CONSUMER SEARCH: NOT ENOUGH OR TOO MUCH?

Rami Zwick
The Hong Kong University of Science and Technology

Amnon Rapoport
University of Arizona

Alison King Chung Lo
Duke University

A. V. Muthukrishnan
The Hong Kong University of Science and Technology

September 20, 2001

Please address all correspondence to:
Rami Zwick
Department of Marketing
The Hong Kong University of Science and Technology
Clear Water Bay, Kowloon
HONG KONG

Tel: 852-2358-7695
Fax: 852-2358-2429
Email: mkzwick@ust.hk

ACKNOWLEDGEMENT. This work has been supported by Research Grants Council of the Hong Kong Special Administration Region, China (Project No. CA98/99.BM01). The second author wishes to thank the Marketing Department at the Hong Kong University of Science and Technology for its support and hospitality

CONSUMER SEARCH: NOT ENOUGH OR TOO MUCH?

ABSTRACT

We study search behavior in a generalized "secretary problem" environment in which consumers search sequentially for the best alternative from a known and finite set of multi-attribute alternatives. In contrast to most previous studies, we make no distributional assumptions about the quality of the alternatives. Rather, at each stage of the search the consumers are only assumed to be able to *rank order* the alternatives they have already inspected in terms of their overall quality. Our study departs from previous experimental investigations of the secretary problem by including search costs and allowing for recall (backward solicitation) of previously inspected alternatives. Both the number of alternatives and the cost of searching are manipulated experimentally in a factorial design. In the no-cost condition, we find that subjects do not search enough, whereas in the cost condition they search too much. We propose a simple behavioral decision model that incorporates both local and global patterns of the sequence--patterns that should have been ignored by a rational consumer--and then show that it can account for the major patterns of the observed results.

Keywords: Search Behavior, Sequential Choice Models, Information Processing

1. INTRODUCTION

There is a rich theoretical literature on sequential decision-making that has been invoked to account for consumer search (e.g., Ratchford and Srinivasan 1993). Framing the problem as a sequential search for lower prices (e.g., McCall and Lippman 1976), the classical model assumes that time is discrete and that, in each period, the consumer observes a price randomly drawn from a known distribution function. The consumer can accept the current price and terminate the search or reject it and search for a lower price. If she later decides to accept a price that she has already rejected, this price is available with some probability that decreases in the number of periods elapsing since this price was rejected. When the number of prices is unlimited, the optimal policy is to set a critical, time-invariant price (threshold) and accept the first price below this threshold. Versions of this model have been studied in the marketing and job search literatures. A survey of this literature reveals that the degree to which consumers adhere to the amount of search prescribed by the optimal solution remains an open question.

1.1. Is Consumer Search Optimal?

Several studies report *optimal amount of searching*. For example, Ratchford and Srinivasn (1993) and Urbany (1986) found that the marginal returns to searching are broadly consistent with what one might expect if consumers balance the cost and benefits of their search. Avery (1996) reported that the cost-benefit search model provides a good fit for pre-store search activity (but not for the in-store search activity), and Kogut (1992) reported that the optimal model is a good predictor of individual search behavior. Other studies report that consumers engage in *considerably less searching* than is predicted by the economics of information theory (Claxton, Fry and Portis 1974; Furse, Punj and Stewart 1984; Newman 1977; Dickson and Sawyer 1990; Grewal and Marmorstein 1994; Moorthy, Ratchford, and Talukdar 1997). There is a smaller stream of recent experimental research reporting that, under certain conditions, people search *more than prescribed* by the optimal models.

For example, in an experimental study of bargaining and search, Zwick and Lee (1999) reported that, under certain conditions, buyers were willing to search too often while rejecting prices that should have been accepted according to perfect rationality. Similarly, Zwick, Weg, and Rapoport (2000) found that buyers were too willing to search for prices from an additional seller even though perfect rationality dictated that they should stick with the first price.

The two general methodologies for studying sequential search are field studies and controlled laboratory experiments. The former methodology, most common in marketing research, has its obvious advantage of ecological validity. However, it is beset by two fundamental difficulties. First, it is not always clear how to measure the amount of searching. Should a field study rely on the self-reported usage of various sources (e.g., Goldman and Johansson 1978; Urbany, Dickson and Kalapurakal 1996, Moorthy, Ratchford and Talukdar 1997), or on a combination of self-reports and various observable behaviors such as dealers' visits? Additionally, how should internal search be integrated with the amount of external search? The other difficulty with field studies is that a decision policy is only optimal with respect to a model that characterizes the search environment and makes precise assumptions about the consumers' beliefs, goals, and preferences. What seems to be sub-optimal search behavior may simply be due to miss-specification of any of these elements. For example, the search cost plays a crucial role in the theoretical models. How should search cost be measured or inferred in a field study? Furthermore, the classical model and many of its variants assume that consumers know the underlying distribution of prices or any other attribute that is being considered. However, the assumption that the *complete distribution* is known with precision—which is critical for the calculation of the reservation price—is difficult to justify. Moreover, what if the search is for an item that cannot be simply characterized along a single dimension such as most high-ticket durable goods (e.g., a house, vacation, university to attend, etc.)? Would it be reasonable to assume that consumers know the *joint* distributions of all the significant underlying attributes?

In a partial attempt to overcome these difficulties, laboratory experiments have been designed to study sequential search behavior (for a brief review see Camerer 1995). A consistent finding that seems to be repeated in the literature is that people search too little compared with the amount of search prescribed under risk-neutrality (e.g., Schotter and Braunstein 1981; Braunstein and Schotter 1982; Cox and Oaxaca 1989; Hey 1987). The reasons for this behavior are not entirely clear. Kogut (1990) reported that subjects stopped searching too early even under the assumption of risk-aversion, and Sonnemans (1998) found that risk-aversion could not explain the less-than-optimal amount of search. In his survey of the literature, Camerer (1995) suggests that heuristic rules might explain the persistent tendency to under-search. The ability to approximate the optimal policy under heuristic rules is probably due to the fact that, in many search environments, the optimal policy is not very sensitive to deviations ("flat maxima") and that several heuristic rules can approximate the optimal solution quite well (Moon and Martin 1990; Seale and Rapoport 1997).

Statements about "too much searching" or "considerably less searching" are obviously only meaningful with respect to an optimal policy. We argue that optimal policies serving as benchmarks in the previous investigations are based on assumptions that are too strict to be met in practice. In the present study, we propose an alternative search model that is based on more plausible assumptions and allows for more generality. Unlike the classical search models, our model does not assume that consumers possess precise knowledge about the characteristics of the underlying distribution of prices or any attribute under consideration. Moreover, and also unlike the classical search models, we do not assume that the search for an item is characterized along a single dimension. Rather, the model that we propose and test below incorporates only information about the rank ordering of the items being considered. This allows generalization of the model to multi-attribute items while bypassing the thorny issue of how the attributes are integrated by the consumer.

1.2. The Search Environment

We study individual search behavior in a class of sequential decision problems in which a decision-maker (DM) is faced with a set of n objects that are to be inspected sequentially. These "objects" may be applicants for a job, apartments for rent, or offers for purchasing an object for sale. In all these cases, during each period of the search process, the DM is faced with a dilemma whether to accept the current object and thereby terminate the search or wait in the hope of obtaining a better object later. There are a number of real-world consumer decisions that are characterized by such environment. For example, when a graduate student who has just arrived to a new college town constructs or otherwise receives a list of apartments, she must decide at each stage of her inspection whether to rent the current apartment or inspect another apartment on her list. In doing so, she runs the risk that an apartment that has already been inspected but not chosen may not be available later. Likewise, when a consumer shops for a used car, he must decide whether to purchase the car he has just inspected or visit another seller and inspect another, hopefully better car. If the delay is long, a previously inspected car may no longer remain available. Finally, when a seller offers a certain item for sale in the market, he must decide upon receiving an offer whether to accept it or reject it and wait for another, possibly better offer.

Sequential observation and selection (search) processes of this type have been studied both theoretically and experimentally in economics (e.g., Cox and Oaxaca 1989 on job searches in labor economics), marketing (e.g., Kogut 1992 on consumer search), and psychology (Rapoport and Tversky 1970; Shapira and Venezia 1981). Our study differs from these studies in the assumptions we make about the DM's information structure.

With few exceptions, the experimental literature on sequential observation and selection has assumed that the objects are drawn from a distribution whose parameters are known with precision.

As mentioned earlier, precise knowledge of the distribution is critical, because without it the optimal reservation values cannot be calculated. These values are often at the right tail of the distribution and are very sensitive to minor changes in the distribution parameters. We contend that this is a strict assumption seldom met in practice. Moreover, the validity of this assumption is difficult to ascertain. The subjects in these experiments are typically instructed that objects are to be independently and randomly drawn from a distribution (e.g., normal) with known parameters. It is unreasonable to assume that inexperienced subjects, even if presented with random samples for a distribution, can accurately measure the probability of each observation and then compare it to some cutoff point, which they are supposed to know how to compute.

We replace this assumption with a considerably weaker assumption that, at each period of the observation and selection process, the DM can only determine the *relative rank* of the current object with respect to the objects that have already been inspected. This assumption allows application of the model to multi-attribute alternatives. With reference to the previous examples, under our assumption the graduate student who inspects potential apartments can only conclude at the end of each inspection that the current apartment is the best, second best, third best, or worst apartment she has seen so far. After driving a few cars, one at a time, and evaluating them in terms of price, comfort, size, color, etc., the car buyer can typically only rank order the most recently tested car in terms of its overall quality with respect to the ones he has already inspected.

