

Consumer Sequential Search: Not Enough or Too Much?

Rami Zwick • Amnon Rapoport • Alison King Chung Lo • A. V. Muthukrishnan

Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong

University of Arizona, Management and Policy, P.O. Box 210108, Tucson, Arizona 85721-0108

Duke University, Fuqua School of Business, Durham, North Carolina 27708-0120

Hong Kong University of Science and Technology, Clear Water Bay, Kowloon Hong Kong

mkzwick@ust.hk • amnon@u.arizona.edu • kcl5@duke.edu • mkmuthu@ust.hk

We study sequential search behavior in a generalized “secretary problem” in which a single object is to be selected from a set of n alternatives. Alternatives are inspected in a random order, one at a time, and only the rank order of the current alternative relative to the ones that have already been observed can be ascertained. At each period, the consumer may either accept the current alternative, continue to search and pay a fixed cost, or recall an alternative that has already been inspected. A recalled alternative is assumed to be available with a known probability. The consumer’s goal is to select the overall best alternative from the fixed set.

We describe the results of an experiment designed to test the optimal model and compare it to a behavioral decision model that incorporates local patterns of the observed sequence. Both set size and search cost are manipulated experimentally in a 2×2 factorial design. Our results show that cost and set size affect the amount of search in the predicted direction. However, in the two no-cost conditions subjects search too little in comparison to the optimal model, whereas in the two cost conditions they search too much. The behavioral decision rule that we propose provides a possible account for the observed pattern of the behavioral regularities.

(Search Behavior; Sequential Choice Models; Behavioral Decision Rules; Information Processing)

1. Introduction

We attempt to answer three questions about consumer decisions in a class of sequential search problems: Do consumers search optimally? If not, under what conditions do they search too much and under what conditions do they not search enough? What behavioral decision rules, if any, govern their sequential search for the best alternative? Whereas the first question has received much attention in the economics and marketing literature, the second, and especially the third, questions have received scant attention. To answer these questions, we describe the optimal decision rule for maximizing expected gain of search. We

then propose an alternative behavioral model that considers feelings of anticipation and regret based on observed characteristics of the sequence of the inspected alternatives, and then derive predictions that may account for deviations from optimality. To answer our three questions, we proceed as follows. First, we briefly review theoretical and experimental findings on consumer search that address the first question. Second, we describe the optimal model and outline a behavioral decision model. Finally, we present an experiment designed to competitively test these two models, and discuss the marketing implications of our findings.

1.1. Is Consumer Search Optimal?

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 a lower price (e.g., Lippman and McCall 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 on the next period. If she later decides to accept a price that she has passed over, 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, no cost of search is assumed, and prices that are passed over can be recalled with probability one; the optimal policy is to set a critical, time-invariant price (threshold) and accept the first price below this threshold. Versions of this model that allow for recall or finite search horizon have been studied in the marketing and job search literatures.

The two general methodologies for testing the implications of these models are field studies and controlled laboratory experiments. The former methodology, most common in marketing research, has the 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 et al. 1996, Moorthy et al. 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 second difficulty with field studies is that a decision policy can only be meaningfully evaluated with respect to a model that fully characterizes the search environment and makes precise assumptions about the consumers' beliefs, goals, and preferences. What seems to be suboptimal search behavior may simply be due to misspecification of any of these elements. For example, search cost plays a crucial role in the theoretical models. How should search costs be measured or inferred in a field study?

Laboratory experiments have been designed to study sequential search behavior (for a brief review, see Camerer 1995). Although controversial on ecological validity, they are ideally suited to overcome the two difficulties mentioned above. A consistent finding that seems to be repeated in the literature is that people search too little compared to 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, 1992) 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 undersearch. However, our study shows that under different conditions the application of heuristic rules can result in both under- and oversearch.

There is also a smaller stream of recent experimental research reporting that, under certain conditions, people search more than prescribed by the optimal models. 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 et al. (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.

Statements about "searching too much" or "searching too little" are obviously only meaningful with respect to an optimal policy. We contend that optimal policies serving as benchmarks in the previous investigations are based on two assumptions that are much too strict to be met in practice. First, the classical model and many of its variants assume that consumers know the underlying distribution of prices with precision. However, the assumption that the entire distribution is known with precision—which is critical for the calculation of the reservation price—is difficult to justify. Second, what if the search is for an item that cannot be simply characterized along a single dimension, such as most big-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 the present study, we propose an alternative model for sequential search that is based on what we believe to be more plausible assumptions about the search environment. Unlike the classical models, our model does not assume that consumers possess precise knowledge of the distribution of prices (or other attributes) or can learn them in the process of searching. Moreover, and also unlike the classical search models, we do not assume that the search for an item is restricted to a single dimension (e.g., price value). Rather, the search environment that we examine incorporates only information about the *overall rank ordering* of the items being observed.

Sequential observation and selection problems of this type, referred to collectively as the “secretary problem,” have attracted considerable attention from applied mathematicians whose major focus has been on characterizing the optimal decision policies for a rich class of problems (Freeman 1983). In contrast, with only a few exceptions (Corbin et al. 1975; Seale and Rapoport 1997, 2000), there have been no attempts to study the decision rules (heuristics) that subjects use to select the best alternative in this class of sequential decision problems. Before proposing behavioral decision rules, we first state a set of assumptions that capture many of the features of consumer search for multiattribute products like cars, apartments, gifts, children’s toys, etc., and then describe what would be an optimal search strategy in such contexts.

