(E) + s(\neg E)}, \tag{10}$$

where $s(\cdot) > 0$ is a real-valued measure of the strength of the evidence one can find for E. The evidence can be arrived at through a process of memory search or by other means [7]. Under the theory, the *focal* event E receives more support than the *non-focal* event $\neg E$. This aspect of the theory is consistent with considerable evidence from psychology that shows that human information processing is biased toward confirmation seeking (i.e., toward searching for evidence for E, rather than for evidence of $\neg E$, which is also normatively relevant). In short, according to the theory, when evaluating the probability that an applicant has a particular absolute rank a, one is likely to give insufficient weight to the possibility that the applicant *does not* have an absolute rank a, focusing (disproportionately) instead on the event that the applicant's absolute rank *is* a.

Perhaps the probability estimation results—at least in part—explain the early stopping finding. If a DM believes it certain that stopping early will produce a good payoff, then stopping early is sensible. Put differently, one possible reason for subjects not behaving in accord with the optimal policy is that their representation of the search problem is distorted. Constructing the optimal policy requires the computation of Equation 3; but the estimation data show that subjects' intuitive probability judgments depart quite radically from the true probabilities. This demonstrates the strength of the assumptions required in order for one to presume that the optimal policy is *a priori* a reasonable predictor of human decision making in this class of problems.

What, then, of questions regarding the rationality of decision making in sequential search problems? Can we say that people are irrational if they cannot intuit the probabilities that fall out of Equation 3? Of course not. But, though it is implausible to presume that people should behave in accord with the optimal policy, the optimal policy still provides some basis for evaluating the experimental decision data. While it may be important in economic contexts to anticipate that people will not search optimally, we are still left wondering what underlying psychology drives sequential search decisions. Fortunately, the optimal policy can and has served as a starting point for understanding *how it is* that people make these decisions. As described above, the cutoff rule—of which the optimal policy is a special case—best accounts for the stopping results. Likewise, comparing the subjects' probability estimates to the true probabilities was informative. Hence, knowledge of the optimal policy can be useful for explanatory purposes.

The GSP captures important features of a number of dynamic decision problems likely to be encountered in the wild (i.e., the so-called "real world"); however, there are, of course, many problems that it does not capture. Quite often we are faced with decision problems in which we must make trade-offs among the attributes of decision alternatives. An academic job may offer a good salary but an undesirable teaching load; a house may be close to one's office (minimizing commute time) but in a poor school district; etc. One can argue that the GSP side-steps these kinds of problems by collapsing the multi-attribute utility of options into a single ranking. Since we are primarily interested in how people actually make decisions, we do not want to assume away this interesting problem of making trade-offs among attributes. We would like to know how people do so, particulary in situations in which options are encountered sequentially. Next, we describe a multi-attribute extension of the GSP, describe its solution, and present some experimental findings.

The Multi-attribute Secretary Problem

The *Multi-attribute Secretary Problem* (MASP) further generalizes the GSP to applicants with multiple features or attributes. Formally, it is defined as follows:

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

5. For each attribute i of the selected applicant, the DM earns a payoff of $\pi^i(a^i)$, where a^i is the selected applicant's absolute rank on attribute i and $\pi^i(1) \geq \ldots \geq \pi^i(n)$.

Before describing the optimal policy for the MASP, we must introduce some notation. The *absolute rank* of the jth applicant on the ith attribute, denoted a_j^i, is simply the number of applicants in the applicant pool, including j, whose ith attribute is at least as good as the jth applicant's. The jth applicant's set of absolute ranks can therefore be represented by a vector $\mathbf{a}_j = (a_j^1, \ldots, a_j^k)$. The *relative rank* of the jth applicant on the ith attribute, r_j^i, is the number of applicants from 1 to j whose ith attribute is at least as good as the jth's. Similar to the GSP, when interviewing an applicant, the DM observes $\mathbf{r}_j = (r_j^1, \ldots, r_j^k)$, and must make her selection decision on the basis of this information.

Though she only observes relative ranks $\mathbf{r}_j$, the DM's payoff for selecting the jth applicant, denoted π_j, is based on the applicant's absolute ranks $\mathbf{a}_j$; specifically,

$$\pi_j = \sum_{i=1}^{k} \pi^i \left(a_j^i\right). \tag{11}$$

An optimal policy for the MASP is one that maximizes the expected value of the selected applicant.

Some related problems have appeared in the OR literature. Gnedin [24] presented the solution to a multi-attribute CSP in which the attributes are independent, and the DM's objective is to select an applicant who is best on at least one attribute. Ferguson [21] generalized the problem presented by Gnedin by allowing dependencies between the attributes, and showed that the optimal policy has the same threshold form as the standard single attribute CSP. Samuels and Chotlos [53] extended the rank minimization problem of Chow et al. [13]. They sought an optimal policy for minimizing the sum of two ranks for independent attributes. The rank sum minimization problem they studied is equivalent to the MASP in which $\pi^1(a) = \pi^2(a) = n - a$. The MASP is more general than these previous problems, as it only constrains the payoff functions to be nondecreasing in the quality of the selected applicant's attributes.