Sequential observation and selection problems of this type have attracted considerable attention from applied mathematicians whose major focus has been on characterizing the optimal policies for a rich class of problems (Freeman 1983). In contrast, with only a few exceptions (Corbin, Olson, and Abbondanza, 1975; Seale and Rapoport 1997, 2000), there have been no attempts to study the decision rules that subjects use to select the best objects in this class of sequential decision tasks.

In what follows, we state a set of assumptions that capture many of the features of consumer search for multi-attribute products like cars, apartments, gifts, children's toys, etc., where an underlying distribution of overall quality cannot be reasonably presumed.

1.3. The Model Assumptions

- There is only a single object to be selected.
- The number of objects to be observed sequentially (which constitute the so-called "consideration set"), denoted by n, is known.
- The n objects are presented to the DM one at a time, in a random order. Therefore, before the search commences, each of the n! orderings is assumed to be equally likely.
- There is a fixed per-object cost for searching[1].
- The DM can determine only the rank (in terms of preference, attractiveness, quality) of the current object relative to the ones she has already observed (with no ties).
- At each period, the DM can either accept the current object, continue to search for the next object, or recall[2] any of the objects she has already observed. If she recalls an object, it is assumed to be available with a certain probability that is known by the DM. If the object is unavailable, the DM can recall another object or continue the search. An object that is not available remains unavailable forever.
- The DM's objective is to select the overall best object from the consideration set.

[1] This can be set to zero. The model rules out the possibility that consumers derive pleasure from the sheer act of searching.

[2] "Recall" in our paper (also called "backward solicitation") does not refer to the act of recalling information from memory but rather to the ability to "call back" an item after it has been passed over.

1.4. The Search for a New Apartment Scenario

To illustrate the sequential search problem, consider the following scenario used in our experiment. Imagine that you intend to move to a new apartment. Your real estate agent has constructed a list of potential apartments that seem to satisfy your needs. Were you to visit all the apartments on the list, you would be able to rank them from best to worst (with no ties). However, visiting all the apartments on the list, besides being costly in terms of your time and the cost of extending your stay where you now live, does not guarantee success. First, the best apartment may no longer be available after you have visited all the apartments (unless it is the last on the list). Second, visiting and inspecting all the apartments on the list is costly and time consuming. Therefore, as is typically the case, you decide on the following plan of action: You visit the first apartment on the list (bearing the cost of one visit). If you believe that this apartment can satisfy your needs, you rent it; otherwise, you inspect the second apartment on the list (bearing the cost of two visits). At this point, you compare the two apartments to each other. If you believe that the second apartment satisfies your needs, you rent it. If not, you either decide to inspect the third apartment on the list, or call back the owner of the first apartment to find out if it is still available. (Given that the second apartment is inferior, you may now believe that the first apartment can satisfy your needs and reduce your total search costs). If the first apartment is available, you rent it. If not, you visit the third apartment on the list (bearing the cost of three visits). This search process continues until you choose an apartment that satisfies you.

Assume that the list contains twenty apartments and that you will be satisfied only if you rent the best apartment on the list ("nothing but the best")[3]. Further, assume that the probability that an apartment that was inspected but not selected is still available in a later period declines geometrically

[3] This assumption is restrictive. We discuss possible extensions in the Discussion section.

with the number of periods since this apartment was inspected (p^r, where r is the lag time[4]). How many apartments would you visit before renting or deciding to recall a previously visited apartment? Will you visit too few apartments in agreement with the statements one finds in the marketing literature on search, will you search too much, or will your behavior be accounted for by the optimal decision policy outlined below?

1.5. Optimal Search Policy

The following search policy maximizes the probability of selecting the overall best alternative, given the above scenario with no search cost: Skip the first s_0-1 items. In period s_0 choose the relative best object (no matter in which period it has been observed[5]). If the relative best (out of the first s_0 items) is not available, continue to search and then choose the next relatively best object in the sequence (Yang, 1974).

The value of s_0 -- the first stopping period -- is the largest s that satisfies the inequality

$$p > 1 - [(s-1)(1-c(s))]^{-1}$$

where

$$c(s) = \begin{cases} \sum_{j=s+1}^{n} \frac{1}{j-1} & if \quad s \neq n \\ 0 & if \quad s = n. \end{cases}$$

The optimal cutoff point for the case with a per unit search cost cannot be expressed by a simple formula. Hence, it is determined numerically.

---

[4] The intuition behind this assumption is as follows. Suppose it takes one day to inspect an apartment and that, if it is not rented immediately, it has the probability of (1-p) of being rented by someone else each day. Hence, the probability that the apartment is still available after r days is p^r.

[5] This, of course, might require recall.

Figure 1 presents the cutoff points (periods) for various lists' length ranging from 1 to 90 for three cost values ($0, $0.1 and $0.3). It assumes a reward of $10 for selecting the best overall object and a probability parameter of r=0.9 for determining the availability of previously inspected objects. To generate this figure, each sequence (length) was replicated 10,000 times, and the expected gain was computed for all possible cutoff values (1 to the end of the list). The figure presents the cutoff values with the highest mean expected payoff for each list's length (i.e., the optimal cutoff value).

Figure 1

CUTOFF PERIOD (S_0) AS A FUNCTION OF LIST LENGTH AND SEARCH COST*

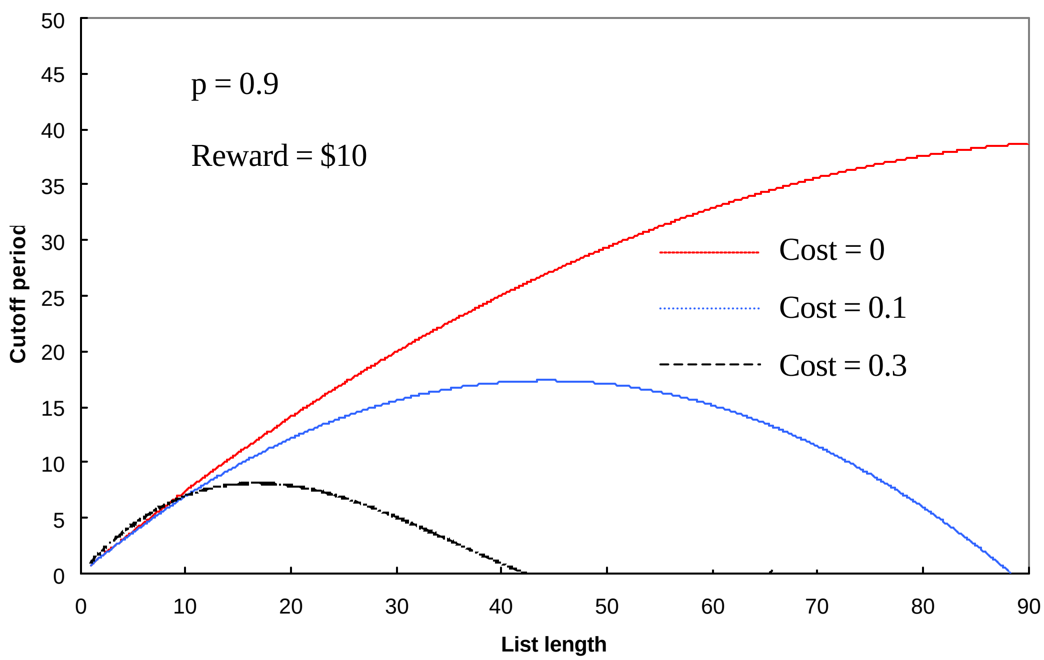

* Based on simulation results with 10,000 sequences for each list's length

The optimal policy dictates that the decision either to accept or reject an object depends only on the currently observed relative rank of this object and on the period of observation, but not on the history of the process.

The remainder of this paper is organized as follows. In the next section, we describe an experiment designed to test the amount of searching by consumers in the search environment described above. The third section presents the major result of the study that, in comparison to the optimal policy, most subjects in our experiment did not search enough in the two no-cost conditions but searched too much in the two cost conditions. We next attempt to account for the observed regularities by proposing a simple behavioral search model that explains simultaneously both under- and over-searching. The paper concludes with a discussion of the limitations of our study and the significance of our findings from both managerial and consumer perspectives.

2. METHOD

2.1. Subjects

Ninety-seven subjects, all undergraduate business students from the Hong Kong University of Science and Technology, participated in several sessions each lasting about 90 minutes. Subjects were recruited through advertisements placed on bulletin boards on campus and class announcements. The announcements promised monetary reward contingent on performance in a marketing study.

2.2. Procedure

The search task was implemented as a Java program on a PC[6], and was introduced as a search for a rental flat[7]. The subjects read online instructions at their own pace and were required to answer correctly several questions before continuing with the experimental task[8]. Subjects were instructed to imagine that their real estate agent had constructed a list of potential flats and that their task was to

[6] The Program is available from http://cebr.ust.hk/software/search/s.zip

[7] A product class frequently studied in conjoint research (cf. Green, Helsen, and Shandler 1988; Johnson, Meyer, and Ghose 1989).

[8] The instructions were delivered via a PowerPoint slide show through the "browsing at a kiosk (full screen)" mode. The instructions (in PowerPoint format) for the (20, cost) condition are available from http://home.ust.hk/~mkzwick/instr_20_COST.ppt

select ("rent") the best flat from the list by sequentially inspecting the flats. The instructions emphasized that:

- on paper all the flats are equally attractive;
- the order in which flats are visited implies nothing about the desirability of the flat;
- only the (relative) rank of the flat, relative to all other flats that have already been visited, is revealed by visiting and inspecting a flat;
- visiting a flat says nothing about other flats that have not been visited;
- the chances that a flat that has been visited in the past is still available decrease with the time since the visit.

