Journal of Economic Behavior & Organization 75 (2010) 338–347

Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/jebo

The devil is in the details: Incorrect intuitions in optimal search

Kimmo Eriksson[a,b,*], Pontus Strimling[b,c]

[a] Mälardalen University, School of Education, Culture and Communication, Box 883, SE-72123 Västerås, Sweden
[b] Stockholm University, Centre for the Study of Cultural Evolution, SE-10691 Stockholm, Sweden
[c] Indiana University, Workshop in Political Theory and Policy Analysis, 513 N. Park Avenue, Bloomington, IN 47408-3895, United States

ARTICLE INFO

Article history:
Received 14 October 2009
Received in revised form 8 April 2010
Accepted 16 April 2010
Available online 24 April 2010

JEL classification:
D83
D84

Keywords:
Optimal search
Heuristics
Secretary Problem
Intuitions
Learning

ABSTRACT

In the classic *Secretary Problem* it has been established that people tend to search somewhat less than is optimal, and a number of explanations have been suggested. Here we propose a new explanation, the Similar-But-Incorrect Intuitions Hypothesis, which says that suboptimal search behavior is to be expected because optimal strategies vary disproportionately with subtle details of the search problem setup, whereas people seem to entertain general intuitions about optimal search. We find support for this hypothesis in experiments on a new search problem, the *Explore-and-Collect Problem*, where the player collects utility from an option every time it is tried and options can be recalled. Although the optimal search effort in this problem is much smaller than for the Secretary Problem, people tend to search only marginally less. This is not predicted by previous explanations for suboptimal search.

1. Introduction

A sequential search problem can be abstractly defined as any problem where a decision maker (DM) has a number of options and can evaluate only one at a time. Depending on the specifics of the problem (such as the underlying distribution of payoffs, the cost of search, the kind of information gained by evaluation, and whether or not recall of options is possible) some search strategy is optimal. It is mathematically demanding to determine the optimal search strategy of a given problem, so it is not surprising that people who act as decision makers in such situations will typically behave suboptimally.

The most famous mathematical formulation of a sequential search problem is the classic *Secretary Problem* (SP), defined as follows. N secretarial applicants, all of different quality, are lined up in random order. DM only cares about finding the best one of all applicants (so we can say that DM's utility is 1 if the best applicant is hired and 0 otherwise). The quality of applicants can be explored one applicant at a time, at no cost, but recall of previous applicants is impossible and DM observes only his relative rank. An applicant who ranks better than all the previous ones is named a *candidate*. It is clearly pointless for DM to stop at an applicant who is not a candidate. The optimal strategy is to stop at the first candidate who appears after the first $N/e \approx 0.37N$ applicants have been evaluated (Ferguson, 1989).

A substantial body of experiments shows that people's behavior in SP is suboptimal in a systematic way. Participants tend to stop searching a little earlier than optimal, i.e., undersearch (Rapoport and Tversky, 1970; Seale and Rapoport, 1997,

* Corresponding author at: Mälardalen University, School of Education, Culture and Communication, Box 883, SE-72123 Västerås, Sweden. Tel.: +46 21101533; fax: +46 21101330.
E-mail addresses: kimmo.eriksson@mdh.se (K. Eriksson), pontusstrimling@gmail.com (P. Strimling).

0167-2681/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.jebo.2010.04.006

2000; Zwick et al., 2003; Bearden et al., 2005, 2006; Schunk and Winter, 2009). Several explanations of such deviations from the optimal strategy have been suggested; see Bearden and Rapoport (2005) for a comprehensive review. Many of these proposed explanations are of a strategic nature, i.e., they explain the systematic deviation from optimal behavior by positing that players are actually optimizing something else than expected payoff (cf. Schunk, 2009). Zwick et al. (2003) conducted an experiment on (a version of) SP where payoff functions explicitly included a search cost and found that with sufficiently high search costs people's behavior switched from undersearch to oversearch. This switch is not predicted by strategic explanations, but Zwick et al. found it to be consistent with people using an optimally chosen "*k* bounce rule" (Hey, 1982, 1987), i.e., a heuristic of the type "stop at the next candidate if there have been at least *k* applicants since the last candidate". This is an example of an order-based explanation.

Here we aim to show that neither strategic nor order-based explanations are sufficient to explain all systematic deviations from optimal strategies in sequential search problems. We will offer a complementary explanation based on erroneous intuitions about such problems. By an *intuition* we here mean a preconceived idea about how much to search.

To see how intuitions might cause systematically biased behavior we will define a new sequential search problem, the *E* xplore-and-Collect Problem (EC). In this problem DM has *N* rounds in which to explore a set of *n* options, each of a fixed utility drawn uniformly from [0,1]. In each round DM picks one of the options, either by exploring a new option or recalling any previously explored option. When an option is picked, DM collects its utility, but observes only its relative rank compared with previously explored options. The crucial decision in EC is when to stop exploring new options and stay with the best of the options explored thus far.