Finding Optimal Policies for the MASP

The probability that the ith attribute of the jth applicant whose relative rank on that attribute is r has an absolute (overall) rank of a is given by Equation 3. To simplify matters, we assume that the k attributes are pairwise independent; that is, $Pr(a^i = a \wedge a^{i'} = a') = Pr(a^i = a)Pr(a^{i'} = a')$ for any pair of attributes i and i'. (Based on work by Ferguson [21], introducing arbitrary correlations ρ among attributes would likely make the determination of the appropriately corresponding Equation 3 intractable, as for any j it would depend—in complicated ways—on $(\mathbf{r}_1, \ldots, \mathbf{r}_j)$, i.e., on the entire history of relative ranks.) Consequently, the expected payoff for selecting the jth applicant is:

$$E\left(\pi_j|\mathbf{r}_j\right) = \sum_{i=1}^{k} \sum_{a=r_j^i}^{n} Pr\left(A = a|R = r_j^i\right) \pi^i(a). \tag{12}$$

At each stage j of the decision problem, the DM must decide to accept or reject an applicant knowing only the applicant's relative ranks $\mathbf{r}_j$. We represent a decision policy for each stage j as a set of acceptable $\mathbf{r}_j$ for that stage $\mathbf{R}_j$. Under the stage policy $\mathbf{R}_j$, the DM stops on an applicant with relative ranks $\mathbf{r}_j$ if and only if $\mathbf{r}_j \in \mathbf{R}_j$. The global policy is just the collection of stage policies $\mathbf{R} = \{\mathbf{R}_1, \ldots, \mathbf{R}_n\}$. By Bellman's [8] Principle of Optimality, for an optimal (global) policy $\mathbf{R}^*$, each sub-policy $\{\mathbf{R}_j, \ldots, \mathbf{R}_n\}$ from stage j to n must also be optimal. Given this property, we can find the optimal policy using straightforward dynamic programming methods by working backward from stage n to stage 1. A procedure for constructing optimal stage policies $\mathbf{R}_j^*$ follows from Proposition 1, which we present below. To simply exposition, we first make the following assumption:

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

Since the DM must accept the nth applicant, if reached, $V_n^* = n^{-1}\sum_{i=1}^{k}\sum_{a=1}^{n}\pi^i(a)$. And, for stages $j < n$, we have

$$V_j^* = Q\left(\mathbf{R}_j^*\right) E\left(\pi_j|\mathbf{R}_j^*\right) + \left[1 - Q\left(\mathbf{R}_j^*\right)\right] V_{j+1}^*, \tag{13}$$

where $E\left(\pi_j|\mathbf{R}_j^*\right) = |\mathbf{R}_j^*|^{-1}\sum_{\mathbf{r}\in\mathbf{R}_j^*} E(\pi_j|\mathbf{r})$ is the expected payoff for stopping at stage j under the optimal stage j policy, and $Q\left(\mathbf{R}_j^*\right) = |\mathbf{R}_j^*|/k^j$ is the probability of stopping under the optimal stage j policy. Given V_n^*, a method for constructing $\mathbf{R}^*$ is entailed by the following proposition:

Proposition 1. $\mathbf{r} \in \mathbf{R}_j^* \Leftrightarrow E\left(\pi_j|\mathbf{r}\right) \geq V_{j+1}^*$.

Invoking Assumption 1, the proof of Proposition 1, presented in Bearden and Murphy [3], follows directly from the Principle of Optimality [8].

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

Proposition 2 follows from Corollary 2.1b in Mucci [38]. Stated simply, Proposition 2 tells us that if it is optimal to stop at stage j when one observes $\mathbf{r}$, then it is optimal to stop when one observes $\mathbf{r}$ in the next stage; by induction, then, it is optimal to stop given $\mathbf{r}$ in *all* subsequent stages. This property allows us to represent the optimal policies rather compactly by specifying for each feasible $\mathbf{r}$ the smallest j for which $\mathbf{r} \in \mathbf{R}^*$.

An Example of a MASP and the Application of Its Optimal Policy

Let us consider how the optimal policy would be applied in a MASP. Table 3 contains an example of an instance of a MASP with $n = 6$ and $k = 2$. The top panel contains the payoffs for each of the attributes. Absolute and relative ranks for each of the 6 applicants are shown in the center panel. We see that applicant 1 has absolute ranks of 2 and 5 on attributes 1 and 2, respectively; her relative ranks are, of course, 1 for both attributes. Applicant 2 has absolute ranks of 4 and 2, and therefore relative ranks of 2 and 1, for attributes 1 and 2, respectively, etc. The bottom panel displays the value of the optimal policy for each applicant position and the expected payoffs for selecting each applicant. Since applicant 3 is the first applicant for which $E\left(\pi_j|\mathbf{r}_j\right) \geq V_{j+1}^*$, she is selected.

Experimental Studies of the MASP

Bearden, Murphy, and Rapoport [4] tested actual DMs on the MASP in two experiments. Using $n = 30$ (up to 30 applicants) and $k = 2$ (2 attributes), the experiments were identical in all regards except for their payoff schemes. Experiment 1 tested a MASP with symmetric payoffs, where $\pi^1(a) = \pi^2(a)$ for all a. Experiment 2 implemented an asymmetric payoff scheme in which attribute 1 was considerably more important (i.e., contributed more to the payoff) than attribute 2. The actual payoff schemes for both experiments are displayed in Table 4. In each condition, each subject played 100 random instances of the MASP. Payoffs were based on a single randomly selected trial; hence, those in Experiment 1 could earn up to \$50 for the one hour session, whereas those in Experiment 2 could earn up to \$40. As with our previous experiments on the GSP, we used a hiring cover story, and instructed the subjects that the attributes were uncorrelated (in language they could easily understand).