- There is only a single object to be selected.
- The number of alternatives to be observed and inspected sequentially (which constitute the so-called “consideration set”), denoted by n , is known.¹
- The Decision Maker (DM) inspects the n alternatives in a random order, one at a time. Therefore, before the search commences each of the $n!$ orderings is assumed to be equally likely.

¹ We assume here that the consideration set is fixed, and independent of the search process. For recent studies that deal with consideration set formation, see Mehta et al. (2003) and Wu and Rangaswamy (2003).

- There is a fixed per-alternative cost for searching.
- The DM can determine only the rank (in terms of preference, attractiveness, or quality) of the current alternative relative to the ones she has already observed (with no ties).

• At each period, the DM can either accept the current alternative, continue to search for the next alternative, or recall² any of the alternatives she has already observed. If she recalls an alternative, it is assumed to be available with a certain probability that is known by the DM. If the alternative is unavailable, the DM can either recall another alternative or continue the search. An alternative that is not available remains unavailable forever.

- The DM’s objective is to select the overall best alternative from the consideration set.

Several of the assumptions that characterize this search environment and are admittedly too restrictive can and have been relaxed. The assumption that the number of alternatives, n , is finite and known can be replaced by the assumption that n is a random variable whose probability distribution is known. The assumption that the per-alternative cost is fixed can be replaced by other assumptions about the cost structure (e.g., cost of search increases exponentially in the number of alternatives inspected). The assumption about the DM’s objective of only deriving utility from choosing the overall best alternative can also be replaced by the assumption that the DM’s utility is inversely related to the true rank of the alternative she chooses or that his objective is to select one of the h best alternatives, $h < n$.

1.2. The Search for a New Apartment Scenario

To illustrate the search environment, consider the following scenario, a variant of which is used later in our experiment. Imagine that you are searching for a new apartment in a particular neighborhood. Using one of the online real estate websites, you select a set of n apartments that includes only those few alternatives best suited to your taste. This initial screening is

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.

virtually costless and can be conducted almost instantaneously using electronic agents that use information about your specific preferences, the alternatives available, and their known values along the important (to you) "search attributes." The alternatives in the consideration set could be searched in more detail to choose the "best" of this reduced set. Although on paper these n apartments seem equally attractive to you, you know that only a physical visit will reveal the hard-to-describe-and-quantify dimensions of "emotional fit," "special feel," "feng shui," etc. (i.e., "experience attributes"). Experiencing these qualities of one apartment tells you nothing about the qualities of the apartments you have not visited yet along the same dimensions.

Were you to visit and inspect all the apartments in the set, you would be able to choose the best based on your taste. However, visiting all the apartments might not be the best option. First, the best apartment may no longer be available after you have visited all the apartments (unless it is the last one you visit). Second, visiting and inspecting all the apartments in the set is costly and time consuming. Therefore, you decide on the following plan of action: Because you have no a priori preferences, you arrange the apartments in the order you plan to visit based on external criteria that have nothing to do with the possible attractiveness of the apartments. Next, you visit the first apartment (bearing the cost of one visit). If you believe that this apartment is the overall best, you rent it; otherwise, you inspect the second apartment (bearing the cost of two visits). At this point, you can compare the two apartments to each other. If you believe that the second apartment is the overall best, you rent it. If not, you either decide to inspect the third apartment, or call back the owner of the first apartment to find out if it is still available. If the first apartment is available, you rent it. If not, you visit the third apartment (bearing the cost of three visits). This search process continues until you choose an apartment that you believe will satisfy you.

Assume that the set contains $n = 20$ apartments and that you will be satisfied only if you rent the best apartment in the set ("nothing but the best"). Further, assume that the probability that an apartment that was inspected but not selected is still available in

a later period declines geometrically with the number of periods since this apartment was inspected (p^r , where r is the lag time).³ How many apartments would you visit before either renting the presently inspected apartment or recalling 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?

1.3. The Optimal Search Policy

The following search policy maximizes the probability of selecting the overall best alternative (and, therefore, the expected gain of search), given the above scenario with no search cost: Visit the first $s_0 - 1$ apartments. In period s_0 , choose the relative best apartment no matter in which period it has been observed. This, of course, might require recall. If the relative best (out of the first s_0 apartments) is not available, continue to search and then choose the next relatively best apartment in the sequence (Yang 1974).

The value of s_0 is the largest s that satisfies the inequality

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

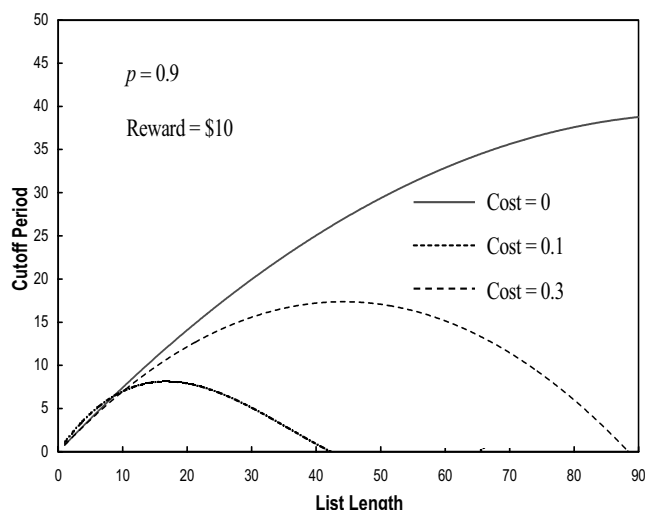
where

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

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.