In Section 2 we will conduct a mathematical analysis of the optimal strategies of the above two problems. For EC we will prove that the optimal strategy is to stop exploring after about $\sqrt{2N}$ rounds (Theorem 1). The difference in the dependence of *N* between the two formulas for optimal amounts of search (linear in *N* for SP, square-root of *N* for EC) implies that for large values of *N* the optimal search effort will be very different in the two problems. For instance, for $N = 80$ the optimal strategy in SP is to stop at the next candidate after the first 30 rounds, whereas the optimal strategy in EC is to stop exploring already after 12 rounds.

Now consider how non-experts will perceive these two problems. It seems reasonable that the Secretary Problem and the Explore-and-Collect Problem will be seen as quite similar; they are both about a DM who has *N* opportunities to evaluate a large set of options one at a time, and who must make a decision about when to stop evaluating new options and instead stick with a known one. The problems differ in several details: whether recall is possible, whether there is payoff at every turn or only at the end, and whether payoffs are uniformly distributed or only the best option pays. However, without a complete mathematical analysis it is impossible to foresee the precise impact that these details have on the optimal strategy, and indeed difficult to foresee even the direction of the impact. Thus, when people are faced with different sequential search problems, we cannot expect them to know how to adapt their strategy from one problem to the next unless they have had substantial previous experience of both problems with appropriate feedback. As we will discuss in Section 6, it seems unlikely that people will get such experience and feedback outside the laboratory. In summary, our main hypothesis reads as follows.

1.1. *The Similar-But-Incorrect Intuitions Hypothesis*

When encountering a sequential search problem, people will have intuitions about the optimal amount of search to be conducted before stopping, and these intuitions (together with the specific information encountered during search) will shape people's actual behavior. People are typically not able to foresee how the optimal amount of search depends on details of the problem. Hence, for different search problems people will tend to have rather similar intuitions about how much to search, although the mathematically optimal strategies may in fact differ greatly. As a consequence, people's intuitions will typically be incorrect for most problems. However, for any specific problem people can learn to adjust their intuitions toward optimality with experience and clear feedback.

To state the key idea in other words: Changing a detail in the problem may change the optimal strategy a lot more than most people will think. This perspective offers another interpretation of the data obtained by Zwick et al. (2003). In their game, the detail of adding a search cost changed the optimal amount of search from 15 to 6 applicants (a 60 percent drop), whereas participants' average amount of search changed from 11.3 to 8.7 applicants (only a 23 percent drop). Zwick et al. showed how such a switch from undersearch to oversearch is predicted from a model where agents use certain optimal order-based heuristics. According to our hypothesis, another possibility is that people's average amount of search is to a large extent determined by intuitions that are not as sensitive to the addition of a search cost as the optimal strategy happens to be.

In this paper we will report three studies that test predictions of our hypothesis for the Secretary Problem and the Explore-and-Collect Problem. As mentioned above, these problems differ in what search effort is optimal, such that for any large value of *N* the optimal amount of search is much greater in SP than in EC. Thus for large *N* the Similar-But-Incorrect Intuitions Hypothesis predicts that people will either drastically oversearch in EC or drastically undersearch in SP. Since previous studies have found only a rather small amount of undersearch in the Secretary Problem, we therefore predict a rather large amount of oversearch in EC. As far as we understand, no other theory on search behavior would make this prediction.

Previous experimental studies of SP have typically found little effect of repeated play. As we will discuss in the next section, EC and SP also differ in the variance in payoff for a fixed strategy. This variance is very large for SP, and hence the feedback on what payoff to expect in the future from a tried strategy is extremely poor. For EC, on the other hand, the payoff variance is rather small, thus providing much clearer feedback. We therefore predict repeated play in EC to result in behavior (and, indeed, intuitions) quickly changing in the direction toward optimality.

In Study 1 we let participants play EC repeatedly with real payoffs and feedback on their earnings. In line with our previous argument, we predict that participants will tend to oversearch compared to the optimal search effort in EC (in contrast to the undersearch typically found for SP) and that they will gradually adjust their behavior toward the optimum, i.e., search less.

Our hypothesis is based on the assumption that people's actual behavior in terms of search effort is shaped not only by the order in which information is encountered during search but also by their intuitions (preconceived ideas) on how much to search. It is difficult to tease these factors apart in people's behavior. However, the types of strategies that are optimal do not depend at all on the order of options but can be fully expressed using only one parameter (for search effort). Therefore, viewed as optimization problems both EC and SP are unchanged if we ask participants only for this parameter instead of having them make one decision for each turn of the game. In this way we can eliminate the component of behavior that is order-based, so that only participants' intuitions play a role. We will use the terminology that asking participants directly for which search effort they want to make in EC or SP (instead of having them play out the entire game) is the *intuition form* of the problem. In the last two studies we use the intuition forms of the problems.

Study 2 is a vignette study (with no actual payoffs) where we ask participants for their intuitions about both EC and SP for varying values of $N = 11$, 80, and 500. Recall that the optimal search effort is proportional to N for SP, but proportional to $\sqrt{N}$ for EC. Thus the Similar-But-Incorrect Intuitions Hypothesis predicts that, for large values of N, undersearch in SP will occur simultaneously with drastic oversearch in EC; in addition, this discrepancy will be more pronounced the larger N gets.

Finally, in Study 3, we use a laboratory experiment with real payoffs to investigate how intuitions in EC and SP change over repeated play. For EC we predict the same pattern as in Study 1, i.e., oversearch that is gradually adjusted toward the optimal amount of search. For SP we predict a replication of the experimental findings in the literature, i.e., slight undersearch and no pronounced adjustment over repeated play.