Threshold representations of the optimal policies for the two MASPs examined in Experiments 1 and 2 are shown in Table 5. The cell entries for a pair of relative ranks (r_1, r_2)

Table 3. MASP Example Problem.

	Payoff Values					
a	1	2	3	4	5	6
$\pi^1(a)$	6	5	4	3	2	1
$\pi^2(a)$	5	4	3	2	0	0
	Example Applicant Sequence					
Applicant (j)	1	2	**3**	4	5	6
a_j^1	2	4	**3**	6	5	1
a_j^2	5	2	**1**	3	6	4
r_j^1	1	2	**2**	4	4	1
r_j^2	1	1	**1**	3	5	4
	Optimal Policy and Payoffs					
Applicant (j)	1	2	**3**	4	5	6
V_{j+1}^*	7.82	7.67	**7.37**	6.83	5.83	—
$E(\Pi_j \vert \mathbf{r}_j)$	5.83	5.73	**7.55**	2.93	1.83	8.00
Π_j	5.00	7.00	**9.00**	4.00	2.00	8.00

correspond to the applicant position at which the optimal DM should begin to accept applicants with those relative ranks. For example, in the symmetric case (Experiment 1), the DM should begin accepting applicants with relative ranks of 1 on both attributes $(1,1)$ at applicant position 7; applicants with a relative rank of 1 on one attribute and 2 on the other $((1,2)$ or $(2,1))$ should be accepted starting at position 12; etc. For the asymmetric payoffs (Experiment 2), the DM should also begin accepting applicants with relative ranks of 1 on both attributes at position 7. Applicants with relative rank 1 on attribute 1 and relative rank 2 on attribute 2 should be accepted starting at position 11. In contrast, when the second attribute (the less important attribute) has relative rank 1 and the first (the more important attribute) has relative rank 2, the DM should wait until applicant position 14 to start accepting applicants with this profile.

In both conditions we observed that the subjects examined fewer options on average than predicted by the optimal policy. For the MASP studied in Experiment 1, the expected length of search is 20.09 (by Equation 9); the empirical average length of search was 15.89. For Experiment 2, the expected search length under the optimal is 19.45; the empirical average was 15.90. These findings are consistent with the previous results on the CSP and GSP that we reported above. Fortunately, the data from the MASP lend themselves to additional analyses that are quite informative.

Recall that Seale and Rapoport [55] tested several heuristics using the number of incorrect predictions ("violations" in their language) as their objective to minimize. To fit the Cutoff Rule, for example, they found the value of t that minimized the number of incorrect stopping decisions for a given subject. To get a deeper understanding of the underlying decision policies subjects employ in the MASP, we used similar methods to estimate subjects' policies. We restricted ourselves to policies of the same form as the optimal policy, viz. to sets of thresholds for each stage j that lead to decisions to stop the search. Given that each of the heuristics tested by Seale and Rapoport could take on at a maximum 80 different values, they could use brute force search without any computational difficulty. Fitting MASP policies is more challenging due to the large feasible set of policies. To get around this problem, we used a heuristic procedure to find (probably) best fitting policies for each subject. Specifically, we used Dueck and Scheuer's [18] *Threshold Accepting* algorithm, which is an easy-to-implement

Table 4. MASP payoffs for Experiments 1 and 2 in Bearden, Murphy, and Rapoport [4].

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

cousin of simulated annealing. The details of our treatment of the optimization problem can be found in [4].

Using the median cutoff (taken over subjects) of the estimated policies for each pair of relative ranks, we examined the difference between the optimal and empirical cutoffs to look for any systematic departures of the empirical policies from the optimal ones. Thus, when the difference is negative for a pair of relative ranks (r_1, r_2), the empirical policy tends to stop earlier than the optimal one. The difference values for policies derived from both experiments are shown in Table 6. In most cases, for both experiments, we observe that the differences are negative, revealing a tendency (or bias) to accept applicants earlier than is optimal. More interesting, the bias is strongest for small relative rank pairs (e.g., $(1,2)$, $(2,2)$, $(2,3)$, etc.). There is also a bias to stop later on pairs for which one attribute guarantees 0 payoff (i.e., $r \geq 6$), and the other does not (i.e., $r < 6$). For example, in the symmetric case, the subjects tended to pass up applicants with relative ranks 1 and 6 even when stopping had a greater expectation under the optimal policy.

How should we interpret the policy results in Table 6? One possibility is that the subjects were using a policy consistent with Herbert Simon's notion of *satisficing* [58]. According to Simon, given the bounds on the capacities of actual agents—in contrast to ideal agents—we should not expect optimizing behavior. Instead, he suggested that agents might search for options that are "good enough," rather than those that are optimal. These good enough options are the ones that satisfice, i.e., that meet the agent's aspirations on *all* (relevant) attributes. Since our subjects tended to stop quite early on applicants with small pairs of relative ranks and to avoid stopping on those with one good relative rank (e.g., 1) and one poor one (e.g., 16), we might suppose that their policy is of the satisficing sort: They want options that are sufficiently good on both attributes, and do not want an option that is too poor on any single attribute, as these do not satisfice.