In two of the four conditions, subjects were instructed that each visit to a flat bears a fixed cost.

Table 1 presents an example of the decision screen. The parameters of the game (number of flats on the list, cost per visit (if any), objective of the DM, and the reward for selecting the best flat) are presented in the top right section of the screen. On the left of the screen a table presents the history of the search up to the current period. It displays the relative ranks of the flats that have already been visited (column 2) and the probability that a previous flat is available if an attempt is made to rent it (column 3). The current period (10 in this example) is highlighted in red, and the best flat from the ones that were inspected is highlighted in blue (flat no. 2 in this example).

Table 1

EXAMPLE FOR THE DECISION SCREEN PRESENTED TO THE SUBJECTS*

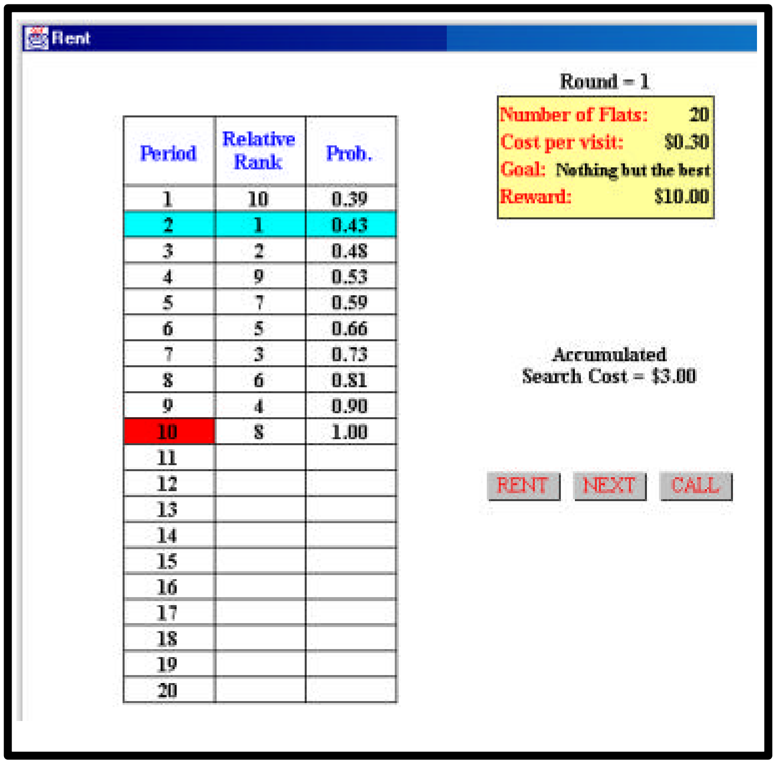

Period	Relative Rank	Prob.
1	10	0.39
2	1	0.43
3	2	0.48
4	9	0.53
5	7	0.59
6	5	0.66
7	3	0.73
8	6	0.81
9	4	0.90
10	8	1.00
11		
12		
13		
14		
15		
16		
17		
18		
19		
20		

* This is for the 20 flats case and positive search cost of $0.30. The statement on the screen "Goal: Nothing but the best" indicates that the reward (of $10) is awarded only if the best flat overall is selected. In the above example, the subject is currently visiting the 10th flat. Its relative rank is 8 and the best flat so far was visited in period 2. The chances that the flat from period 2 is still available are 0.43.

After visiting a flat (that is not the first on the list), subjects could choose one of three actions (see Table 1):

- RENT the current flat, thereby terminating the search,
- Visit the NEXT flat, or
- CALL a previously visited flat[9]. If this flat is available, it is rented and the search is terminated; otherwise, the same three options are again available to the subject[10].

[9] This option, of course, was not available in the first period.

[10] As mentioned earlier, if recalling a flat reveals that it is not available, it remains unavailable for the duration of the round.

If a flat is rented, or after being recalled it is available, the search is terminated and the (absolute) ranks of all the flats on the list are unveiled. If the selected flat is the best, the subject is provided with the reward minus his or her accumulated search cost (if any). If the selected flat is not the best, the subject is charged an amount equivalent to his or her accumulated search cost (if any).

Each subject participated in 100 rounds of the search game with the same parameter values. The number of rounds was not disclosed. Subjects were informed that they would be paid a certain percentage of their cumulative earnings from all 100 rounds plus $10 for completing the session[11]. On average, the subjects earned $159.43[12].

2.3. Experimental Design

A 2 (number of flats on the list) × 2 (search cost values) × 100 (replications) factorial design was used. The first two factors were between subjects, and the last within subjects. The number of flats, n, was either 20 or 60. The search cost was either $0 (zero) or positive ($0.1 and $0.3 in the 60 and 20 flats conditions, respectively). The reward for selecting the best flat was fixed at $10 and the probability that a flat was still available upon recall was set at 0.9^r, where r is the lag time from the current period (r=0,1,2, …)[13]. To render the task to be about equally profitable in all four conditions (based on the optimal policy), subjects were paid different percentages of their cumulative earnings in the four conditions, and the search cost (when it was positive) also varied as a function of the list length. Table 2 presents the parameter values, number of subjects in each condition, cutoff period of the optimum policy, probability of selecting the overall best flat, and the associated expected gain if the optimal policy is followed.

---

[11] All amounts are in Hong Kong dollars. The exchange rate at the time of the study was US$1=HK$7.8.

[12] $50 per hour was the hourly wage for an on-campus job at the time of the study.

[13] The formula was explained in the instructions and a hard copy table with probabilities associated with all possible lag times was available to the subjects at all times.

Table 2

PARAMETER VALUES IN THE FOUR EXPERIMENTAL CONDITIONS AND OPTIMAL STRATEGIES

	20 flats		60 flats	
Search cost	\$0.0	\$0.3	\$0.0	\$0.1
Cutoff (s_0)	15	6	32	13
Prob.	0.502	0.320	0.400	0.269
Expected gain	5.020	0.897	4.000	0.525
Payment	40%	300%	50%	500%
N	23	22	28	24

Cutoff	The cutoff period based on the optimal policy
Prob.	The probability of selecting the best flat if the optimal policy is followed
Expected gain	The expected gain per round if the optimal policy is followed
Payment	The percentage of the accumulated gains (in all 100 rounds) paid to subjects at the end of the session
N	Number of subjects

3. RESULTS

Figure 2 illustrates the efficiency of searching in comparison with the optimal policy (100% efficiency) for various cutoff values in the neighborhood of the optimal cutoff value (±6). As might be expected, the optimal policy is more sensitive to deviations when the number of items is smaller and the cost is higher. The efficiency loss for sub-optimal behavior can be substantial. For example, setting a cutoff value of 6 periods too early in the shorter list (n=20) results in an efficiency loss of 21.4% and 77.8% for the no-cost and cost conditions, respectively. Such significant losses indicate that the optimal policy is not only theoretically important but also has practical significance.

Figure 2

EFFICIENCY OF SEARCH FOR VARIOUS CUTOFF VALUES IN THE NEIGHBORHOOD OF THE OPTIMAL VALUE

N = 20

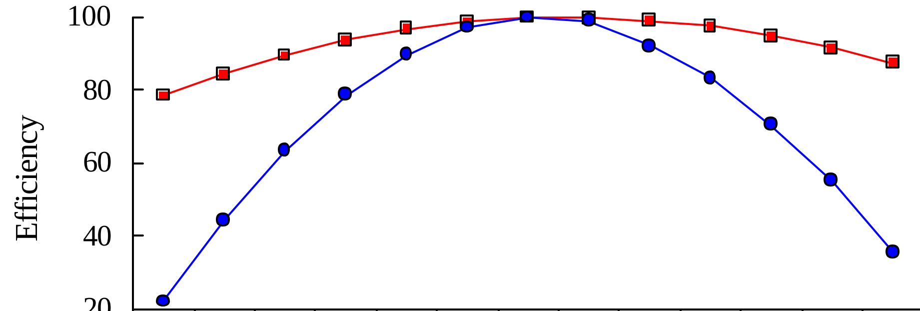

	-6	-5	-4	-3	-2	-1	0	+1	+2	+3	+4	+5	+6
No Cost	78.6	84.5	89.6	93.7	96.8	98.7	100	99.8	99.1	97.5	94.9	91.5	87.5
Cost	22.2	44.4	63.4	78.8	89.8	97.2	100	99.0	92.0	83.3	70.6	55.1	35.5

N = 60

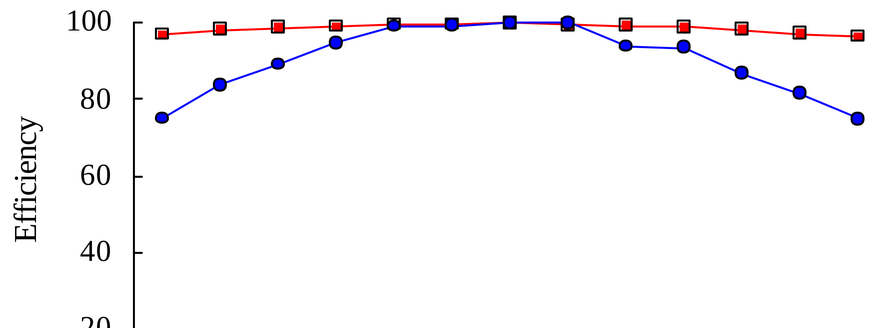

	-6	-5	-4	-3	-2	-1	0	+1	+2	+3	+4	+5	+6
No Cost	97.1	98.1	98.8	99.2	99.6	99.7	100	99.4	99.2	98.9	98.2	97.2	96.4
Cost	75.0	83.8	89.2	94.9	99.1	99.2	100	99.9	93.8	93.4	86.6	81.5	74.7

3.1. Period of First Action

Figure 3 displays the major findings, namely, the frequency distributions of the **P**eriod of **F**irst **A**ction (PFA) that can be either to recall a previous flat or accept the present flat. The distributions are shown by list length (row), cost/no-cost condition (column), and subject (box plot). The bottom and top edges of each box plot are located at 25th and 75th percentiles of the sample. The vertical lines extend from the box as far as the data extend, to a distance of at most 1.5 interquartile range. Any value more extreme than this is marked by a dot. A reference line is drawn horizontally at the optimal PFA. Subjects are ordered from left to right by the value of the 25th percentile of their distribution.