2. The optimal strategies of the Secretary and Explore-and-Collect Problems

2.1. The Secretary Problem

The optimal strategy of the Secretary Problem is well-known (Ferguson, 1989). Let s_k^{SP} be the random binary variable indicating success in obtaining the best applicant out of N if one follows the strategy to stop at the first candidate after passing up the first $k-1$ applicants. The value of k is called the threshold. Then

$$E(s_k^{\mathrm{SP}}) = \frac{k-1}{N}\sum_{i=k}^{N}\frac{1}{i-1}, \quad k > 1. \tag{1}$$

This is maximized for threshold value

$$k' = \min\left\{k \geq 1 : \sum_{i=k+1}^{N}\frac{1}{i-1} \leq 1\right\},$$

which quickly tends to $1/e \approx 0.37$ as N grows large. The expected success at the optimal strategy, $E(s_{k'}^{\mathrm{SP}})$, also tends to $1/e$.

Here we will also be interested in how precisely s_k^{SP} measures the expected success, i.e., we want the relative standard deviation of s_k^{SP} for any threshold k:

$$\mathrm{RSD}(s_k^{\mathrm{SP}}) = \frac{\sqrt{\mathrm{Var}(s_k^{\mathrm{SP}})}}{E(s_k^{\mathrm{SP}})} = \frac{\sqrt{E(s_k^{\mathrm{SP}}) - E(s_k^{\mathrm{SP}})^2}}{E(s_k^{\mathrm{SP}})},$$

into which we can plug in (1). For the optimal strategy this value tends to $\sqrt{e-1} \approx 131$ percent. Thus, the precision is very bad.

2.2. The Explore-and-Collect Problem

The Explore-and-Collect Problem seems to capture in idealized form some aspects of sequential search in everyday life. For instance, consider trying a new brand of washing powder among the dozen of brands on the shelf in your shop. When trying this washing powder you do not only gain some information about its quality, but you actually get the benefit of using it. Or, say that you arrive in a new city in order to stay for a week. Each day you will go to a café for breakfast. There are at least 20 cafés in the city and you want to maximize your breakfast experience during your stay. How many breakfast places should you try before you stick to the best of those you have already been to?

Our analysis of EC starts with the simple observation that a strategy can never be optimal if it picks some option X twice before another option Y is picked for the first time. The reason is that if Y is to be picked at all, it must be better to do it before picking X a second time because if Y turns out to be better than X then you would do better to pick Y instead of X a second time. Therefore the optimal strategy must be to pick different options for some number (say k) of rounds, and then consistently pick the best option known to you for the remainder of the game. We can think of k as the stopping time of the exploration phase, or the *search effort*. The above observation reduces the problem to optimizing the search effort k.

Define three random variables as follows: U_i is the utility of the ith randomly picked option; $B_k := \max\{U_1, \ldots, U_k\}$ is the utility of the best of k randomly picked options; $s_k^{\mathrm{EC}} := (U_1 + \cdots + U_k) + (N-k)B_k$ is the total utility obtained from following the strategy of making k picks of different options followed by $N-k$ picks of the best option found during the previous phase. The expected utility of this strategy can then be expressed as

$$E(s_k^{\mathrm{EC}}) = E(U_1 + \cdots + U_k) + (N-k)E(B_k). \tag{2}$$

Given some probability distribution of utilities, Eq. (2) allows us to calculate the utility of the strategy of making search effort k. Observe that Eq. (2) does not contain n, the number of options. Consequently the optimal number of picks does not depend on the number of options, as long as the number of options is large enough to allow the theoretically optimal number of picks. We will now determine a formula for the optimal search effort when utilities of options are i.i.d. random variables uniformly distributed on $[0, 1]$. In this case the cumulative distribution function of B_k is $F_{B_k}(x) = x^k$ for $0 \le x \le 1$. Hence the density function is $f_{B_k}(x) = kx^{k-1}$, and it is then elementary to derive expected value and variance:

$$E(B_k) = \frac{k}{k+1}, \quad \mathrm{Var}(B_k) = \frac{k}{(k+2)(k+1)^2}.$$

Every U_i is distributed as $U_1 = B_1$, so we have

$$E(U_1 + \cdots + U_k) = \frac{k}{2}, \quad \mathrm{Var}(U_1 + \cdots + U_k) = \frac{k}{12}.$$

Eq. (2) for the expected utility of the strategy with stopping time k now takes the form:

$$u(k) := E[s_k^{\mathrm{EC}}] = \frac{k}{2} + \frac{(N-k)k}{k+1} = N + 1 - \frac{k}{2} - \frac{N+1}{k+1}. \tag{3}$$

Differentiation of the expected utility $u(k)$ with respect to k establishes that it is unimodal for $k > 0$ (see Fig. 1). Define the ceiling function $\lceil x \rceil$ as the smallest integer $\ge x$.