One possible line of future inquiry is to look at optimal satisficing strategies in the MASP. Suppose the DM wants to guarantee (or maximize the probability) that she selects an applicant that is acceptable on all attributes. Perhaps problems with this objective more closely match those that actual DMs face. Thus, optimal satisficing MASP policies may have both theoretical and practical importance.

The problems we have focused on up to now have involved selecting an option from a set of options presented sequentially. Next, we look at a class of problems in which the DM must assign each of the sequentially encountered options to open positions.

Sequential Assignment

Derman, Leiberman, and Ross [16] considered the following problem: A DM observes a sequence of n jobs for which there are m machines available. Each job j has a *value*, which is a random variable X_j that takes on the value x_j. The cost of assigning job j to machine

Table 5. Optimal policies for Experiments 1 and 2 from Bearden, Murphy, and Rapoport [4]. For a given pair of relative ranks (r_1, r_2), the table entry is the smallest j for which applicants with relative ranks (r_1, r_2) are selected. This representation is analogous to the the cutoff representation $\mathbf{t}$ of optimal policies for the GSP.

	Experiment 1 (Symmetric)					
	$r_2=1$	$r_2=2$	$r_2=3$	$r_2=4$	$r_2=5$	$r_2=6$
$r_1=1$	7	12	14	15	16	16
$r_1=2$	12	19	22	24	25	26
$r_1=3$	14	22	25	27	27	28
$r_1=4$	15	24	27	28	29	29
$r_1=5$	16	25	27	29	30	30
$r_1=6$	16	26	28	29	30	30

	Experiment 2 (Asymmetric)					
	$r_2=1$	$r_2=2$	$r_2=3$	$r_2=4$	$r_2=5$	$r_2=6$
$r_1=1$	7	11	12	13	13	13
$r_1=2$	14	19	22	23	24	24
$r_1=3$	17	23	25	26	27	27
$r_1=4$	19	25	27	28	29	29
$r_1=5$	20	26	28	29	30	30
$r_1=5$	21	27	29	30	30	30

Table 6. Difference in location of the median empirical and the optimal cutoff locations for Experiments 1 and 2 in Bearden, Murphy, and Rapoport [4]. For a given pair of relative ranks (r_1, r_2), the table entry is the smallest j for which applicants with relative ranks (r_1, r_2) are selected.

	Experiment 1 (Symmetric)					
	$r_2=1$	$r_2=2$	$r_2=3$	$r_2=4$	$r_2=5$	$r_2=6$
$r_1=1$	0	-4	-5	-3	-3	2
$r_1=2$	-4	-9	-10	-9	-7	-1
$r_1=3$	-5	-10	-9	-8	-5	-2
$r_1=4$	-3	-9	-8	-6	-4	0
$r_1=5$	-3	-7	-5	-4	-1	0
$r_1=6$	2	-1	-2	0	0	0
	Experiment 2 (Asymmetric)					
	$r_2=1$	$r_2=2$	$r_2=3$	$r_2=4$	$r_2=5$	$r_2=6$
$r_1=1$	0	-4	0	-1	8	8
$r_1=2$	-4	-6	-5	-3	1	1
$r_1=3$	-3	-3	-3	-1	-2	0
$r_1=4$	-2	-1	-2	-3	-1	-1
$r_1=5$	2	0	0	0	0	0
$r_1=6$	4	3	1	0	0	0

Note. Negative differences indicate that cutoffs are shifted toward stopping too early.

i is $c_i x_j$, and the DM's objective is minimize her expected costs. In the simplest case, the jobs are sampled i.i.d. according to the distribution function $f(x)$, which is known by the DM.

Chun and Sumichrast [14] extended the problem presented by Derman et al. to scenarios in which the payoffs are determined only on the basis of the ranks of the jobs. Under this formulation, we need not assume that *a priori* the DM has full distributional information on the jobs. Chun and Sumichrast's sequential selection and assignment problem (SSAP) can be described as follows:

1. There are n applicants for m positions. Each applicant can be ranked in terms of quality with no ties. (For brevity and with no loss of generality (see [14]), we will only consider cases in which $n = m$.) Associated with each position i is a cost c_i, where $c_1 \leq c_2 \leq \ldots \leq c_m$.
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. Each applicant must be assigned to an open position. Once assigned to a position i, the applicant cannot be re-assigned to another position.
5. The total cost for assigning the n applicants is $\sum_j a_j c_i$, where a_j is the absolute rank of the jth applicant and c_i is the cost of the position to which j is assigned. The DM's objective is to minimize her total assignment costs.