Inspection of Fig. 3 reveals that for both lists (n=20 or n=60) most subjects took the first action earlier than prescribed by the optimal policy for the no-cost condition, thereby supporting the common claim that consumers do not search enough. However, the opposite result is depicted in the two cost conditions, where, for most subjects, the PFA exceeds the value prescribed by the optimal policy on most trials, thereby indicating excessive searching. A statistical analysis is presented next to support these claims.

To set our analysis in the context of previous research, we refer to each item in the sequence as "applicant." An applicant who is best relative to the ones that have already been inspected is called a "candidate." Note that being a candidate is a necessary but not sufficient condition for being the best overall. Table 3 presents the mean and standard deviation of the PFA by the type of the first action (accept or recall) and by the status of the applicant that has been inspected at the period of the first action (candidate or not). The boxed cells indicate the cutoff period for the optimal policy, and the shaded entries indicate errors in either accepting a non-candidate or recalling a candidate in a period when another candidate is encountered (clearly the current candidate is the highest ranked).

Figure 3

PERIOD OF FIRST ACTION

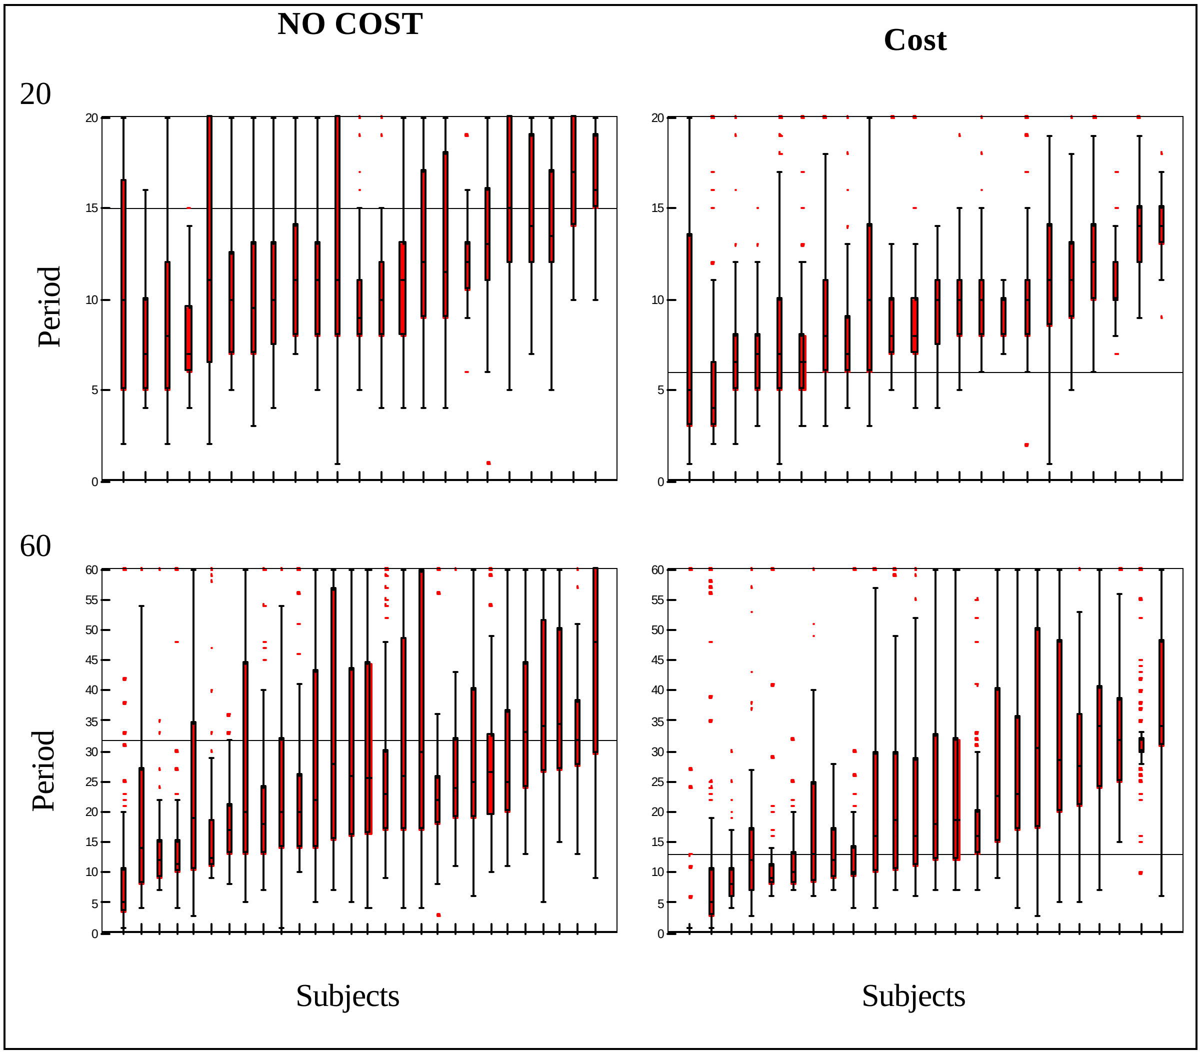

Table 3

MEAN AND STANDARD DEVIATION OF PERIOD OF FIRST ACTION

	Applicant in the PFA		20 Accept		20 Recall	60 Accept		60 Recall
No cost	Candidate	MEAN	11.26	15	12.29	23.15	32	20.74
		STD	5.39		4.11	15.54		17.28
		%N	6		2	8		1
		t	8.04			8.48		
	Not a candidate	MEAN	19.73	15	11.77	50.47	32	25.37
		STD	1.79		4.59	19.71		14.54
		%N	2		90	7		84
		t			32.04			22.18
Cost	Candidate	MEAN	8.69	6	10.33*	17.77	13	27**
		STD	4.44		7.37	16.03		15.56
		%N	14		0	13		0
		t	10.45			5.34		
	Not a candidate	MEAN	7.31	6	9.79	52.22	13	21.2
		STD	6.39		3.76	13.55		14.35
		%N	2		84	4		83
		t			43.29			25.48

* Based on 3 observations
** Based on 2 observations
%N: % of time the relevant category occurred in the period of first action.

t: t value for testing the hypothesis that the observed mean is the same as expected (appears in the box to the right or the left of the mean).

3.2. The Effect of Experience, Search Cost, and Set Size on the Period of First Action

To test for the effect of experience we divided each subject's responses into two blocks of 50 trials each. We than subjected the PFA scores to a $2 \times 2 \times 2$ (search cost by number of items by block) ANOVA (with repeated measure on the block variable). Block and the interactions of block with the other two variables, as well as the triple interaction were not significant (F=3.01, 3.14, 2.86 and 2.72, for block, the two interactions of block with cost and block with number of items, and for the triple interaction, respectively, p>0.05). Search cost and number of items affected the amount of search before the first action. Both main effects and the interaction were significant (F=3026.63, 215.19 and 27.69, for number of items, search cost, and the two-way interaction, respectively, p<0.0001). On average, the number of skipped items before the first action was higher for the no-cost than for the cost condition (20.10 vs. 16.00), and higher for the 60 than for the 20 items lists (24.55 vs. 10.77). The interaction is due to the fact that the cost reduces the number of skipped items by about 5 when n=60 but only by 2 when n=20.

The first action was taken earlier than predicted by the optimal policy in the no-cost condition and later than predicted in the cost condition for both list lengths and when actions were not classified as errors[14].

For the no-cost 20-item list, the mean PFAs are 11.26 and 11.77 for the accept and recall first actions, respectively, whereas the predicted cutoff period is 15. For the no-cost 60 items list, the mean PFAs are 23.15 and 25.37 for the accept and recall first actions, respectively, whereas the predicted cutoff period is 32. All of these differences are significant at the 0.001 level (see t values in the table). For the cost 20-item list, the means PFAs are 8.69 and 9.79 for the accept and recall first action, respectively, whereas the predicted cutoff period is 6. For the cost 60-item list, the means PFAs are 17.77 and 21.2 for the accept and recall first actions, whereas the predicted cutoff period is 13. All of these differences are also significant at the 0.001 level (see t values in the table).

[14] The highest error level is 8% in the 60 no-cost cell. The corresponding error levels for the other cells are: 4%, 2% and 4% for the (20, no-cost), (20, cost) and (60, cost) conditions, respectively. Many of the errors of accepting a non-candidate are due to accepting the last applicant without an attempt to recall the last candidate. For example, the mean period in the (20, no-cost) condition for such an action is 19.73, and 50.47 in the (60, no-cost) case, and 52.22 in the (60, cost) case.