Theorem 1. *The optimal search effort in the Explore-and-Collect Problem is*

$$\hat{k} = \min\{\lceil \sqrt{2N + 9/4} - 3/2 \rceil, n\}.$$

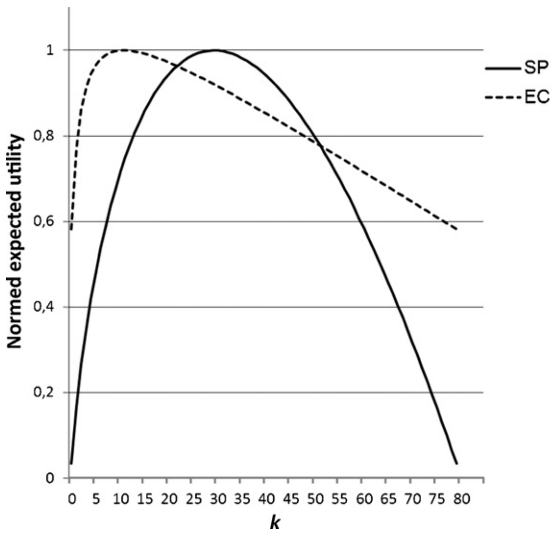

Fig. 1. For varying search effort k, the curves show normed expected utility $u(k)/u(\hat{k})$ in the Explore-and-Collect Problem with $N = 80$ picks and normed expected success $E(s_k^{\mathrm{SP}})/E(s_{k'}^{\mathrm{SP}})$ in the Secretary Problem with $N = 80$ applicants.

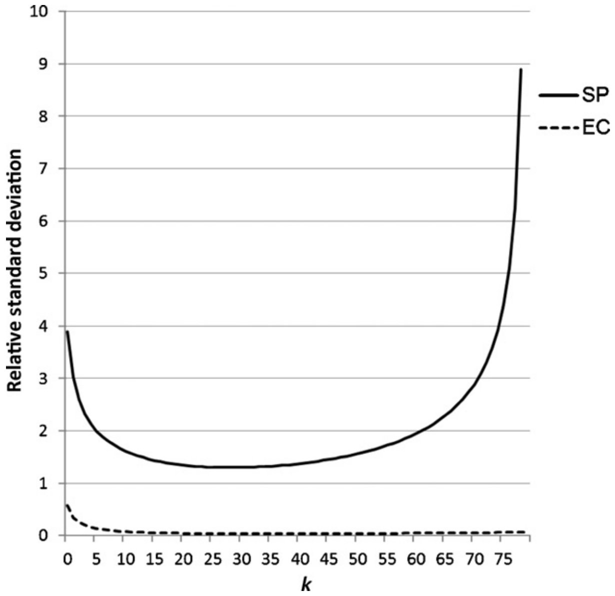

Fig. 2. For $N = 80$ and varying k, the solid curve shows the relative standard deviation of s_k^{SP}. The dashed curve shows the upper bound from Eq. (4) of the relative standard deviation of s_k^{EC}.

Proof. Given (3) the equation $u(k^*) = u(k^* + 1)$ has the solution

$$k^* = \sqrt{2N + 9/4} - 3/2.$$

Thanks to unimodality, the integer value that maximizes $u(k)$ must be $\lceil k^* \rceil$, because this is the unique integer in the interval $[k^*, k^* + 1]$ except in the degenerate case when k^* happens to be an integer, in which case both k^* and $k^* + 1$ are optimal integer values of k. However, if $k^* > n$ then the best one can achieve with n picks is to use them all in the search phase. □

In order to obtain the relative standard deviation of s_k^{EC}, first observe that the sum $U_1 + \cdots + U_k$ must be positively correlated to B_k. Therefore we have the inequality $\mathrm{Var}(s_k^{\mathrm{EC}}) \leq \mathrm{Var}(U_1 + \cdots + U_k) + \mathrm{Var}((N-k)B_k)$. Consequently,

$$\mathrm{RSD}(s_k^{\mathrm{EC}}) = \frac{\sqrt{\mathrm{Var}(s_k^{\mathrm{EC}})}}{E(s_k^{\mathrm{EC}})} \leq \frac{\sqrt{\frac{k}{12} + \frac{(N-k)^2 k}{(k+2)(k+1)^2}}}{N + 1 - \frac{k}{2} - \frac{N+1}{k+1}}. \tag{4}$$

2.3. Comparison between SP and EC

Fig. 1 shows how the expected utility (normed to have maximum 1) varies with search effort k for both problems with $N = 80$ (the value used in Study 3). The optimum for EC is clearly much lower than the optimum for SP. There is no obvious difference in how peaked the curves are around the optimum. i.e., the expected loss from suboptimal play is about the same in both problems.

Fig. 2 shows how the relative standard deviation varies with k for both problems. The relative standard deviation is consistently smaller in EC than in SP by several orders of magnitude, which means that in the former problem a single outcome contains much more information about the expected outcome.

Now consider an agent who plays these games repeatedly with feedback on earnings (s_k^{SP} and s_k^{EC}, respectively). Assuming the agent desires to play optimally, the possibilities of learning from feedback to adjust the search effort k toward optimality are much greater in EC than in SP, since the difference in outcomes between different strategies are much more precisely measured.

3. Study 1

Here we test the prediction that people will tend to oversearch in the Explore-and-Collect Problem but adjust their behavior toward more optimal strategies over repeated play.