Computing Optimal Policies for the SSAP

Chun and Sumichrast [14] presented a procedure for determining optimal assignment policies for the SSAP. Interestingly, the optimal policy does not depend on the values of the position costs c_i; only the rank ordering of the costs matters. The optimal policy is expressed as sets of *critical relative ranks* $r^*_{i,k}$ for each stage $k = n - j + 1$ (where k is simply the number of remaining to-be-interviewed applicants, including the current one). The critical ranks work as follows: Assign an applicant with relative rank r_k at stage k to the ith position if $r^*_{i-1,k} < r_k < r^*_{i,k}$. The critical ranks are computed by recursively solving:

$$r^*_{i,k} = \begin{cases} 0 & \text{for } i = 0 \\ \frac{1}{n-k+3} \sum_{r_{k-1}}^{n-k+2} \min\left\{\max\left\{r_{k-1}, r^*_{i-1,k-1}\right\}, r^*_{i,k-1}\right\} & \text{for } 1 \leq i < k \\ \infty & \text{for } i = k \end{cases} \tag{14}$$

Critical ranks for a problem with $n = 9$ are shown in Table 7. The first applicant ($j = 1$ or $k = 9$) should always be assigned to position 5, since she will always have a relative rank 1, which is between $r^*_{4,9}$ and $r^*_{5,9}$. If the second applicant's ($j = 2$ or $k = 8$) relative rank is 1, she should be assigned to position 3; otherwise, she should be assigned to position 6 (since her relative rank will therefore be 2).

An Example of an SSAP and the Application of Its Optimal Policy

To fully illustrate this implementation, let us consider a problem with $n = 9$ in which the applicants are patients who must be assigned to hospitals beds. The absolute and relative ranks of the 9 patients are displayed in Table 8. We assume that the absolute ranks correspond to the severity of the patient's malady, with lower ranks representing more serious cases. Each patient must be assigned to a bed in one of three hospitals that differ in terms of their costs. Table 9 shows the configuration of 9 beds that are distributed across three hospitals. Hospital A is a high-cost hospital that should be used for severely injured patients; Hospital B is for intermediate cases; and Hospital C is for the least severe cases. Further,

TABLE 7. Critical relative ranks for the SSAP with $n = 9$.

State	$k=9$	$k=8$	$k=7$	$k=6$	$k=5$	$k=4$	$k=3$	$k=2$	$k=1$
$i=0$	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
$i=1$	0.50	0.75	1.00	1.33	1.75	2.31	3.11	4.50	∞
$i=2$	0.64	0.96	1.42	1.92	2.59	3.50	4.89	∞	
$i=3$	0.75	1.25	1.78	2.50	3.41	4.69	∞		
$i=4$	0.92	1.50	2.22	3.08	4.25	∞			
$i=5$	1.08	1.75	2.58	3.67	∞				
$i=6$	1.25	2.04	3.00	∞					
$i=7$	1.36	2.25	∞						
$i=8$	1.50	∞							
$i=9$	∞								

TABLE 8. Absolute and relative ranks for 9 patients.

Patient Number	1	2	3	4	5	6	7	8	9
a_j	6	2	5	3	8	1	9	7	4
r_j	1	1	2	2	5	1	7	6	4

TABLE 9. Hospital bed positions for example. Bed numbers are shown in parentheses.

Hospital A (High Cost)	Hospital B (Med. Cost)	Hospital C (Low Cost)
Bed A_1 (1)	Bed B_1 (3)	Bed C_1 (6)
Bed A_2 (2)	Bed B_2 (4)	Bed C_2 (7)
–	Bed B_3 (5)	Bed C_3 (8)
–	–	Bed C_4 (9)

TABLE 10. Optimal assignments for the example based on the ranks in Table 8 for the positions in Table 9.

Hospital A (High Cost)	Hospital B (Med. Cost)	Hospital C (Low Cost)
1	2	5
4	3	7
–	6	9
–	–	8

within each hospital the beds can be ranked in terms of costs. The most costly bed is Bed A_1 in Hospital A, and the least costly bed is Bed C_4 in Hospital C. For purposes of optimal assignment all that matters is the cost of a particular bed c_i, where $c_1 \geq c_2 \geq \ldots \geq c_9$. Applying the assignment dictated by the critical ranks shown in Table 7, we determine the assignments shown in Table 10.

Experimental Studies of the SSAP

Bearden, Rapoport, and Murphy [5] conducted three experiments on the SSAP. In each experiment, the subjects were asked to make triage decisions in a computer-controlled task. The SSAP was described to the subjects in simple language. They were asked to imagine that they had to assign patients to hospital beds after a mass casualty event, and were told that they would be paid based on the quality of their assignments. The costs within a hospital were constant, but hospitals differed from one another in their costs.

Table 11. Proportion of times patients were assigned to each hospital with probability greater than predicted by the optimal assignment policy. Hospitals are ordered with respect to their costs. Hospital A is the highest cost hospital and D is the lowest cost one. Data taken from Bearden, Rapoport, and Murphy [5].

	Hospital A	Hospital B	Hospital C	Hospital D
Experiment 1	0.50	0.92	0.92	0.25
Experiment 2	0.25	0.92	0.92	0.20
Experiment 3	0.31	0.93	0.78	0.20

The three experiments differed in the number of to-be-assigned patients (12 in Experiment 1, and 24 in Experiments 2 and 3), and the number of bed positions in each of four hospitals. (For our $n = 9$ example, we used the configuration in Table 9; however, we could have used other configurations, such as 3 beds per hospital. The particular configurations used do not affect the optimal assignment, but we suspected they might make a psychological difference.) In each experiment, each subject played a total of 60 instances of the SSAP.