3.3. The Effect of Search Cost and Set Size on the Search Termination Period

The majority of first actions were to recall in all four conditions. The percentages were 90.02, 84.01, 84.00 and 83.00 in the (no-cost, 20), (no-cost, 60), (cost, 20) and (cost, 60) conditions, respectively. The predicted percentages derived from the optimal policy are 93.3, 96.87, 83.33 and 92.31[15]. Since accepting a candidate terminates the search with certainty, the lower observed rate of recall as a first action in the no-cost condition coincides with the above findings that less searching than expected is taking place. For the cost condition, the rate of recall as the first action is close to that expected when n=20, and below expectation when n=60. This could imply that although the first action is taken later than predicted and because the action is more likely than expected to be acceptance rather than recall, the overall number of inspected items is not necessarily higher than predicted. However, our next analysis shows that this is not the case.

Figure 4 displays the observed and predicted cumulative probability distributions of the search being terminated at any particular period. The theoretical distributions are depicted with a dashed line and the observed distributions with a solid line. The no-cost conditions are represented by thick lines and the cost conditions by thin lines. The findings regarding the PFA extend to the overall number of items inspected before the search is terminated. For the no-cost case, for both list lengths and any given period, the observed probability of search termination is higher than expected, emphasizing the fact that searching is less extensive than predicted. On the other hand, for the two cost conditions, except for the few periods before the predicted cutoff period (where termination is never expected), the observed probability of termination is lower than expected, emphasizing the fact that searching is more extensive than predicted.

[15] The probability that the j^{th} applicant is a candidate is 1/j. Hence, the probability that the first action is recall is given by $1–1/s^*$, where s^* is the cutoff period.

Figure 4

OBSERVED AND EXPECTED CUMULATIVE PROBABILITY DISTRIBUTIONS OF SEARCH TERMINATION BY PERIOD

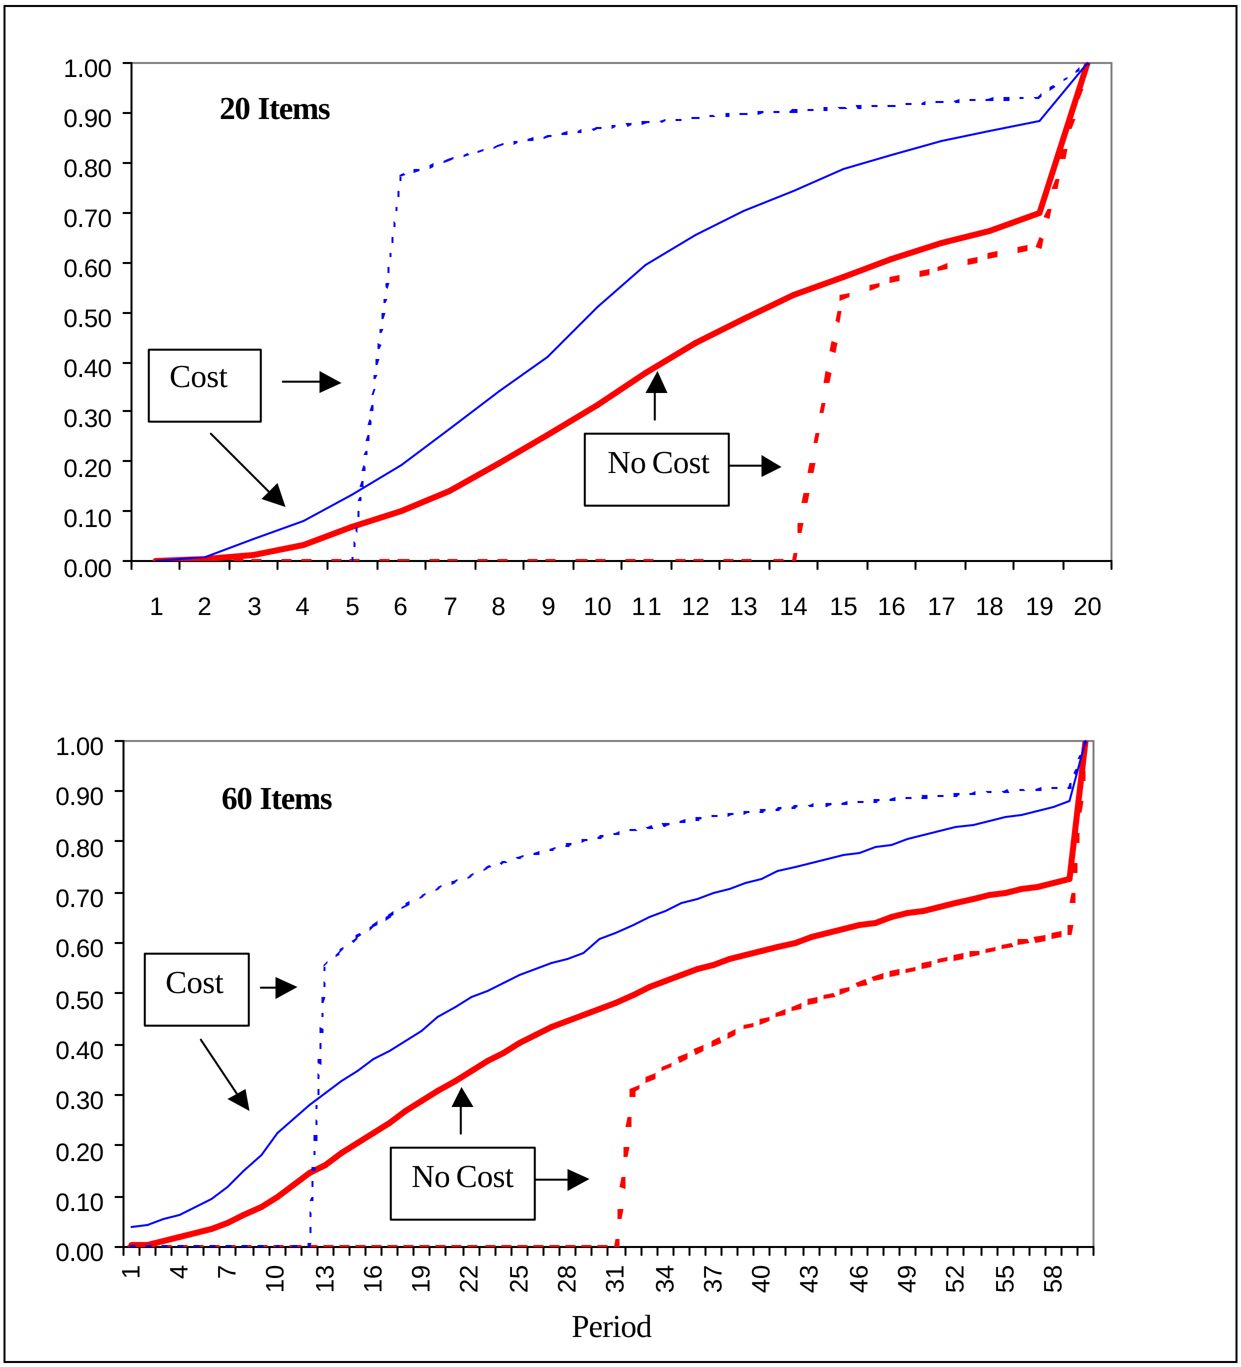

3.4. Cost of Sub-optimality

Table 4 compares the observed vs. predicted payoffs of individual subjects. Predicted payoffs are derived from the optimal policy; they are calculated on the basis of the *actual sequences* observed by each subject. Each sequence was simulated 100 times to account for the probability that a recalled item might not be available if the optimal policy required recall.

By not searching enough in the two no-cost conditions, subjects lost 15.67% and 21.28% of the total potential accumulated payoffs in comparison with the optimal policy for the cases n=20 and n=60, respectively. By searching too much in the two cost conditions, the loss was much higher at 40.69% when n=20 and more than 100%[16] when n=60. Clearly, sub-optimal behavior, especially in the cost condition, is significant and deserves attention. Notwithstanding the previous statement, a few subjects did better than expected (there were, respectively, 2, 8, 2 and 3 in the [20, 0], [20, 0.3], [60, 0] and [60, 0.1] conditions).

We have shown that most subjects did not follow the optimal policy and that the deviation from optimality was substantial and systematic. In the no-cost condition, subjects did not search enough, but in the cost condition they searched too much. Additionally, the sub-optimal search behavior was rather costly. How can we explain such behavior, and in particular can a descriptive model account for both the under- and over-search as a function of search cost?

---

[16] It is not quite clear how to measure the loss in efficiency here since the predicted value is positive and the observed value is negative.