Figure 1 exhibits the cutoff periods for various sets, sizes ranging from 1 to 90 for three different per-period cost values (\$0, \$0.1, and \$0.3). It assumes a reward of \$10 for selecting the best overall alternative and a probability of $p = 0.9$ for determining the availability of previously inspected alternatives. To generate this figure, each sequence was replicated

³The 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 .

Figure 1 Cutoff Period (S_0) as a Function of Set Size and Search Cost*

Note. *Based on simulation results with 10,000 sequences for each set's size.

10,000 times, and the expected gain was computed for all possible cutoff values (1 to the end of the list). The figure shows the cutoff values with the highest mean expected payoff for each set size (i.e., the optimal cutoff value).

1.4. Behavioral Decision Rules

To set our discussion in the context of previous theoretical research, we refer to each member of the consideration set as an *applicant*. In each period, the best applicant among the ones that have already been inspected (including the current one) is called a *candidate*. Being a candidate is a necessary but not sufficient condition for an applicant to be the overall best.

The optimal policy dictates that a noncandidate is never accepted or recalled and that the decision to either accept or reject a candidate or recall the last candidate that has been encountered depends only on the period of the observation but not on the history of the process (e.g., local patterns of the sequence). All candidates before the cutoff period are initially rejected. If the applicant at the cutoff period is a candidate, it is accepted; otherwise, the last encountered candidate is recalled. If it is not available, the first candidate after the cutoff period is accepted.

Experimental results reported by Seale and Rapoport (1997, 2000), who studied a similar sequen-

tial search problem with zero search cost and no recall, indicate that the second implication of the optimal model is violated; the perceived patterns in the sequence, that otherwise is known to be random, do affect the DM's decisions.

Based jointly on behavioral and self-report data,⁴ Seale and Rapoport (1997, 2000) found that three factors dominated subjects' decisions to accept a candidate. First, the period in which the candidate appears; second, the number of previously encountered candidates; and third, the number of periods since the last candidate was observed. Whereas only the first factor is relevant for the optimal decision, almost all the subjects reported using decision rules combining in one way or another all three factors. No other decision rules were reported by any significant number of subjects.

In the present study we first formulate these factors precisely and study their qualitative effects on the length of search in our task. In the result section we propose a simple combination rule of the above three factors that can explain the observed behavioral regularities that we document.

We propose that both regret and anticipation influence subjects' actions. Anticipation that the next candidate will be the overall best builds up as a function of the lag (number of periods) since the last candidate was encountered. However, anticipation is reduced if the consumer encountered many candidates in the previous rounds. That is, on observing many candidates in the earlier rounds the consumer assumes that the likelihood of encountering such candidates in subsequent rounds would also be the same. Hence, he could wait for an "even better" candidate. Thus, the longer the time lag since the last candidate was encountered and the lower the number of candidates encountered thus far, the higher would be the probability of the consumer accepting the current applicant if it is a candidate.

⁴ After completing 100 trials, subjects in Seale and Rapoport (1997, 2000) were asked to complete a questionnaire in which they had to describe in detail their strategy to stop on a particular applicant and rate the importance of five factors in their stopping decisions. To motivate thoughtful answers, the subjects were informed that they would play one final trial, worth US\$3 for a correct selection, following the decision rule they documented on their questionnaire.

We also assume that consumers are governed by feelings of regret that an earlier encountered candidate that might have been the overall best was passed over. As it builds up, regret increases the probability that a previously rejected option will be recalled. We propose that the degree of regret is a function of the time lag since the last candidate was encountered.

Specifically, given a specific period, t , let $AROCA_t$ be the *Average Rate of Candidate Arrival* computed as the number of previously encountered candidates divided by $t - 1$, and let $PSLC_t$ be the number of *Periods Since the Last Candidate* was encountered.⁵ $AROCA$ can be interpreted as a global measure of arrival rate computed over the entire sequence up to the current period, and $PSLC$ can be interpreted as a local measure of arrival rate that is relevant in the vicinity of the current period. We formulate below several qualitative predictions with regard to the effects of $AROCA$ and $PSLC$ on the decisions to accept the current applicant if it is a candidate or recall the last candidate if the current applicant is not a candidate. These predictions are in line with the well-known findings that people tend to perceive local patterns in an otherwise global random sequence such as the law of small numbers and the misperception of randomness (e.g., the “gambler’s fallacy,” “magical thinking,” illusion of control, and the belief in “hot hand”).

Decision to Accept. Given a specific period, t , in which a candidate was encountered, if the sequence up to t had relatively many candidates (high $AROCA$), then the probability of immediately accepting that candidate is lower compared to experiencing a more sparse sequence before period t (low $AROCA$). This behavior follows the (erroneous) belief that the past rate of arrivals will continue in the future; hence, given a past sequence “rich” with candidates, subjects believe that many more candidates are likely to appear in the remaining sequence. 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. Similarly, a recent shortage of

candidates (high $PSLC$) increases the likelihood that the current candidate will be accepted.

Decision to Recall. Given a specific period, t , in which a noncandidate has been encountered, if the sequence up to t had relatively many candidates (high $AROCA$), then the probability of recalling the last candidate is lower compared to experiencing a more sparse sequence before period t (low $AROCA$). Conversely, a sequence in which candidates appear infrequently increases the chances that the last candidate will be recalled. Similarly, a recent shortage of candidates (high $PSLC$) increases the likelihood that the last candidate will be recalled.