Table 1
Maximum likelihood estimation of the fixed effects of model (5) using the R package.

Effect	Estimate	S.E.	DF	t-value	Pr > $\|t\|$
intercept	10.97	0.70	656	15.73	<0.001
round (t)	−0.08	0.03	656	−2.41	0.016
gender$_i$	3.70	1.10	45	3.37	0.002
gender$_i \times t$	−0.05	0.54	656	−1.07	0.286

3.1. Subjects

48 participants (19 women, 28 men, 1 missing data on gender; median age 26 years) were recruited from a pool of volunteering college students from miscellaneous study programs at a Swedish university. Participants received a cinema ticket as a show-up fee and were informed that they would also be paid according to how many points they earned in the game they were to play. Average payment was around 100 SEK.

3.2. Procedure

Each participant sat at a desk equipped with a computer, running 15 rounds of the Explore-and-Collect Problem, see appendix in Supplementary material for complete instructions. Every round had $N = 22$ picks and $n = 22$ options with uniformly distributed underlying values. After every round each participant was told the number of points he or she earned in that round (and no other feedback).

3.3. Analysis

We will define a participant's search effort in a given round as the number of different options he or she explored in that round. We will compare this to the optimal search effort, which according to Theorem 1 is $\lceil\sqrt{2\cdot 22+9/4}-3/2\rceil = 6$.

To study learning during repeated play we set up a random-effects regression model, an established approach to study sequential behavior in games (cf. Croson, 2007; Eriksson and Strimling, 2009). The dependent measure is the search effort for an individual i in round t. The key independent variable is the round t; we predict that search effort will decrease with t. We also control for gender and interaction between gender and t, because in a parallel study of EC in a social learning context we have found that women tend to search more than men and that this effect varies with time (Eriksson and Strimling, 2009). Our model is

$$\text{search effort}_{t,i} = a + c_1 t + c_2 \text{gender}_i + c_3 \text{gender}_i \times i + e_i + e_{t,i}, \tag{5}$$

for each participant i in the sample; gender$_i$ is coded 1 if i is a woman, 0 if man; e_i is the error term for between-individual variation, and $e_{t,i}$ is the residual error term.

3.4. Results

The average search effort per round was 11.6, which is greater than the optimal value of 6 ($t(47) = 9.55$, $p < 0.001$). Table 1 shows the result of the regression model. As predicted, the search effort decreases with repeated play. Further, there is an effect of gender, with women searching more than men (i.e., they are even farther from optimal levels of search). The interaction between gender and round of the game is not significant.

3.5. Discussion

As predicted we found that people on average searched much more than optimal in the Explore-and-Collect Problem and improved over repeated play. We also found a gender difference in search behavior that, although we have no explanation for it, is consistent with findings in a parallel study on EC.

4. Study 2

We conducted a vignette study in order to investigate people's intuitions about how much to search in the Secretary Problem and the Explore-and-Collect Problem. Our prediction is that intuitions are pretty similar for the two problems, and very inaccurate for EC when N is large.

4.1. Subjects and procedure

66 participants (31 women and 35 men; median age 23 years) were recruited from the same pool of volunteers as in Study 1, with no overlap between the samples. Participants sat in separate cubicles and answered a survey where they were

Table 2
Optimal (opt) and mean estimated (mean est) search efforts in the Secretary Problem and the Explore-and-Collect Problem. Relative deviation (rel dev) is the difference between mean est and opt divided by opt. Stars indicate statistical significance of relative deviations being different from zero (t-tests).

N	Secretary Problem			Explore-and-Collect Problem		
	opt	mean est (SD)	rel dev	opt	mean est (SD)	rel dev
11	4	4.8 (1.9)	+20% **	4	5.1 (2.6)	+28% **
80	30	26.1 (17.4)	−13%	12	20.2 (13.8)	+68% ***
500	185	122.6 (115.6)	−34% ***	31	104.7 (105.4)	+238% ***

*$p < 0.05$.
** $p < 0.01$.
*** $p < 0.001$.

presented with two scenarios describing the Secretary Problem and the Explore-and-Collect Problem, both framed in terms of hiring applicants (details in appendix in Supplementary material). For both problems they were asked to consider three cases ($N = 11$, 80 or 500). For each case participants estimated how many applicants they would interview before being prepared to hire (in the SP scenario), and how many applicants they would try before returning to the best one they had tried (in the EC scenario). Participants received a cinema ticket as a show-up fee. There were no other monetary incentives connected with the survey.

4.2. Analysis

The optimal search effort for each vignette (computed from the formulas presented earlier) is presented in Table 2. For $N = 11$ the optimal search effort is the same (4) for both SP and EC. For $N = 80$, on the other hand, the optimal search effort is 2.5 times larger in SP than in EC. For $N = 500$ the optimal search effort is 6 times larger in SP than in EC.

4.3. Results

Results on participants' estimated search efforts are shown in Table 2. The differences in the estimations between SP and EC are consistently small (all differences are less than 30 percent, and statistically significant only in the case $N = 80$).

For the case of the smallest value of N, the estimated search efforts in the two scenarios are close to each other and both are larger than the optimal search effort.