Table 4

OBSERVED AND EXPECTED ACCUMULATED PAYOFFS

	20				60			
	No cost		Cost		No cost		Cost	
Subject*	OBS	EXP	OBS	EXP	OBS	EXP	OBS	EXP
1	440	483.8	-78.1	89.1	30	414.5	-19.4	32.7
2	420	524.0	-19.9	51.0	210	366.7	3.1	94.7
3	320	494.9	124.5	93.3	270	359.3	88.3	48.4
4	440	458.3	34.8	82.6	360	405.1	31.5	50.7
5	410	477.7	59.2	149.9	260	394.5	33.7	52.5
6	370	494.5	90.3	100.1	310	384.9	-3.1	-5.4
7	360	474.1	30.0	96.3	330	452.1	7.5	76.5
8	390	538.1	15.2	78.8	280	360.8	79.1	71.6
9	470	532.7	1.8	133.0	230	356.8	5.3	92.3
10	430	543.8	47.7	142.6	280	362.6	65.8	94.6
11	350	503.2	133.7	75.3	430	458.0	-48.5	16.2
12	360	532.1	66.3	128.3	310	408.2	-174.3	69.3
13	440	510.8	70.4	72.2	330	373.2	-1.2	37.6
14	330	477.9	151.3	130.4	340	379.4	87.8	100.8
15	430	517.9	27.8	52.7	330	409.9	17.0	84.1
16	330	453.4	56.8	34.1	360	403.2	-57.5	33.8
17	460	460.7	53.1	52.3	330	360.8	19.7	73.2
18	450	477.2	62.5	45.5	310	411.3	-22.5	74.2
19	570	562.5	1.7	118.5	320	412.2	-134.8	9.1
20	480	508.7	115.1	81.5	400	451.4	-30.9	14.2
21	470	523.9	116.4	77.4	290	416.7	-175.4	-2.7
22	540	490.3	13.9	95.3	410	400.7	-149.7	71.5
23	440	462.2			340	363.8	-42.3	82.3
24					280	386.7	-50.3	6.2
25					420	413.0		
26					320	365.7		
27					270	348.6		
28					360	444.6		
MEAN	421.7	500.1	53.4	90.0	311.0	395.2	-19.6	53.3
STD	64.5	30.5	55.0	32.9	76.8	31.6	76.1	33.6

* - Subjects' ID is based on the ranking presented in Figure 3

3.5. Behavioral Decision Rules

The findings reported above indicate that the stopping and recall criteria used by the subjects do not remain fixed as prescribed by the optimal policy but are sensitive not only to the search costs and list length but also to the characteristics of the observed sequence. The optimal policy dictates that the decision to accept or reject a candidate or to recall a previously encountered candidate depends only on the currently observed relative rank (which must be 1 in the accept case and different from 1 in the recall case) and on the period of the observation, but not on the history of the process. All candidates before the cutoff period are initially rejected, and the first candidate after the cutoff period is accepted, if it is reached. Based on previous studies of sequential search (Seale and Rapoport 1997, 2000; Saad and Russo 1996; Kraft and Lee 1979), we investigated the effects of the following two factors on the decision to accept or reject a candidate and on the decision to recall a previously encountered candidate: (1) the number of **P**eriods **S**ince the **L**ast **C**andidate was encountered (PSLC)[17], and (2) the **A**verage **R**ate **O**f **C**andidate **A**rrival (AROCA) computed at each period as the number of previously encountered candidates divided by the period - 1.

3.5.1. Decisions to Accept

The decisions in periods in which a candidate was encountered to either accept the candidate or continue the search[18] (as the dependent variables) were subjected to a multivariate dynamic ordered probit analysis (Dueker 1999) with period, average rate of candidate arrival (AROCA), and number of periods since the last candidate was encountered (PSLC) as the independent variables. Table 5 (upper panel) presents the maximum-likelihood estimates of the regression coefficients where the parameters

[17] PSLC and AROCA are not defined for t=1; for the present analysis we set PSLC=0 and AROCA=0 for t=1.

[18] It is a strategic error to recall a previously encountered candidate in a period in which a candidate is encountered. For the purpose of this analysis, errors of this type were omitted. Table 3 shows that these errors occurred very infrequently.

correspond to the Accept decision. In all conditions, a Pearson chi-square overall goodness of fit test cannot be rejected[19].

As mentioned earlier, the optimal policy implies that the average rate of candidate arrival and the length of time since the last candidate was encountered are irrelevant. The observed behavioral patterns clearly violate these implications. Although in all conditions the rate of accepting a candidate increases with the period (see the positive significant coefficients for the period in all conditions), these rates are also a function of the number of periods since the last candidate was encountered (see the positive significant coefficients for PSLC in all conditions). In the shorter lists (n=20) it is a function of the average rate of candidate arrival (see the negative significant coefficients for AROCA).

AROCA can be interpreted as a global measure of arrival rate computed over the entire sequence up to the current period. PSLC can be interpreted as a local measure of arrival rate, relevant in the vicinity of the current period. In the shorter lists (n=20) for both the cost and no-cost conditions, subjects are sensitive to both measures. If, given a specific period, *v*, in which a candidate has been encountered, the sequence up to *v* had relatively many candidates, then the probability of immediately accepting that candidate is lower compared to experiencing a more sparse sequence before period *v*. This behavior follows the (erroneous) belief that the past rate of arrival will continue to the future, hence given a past sequence "rich" with candidates, subjects believe that many more candidates are

[19] We have used the multivariate dynamic ordered probit analysis proposed by Dueker (1999) because it allows for conditional heteroscedasticity to exist in a qualitative response model of time series data. The procedure addresses this issue by adding Markov-switching heteroscedasticity to a dynamic ordered probit model. Note that in our data period and PSLC are structurally correlated: PSLC is always less than or equal to the period, however, given a period, except for the upper bound on PSLC it is free to vary from 1 to this bound. Moreover, given the random sequences, we can compute the probabilities of the various switching stages. The analysis suffers from a censuring problem. Given an acceptance, the sequence is terminated. Further, the probability that a random sequence contains, for example, only one candidate before the 20th period is very low; hence, the deeper we get into the sequence the fewer actual observations we have to rely upon for low values of AROCA.

likely to appear in the remaining sequence[20]. Conversely, a sequence in which candidates appear infrequently (an extreme case is if, by chance, the applicants appear in a descending order) increases the rate of acceptance of the current candidate. This is in line with the finding that people tend to perceive local patterns in an otherwise global random sequence. Similarly, a recent shortage of candidates increases the likelihood that the current candidate will be accepted.

In the longer lists (n=60) for both cost and no-cost conditions subjects were only sensitive to the period and the local measure of arrival rate. Note that AROCA is a much more difficult index to process cognitively. It was not immediately available on the screen as such[21]. On the other hand, PSLC was immediately visible on the screen since the last candidate was always highlighted in blue (see Table 1). Consequently, we conjecture that the recency effect was more pronounced in a condition where the derivation of the global measure of arrival rate was harder to process.

---

[20] For every period, we can compute the probability that n candidates were encountered previously. While it is true, for example, that it is more likely that a candidate in period 9 is the 4th rather than the 2nd candidate, the probability that this candidate is the overall best is 1/(n-9) in both cases. That is, the probability that a candidate is the overall best is independent of the number of previously encountered candidates.

[21] All the information needed to compute this index was available on the screen at all times. However, it required some processing of determining, for example, if an applicant was a candidate at the time when it was presented.

Table 5

PROBIT ANALYSIS

Decision to accept									
		No Cost				Cost			
		EST	STD	χ^2	Pr > χ^2	EST	STD	χ^2	Pr > χ^2
20	Intercept	-1.682	0.188	48.15	0.000	-1.555	0.055	23.38	0.000
	Period	1.287	0.148	22.01	0.000	1.674	0.260	23.19	0.000
	PSLC	1.889	0.231	8.23	0.006	1.192	0.092	26.11	0.000
	AROCA	-1.706	0.363	3.09	0.037	1.746	0.263	2.82	0.044
60	INTER	-1.594	0.058	126.75	0.000	-1.234	0.325	218.89	0.000
	Period	1.582	0.135	18.06	0.000	1.912	0.331	15.94	0.000
	PSLC	1.502	0.207	7.35	0.000	1.714	0.053	18.71	0.000
	AROCA	1.448	0.685	0.32	0.491	1.055	0.255	1.20	0.236
Decision to recall									
20	Intercept	-1.171	0.075	220.65	0.000	-1.796	0.419	129.96	0.000
	Period	1.658	0.067	942.76	0.000	1.884	0.117	299.55	0.000
	PSLC	1.212	0.046	258.62	0.000	1.767	0.146	18.80	0.000
	AROCA	-0.881	0.187	3.61	0.059	-1.406	0.276	3.97	0.042
60	INTER	-1.747	0.247	52.55	0.000	-1.431	0.119	439.38	0.000
	Period	1.102	0.017	135.34	0.000	1.682	0.101	513.13	0.000
	PSLC	1.010	0.027	163.78	0.000	1.037	0.048	1103.84	0.000
	AROCA	0.588	0.058	0.48	0.388	-0.933	0.277	1.96	0.083

3.5.2. Decisions to Recall

The decisions in periods in which a candidate was not encountered to either recall a previously encountered candidate or to continue the search were subjected to a similar Probit analysis with period, AROCA, and PSLC as the independent variables. Table 5 (lower panel) presents the maximum-likelihood estimates of the regression coefficients where the parameters correspond to the Recall decision. In all conditions, a Pearson chi-square overall goodness of fit test cannot be rejected.

The optimal policy dictates that recall decisions will occur not more than once at the period just after the threshold, if the applicant in this period is not a candidate. The characteristics of the sequence up to this period should be entirely irrelevant. The observed behavioral patterns clearly violate this prediction. Although in all conditions the rate of recall increases with the period (see the positive significant coefficients for Period in all conditions), these rates are also positively affected by the PSLC. In the shorter lists (n=20) the rate of recall decisions decreases as AROCA increases.

3.6. Summary of Behavioral Findings

Our results indicate that search decisions are influenced by two components that we term looking backward with *regret* and looking forward with *anticipation*. Anticipation that the next candidate will be the overall best builds up as a function of the lag since the last candidate was encountered. It decreases as a function of the average rate of candidate arrival. The longer the lag since the last candidate and the lower the average rate of candidate arrival, the higher the probability that the next candidate will be accepted. Regret that the last encountered candidate was the overall best affects the decision to recall and it increases with the lag since the last candidate was encountered. At a point when a candidate is observed any regret disappears.