Note that $AROCA$ is a much more difficult index to process cognitively than is $PSLC$. Whereas for $PSLC$ the consumer only has to recall how long it has been since the last candidate was observed, for $AROCA$ the consumer needs to take into account the status of an applicant (candidate or noncandidate) *at the time that it was encountered*, even if from the perspective of the current period only one (or none) of the previously encountered candidates can still be classified as a candidate. Consequently, although we expect both factors to affect decisions, we conjecture that the cognitively simpler index of $PSLC$ would dominate.

We described above some qualitative predictions as to the effects of local observed patterns in the sequence, patterns that the rational consumer should ignore. A quantitative decision rule based on these cognitive processes is formulated and tested in §3.6.

2. Method

2.1. Subjects

Ninety-seven subjects, all undergraduate business students at 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 through 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 and was introduced as a search for a rental

⁵ $PSLC$ and $AROCA$ are not defined for $t = 1$; for the present analysis we set $PSLC = 0$ and $AROCA = 0$ for $t = 1$.

apartment. The subjects read online instructions at their own pace and were required to answer several questions correctly to verify their understanding of the task before continuing with the experimental task. Similarly to the scenario described above, the subjects were instructed to imagine that their real estate agent had constructed a finite list of potential apartments and that their task was to select (“rent”) the best apartment from the list by sequentially inspecting the apartments. The instructions emphasized that:

- on paper all the apartments are equally attractive;
- the order in which apartments are visited implies nothing about the desirability of the apartment;
- only the rank of the apartment, relative to all other apartments that have already been visited, is revealed by visiting and inspecting that apartment;
- visiting an apartment implies nothing about other apartments that have not been visited;
- the probability that an apartment that has been visited in the past is still available decreases with the time since the visit.

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

Figure 2 presents an example of the decision screen. The parameters of the game are presented in the right section of the screen. The left part of the screen presents the history of the search up to the current period; it displays the relative ranks of the apartments that have already been visited (Column 2) and the probability that a previous apartment is available if an attempt is made to recall it (Column 3). The current period (footnote a, 10 in this example) and the best apartment from the ones that were inspected (footnote a, Apartment 2 in this example) are shown in Column 1.

After visiting an apartment (that is not the first on the list), subjects could choose one of three actions:

- RENT the current apartment, thereby terminating the search,
- Visit the NEXT apartment, or

Figure 2 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 | | |

| Round = 1 | |
|------------------|----------------------|
| Number of flats: | 20 |
| Cost per visit: | \$0.30 |
| Goal: | Nothing but the best |
| Reward: | \$10 |

Accumulated
Search Cost = \$3.00



*This is an example for the 20 apartments 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 apartment overall is selected. In the above example, the subject is currently visiting the 10th apartment. Its relative rank is 8, and the best apartment so far was visited in period 2. The probability that the apartment from period 2 is still available is 0.43.

• CALL a previously visited apartment. If this apartment is available, it is rented and the search is terminated; otherwise, the same three options are again available to the subject.

If an apartment is rented, or a recalled apartment is available, the search is terminated and the (absolute) ranks of all the apartments on the list are unveiled. If the selected apartment is the best, the subject is provided with the reward (\$10.00) minus accumulated search cost (if any). If the selected apartment is not the best, the subject is charged an amount equivalent to the 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 the rounds plus \$10 for completing the session.⁶ On average, the subjects earned \$159.43, which was equivalent to approximately three hours' wages paid on campus.

2.3. Experimental Design

A 2 (number of apartments 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 apartments, n , was either 20 or 60. The search cost, c , was either \$0 (zero) or positive (\$0.1 for $n = 20$ and \$0.3 for $n = 60$). The reward for selecting the best apartment was fixed at \$10 and the probability that an apartment was still available was set at 0.9^r , where r is the lag time from the current period ($r = 0, 1, 2, \dots$). To render the task 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 set size. Table 1 presents the parameter values, number of subjects in each condition, cutoff period of the optimal policy, probability of selecting the overall best apartment, and the associated expected gain if the optimal policy is followed.

⁶ All amounts are in Hong Kong dollars. The exchange rate at the time of the study was US\$1=\$HK7.80.

Table 1 Parameter Values in the Four Experimental Conditions and Optimal Strategies


| | $n = 20$ Apartments | | $n = 60$ Apartments | |
|------------------------|---------------------|-------|---------------------|-------|
| Search cost (\$) | 0.0 | 0.3 | 0.0 | 0.1 |
| Cutoff value (s_0) | 15 | 6 | 32 | 13 |
| Probability | 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 |

Notes. Cutoff value: the cutoff period based on the optimal policy; probability: the probability of selecting the best apartment 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 3 depicts the efficiency of searching in comparison with the optimal policy for various cutoff values in the neighborhood of the optimal value (± 6). As

Figure 3 Efficiency of Search for Various Cutoff Values in the Neighborhood of the Optimal Value