We now focus on the two cases where N is large (80 or 500). As predicted, participants' estimated search efforts deviate from optimal in opposite directions for SP (undersearch) and EC (oversearch), and the relative deviations are much larger for EC than for SP.

4.4. Discussion

As a check of the validity of our results, we can compare our findings with those of Seale and Rapoport (1997) for SP with $N = 80$ applicants, where they fitted threshold models to data and found the modal value of inferred thresholds to be 21. Here was asked explicitly for thresholds and received an average response of 26. These results from different methods seem roughly consistent.

Our predictions for large values of N were confirmed: People seem to intuitively make a much smaller distinction between how to behave in these two search problems than is optimal.

Since these vignettes ask only about the threshold, all potential effects of heuristics based on order have been eliminated. Nonetheless, we replicated previous experimental findings for both problems. For EC we found oversearch as we did in Study 1. Similarly, for SP we found the well-known undersearch for large values of N (but significant only for $N = 500$). However, for the small value of $N = 11$ we found that the intuition for the Secretary Problem was to oversearch. We do not know of any experiments carried out on SP for such a small N; it would be quite interesting to see whether this is a robust result.

5. Study 3

Here we run an experiment where participants play both SP and EC repeatedly. As in the vignettes in Study 2, we use the "intuition form" of the problems, i.e., participants must decide from the outset how much to search.

5.1. Subjects

70 participants (30 women, 40 men; median age 24 years) were recruited from the same pool of volunteers as in the previous studies, with no overlap between the samples. Participants received a cinema ticket as a show-up fee and were informed that they would also be paid according to how many points they earned in the game they were to play, according to a certain exchange rate (1000 points = 1SEK). Average payment was around 90 SEK.

5.2. Procedure

Each participant sat at a desk equipped with a computer. First they were presented with the rules for both the Secretary Problem and the Explore-and-Collect Problem. In both problems the screen showed 80 boxes.

The SP instructions read: “Try to find the best box! The boxes contain the integer values from 1 to 80 in a random order. You must decide how many boxes to begin to explore. The computer will then continue exploring until it finds a box that has a higher value than all previous ones (if there is such a box). If the box selected in this way has the highest value of all 80 boxes, you win the game and gain 10000 points.” Once the player decided how many boxes to explore, the computer revealed the values of all the boxes in a sequence of three steps: first, the boxes the player decided to explore; second, the boxes explored by the computer up to the box that got selected (clearly marked); last, all remaining boxes (clearly marking the best of all boxes, if different from the selected box). Finally, the computer stated either “The selected box was the best one — you win” or “The selected box was not the best one — you did not win.”

The EC instructions read: “Try to maximize your earnings in 80 picks from the boxes! The boxes contain the integer values from 1 to 80 in a random order. You must decide how many different boxes to pick. For each pick you will receive the value of that box. The computer will then continue to pick for you the best of these boxes for the remaining picks. After the player decided how many different boxes to pick, the computer revealed the values of all the boxes in a sequence of two steps: first, the boxes the player decided to pick (with the best one clearly marked); last, all remaining boxes. Finally, the computer stated the value of all picks, e.g. “In your 5 different picks you received $41 + 65 + 7 + 59 + 34$ points and in the following 75 picks you received 65 points each time, for a total of 5081 points.”

Every participant had one test run of each game and then played 10 rounds of one game followed by 10 rounds of the other game. The order of the games were randomized, with 34 resp. 36 participants playing SP resp. EC first.

5.3. Analysis

Recall from Study 2 that for $N = 80$ the optimal search effort is 30 in SP and 12 in EC. For each of the games we carry out the same statistical analysis as in Study 1.

5.4. Results

First we analyzed whether the order of the games had any effect. There was a tendency that the search effort was lower in the first round of SP if SP was played as the second game (mean 31.4 vs 36.0 if SP was played first) but not statistically significant ($t(68) = 0.854$, $p = 0.396$). For the remaining analysis we therefore pool the data for both orders of the games.

Fig. 3 shows for each game how the mean search effort changed over the 10 rounds. For SP the average search effort per round was 30.9, which is not significantly different from the optimal value of 30 ($t(69) = 0.75$, $p = 0.456$). For EC, in contrast, the average search effort per round of 24.6 was much larger than the optimal value of 12 ($t(69) = 7.03$, $p < 0.001$).

The results from the regressions are presented in Tables 3 and 4. For EC we replicate the findings from Study 1 of a decreasing search effort with repeated play, and a greater search effort by women. Here we also find a significant interaction, with women decreasing their search effort at a higher rate (from their higher initial levels). For SP all the tendencies are the same but not statistically significant.

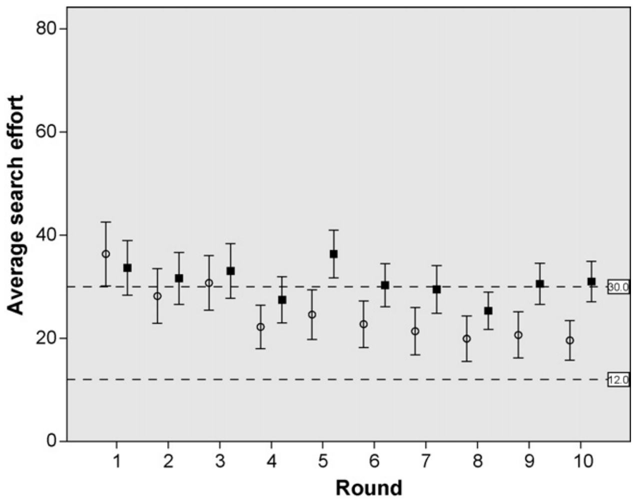

Fig. 3. Mean search effort per round for SP (squares) and EC (circles). Error bars signify ±1 standard error. The dashed lines mark the optimal search effort in SP (upper line) resp. EC (lower line).