3.7. Implications of the Behavioral Regularities

Can a decision rule (policy) that focuses on counting candidates and non-candidates explain our basic finding that most subjects did not search enough in the no-cost condition and searched too much in the cost condition? We show below that the answer is positive. Consider the following decision heuristic (behavioral model) termed **C**andidate/**N**o-**C**andidate **C**ounting **P**olicy (CNCCP): a threshold value j is established, and the following counter is set:

$$\text{counter}_t = t + \text{PSLC}_t - \text{AROCA}_t.$$

If $\text{counter}_t > j$ and the applicant in period t+1 is a candidate, then it is accepted; otherwise, the last candidate is recalled. If the last candidate is available, the search process ends; otherwise, the search continues to the next candidate that is immediately accepted. If no candidate is encountered after an unsuccessful recall, then the last applicant is accepted. If $\text{counter}_t \leq j$ for all $t<T$ (the number of applicants), then in period T, the last candidate is recalled. If it is available, the process terminates; otherwise, the process ends by accepting the last applicant. This behavioral decision model is motivated by our findings regarding the effects of Period, PSLC, and AROCA on the decision to accept, reject, or recall. Note that the effect of AROCA ($0< \text{AROCA} \leq 1$) is much smaller than Period and PSLC, reflecting our experimental findings.

Given the specific parameter values of the game (number of applicants, reward and search cost), we can determine the optimal j that maximizes the expected value of the search, given the behavioral model proposed above. Given j, we can compute the **E**xpected **P**eriod of **F**irst **A**ction (EPFA). Note that the optimal policy only counts the number of periods and hence the PFA is fixed. In contrast, CNCCP counts the number of candidates and non-candidates. Consequently, PFA depends on the actual realization of the random sequence. Figure 5 displays the PFA (solid line) and EPFA (dashed line) values for either the optimal policy or the CNCCP as a function of the number of applicants (upper and lower figures) and the search cost (on the horizontal axis). A bold vertical line above each cost level represents the range of EPFA corresponding to an 80% efficiency level with respect to the optimal CNCCP.

The following pattern emerges: for both the no-cost and low-cost environments, the CNCCP model with optimal value of j prescribes a PFA that is, on average, earlier than the cutoff period prescribed by the optimal policy. On the other hand, in the high-cost environment, the CNCCP model with the optimal value of j prescribes a PFA that is, on average, later than the cutoff period prescribed by the optimal policy[22]. Further, even if the subjects follow the CNCCP but do not adhere to the optimal threshold, the pattern of over- and under-search as a function of cost is expected to emerge if the deviation from the optimal threshold value is not too extreme. This is shown by the ranges that represent the 80% efficiency levels. Almost the entire range lies below the PFA for the no- and low-cost levels and above PFA for high-cost levels.

[22] For both strategies, the first action is expected to occur sooner as the cost increases up to such a cost that no matter how much higher the cost is, the best option is to accept the first applicant. Hence, the discrepancy between the two strategies disappears at this cost level.

Figure 5

(E)PFA FOR THE OPTIMAL POLICY AND CNCCP AS A FUNCTION OF COST

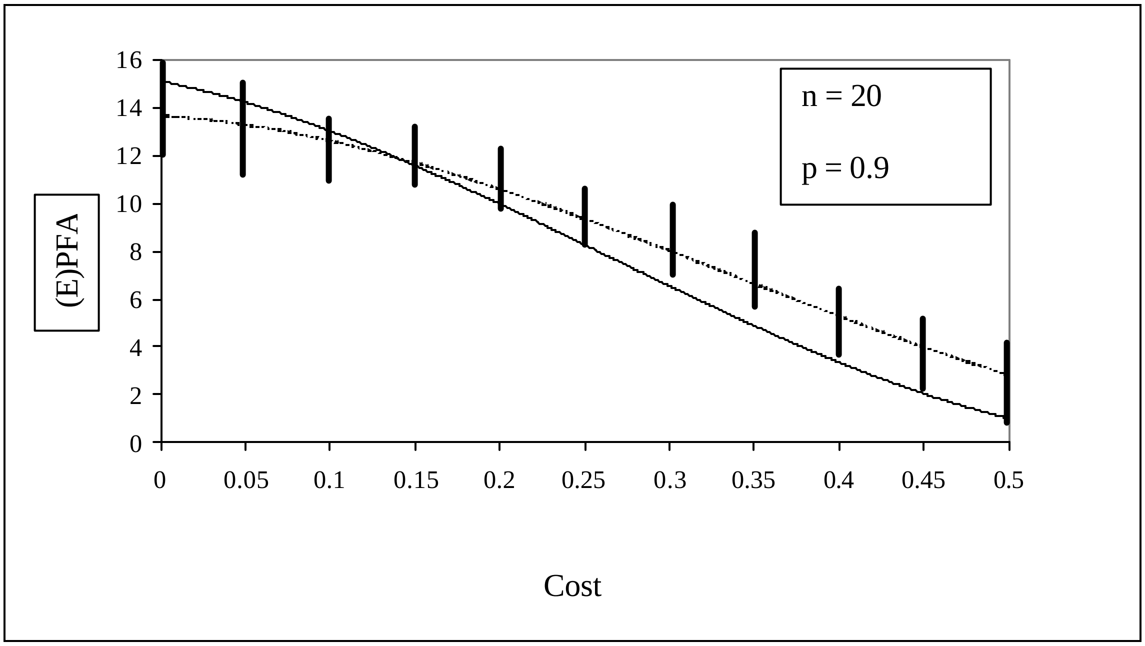

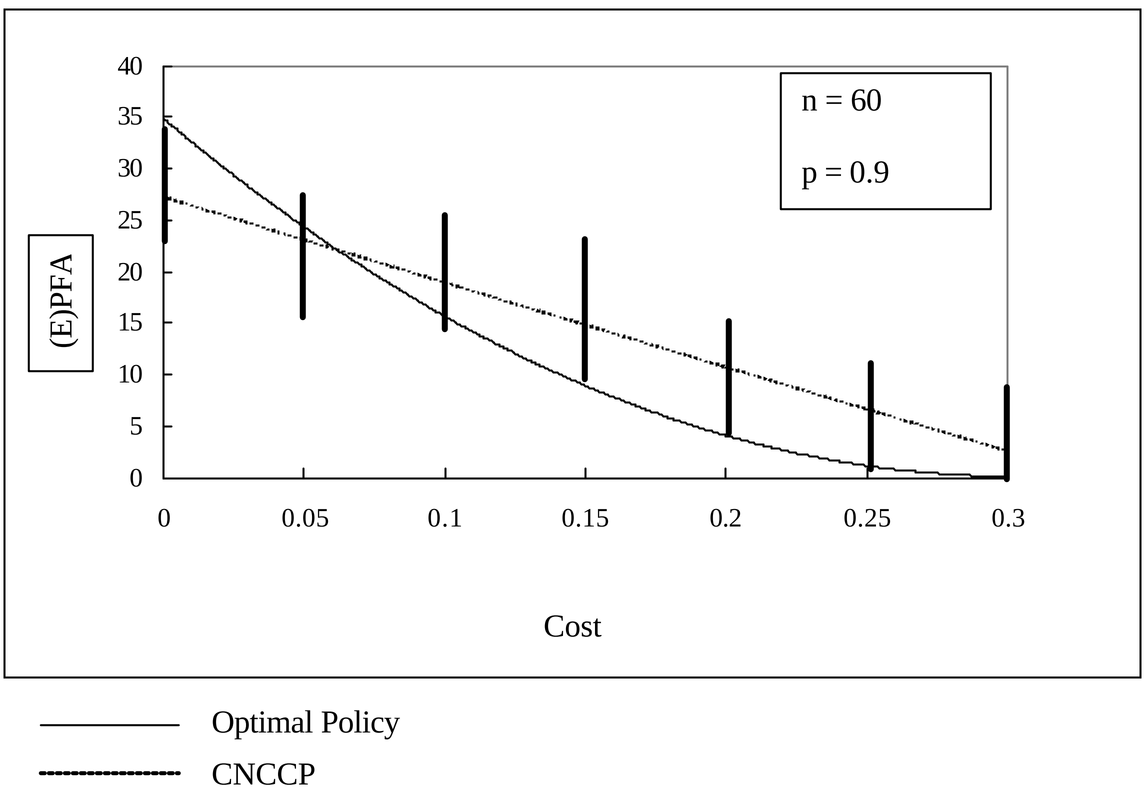

4. DISCUSSION

A common claim is made that consumers engage in insufficient external searching prior to purchasing, even for major purchases involving furniture, appliances, and automobiles (Claxton, Fry and Portis 1974; Furse, Punj and Stewart 1984; Newman 1977; Dickson and Sawyer 1990). These studies have found that a significant number of consumers make major decisions after shopping at a single retailer and/or considering only one brand. The general claim that seems to be drawn from these studies, namely, that consumers do not search enough, is questionable for three reasons. First, the finding that consumers engage in insufficient external searching does not necessarily mean that the amount of overall searching is low because this finding fails to consider the possibility of internal searching. Consumers do not approach the purchase of major items completely uninformed. Typically, and to varying degrees, they gather relevant information and deliberate on it—often at some length--prior to starting the external search. Second, it is not always obvious how to measure the amount of external searching. Thirdly, given a correctly measured amount of searching, the "not enough" statement can only derive meaning with respect to the known optimal length of searching that, in turn, must be based on modeling the search environment from the consumer's perspective.