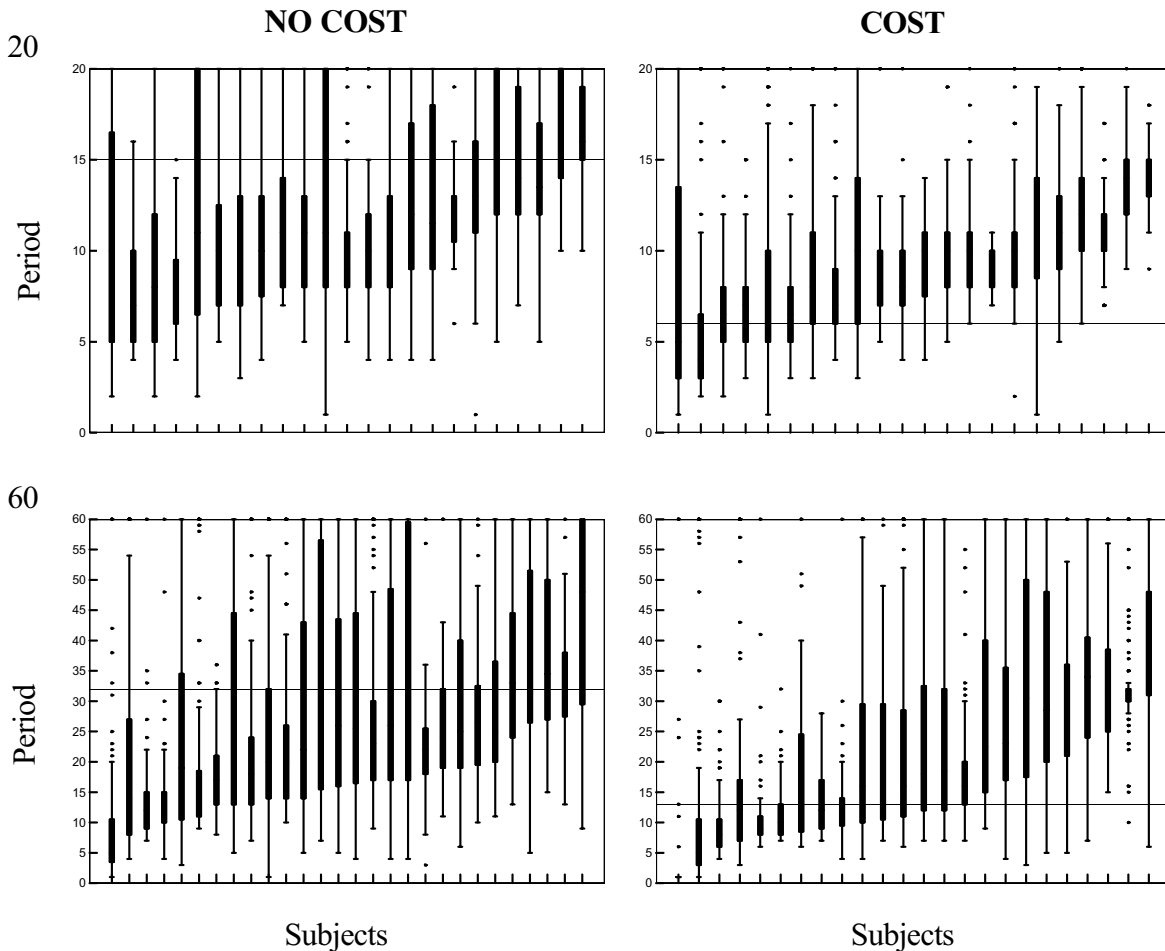
might be expected, the optimal policy is more sensitive to deviations from it when the number of items is smaller and the cost of search is higher. The efficiency loss for suboptimal behavior can be substantial. For example, setting a cutoff value of six periods too early in the shorter list ($n = 20$) results in efficiency losses 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 implications.

3.1. Period of First Attempt

Give the options to “rent,” “recall,” or “next,” the optimal policy identifies the period at which the first attempt at renting (either by renting the current

apartment or by recalling a previously inspected apartment) should take place (i.e., the cutoff period). Figure 4 displays the major findings with regard to the optimal strategy, namely, the frequency distributions of the *Period of First Attempt* (PFA). The distributions are shown by set size (row), cost/no-cost condition (column), and subject (box plot). The bottom and top edges of each box 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 that is 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.

Figure 4 Period of First Attempt



Inspection of Figure 4 reveals that for both lists ($n = 20$ or $n = 60$), most subjects made the first attempt 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.

Table 2 presents the mean and standard deviation of the PFA scores by the type of the first attempt (accept or recall) and the status of the applicant that has been inspected at the period of the first attempt (candidate or not). The boxed cells indicate the cutoff

period for the optimal policy. The shaded entries indicate errors in either accepting a noncandidate or recalling a candidate in a period when another candidate is encountered (clearly, the current candidate is the highest ranked).

3.2. The Effect of Experience, Search Cost, and Set Size on the PFA

To test for the effect of experience, we divided each subject's responses into two blocks of 50 trials each. Then, the PFA scores as well as subjects' payoffs were subjected to a $2 \times 2 \times 2$ (search cost by number of items by block) ANOVA (with repeated measures on the block variable). The main block effect, the two-way interactions of block with the other two variables, and the triple interaction, were not significant for both dependent variables. Further analysis of the payoff measure is reported in §3.4. Search cost (c) and consideration set size (n) affected the amount of search before the first attempt. Both main effects and the $c \times n$ two-way interaction effects were significant ($F = 3,026.63, 215.19, \text{ and } 27.69$, for number of items, search cost, and the two-way interaction, respectively, $p < 0.0001$). On average, the number of items that were inspected before the first attempt was higher for the no-cost than the cost condition (20.10 versus 16.00), and higher for the 60 than the 20 item lists (24.55 versus 10.77). The $c \times n$ interaction is due to the fact that the cost reduces the number of alternatives that were passed over by about 5 when $n = 60$, but only by 2 when $n = 20$.