Table 3
The fixed effects of model (5) estimated for the search effort in the Explore-and-Collect Problem (Study 3).

Effect	Estimate	S.E.	DF	t-value	Pr > |t|
intercept	28.56	2.69	628	10.60	<0.001
round (t)	−1.16	0.25	628	−4.67	<0.001
gender$_i$	11.20	4.12	68	2.72	0.008
gender$_i$ × t	−1.00	0.38	628	−2.62	0.009

Table 4
The fixed effects of model (5) estimated for the search effort in the Secretary Problem (Study 3).

Effect	Estimate	S.E.	DF	t-value	Pr > |t|
intercept	32.03	2.27	628	14.12	<0.001
round (t)	−0.28	0.30	628	−0.94	0.348
gender$_i$	2.76	3.47	68	0.80	0.429
gender$_i$ × t	−0.33	0.46	628	−0.71	0.478

5.5. Discussion

For the Explore-and-Collect Problem we replicated the findings from our previous studies: people on average search much more than optimal improved over repeated play. The gender difference from Study 1 showed up again in this study.

For the Secretary Problem we did *not* replicate the undersearch typically found. However, recall that in our vignette study the undersearch was not significant for $N = 80$ (the case studied here) but only in the case of $N = 500$.

6. Conclusions

We believe our three studies together support the Similar-But-Incorrect Intuitions Hypothesis about people's behavior in optimal search. In other words, people tend to have an initial idea about how much to search, and this intuition is not very sensitive to the details of the search situation, such as precisely what parameter is optimized or whether payoffs are obtained in every step or only in the end. As mathematical analysis shows, the optimal behavior may be very sensitive to these details. If two problems have very different optimal strategies but an individual has similar intuitions for both problems, then the intuition must be incorrect for at least one of the problems. Given clear feedback on behavior in a specific search problem, people can gradually learn to correct incorrect intuitions for this specific problem.

As we discussed in Section 1, both incorrect intuitions and order-based heuristics may be major contributing factors to systematic suboptimal behavior in sequential search problems in general. Our studies focused on the Secretary Problem and the Explore-and-Collect Problem. In Study 1 we found systematic oversearch in EC as predicted by the Similar-But-Incorrect Intuitions Hypothesis. This finding was replicated in the later studies where the intuition form of the problem was used, consistent with the idea that this result is driven by incorrect intuitions. For SP, on the other hand, the established finding is that people have a tendency to slight undersearch. In Study 2 we found a similar tendency also in people's intuitions for SP, but this result was not replicated in Study 3. Thus, for the Secretary Problem the effect of incorrect intuitions remains unclear.

We have not seen our hypothesis explicitly treated before, but we wish to point out that some thoughts along the same line can be found in a short note of Bearden (2006). He analyzed a version of SP where the quality of applicants are drawn uniformly from [0,1]. DM can observe only whether or not the current applicant is a candidate (i.e., best among those seen so far), and tries to optimize the expected quality of the hired candidate.[1] Bearden proved that the optimal strategy in this problem is to stop at the first candidate after the first $\sqrt{N}$ applicants, i.e., about the same as the optimal search effort for the Explore-and-Collect Problem. Bearden suggested that the payoff scheme in his problem is more natural than in the classic SP and speculated that early stopping in the latter problem may be just "an artifact of the problem's unusual payoff scheme" (Bearden, 2006, p. 59). However, intuitions for SP are pretty close to optimal, whereas the Similar-But-Incorrect Intuitions Hypothesis would predict drastic oversearch also in Bearden's problem as the optimal search effort in this problem is even smaller than in EC.

So, how can it be that people's intuitions for sequential search problems such as EC can be so bad? While there has been much study of the use of simple heuristics for difficult problems (cf. Kahneman et al., 1982; Gigerenzer et al., 1999; Gigerenzer and Selten, 2001), little attention has been paid to how the use of heuristics change over an individual's lifetime. For the Explore-and-Collect Problem in an experimental setting we found that people's intuitions and behavior change toward the optimal strategy over repeated play with detailed feedback. As we discussed earlier, the feedback is inherently of higher quality for the Explore-and-Collect Problem than for the Secretary Problem, and hence learning is more likely to occur in the

[1] Readers versed in the theory of optimal stopping should take note that Bearden's version is not equivalent to the Rank-Optimizing Secretary Problem solved by Chow et al. (1964), in which DM had access to the relative rank of every explored applicant. As the mathematical analysis turns out to be much easier to do for Bearden's version of the problem, it might be worthwhile to consider a two-player version of Bearden's problem (cf. Eriksson et al., 2007).