The SSAP is a relatively complex task and produces data that can be analyzed in a number of ways. Due to space considerations, we focus on what we consider to be the most interesting finding from these experiments. The original paper [5] can be consulted for a fuller treatment

For a given instance of a SSAP, we can determine the god's-eye (or *a priori*) optimal assignment. This, however, is of little use in evaluating the assignments of an experimental subject. What we need to do for a given subject and a given problem instance is determine what is *conditionally optimal*, that is, what the subject should do—were she wishing to behave optimally—given what she has done up to that point. For example, the god's-eye optimal policy might dictate that patient j be assigned to Hospital B; but if B is full by the time this patient is observed, then this assignment is impossible. What we need to determine for j is how she should be assigned given the assignments that have previously been made (from 1 to $j-1$). What we report next is based on this notion of conditionally optimal assignment.

Using the experimental data, for each applicant position we can determine the probability that a patient will be assigned to each of the four hospitals under conditionally optimal assignment. Likewise, for each of the n patient positions we can get the empirical probabilities (the proportions) of assignments to each of the four hospitals. When we do so, we find a systematic difference between the empirical and optimal assignment probabilities. Table 11 shows the proportion of times that the empirical probabilities exceeded the optimal ones.

Across the three experiments, the results are unambiguous: The subjects tended to assign patients to the intermediate hospitals with greater probability than was optimal. The results reveal a tendency to try to keep open beds in the extreme (really good (highest cost) and really poor (lowest cost)) hospitals. It seems that the subjects wished to reserve positions for the really injured and the not-so-injured patients, when doing so was suboptimal.

The triage task used in the experiments is quite artificial. However, it is not implausible to suppose that civilians who might be called on to make triage-type decisions in times of truly mass casualty situations (say a low-yield nuclear explosion in NYC) might demonstrate similar types of biases. (There are well-defined triage procedures for trained professionals, but, unfortunately, one can easily imagine situations in which these relatively small group of individuals would be overwhelmed and might have to rely on civilians for help.)

Conclusion

Our primary goal in writing this chapter is to stir the interests of OR researchers in the way *actual* DMs tend to solve the sort of problems faced by *ideal* agents in the Platonic OR world.

Not surprisingly, actual DMs do not always behave in accord with the dictates of optimal decision policies, which (quite often) must be determined by computationally intensive (at least with respect to the capacities of normal humans) procedures. *That* humans do not make optimal decisions is not particularly interesting. *How* it is that they do, in fact, make decisions should, we think, be of considerable interest to OR researchers. There are a number of reasons for this belief.

The first reason is that a consideration of actual human cognition can lead to new research questions. Given what we know about the bounds of human cognition—or simply making reasonable assumptions regarding these bounds—we can formulate new sets of optimization problems. We could then ask: How well do humans perform relative to the appropriate *constrained optimal*? This is different from the standard question regarding optimality, which is: How well do humans perform relative to an *unconstrained optimal*? Determining the appropriate constraints on particular problems and studying the resulting constrained version can, we think, lead to interesting problems. See Shuford [57] for a nice example of this approach.

Bearden and Connolly [2], for example, compared the behavior of subjects in a full-information multi-attribute optimal stopping task under two different conditions. In the *unconstrained* condition, subjects observed the actual values of the attributes of the encountered options. (A complete description of the problem faced by the subjects and the design of the experimental procedure can be found in the original paper. A procedure for computing optimal policies for the problem can be found in Lim, Bearden, and Smith [35].) The *constrained* condition forced the subjects to use a policy of the same form as satisficing, as proposed by Simon [58]. Specifically, the subjects set aspiration levels or cutoffs for each attribute, and were then simply shown whether the attribute values were above (satisfactory) or below (unsatisfactory) their aspiration levels. Based on the procedures developed by Lim et al., Bearden and Connolly found optimal aspiration levels for the problem faced by the subjects; that is, they found optimal satisficing decision policies. One of the more interesting results from this study is that the subjects who were forced to satisfice (or to play the constrained problem) tended to set their aspiration levels too high, which caused them to accumulate higher than expected search costs and, consequently, to obtain lower than expected net earnings. Going beyond what can be concluded legitimately, one might say that even if people do tend to satisfice when they search, they do not do so optimally because they set their aspiration levels too high. Surprisingly, little has been said (in psychology or elsewhere) about the possibility of optimal satisficing. Research that examines the theoretical performance of relatively easy-to-implement heuristics, such as those studied by Stein et al. [60], can be of considerable importance to experimental psychologists and experimental economists, in additional to the typical readers of OR journals (see [22], particularly p. 287-308, for some interesting examples of work along these lines).

Another reason why OR researchers should consider behavioral studies of decision making is that this work may impact more traditional OR problems. Perhaps by understanding the ways in which human behavior compares to optimality, one can design solutions to problems that are robust to human shortcomings. Whether one is determining optimal facility layouts or how fires should best be fought, understanding the properties of the actual agents who will be on the ground working in these facilities and making decisions when fires erupt must have some utility.

There are a number of areas in experimental psychology that we have not discussed in which OR methods play a valuable role (such as the study of problems involving optimal control; see [11] for an overview). And there are many more areas in which OR methods may serve a useful purpose. Unfortunately, much lip service is given to the need for "interdisciplinary research," but it seems that little is done. Very few of our colleagues in psychology have ever heard of INFORMS. Likewise, few of our colleagues in OR have heard of the Society for Judgment and Decision Making, the main research society for behavioral decision research. We can vouch for the fruitfulness of cross-discipline collaboration, and hope that in the future there is more cross-fertilization between experimental psychology and OR.