The first attempt was taken earlier than predicted in the no-cost condition and later than predicted in the cost condition for both set sizes and when actions were not classified as errors (shaded cells in Table 2). For condition ($n = 20, c = 0$), the mean PFA scores are 11.26 and 11.77 for the first attempts to accept and recall, respectively, compared to the predicted cutoff period of 15. For condition ($n = 60, c = 0$), the mean PFAs are 23.15 and 25.37 for the first decisions to accept and recall, respectively, in comparison with the predicted cutoff period of 32. All of these differences are significant by a t -test at the 0.001 level (see t values in Table 2). For condition ($n = 20, c = 0.1$), the mean PFAs are 8.69 and 9.79 for the first decisions to accept and recall, respectively, whereas the predicted

Table 2 Mean and Standard Deviation of Period of First Attempt

| Applicant in the PFA | | $n=20$ | | $n=60$ | | | | |
|----------------------|-----------------|--------|--------|--------|--------|-------|----|-------|
| | | Accept | Recall | Accept | 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 |

Notes. %N: % of time the relevant category occurred in the Period of First Attempt. 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).

*Based on 3 observations.

**Based on 2 observations.


cutoff period is 6. For condition ($n = 60, c = 0.3$), the mean PFAs are 17.77 and 21.2 for the first attempts to accept and recall, whereas the predicted cutoff period is 13. All of these differences are also significant at the 0.001 level (see t values in Table 2).

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

In all four conditions, the majority of first attempts were to recall a candidate that was passed over. The percentages were 90.02, 84.01, 84.00, and 83.00 in conditions ($n = 20, c = 0$), ($n = 60, c = 0$), ($n = 20, c = 0.1$), and ($n = 60, c = 0.3$), respectively. The predicted percentages derived from the optimal policy are 93.3, 96.87, 83.33, and 92.31.⁷ Because accepting a candidate terminates the search with certainty, the lower observed rate of recall as a first attempt in the no-cost conditions coincides with the above findings that less searching than expected was taking place. For the cost conditions, the rate of recall as the first attempt was close to that expected when $n = 20$, and below expectation when $n = 60$. This could imply that although the first attempt 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 5 displays the observed and predicted cumulative probability distributions of the search being terminated at any particular period. The theoretical distributions are depicted with dashed lines and the observed distributions with solid lines. 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 alternatives inspected before the search is terminated. In the no-cost case, for both set sizes and any given period, the observed probability of search termination is higher than expected, emphasizing the fact that searching is less extensive than predicted. In the two cost conditions, except for the few periods before the predicted

Figure 5 Observed and Expected Cumulative Probability Distributions of Search Termination by Period



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.

3.4. Cost of Suboptimal Decisions

We compared the observed and 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 on average 15.67% and 21.28% of

⁷ The probability that the j th applicant is a candidate is $1/j$. Hence, the probability that the first attempt is recall is given by $1 - 1/s^*$, where s^* is the cutoff period.

the potential cumulative payoffs in comparison with the optimal policy for the list lengths $n = 20$ and $n = 60$, respectively. These losses are substantive. By searching too much in the two cost conditions, the loss was much higher at an average of 40.69% when $n = 20$ and more than⁸ 100% when $n = 60$. Clearly, suboptimal behavior, especially in the two cost conditions, is significant and deserves attention.

To recapitulate, our analyses show that most subjects did not follow the optimal policy and that the deviations from optimality were substantial, systematic, and costly. In the two no-cost conditions, subjects did not search enough, whereas in the two cost conditions they searched too much. How can one explain such behavioral patterns and, in particular, can a descriptive model account *simultaneously* for both the under- and oversearch as a function of search cost? We move next to discuss the results with regard to the behavioral decision heuristics proposed in the introduction section.

3.5. Tests of the Behavioral Decision Rules

3.5.1. The Decision to Accept. The decisions in periods in which a candidate was encountered to either accept this candidate or continue the search (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. This analysis was conducted at both the individual and aggregate levels. Table 3 (upper panel) presents the maximum likelihood estimates of the regression coefficients at the aggregate level where the parameters correspond to the decision to accept. In all conditions, a Pearson chi-square overall goodness-of-fit test cannot be rejected.⁹ Ninety-four percent of

subjects across conditions exhibited the same pattern of significant factors at the individual level.

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 entirely 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 increasing with the number of periods since the last candidate was encountered (see the positive significant coefficients for PSLC in all conditions) and decreasing with the average rate of candidate arrival in the smaller set sizes (see the negative significant coefficients for AROCA in $n = 20$). These findings support our predictions regarding the effects of PSLC and AROCA on the decisions to accept, including the fact that the recency effect of PSLC is more pronounced in a condition where the derivation of the global measure of arrival rate is harder to process (in the larger set sizes).

3.5.2. The Decision to Recall. The decisions in periods in which a candidate was not encountered to either recall the last encountered candidate or continue the search were subjected to a similar probit analysis with period, AROCA, and PSLC as the independent variables. This analysis was conducted at both the individual and aggregate levels. Table 3 (lower panel) presents the maximum-likelihood estimates of the regression coefficients where the parameters correspond to the decision to recall. In all conditions, a Pearson chi-square overall goodness-of-fit test cannot be rejected. Eighty-nine percent of the subjects across conditions exhibited

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 censoring 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.

⁸ It is not quite clear how to measure the loss in efficiency here, since the predicted value is positive (\$53.3), and the observed value is negative (-\$19.6).

⁹ 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.

Table 3 Probit Analysis

| | No Cost | | | | Cost | | | |
|---------------------------|---------|-------|----------|---------------|--------|-------|----------|---------------|
| | EST | STD | χ^2 | Pr > χ^2 | EST | STD | χ^2 | Pr > χ^2 |
| <i>Decision to Accept</i> | | | | | | | | |
| <i>n = 20</i> | | | | | | | | |
| 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 |
| <i>n = 60</i> | | | | | | | | |
| Intercept | -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 |
| <i>Decision to Recall</i> | | | | | | | | |
| <i>n = 20</i> | | | | | | | | |
| 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 |
| <i>n = 60</i> | | | | | | | | |
| Intercept | -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 | 1,103.84 | 0.000 |
| AROCA | 0.588 | 0.058 | 0.48 | 0.388 | -0.933 | 0.277 | 1.96 | 0.083 |

virtually the same pattern of significant factors at the individual level.