former game. However, for sequential search problems in natural settings it seems to us very unlikely that a similar learning process would occur. There are several reasons for this. To begin with, there is the phenomenon that situations differ in details that have great impact on the optimal behavior. It seems unlikely that the exact same situation arises repeatedly and equally unlikely that the effect of changing just one detail is ever observable. Even more problematic is that feedback on one's choice of strategy will typically be extremely poor, because one does not get to see what could have happened if one had used a different strategy. For instance, in order to get proper feedback in a secretarial hiring scenario one would have to fill a large number of secretarial positions, varying the number of interviews one conducts before potentially stopping (with each of the numbers used several times to obtain any kind of reliable data). One would then have to find a way of comparing the secretaries hired through each strategy and also somehow compare them to the secretaries one did not even interview. Given the infeasibility of this kind of learning in natural settings it is not surprising that people have similar but incorrect intuitions about games that actually have very different optimal strategies, nor is it surprising that people start to develop better strategies when detailed feedback becomes available in an experimental situation.

Sequential search problems occur in many different areas studied by economists, including search for labor, goods, or marriage partners. The above argument for the infeasibility of learning optimal strategies through feedback seems to be generally applicable. For instance, it would not really matter whether only ranks or whatever counts as payoffs are observable, nor whether the distribution of payoffs is known or unknown. Regardless of the area in which sequential search problems arise, we predict that agents will typically have poor ideas about the optimal amount of search.

Finally, we must mention the gender difference we found in the Explore-and-Collect Problem both here and in another study (Eriksson and Strimling, 2009): women tended to oversearch more than men. Intriguingly, this phenomenon does not seem to be predicted by any established theory on cognitive sex differences (cf. Halpern, 2000). We leave it to future research to determine the scope of this phenomenon and to offer an explanation.

Acknowledgments

This research was supported by the CULTAPTATION project (European Commission contract FP6-2004-NEST-PATH-043434) and the Swedish Research Council.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.jebo.2010.04.006.

References

Bearden, J.N., 2006. A new secretary problem with rank-based selection and cardinal payoffs. Journal of Mathematical Psychology 50, 58–59.
Bearden, J.N., Rapoport, A., 2005. Operations research in experimental psychology. In: Smith, J.C. (Ed.), Tutorials in Operations Research. Informs, Baltimore, MD, pp. 213–236.
Bearden, J.N., Murphy, R.O., Rapoport, A., 2005. A multi-attribute extension of the secretary problem: theory and experiments. Journal of Mathematical Psychology 49, 410–422.
Bearden, J.N., Rapoport, A., Murphy, R.O., 2006. Experimental studies of sequential selection and assignment with relative ranks. Journal of Behavioral Decision Making 19, 229–250.
Chow, Y.S., Moriguti, S., Robbins, H., Samuels, S.M., 1964. Optimal selection based on relative rank (the "Secretary problem"). Israel Journal of Mathematics 2, 81–90.
Croson, R., 2007. Theories of commitment, altruism and reciprocity: evidence from linear public goods games. Economic Inquiry 45, 199–216.
Eriksson, K., Sjöstrand, J., Strimling, P., 2007. Optimal expected rank in a two-sided secretary problem. Operations Research 55, 921–931.
Eriksson, K., Strimling, P., 2009. Biases for acquiring information individually rather than socially. Journal of Evolutionary Psychology 7, 309–329.
Ferguson, T.S., 1989. Who solved the secretary problem? Statistical Science 4, 282–289.
Gigerenzer, G., Todd, P., the ABC research group, 1999. Simple Heuristics that Make Us Smart. Oxford University Press, Oxford.
Gigerenzer, G., Selten, R., 2001. Bounded Rationality: The Adaptive Toolbox. MIT Press, Cambridge, MA.
Halpern, D.F., 2000. Sex Differences in Cognitive Abilities, 3rd ed. Lawrence Erlbaum Associates, Mahwah, NJ.
Hey, J., 1982. Search for rules for search. Journal of Economic Behavior and Organization 3, 65–81.
Hey, J., 1987. Still searching. Journal of Economic Behavior and Organization 8, 137–144.
Kahneman, D., Slovic, P., Tversky, A., 1982. Judgment under Uncertainty: Heuristics and Biases. Cambridge University Press, Cambridge.
Rapoport, A., Tversky, A., 1970. Choice behavior in an optimal stopping task. Organizational Behavior and Human Performance 5, 105–120.
Schunk, D., 2009. Behavioral heterogeneity in dynamic search situations: theory and experimental evidence. Journal of Economic Dynamics and Control 33, 1719–1738.
Schunk, D., Winter, J., 2009. The relationship between risk attitudes and heuristics in search tasks: a laboratory experiment. Journal of Economic Behavior and Organization 71, 347–360.
Seale, D.A., Rapoport, A., 1997. Sequential decision making with relative ranks: an experimental investigation of the "Secretary Problem". Organizational Behavior and Human Performance 69, 221–236.
Seale, D.A., Rapoport, A., 2000. Optimal stopping behavior with relative ranks: the secretary problem with unknown population size. Journal of Behavioral Decision Making 13, 391–411.
Zwick, R., Rapoport, A., Lo, A.K.C., Muthukrishnan, A.V., 2003. Consumer sequential search: not enough or too much? Marketing Science 22, 503–519.