Acknowledgements

We gratefully acknowledge financial support by a contract F49620-03-1-0377 from the AFOSR/MURI to the Department of Systems and Industrial Engineering and the Department of Management and Policy at the University of Arizona. We would also like to thank J. Cole Smith and Ryan O. Murphy for their feedback.

References

[1] Y. Anzai. Cognitive control of real-time event driven systems. *Cognitive Science*, 8:221–254, 1984.

[2] J. N. Bearden and T. Connolly. Satisficing in sequential search. *Organizational Behavior and Human Decision Processes*, submitted 2005.

[3] J. N. Bearden and R. O. Murphy. On generalized secretary problems. *Theory and Decision*, submitted 2004.

[4] J. N. Bearden, R. O. Murphy, and A. Rapoport. A multi-attribute extension of the secretary problem: Theory and experiments. *Journal of Mathematical Psychology*, submitted 2004.

[5] J. N. Bearden, A. Rapoport, and R. O. Murphy. Assigning patients to hospitals in times of disaster: An experimental test of sequential selection and assignment. 2004.

[6] J. N. Bearden, A. Rapoport, and R. O. Murphy. Sequential observation and selection with rank-dependent payoffs: An experimental test. *Management Science*, submitted 2004.

[7] J. N. Bearden and T. S. Wallsten. MINERVA-DM and subadditive frequency judgments. *Journal of Behavioral Decision Making*, 26:349–363, 2004.

[8] R. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, 1957.

[9] B. Brehmer. Dynamic decision making: Human control of complex systems. *Acta Psychologica*, 81:211–241, 1992.

[10] B. Brehmer and R. Allard. Real-time dynamic decision making: Effects of task complexity and feedback delays. In J. Rasumussen, B. Brehmer, and J. Leplat, editors, *Distributed Decision Making: Cognitive Models for Cooperative Work*. Wiley, Chinchester, UK, 1991.

[11] J. R. Busemeyer. Dynamic decision making. In N. J. Smelser and P. B. Baltes, editors, *International Encyclopedia of the Social & Behavioral Sciences*, pages 3903–3908. Elsevier, Oxford, 2001.

[12] C. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, Princeton, NJ, 2003.

[13] Y. S. Chow, S. Moriguti, H. Robbins, and S. M. Samuels. Optimal selection based on relative rank (the "secretary problem"). *Israel Journal of Mathematics*, 2:81–90, 1964.

[14] Y. H. Chun and R. T. Sumichrast. A rank-based approach to the sequential selection and assignment problem. *European Journal of Operational Research*, forthcoming.

[15] R. M. Corbin. The secretary problem as a model of choice. *Journal of Mathematical Psychology*, 21:1–29, 1980.

[16] C. Derman, G. J. Lieberman, and S. M. Ross. A sequential stochastic assignment problem. *Management Science*, 18:349–355, 1972.

[17] A. Diederich. Sequential decision making. In N. J. Smelser and P. B. Baltes, editors, *International Encyclopedia of the Social & Behavioral Sciences*, pages 13917—13922. Elsevier, Oxford, 2001.

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

[19] R. J. Ebert. Human control of a two-variable decision system. *Organizational Behavior and Human Performance*, 7:237–264, 1972.

[20] T. S. Ferguson. Who solved the secretary problem? *Statistical Science*, 4:282–296, 1989.

[21] T. S. Ferguson. Best-choice problems with dependent criteria. In *Strategies for Sequential Search and Selection in Real Time*. American Mathematical Society, Providence, RI, 1992.

[22] G. Gigerenzer, P. M. Todd, and ABC Group. *Simple Heuristics that Make Us Smart*. Oxford University Press, Oxford, England, 1999.

[23] J. Gilbert and F. Mosteller. Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, 61:35–73, 1966.

[24] A. V. Gnedin. A multicriteria problem of optimal stopping of a selection process. *Automation and Remote Control*, 42:981–986, 1981.

[25] C. Hempel. *Philosophy of Natural Science*. Prentice-Hall, Englewood Cliffs, N.J., 1966.

[26] R. Hertwig and A. Ortmann. Experimental practices in economics: A methodological challenge for psychologists? *Behavioral and Brain Sciences*, 23:383–403, 2001.

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

[28] J. D. Hey. Search for rules of search. *Journal of Economic Behavior and Organization*, 3:65–81, 1982.

[29] J. D. Hey. Still searching. *Journal of Economic Behavior and Organization*, 8:137–144, 1987.

[30] J. M. C. Hutchinson and J. M. McNamara. Ways to test stochastic dynamic programming models empirically. *Animal Behavior*, 59:656–676, 2000.

[31] R. J. Jagacinski and R. A. Miller. Describing the human operator's internal model of a dynamic system. *Human Factors*, 20:425–433, 1978.

[32] J. H. Kerstholt and J. G. W. Raaijmakers. Decision making in dynamic task environments. 1997.

[33] D. Kleinmuntz and J. Thomas. The value of action and inference in dynamic decision making. *Organizational Behavior and Human Decision Processes*, 39:341–364, 1987.

[34] O. Larichev. Normative and descriptive aspects of decision making. In T. Gal, T. Stewart, and T. Hanne, editors, *Multicritia Decision Making: Advances in MCDM Models, Algorithms, Theory and Applications*, pages 5.1—5.24. Kluwer Academic Publishing, Boston, 1999.