The optimal policy dictates that recall decisions will occur no more than once at the threshold period 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 but are in line with our proposed heuristic rules. Although in all conditions the rate of recall increases with the period (see the positive significant coefficients for the period in all conditions), these rates are also positively affected by PSLC (in both set sizes) and negatively affected by AROCA in the smaller set sizes ($n = 20$).

3.6. Implications of the Behavioral Regularities

Can a decision rule that focuses on counting candidates and noncandidates explain our basic finding that most subjects did not search enough in the

no-cost conditions and searched too much in the cost conditions? We answer this question affirmatively.

When asked by Seale and Rapoport (1997, 2000) to rate the importance of period, PSLC, and AROCA¹⁰ in their selection decisions, most subjects rated them about equal (about 5–6 on a 7-point scale). Following this finding, we propose a very simple decision rule, termed *Candidate/No-Candidate Counting Policy* (CNCCP), that assigns equal weights to all three factors: the DM establishes a threshold value j , 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

¹⁰ In Seale and Rapoport (1997, 2000), the question was simply the number of previous candidates.


and the next candidate is immediately accepted. If no candidate is encountered after an unsuccessful recall, then the last applicant is accepted. If counter, $t \leq j$ for all $t < n$ (the number of applicants), then in period n the last candidate is recalled. If it is available, the process terminates; otherwise, the process ends by accepting the last applicant. Because the AROCA scores ($0 < \text{AROCA} \leq 1$) are constrained, their effect is smaller than the effect of Period and PSLC, reflecting our experimental findings.

This strategy shares with the optimal strategy the assumption that the mental process of reaching a decision takes the form of counting. It differs from the optimal strategy in the events that are supposed to be counted. In both cases, the effect of set size, cost, reward, and recall probabilities on the amount of search is via the effect on the selected threshold, but not on events that are supposed to be counted.

Given the specific parameter values of the game (number of applicants, reward, and search cost), we determined (by a numerical algorithm) the optimal value of j that maximizes the expected value of the search, given the decision rule proposed above. Given j , we can then compute the *Expected Period of First Attempt* (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 noncandidates. Consequently, PFA depends on the *actual realization* of the random sequence. Figure 6 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 the 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,

Figure 6 (E)PFA for the Optimal Policy and CNCCP as a Function of Cost



later than the cutoff period prescribed by the optimal policy.¹¹ Further, even if the subjects follow the CNCCP but do not adhere to the optimal threshold, the pattern of over- and undersearch 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.

¹¹ For both strategies, the first attempt 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.

The relationship between the optimal strategy and CNCCP is based on simulations with parameters used in our experiment (n and p) and for various cost levels. We repeated these simulations with n ranging from 10 to 100 (by 10 period steps) and p ranging from 0.1 to 0.9 (by 0.1 steps) and found the same pattern of crossover. That is, in all these cases CNCCP predicts undersearch for no-cost and low-cost conditions and oversearch for high-cost conditions. We cannot prove, however, that this would be true for all possible parameter combinations.

4. Discussion

A common claim is made that consumers engage in insufficient external search prior to purchasing, even for major purchases involving furniture, appliances, and automobiles (Claxton et al. 1974, Dickson and Sawyer 1990, Furse et al. 1984, Newman 1977). Using an experimental methodology for testing this claim, we find that, on average, our subjects did not search enough in the two no-cost conditions in agreement with most previous claims. 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 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.

The immediate contribution to the academic marketing literature is obvious. For the first time, an argument is advanced 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 search is expected. The policy implications are also clear. Previously, the common theme among consumer advocates was how to teach consumers to search smartly with the understanding that smart searching would require more searching than commonly takes place. Our results question this policy position. We suggest that smart searching sometimes means less searching than our intuition leads us to perform.

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 is not realistic, we have opted to study it for practical reasons. The optimal policies for alternative objectives under the same search environment, such as a “satisficing” goal in which the reward is awarded if the selected item is one of the best h items ($1 \leq h < n$) (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 considerably more complicated. Manipulating the DM’s objective is an important future research opportunity to pursue.

Note that we have assigned a goal to the subjects. That is, they were rewarded for only renting the best apartment in the set. It is clear that subjects understood this goal and worked to achieve it, as reflected in the relatively few cases where subjects rented or recalled an apartment that was known for sure not to be the overall best (see Table 2 shaded cells). Consequently, our results cannot be explained by assuming that subjects had adopted a different goal, such as satisfying rather than “nothing but the best.” Furthermore, if indeed subjects follow a satisfying strategy that by its very definition would imply undersearch vis-à-vis the optimal strategy for the “nothing but the best” goal, it is not clear why oversearch would be observed in the high-cost condition. We argue that the explanation is based on heuristic rules given the exogenously determined goal rather than assuming distortion of the goal itself.

Second, our statistical analysis is mostly conducted at the group level; only a few descriptive statistics are presented at the individual level (see Figure 4). Note, however, that our major findings of under- and oversearch as well as the sensitive to PSLC and AROCA hold for most subjects (see Figure 4).

Thirdly, one may argue that the model we propose for consumer sequential search 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 items 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 postdecisional regret). Our results speak to these cases.