The experimental methodology used in the present study alleviates many of the problems concerning the measurement of length of searching, assessment of the search cost, and specification of the consumer's objective. Moreover, this methodology allows the study of sequential searches for objects with multiple dimensions without the necessity of specifying the trade-off the consumer has to make among the various dimensions. As mentioned briefly earlier, previous experiments in the marketing literature commonly examined consumers' sequential searching behavior in a single-attribute space, most often focussing on price or quality. The optimal policy in this case is based on the assumption that consumers know the distribution of the single attribute or that they learn the shape of the distribution by sequential sampling. We contend that both assumptions are unrealistic, given the typical level of knowledge of the average consumer (Dickson and Sawyer 1990).

We find that, on average, our subjects did not search enough in the two no-cost conditions in agreement with previous findings of insufficient search. We also find that, on average, our subjects searched too much in the two cost conditions. The statement that consumers do not search enough is challenged not by arguing that consumers search just the right amount but by demonstrating that under certain circumstances too much searching is taking place. We also find that, in contrast to the optimal policy, local features of the observed sequences--features that rational consumers should disregard--influence the length of the search.

Five issues warrant further discussion. First, our results are limited to the search environment implemented in our study: a generalized secretary problem with fixed cost per inspection, recall with a geometric decay function, and an objective of selecting nothing but the best. The cost and recall factors are easily justifiable as they bring the classical secretary problem much closer to real-world consumers' search environments. Although the objective criterion to be maximized is more difficult to justify, we have opted to study it for practical reasons. The optimal policies for alternative objectives under the same search environment, such as a "satisfysing" goal in which the reward is awarded if the selected item is one of the best r items ($1 \leq r < n$) (Yeo 1998) or in which the award is proportional to the absolute rank of the selected item, call for a multiple-threshold search policy that renders the experimental investigation more complicated. Manipulating the DM's objective is an important future research opportunity to pursue.

Second, our statistical analysis is mostly conducted at the group level; only a few descriptive statistics are presented at the individual level (see Figure 3 and Table 4). Note that our major findings of under- and over-search hold for most subjects (see Figure 3). Furthermore, an individual level

Probit analysis on the decision to accept and recall yielded virtually the same results as reported on the group level for almost all subjects. These are very robust findings that seem to indicate that almost all subjects are sensitive to PSLC and AROCA in such a way that explains their tendency to under- or over-search as a function of cost.

Thirdly, Seale and Rapoport (1997, 2000) conducted an experimental investigation of the secretary problem without recall and without an exogenously imposed cost. They accounted for their major finding of insufficient searching by a cutoff model that postulates an *endogenous* cost of searching. Assuming that the endogenous search cost (i.e., the subject's personal cost of time) is independent of the experimentally imposed exogenous cost, such an additional cost should have contributed to under- search in *both* the no-cost and cost conditions. Our results indicate that the main forces behind the length of search are the local features of the sequences observed by the subjects and their sensitivity to various patterns such as PSLC and AROCA. In Seale and Rapoport (1997, 2000) both endogenous costs and the above factors (that have been identified as playing a major role in their studies) decreased the length of searching in the same fashion. However, when endogenous cost and local characteristics of the sequence affect the amount of search differently, the local characteristics seem to dominate, resulting in too much searching for the two cost conditions.

Fourth, one may argue that the model we propose for consumer sequential searching is not realistic because it assumes that after a selection is made, thereby terminating the search, the consumer is informed whether or not the selected item is, indeed, the overall best. This is the case even though not all the alternatives have been inspected. Clearly, this information is often not available in realistic search environments. Nevertheless, in many cases consumers continue to search passively (with no cost) even after making a purchase in order to evaluate the quality of the actual selection (or to minimize post-decisional regret). Our results speak to these cases.

Finally, the proposed behavioral model (CNCCP) does not predict too much searching when the search is costly. Rather, it predicts that at some point that is a function of the number of objects, n, and the reward-to-cost ratio, the under-search will change to over-search. In our study, the ratio of reward to cost was sufficiently high to reverse the common finding of under-search. It may be interesting to experimentally investigate the implications of CNCCP by systematically manipulating the reward-to-cost ratio.

The immediate contribution to the academic marketing literature is obvious. For the first time an argument is made that the common observation that consumers either do not search enough or search just the right amount should be extended to accommodate circumstances when too much searching is expected. Further, we propose and test a behavioral decision model that can explain simultaneously all the three major findings of the present study as a function of the cost in the search environment examined in this study.

The managerial implications are less obvious because some managers would like to encourage consumers to search whereas others would be happy with very limited search activities. However, the policy implications are again clear. Previously, the common theme among consumer advocates was how to teach consumers to search smartly with the understanding that smart searching would require, under most circumstances, more searching than commonly takes place. Our results question this policy position. We suggest that smart searching sometimes means less searching than our intuitions lead us to perform.

REFERENCES

Avery, Rosemary J. (1996). Determinants of search for nondurable goods: An empirical assessment of the economics of information theory. *Journal of Consumer Affairs*, 30, 390-420.

Braunstein, Yale, Andrew Schotter (1982). Labor market search: An experimental study. *Economic Inquiry*, 20, 133-144.

Camerer, Colin (1995). Individual Decision Making. in John H. Kagel and Alvin E. Roth, *Handbook of Experimental Economics*, Princeton: Princeton University Press, pp. 587-703.

Claxton, John D., Joseph N. Fry, Bernard Portis (1974). A taxonomy of prepurchase information gathering patterns. *Journal of Consumer Research*, 1, 35-42.

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

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

Dueker, Michael (1999). Conditional Heteroscedasticity in Qualitative Response Models of Time Series: A Gibbs-Sampling Approach to the Bank Prime Rate. *Journal of Business and Economic Statistics*, 17, 466-72.

Dickson, Peter R., Alan G. Sawyer (1990). The price knowledge and search of supermarket shoppers. *Journal of Marketing*, 54, 42-53.

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

Furse, David H., Girish N. Punj, David W. Stewart (1984). A typology of individual search strategies among purchasers of new automobiles. *Journal of Consumer Research*, 10, 417-431.

Goldman, Arieh, J. K Johansson (1978). Determinants of search for lower prices: An empirical assessment of the economics of information theory. *Journal of Consumer Research*, 5, 176-186.

Green, Paul, K. Helsen, B. Shandler (1988). Conjoint internal validity under alternative profile presentations. *Journal of Consumer Research*, 15, 392-397.

Grewal, Dhruv, Howard Marmorstein (1994). Market price variation, perceived price variation, and consumers' price search decisions for durable goods. *Journal of Consumer Research*, 21, 453-460.

Hey, John D. (1987). Still Searching, *Journal of Economic Behavior and Organization*, 8, 137-144.

Johnson, Eric J., Robert J. Meyer, Sanjoy Ghose (1989). When Choice Models Fail: Compensatory Representations in Negatively Correlated Environments. *Journal of Marketing Research*, 26, 255-270.

Kogut, Carl A. (1990). Consumer search behavior and sunk cost. *Journal of Economic Behavior and Organization*, 14, 381-392

Kogut, Carl A. (1992). Recall in Consumer Search. *Journal of Economic Behavior and Organization*, 17(1) 141-151.

Kraft, D. H., T. Lee (1979). Stopping rules and their effect on expected search length. *Information Processing & Management*, 15, 47-58.

Lippman, S. A., J. J. McCall (1976). The economics of job search: A survey (I and II). *Economic Inquiry*, 14, 155-189, 347-368.

Moon P., A. Martin (1990). Better heuristics for economic search: Experimental and simulation evidence. *Journal of Behavioral Decision Making*, 3, 175-193

Moorthy, Sridhar, Brain T. Ratchford, Debabrata Talukdar (1997). Consumer information search revisited: Theory and empirical analysis. *Journal of Consumer Research*, 23, 263-277.

Newman, Joseph W. (1977). Consumer external search" amount and determinations. in Arch G. Woodside, Jagdish N. Sheth and Peter D. Bennett, eds., *Consumer and Industrial Buyer Behavior*. New York: North-Holland, 79-94.

Rapoport, Amnon, Amos Tversky (1970). Choice behavior in an optimal stopping task. *Organizational Behavior and Human Performance*, 5, 105-120.

Ratchford, Brain T., Narasimhan Srinivasan (1993). An empirical investigation of return to search. *Marketing Science*, 12, 73-87.

Saad, Gad, Jay Edward Russo (1996). Stopping criteria in sequential choice. *Organizational Behavior and Human Decision Processes*, 67, 258-270.

Schotter, Andrew, Yale Braunstein (1981). Economic Search: An experimental study. *Economic Inquiry*, 19, 1-25.

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

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

Shapira, Zur, Itzhak Venezia (1981). Optional stopping on nonstationary series. *Organizational Behavior and Human Performance*, 27, 32-49.

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

Urbany, Joel E. (1986). An experimental examination of the economics of information. *Journal of Consumer Research*, 3(2), 257-271.

Urbany, Joel E., Peter R. Dickson, Rosemary Kalapurakal (1996). Price search in the retail grocery market. *Journal of Marketing*, 60, 91-104.

Yang, Mark C. K. (1974). Recognizing the maximum of a random sequence based on relative rank with backward solicitation. *Journal of Applied Probability*, 11, 504-512.

Yeo, Geoffrey F. (1998). Interview costs in the secretary problem. *Australian & New Zealand Journal of Statistics*, 40, 215-219.

Zwick, Rami, Ching Chyi Lee (1999). Bargaining and search: An experimental study. *Group Decision and Negotiation*, 8, 463-487.

Zwick, Rami, Eythan Weg, Amnon Rapoport (2000). Invariance failure under subgame perfectness in sequential bargaining. *Journal of Economic Psychology*, 21, 517-544.