[35] C. Lim, J. N. Bearden, and J. C. Smith. Sequential search with multi-attribute options. *Decision Analysis*, submitted 2005.

[36] D. V. Lindley. Dynamic programming and decision theory. *Applied Statistics*, 10:39–51, 1961.

[37] S. Moriguti. Basic theory of selection of relative rank with cost. *Journal of Operations Research Society of Japan*, 36:46–61, 1993.

[38] A. G. Mucci. Differential equations and optimal choice problems. *Annals of Statistics*, 1:104–113, 1973.

[39] G. A. Parker and J. M. Smith. Optimality theory in evolutionary biology. *Nature*, 348:27–33, 1990.

[40] E. L. Pressman and I. M. Sonin. The best choice problem for a random number of objects. *Theory of Probability and Its Applications*, 17:657–668, 1972.

[41] W. V. Quine. *From a Logical Point of View*. Harvard University Press, Cambridge, MA, 1952.

[42] H. Raiffa. *Decision Analysis*. Addison-Wesley, Reading, MA, 1968.

[43] A. Rapoport. Sequential decision-making in a computer-controlled task. *Journal of Mathematical Psychology*, 1:351–374, 1964.

[44] A. Rapoport. A study of human control in a stochastic multistage decision task. *Behavioral Science*, 11:18–32, 1966.

[45] A. Rapoport. Dynamic programming models for multistage decision making. *Journal of Mathematical Psychology*, 4:48–71, 1967.

[46] A. Rapoport. Research paradigms for studying dynamic decision behavior. In H. Jungermann and G. De Zeeuw, editors, *Utility, Probability, and Human Decision Making*, pages 349–369. Riedal, Dordrecht, Holland, 1975.

[47] A. Rapoport, L. V. Jones, and J. P. Kahan. Gambling behavior in multiple-choice multistage betting games. *Journal of Mathematical Psychology*, 7:12–36, 1970.

[48] A. Rapoport and A. Tversky. Cost and accessibility of offers as determinants of optional stopping. *Psychonomic Science*, 4:145–146, 1966.

[49] A. Rapoport and A. Tversky. Choice behavior in an optimal stopping task. *Organizational Behavior and Human Performance*, 5:105–120, 1970.

[50] A. Rapoport and T. S. Wallsten. Individual decision behavior. *Annual Review of Psychology*, 23:131–176, 1972.

[51] M. Sakaguchi. Dynamic programming of some sequential sampling design. *Journal of Mathematical Analysis and Applications*, 71:680–683, 1961.

[52] S. M. Samuels. Secretary problems. In B. K. Ghosh and P. K. Sen, editors, *Handbook of Sequential Analysis*, pages 381–405. Marcel Dekker, New York, 1991.

[53] S. M. Samuels and B. Chotlos. A multiple criteria optimal selection problem. In J. Van Ryzin, editor, *Adaptive Statistical Procedures and Related Topics: Proceedings of a Symposium in Honor of Herbert Robbins*. Brookhaven National Laboratory, Upton, NY, 1987.

[54] L. J. Savage. *The Foundations of Statistics*. Wiley, New York, 1954.

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

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

[57] E. H. Shuford Jr. Some bayesian learning processes. In M. W. Shelly and G. L. Bryan, editors, *Human Judgment and Optimality*. Wiley, New York, 1964.

[58] H. A. Simon. A behavioral model of rational choice. *Quarterly Journal of Economics*, 69:99–118, 1955.

[59] P. Slovic, B. Fischoff, and S. Lichtenstein. Behavioral decision theory. *Annual Review of Psychology*, 28:1–39, 1977.

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

[61] J. D. Sterman. Misperceptions of feedback in dynamic decision making. *Organizational Behavior and Human Decision Processes*, 43:301–335, 1989.

[62] J. D. Sterman. Learning in and about complex systems. *System Dynamics and Review*, 10:291–330, 1994.

[63] M. Toda. The design of the fungus eater: A model of human behavior in an unsophisticated environment. *Behavioral Science*, 7:164–183, 1962.

[64] A. Tversky and D. Koehler. Support theory: A non-extensional representation of subjective probability. *Psychological Review*, 101:547–567, 1994.

[65] B. Van Fraassen. *The Scientific Image*. Clarendon Press, Oxford, England, 1980.

[66] J. von Neumann and O. Morgenstern. *The Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1944.

[67] D. von Winterfeldt and W. Edwards. Flat maxima in linear optimization models. Technical report, Engineering Psychology Laboratory, University of Michigan, 1973.

[68] D. von Winterfeldt and W. Edwards. *Decision Analysis and Behavioral Research*. Cambridge University Press, Cambridge, England, 1986.

[69] M. C. K. Yang. Recognizing the maximum of a random sequence based on relative rank with backward solicitation. *Journal of Applied Probability*, 11:504–512, 1974.

[70] A. J. Yeo and G. F. Yeo. Selecting satisfactory secretaries. *Australian Journal of Statistics*, 36:185–198, 1994.

[71] R. Zwick, A. Rapoport, A. K. C. Lo, and A. V. Muthukrishnan. Consumer sequential search: Not enough of too much? *Marketing Science*, 22:503–519, 2003.