Fourthly, together with the studies of Seale and Rapoport (1997, 2000), our study provides strong indication that consumers are sensitive to local patterns in the observed sequence (i.e., AROCA and PSLC), patterns that otherwise should have been ignored by a fully rational consumer. Further, we have demonstrated that under the simplest integration rule of AROCA, PSLC, and period number, and under the assumption that consumers aim to maximize the expected gain from search (given the "nothing but the best" goal), the pattern of under- and oversearch as a function of cost is expected.

Our study was not designed to test the CNCCP model; hence, we do not claim that this is the only model capable of explaining the observed findings. We do claim, however, that CNCCP provides the simplest account of our data, and for parsimonious reasons we tentatively present it for future research. We expect the model to account for the results of future studies designed to better understand the behavioral process of consumer sequential search.

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 alternatives, n , and the reward-to-cost ratio, the undersearch will change to oversearch. In our study, the reward-to-cost ratio was sufficiently high to reverse the common finding of undersearch. It may be instructive to experimentally investigate the implications of CNCCP by systematically manipulating the reward-to-cost ratio.

Acknowledgments

This work has been supported by the Research Grants Council of the Hong Kong Special Administration Region, China (Project

No. CA98/99.BM01) and by a grant from the AFOSR/MURI (Contract No. F49620-03-1-0377).

References

- Braunstein, Yale, Andrew Schotter. 1982. Labor market search: An experimental study. *Econom. Inquiry* 20 133–144.
- Camerer, Colin. 1995. Individual decision making. John H. Kagel, Alvin E. Roth, eds. *Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ, 587–703.
- Claxton, John D., Joseph N. Fry, Bernard Portis. 1974. A taxonomy of prepurchase information gathering patterns. *J. Consumer Res.* 1 35–42.
- Corbin, Ruth M., C. R. Olson, M. Abbondanza. 1975. Context effects in optimal stopping rules. *Organ. Behavior Human Performance* 14 207–216.
- Cox, James C., Ronald L. Oaxaca. 1989. Laboratory experiments with a finite horizon job-search model. *J. Risk 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. *J. Business Econom. Statist.* 17 466–472.
- Dickson, Peter R., Alan G. Sawyer. 1990. The price knowledge and search of supermarket shoppers. *J. Marketing* 54 42–53.
- Freeman, P. R. 1983. The secretary problem and its extensions: A review. *Internat. Statist. Rev.* 51 189–206.
- Furse, David H., Girish N. Punj, David W. Stewart. 1984. A typology of individual search strategies among purchasers of new automobiles. *J. Consumer Res.* 10 417–431.
- Goldman, Arieh, J. K. Johansson. 1978. Determinants of search for lower prices: An empirical assessment of the economics of information theory. *J. Consumer Res.* 5 176–186.
- Hey, John D. 1987. Still searching. *J. Econom. Behavior Organ.* 8 137–144.
- Kogut, Carl A. 1990. Consumer search behavior and sunk cost. *J. Econom. Behavior Organ.* 14 381–392.
- . 1992. Recall in consumer search. *J. Econom. Behavior Organ.* 17(1) 141–151.
- Lippman, S. A., J. J. McCall. 1976. The economics of job search: A survey (I and II). *Econom. Inquiry* 14 155–189, 347–368.
- Mehta, Nitin, Surendra Rajiv, Kannan Srinivasan. 2003. Price uncertainty and consumer search: A structural model of consideration set formation. *Marketing Sci.* 22(1) 58–84.
- Moorthy, Sridhar, Brian T. Ratchford, Debabrata Talukdar. 1997. Consumer information search revisited: Theory and empirical analysis. *J. Consumer Res.* 23 263–277.
- Newman, Joseph W. 1977. Consumer external search: Amount and determinations. Arch G. Woodside, Jagdish N. Sheth, Peter D. Bennett, eds. *Consumer and Industrial Buyer Behavior*. North-Holland, New York, 79–94.
- Ratchford, Brian T., Narasimhan Srinivasan. 1993. An empirical investigation of return to search. *Marketing Sci.* 12 73–87.
- Schotter, Andrew, Yale Braunstein. 1981. Economic search: An experimental study. *Econom. Inquiry* 19 1–25.

- Seale, Darryl A., Amnon Rapoport. 1997. Sequential decision making with relative ranks: An experimental investigation of the secretary problem. *Organ. Behavior 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. *J. Behavioral Decision Making* **13** 391–411.
- Sonnemans, Joep. 1998. Strategies of search. *J. Econom. Behavior Organ.* **35** 309–332.
- Urbany, Joel E., Peter R. Dickson, Rosemary Kalapurakal. 1996. Price search in the retail grocery market. *J. Marketing* **60** 91–104.
- Yang, Mark C. K. 1974. Recognizing the maximum of a random sequence based on relative rank with backward solicitation. *J. Appl. Probab.* **11** 504–512.
- Yeo, Geoffrey F. 1998. Interview costs in the secretary problem. *Australian New Zealand J. Statist.* **40** 215–219.
- Wu, Jianan, Arvind Rangaswamy. 2003. A fuzzy set model of search and consideration with an application to an online market. *Marketing Sci.* **22**(3) 411–434.
- Zwick, Rami, Ching Chyi Lee. 1999. Bargaining and search: An experimental study. *Group Decision Negotiation* **8** 463–487.
- , Eytan Weg, Amnon Rapoport. 2000. Invariance failure under subgame perfectness in sequential bargaining. *J. Econom. Psych.* **21** 517–544.

This paper was received October 5, 2001, and was with the authors 7 months for 3 revisions; processed by Barbara Kahn.